A "CUBIST" DECOMPOSITION OF THE HANDEL-MOSHER AXIS BUNDLE AND THE CONJUGACY PROBLEM FOR $\mathrm{Out}(F_r)$

CATHERINE EVA PFAFF AND CHI CHEUK TSANG

Abstract. We show that the axis bundle of a nongeometric fully irreducible outer automorphism admits a canonical "cubist" decomposition into branched cubes that fit together with special combinatorics. From this structure, we locate a canonical finite collection of periodic fold lines in each axis bundle. This gives a solution to the conjugacy problem in $\mathrm{Out}(F_r)$ for fully irreducible outer automorphisms. This can be considered as an analogue of results of Hamenstädt and Agol from the surface setting, which state that the set of trivalent train tracks carrying the unstable lamination of a pseudo-Anosov map can be given the structure of a CAT(0) cube complex, and that there is a canonical periodic fold line in this cube complex.

1. Introduction

Let F_r be the free group of rank $r \geq 2$ and $\mathrm{Out}(F_r)$ its outer automorphism group. The **outer space** associated to F_r, defined in [CV86] and denoted here by CV_r, is the projectivized space of marked metric graphs with fundamental group isomorphic to F_r (§2.1). Equivalently, CV_r is the projectivized space of free, properly discontinuous F_r-actions on $\mathbb{R}$-trees. One can compactify CV_r by taking $\overline{CV_r}$ to be the projectivized space of very small F_r-actions on $\mathbb{R}$-trees [BF94, CL95].

An outer automorphism of F_r is **fully irreducible** if none of its powers preserve the conjugacy class of a nontrivial proper free factor. Each fully irreducible φ acts on $\overline{CV_r}$ with north-south dynamics [LL03], i.e. there is an attracting tree $T_+ \in \partial CV_r$ and a repelling tree $T_- \in \partial CV_r$ such that $\varphi^k(\overline{CV_r} \backslash T_-)$ converges to T_+, uniformly on compact sets, as $k \to \infty$. We call φ **geometric** if it is induced by a homeomorphism of a surface. In this paper, we will be solely concerned with nongeometric fully irreducible outer automorphisms.

A **fold line** for φ is a bi-infinite path in CV_r which underlies a family of folding maps $h_{rs} : T_r \to T_s$ and which limits to T_+ in forward time and T_- in backward time. Morally, fold lines are a version of axes for the action of φ, but defined purely in terms of the combinatorics of graphs and without reference to a metric on CV_r. The **axis bundle** of φ, introduced in [HM11] and denoted here by $\mathcal{A}_\varphi$, is the union of fold lines for φ.

The coarse topology of $\mathcal{A}_\varphi$ was classified in [HM11], where it is shown that the inclusion of any fold line is a proper homotopy equivalence. The finer combinatorics of $\mathcal{A}_\varphi$ has remained mysterious. For example, the 'tripod fold' and 'singularity merging' in [Pfa24] and [AKKP19] seem to suggest certain non-homogeneity along and among axis bundles respectively, although these phenomena are not yet well understood.

In this paper, we aim to initiate a new framework for studying the axis bundle by showing that it admits a specific type of combinatorial structure (explained in detail in §3-4):

Theorem 1.1. *Let φ be a nongeometric fully irreducible outer automorphism. Then $\mathcal{A}_\varphi$ admits a canonical structure of a cubist complex. From the cubist complex structure, there is a canonically defined directed graph $\mathfrak{c}_{\mathcal{A}_\varphi}$ embedded in $\mathcal{A}_\varphi$, we call the* ***cardiovascular system****, satisfying the following properties:*

- *(i) There is a finite set of bi-infinite directed lines on which φ acts periodically. We call each line in this collection an* ***artery****.*
- *(ii) Each vertex of $\mathfrak{c}_{\mathcal{A}_\varphi}$ has a unique outgoing edge, thus has a well-defined forward trajectory. Each forward trajectory eventually enters an artery.*
- *(iii) Any two arteries are related by sweeping across finitely many 2-dimensional branched cubes.*

Both authors are grateful to Lee Mosher for inspiring conversations and interest in their work. The first author is supported by an NSERC discovery grant and is appreciative of the hospitality of the Institute for Advanced Study (and Bob Moses for funding her 2024-2025 membership) and CIRGET (under the generous direction of Steve Boyer and Vestislav Apostolov). This project was completed while the second author was a CRM-ISM postdoctoral fellow based at CIRGET. He would like to thank the center for its support.

Here, the notion of a cubist complex is based off that of a cube complex, with the key differences being:

(1) A cube complex is a space decomposed into cubes, whereas a cubist complex is a space decomposed into **branched cubes**, which are unions of cubes glued along certain affine slices. See the figure below for a 2-dimensional branched cube formed by gluing three 2-dimensional cubes.

(2) A cubist complex is inherently directed: Each branched cube is equipped with a splitting vertex and an 'antipodal' folding vertex. Intuitively, the **folding direction** refers to paths from the splitting vertex to the folding vertex, whereas the **splitting direction** is the opposite direction. The branched cubes in a cubist complex must intersect in a manner preserving this directionality.

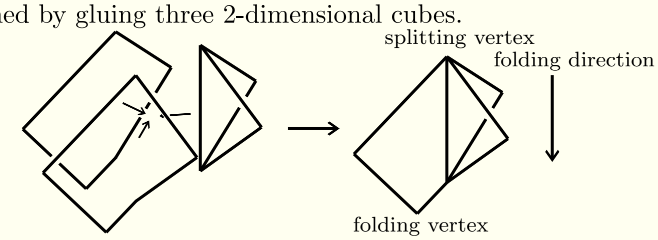

(3) When two cubes intersect in a cube complex, their intersection is a (complete) face of each of the two cubes, whereas when two branched cubes intersect in a cubist complex, their intersection is a (complete) face in the splitting side of one branched cube and a *subset* of a face in the folding side of the other branched cube. Morally, this causes the branched cubes to get 'finer', resembling a 'Zenotic' division, in the folding direction. See Figure 1, the black squares, for an example.

Property (2) follows directly from how we defined the branched cubes. For property (1), the branching of the branched cubes arises when the turns share directions. For example, if three directions lie in a single gate, then there are three edges in the branched cube that come from the three ways of choosing a pair of directions out of the three to fold, but folding two of the pairs will force the remaining pair to be folded as well, resulting in a branched cube depicted in the image accompanying (2) above.

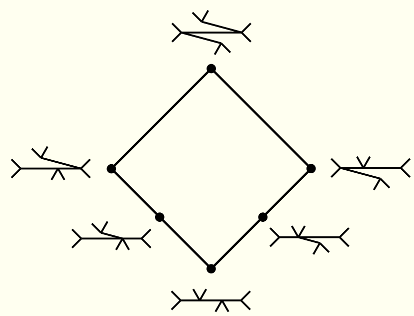

For property (3), the Zenotic division in the folding direction arises if folding some of the pairs of directions causes vertices to appear in the path of the remaining folds (some folds are longer than others), so that in the process of carrying out these remaining folds, one passes through multiple fully preprincipal train tracks. To the right is an example.

We give an outline of how the cubist complex structure arises for an axis bundle $\mathcal{A}_\varphi$: For each graph $T \in \mathcal{A}_\varphi$, there is a map $T \to T_+$ that respects the F_r-actions and restricts to an isometry on each leaf of the **attracting lamination** Λ [BFH97] determined by the repelling tree. In general, there is an interval's worth of such maps, but by taking the rightmost one, we associate a canonical **Λ-isometry** to each T. We say a pair of directions at a vertex of T is an **illegal turn** if they are identified in T_+ under the Λ-isometry. A **gate** is an equivalence class of directions, under the equivalence relation generated by illegal turns.

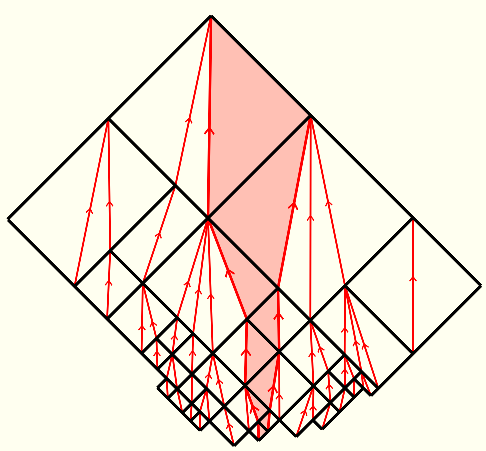

Figure 1. A cubist complex (in black) and its cardiovascular system (in red).

The branched cubes in the cubist complex structure are defined as follows. The nodes of the cubist complex structure, i.e. the 0-dimensional cubes, in $\mathcal{A}_\varphi$ are the **fully preprincipal weak train tracks**. These are the graphs in CV_r that have at least three gates at each vertex, as introduced by the first author in [Pfa24]. The edges, i.e. the 1-dimensional cubes, are obtained by folding illegal turns. The higher dimensional branched cubes are obtained by folding illegal turns by various amounts.

From this cubist structure, the cardiovascular system $\mathfrak{c}_{\mathcal{A}_\varphi}$ is defined by connecting, via directed segments, the nodes of $\mathcal{A}_\varphi$ in the splitting direction. See the Figure 1 red graph, for an example, and see §4.4 for the precise definition. The refinement of X in the folding direction implies that each node in $\mathfrak{c}_{\mathcal{A}_\varphi}$ has precisely one outgoing edge but possibly many incoming edges. Thus each node has a unique forward trajectory, and intuitively these trajectories (which go in the splitting direction) tend to converge together. The portion of a forward trajectory from the point it converges into a translate on then lies on an artery. This shows Theorem 1.1(i)-(ii). Item (iii) follows from a more intricate analysis of the combinatorics of cubist complexes.

The nodes and edges of each artery determine a **periodic folding sequence** for φ, i.e. $P, N \in \mathbb{Z}$ and a sequence of folding maps $h_n : \tau_{n+1} \to \tau_n$ for each $n \in \mathbb{Z}$ satisfying that $\tau_{n+P} = \varphi^N(\tau_n)$, defined up to a translation of the indexing. By varying over the arteries, we obtain as such a canonical finite collection $\mathfrak{s}_\varphi$ of periodic folding sequences for φ. Moreover, Theorem 1.1(iii) states that the different sequences in this collection differ in a highly controlled way.

We can use this collection to provide one solution to the conjugacy problem:

Theorem 1.2. *Two nongeometric fully irreducible $\varphi, \varphi' \in \mathrm{Out}(F_r)$ are conjugate if and only if the collections $\mathfrak{s}_\varphi$ and $\mathfrak{s}_{\varphi'}$ share a common periodic folding sequence (up to a translation on the indexing).*

Finally, we highlight that the cubist complex structure and cardiovascular system on $\mathcal{A}_\varphi$ are very computable objects, as demonstrated in Section 5 via some examples, that also showcase a variety of phenomena.

1.1. **Geometry of the axis bundle.** Handel and Mosher, via a list of questions in [HM11], and Bridson and Vogtmann [BV06, Question 3], more directly, ask: Describe the geometry of $\mathcal{A}_\varphi$ for a fully irreducible $\varphi \in \mathrm{Out}(F_r)$ acting on CV_r. The cubist complex structure gives some new perspective on this question.

Each node of the cardiovascular system $\mathfrak{c}_{\mathcal{A}_\varphi}$ determines a directed ray. These directed rays converge together into a finite collection of periodic bi-infinite lines.

We draw a parallel between this situation and the dynamics of geodesics in hyperbolic space $\mathbb{H}^n$: For a fixed boundary point $\xi \in \partial\mathbb{H}^n$, consider the collection of geodesics that limit onto ξ. The geodesics in this collection converge together in forward time, and if φ is a loxodromic isometry of $\mathbb{H}^n$ fixing ξ, then one can extract from this collection a geodesic that is invariant under φ.

This is, however, not a perfect analogy since the axis bundle can have multiple arteries. In the case of multiple arteries, Theorem 1.1(iii) states that any two arteries are related by sweeping across finitely many 2-dimensional branched cubes. One can morally regard the union of the 2-dimensional branched cubes as a flat strip, see the shaded region of Figure 2 left, and we can revise the analogy by 'blowing up' the invariant geodesic in $\mathbb{H}^n$ by inserting a Euclidean strip, see Figure 2 right.

1.2. **Motivation from surfaces.** Recall that a **train track** on a surface S is an embedded graph τ where the half-edges at each vertex are tangent to a common tangent line. In [Ham09], Hamenstädt showed that the set of trivalent train tracks on S can be realized as the vertex set of a cube complex $\mathcal{TT}_S$ by taking the edges to be splitting moves, and the higher-dimensional cubes to be spanned by commuting splitting moves. We regard $\mathcal{TT}_S$ as an analogue of outer space in the surface setting.

Let f be a pseudo-Anosov map on S, and let ℓ_+ be the attracting lamination of f. The results of [Ham09] imply that the set of trivalent train tracks that carry ℓ_+ determines a CAT(0) subcomplex of $\mathcal{TT}_S$. We denote this subcomplex as $\mathcal{A}_f$ and regard it as an analogue of the axis bundle. This analogy is justified by the fact that the axis bundle $\mathcal{A}_\varphi$ admits a similar description as the set of graphs Γ that carry the expanding lamination Λ, in the sense that there is a Λ-isometry from $\widetilde{\Gamma}$ to the attracting tree T_+.

In this perspective, the first statement in Theorem 1.1 that $\mathcal{A}_\varphi$ is a cubist complex is an analogue of the fact that $\mathcal{A}_f$ is a cube complex. However, the vertices of $\mathcal{A}_f$ are trivalent train tracks, which are the 'most generic' type of train track, in the sense that they form an open subset of the set of all train tracks carrying ℓ_+. On the other hand, the vertices of $\mathcal{A}_\varphi$ are fully preprincipal train tracks, which are the 'least generic'

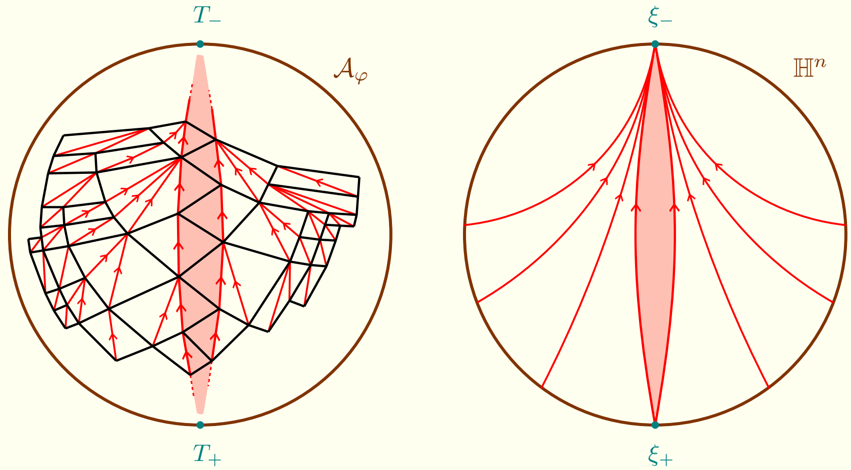

FIGURE 2. The dynamics of the cardiovascular system resembles that of geodesics in hyperbolic space that limit onto a fixed boundary point ξ_-, possibly with the invariant geodesic being blown-up into a flat Euclidean strip (region in red).

type of train track since they are the cubes with the highest possible codimension. Hence it might be more accurate to say that our cubist complex structure on $\mathcal{A}_\varphi$ is an analogue of a sort of 'dual complex' of $\mathcal{A}_f$.

We also mention an unfinished monograph of Mosher [Mos03] that explores similar ideas of constructing complexes out of train tracks and splitting moves (albeit of a more general type than those in [Ham09]).

In [Ago11], Agol showed that there is a canonical axis of the action of f on $\mathcal{A}_f$. The corresponding periodic splitting sequence gives a complete conjugacy invariant of the pseudo-Anosov map f. Theorem 1.1(i) and Theorem 1.2 can be considered as analogues of these facts.

Note that our methods differ from those of Agol. In [Ago11], Agol considers the **maximal splitting** operation, which means splitting the branches of a train track that have maximal weight (measured by the transverse measure on ℓ_+), while the operation underlying our cardiovascular system is, in some sense, performing all possible splits.

Question 1.3. *Is there a meaningful version of Agol's maximal splitting operation in the free group setting?*

There are some inherent difficulties in a naïve generalization. For example, while a natural choice of weights can be assigned to branches via studying the incidence matrices, there is no guarantee that a branch with maximal weight can be split or that branches of maximal weight are disjoint. The latter fact is in part due to our train tracks not being trivalent - some arguments in [Ago11] break down because of this.

Nevertheless, a positive answer to Question 1.3 would suggest a positive answer to the following question.

Question 1.4. *Is it possible to upgrade Theorem 1.1(i) from a finite canonical collection of fold lines to a single canonical fold line?*

One motivation for Question 1.4 is that pseudo-Anosov maps have a unique axis in Teichmüller space. By contrast, fully irreducible outer automorphisms have, in general, uncountably many fold lines. Theorem 1.1(i) extracts a finite canonical collection of fold lines out of these, but it would be more satisfying if one could take things a step further and extract a single canonical fold line.

1.3. **Generalization to cut decomposition axis bundles.** A **(weak) periodic Nielsen path (PNP)** in a $T \in \mathcal{A}_\varphi$ is a homotopically nontrivial path in T whose endpoints are principal points with the same image in T_+. Each PNP can be written as a concatenation of indivisible PNPs, which are paths of the form $\alpha_1^{-1} * \alpha_2$, with α_1 and α_2 mapped to the same interval in T_+. We refer to [BH92, HM11] for details.

Intuitively, a PNP represents some redundancy of T: Given an indivisible periodic Nielsen path $\alpha_1^{-1} * \alpha_2$, one can fold T by identifying α_1 and α_2 and get a 'reduced' element $\overline{T}$ of $\mathcal{A}_\varphi$ with a naturally induced Λ-isometry $\overline{T} \to T_+$.

For this reason, it is sometimes convenient to consider instead the **stable axis bundle** $\mathcal{SA}_\varphi$, which is the subset of $\mathcal{A}_\varphi$ that consists of elements without PNPs. See, for example, [HM11, MP16]. The machinery developed in this paper fully carries over to this stable category.

To be even more general, we define the **cut decomposition axis bundles**: The possible PNPs in trees in $\mathcal{A}_\varphi$ correspond to the cut vertices of $IW(\varphi)$, see [HM11, §4]. Fixing a subset of possible PNPs is equivalent to choosing a cut decomposition $\mathcal{G}$ of $IW(\varphi)$. For each cut decomposition $\mathcal{G}$, we define the **$\mathcal{G}$-axis bundle** $\mathcal{GA}_\varphi$ to be the subset of $\mathcal{A}_\varphi$ that consists of elements whose PNPs determine the cut decomposition $\mathcal{G}$.

Theorem 1.5. *Let φ be a nongeometric fully irreducible outer automorphism. Let $\mathcal{G}$ be a cut decomposition of $IW(\varphi)$. The $\mathcal{G}$-axis bundle $\mathcal{GA}_\varphi$ of φ admits a canonical structure of a cubist complex which makes it a subcubist complex of $\mathcal{A}_\varphi$. From the cubist complex structure, there is a canonically defined directed graph $\mathfrak{c}_{\mathcal{GA}_\varphi}$ embedded in $\mathcal{GA}_\varphi$, which we call the **cardiovascular system**, satisfying the following properties:*

*(i) There is a finite set of bi-infinite directed lines on which φ acts periodically. We call each line in this collection an **artery**.*

(ii) Each vertex of $\mathfrak{c}_{\mathcal{GA}_\varphi}$ has a unique outgoing edge, thus has a well-defined forward trajectory. Each forward trajectory eventually enters an artery.

(iii) Any two arteries are related by sweeping across finitely many 2-dimensional branched cubes.

We expect that the cubist complex machinery can be used to study the following problem.

Question 1.6. *How do the cut decomposition axis bundles $\mathcal{GA}_\varphi$ sit inside the full axis bundle $\mathcal{A}_\varphi$?*

As a final remark, we note that even though $\mathcal{GA}_\varphi$ is a subcubist complex of $\mathcal{A}_\varphi$, the arteries in $\mathcal{GA}_\varphi$ and $\mathcal{A}_\varphi$ can be completely different. We demonstrate an example of this in Example 5.4.

2. Preliminary definitions and notation

Throughout this paper, we write F_r for the rank r free group and $\mathrm{Out}(F_r)$ for its outer automorphism group. We will always assume that $r \geq 3$. Finally, $\varphi \in \mathrm{Out}(F_r)$ will denote a nongeometric fully irreducible outer automorphism. Recall from the introduction that this means no power of φ preserves the conjugacy class of a nontrivial proper free factor and φ is not induced by a homeomorphism of a surface.

2.1. Culler-Vogtmann Outer Space.

Outer space was introduced by Culler and Vogtmann in [CV86] as an $\mathrm{Out}(F_r)$ analogue to Teichmüller space.

Let R_r be the r-petaled rose, i.e. the graph with precisely r edges and one vertex. Recall from [BH92] that a **marked graph** of rank r is a connected finite graph Γ, with no valence 1 or 2 vertices, together with an isomorphism $\pi_1(\Gamma) \cong F_r$ defined via a homotopy equivalence (called the **marking**) $\rho\colon \Gamma \to R_r$. Marked graphs $\rho\colon \Gamma \to R_r$ and $\rho'\colon \Gamma' \to R_r$ are considered equivalent when there exists a homeomorphism $h\colon \Gamma \to \Gamma'$ such that $\rho' \circ h$ is homotopic to ρ. We denote the edge set by $E\Gamma$ and the vertex set by $V\Gamma$.

A **metric** on Γ is the path metric determined by choosing for each edge e of Γ a length $\ell(e)$ and a characteristic map $j_e\colon [0, l(e)] \to e$, in the sense of CW complexes. A metric is determined, up to homeomorphism isotopic to the identity, by the assignment of lengths $\ell(e)$. The **volume** of Γ is defined as $\mathrm{vol}(\Gamma) := \sum\limits_{e \in E(\Gamma)} \ell(e)$.

The **unprojectivized outer space** $\widehat{CV_r}$ is the space of all metric marked graphs of rank r modulo marking-preserving isometry. The **outer space** CV_r itself is the projectivization of $\widehat{CV_r}$, i.e. the quotient of $\widehat{CV_r}$ by homothety. By instead viewing the points in CV_r as those Γ with $\mathrm{vol}(\Gamma) = 1$, one can see its decomposition into disjoint open simplices, one for each marked graph.

Lifting to the universal covers, one obtains an alternative definition of CV_r: Given a marked graph (Γ, m) in CV_r, by lifting to the universal cover, one obtains a simplicial tree with a free $\pi_1(\Gamma) \cong F_r$-action by deck transformations. The **compactified outer space** $\overline{CV_r} = CV_r \cup \partial CV_r$ is the space of minimal, very small F_r-actions on $\mathbb{R}$-trees, known as **F_r-trees**, modulo F_r-equivariant homothety.

Throughout this paper we often pass between the perspective of CV_r as a space of marked graphs or a space of F_r-trees without comment.

A **direction** at a point p in an $\mathbb{R}$-tree T is a connected component of $T \backslash \{p\}$ and if there are ≥ 3 directions at p it is a **branch point**. A **turn** at p is an unordered pair of directions at p. Let $\mathcal{D}_pT$ denote the set of directions at p, or $\mathcal{D}(p)$ if T is clear. Define $\mathcal{D}T := \cup_{p \in T} \mathcal{D}_pT$. Given a locally injective continuous map $f\colon T \to T'$ of $\mathbb{R}$-trees, define a **direction map** $Df\colon \mathcal{D}_pT \to \mathcal{D}_{f(p)}T'$ by sending a direction to its f-image.

2.2. **Train track representatives.** A homotopy equivalence $g\colon \Gamma \to \Gamma$ of a marked graph Γ is a **train track representative** for $\varphi \in Out(F_r)$ if it maps vertices to vertices, $\varphi = g_*\colon \pi_1(\Gamma) \to \pi_1(\Gamma)$, and $g^k\mid_{int(e)}$ is locally injective for each $e \in E\Gamma$ and $k > 0$. Many of the definitions and notation for discussing train track representatives were established in [BH92] and [BFH00]. We recall some here.

We denote the vertex set by $V\Gamma$ and edge set by $E\Gamma$. A **direction** at $v \in V\Gamma$ is a germ of initial segments of directed edges emanating from v. The set of directions at v is denoted $\mathcal{D}(v)$ and Dg will denote the direction map induced by g. We call a point v **periodic** if there exists a $j \geq 1$ such that $g^j(v) = v$ and a direction d at a periodic point v **periodic** if $Dg^k(d) = d$ for some $k > 0$. We call an unordered pair of directions $\{d_i, d_j\}$, based at the same point, a **turn**.

We call a locally injective path **tight**. Recall from [BH92] that a nontrivial tight path ρ in Γ is a **periodic Nielsen path (PNP)** for g if $g^k(\rho) \simeq \rho$ rel endpoints for some $k \in \mathbb{Z}_{>0}$, a **Nielsen path (NP)** if the period is one, and an **indivisible Nielsen path (iNP)** if it further cannot be written as a concatenation $\rho = \rho_1\rho_2$, where ρ_1 and ρ_2 are also NPs for g. [BH92] gives an algorithm for finding a representative with the minimal number of Nielsen paths, such a representative is called a **stable** representative.

As in [HM11], we call a periodic point $v \in \Gamma$ **principal** that either has at least three periodic directions or is an endpoint of a periodic Nielsen path.

A train track representative is called **rotationless** if every principal point is fixed and every periodic direction at each principal point is fixed. We use from [FH11, Corollary 4.43] that rotationless powers exist, depend only on the rank r, and fix all PNPs.

2.3. **Attracting tree T_+ and lamination Λ.** Each fully irreducible $\varphi \in \mathrm{Out}(F_r)$ acts on $\overline{CV_r}$ with an attracting tree $T^+_\varphi \in \partial CV_r$ and a repelling tree $T^-_\varphi \in \partial CV_r$ [LL03]. Dual to T^+_φ and T^-_φ, we have the repelling and attracting laminations respectively. In this paper we only concern ourselves with the attracting tree and lamination, which we thus write succinctly as T_+ and Λ when φ is clear. In this subsection, we provide a description of these objects in terms of train track representatives.

Construction 2.1 (Attracting tree T_+)**.** Let $g\colon \Gamma \to \Gamma$ be a train track representative of φ and $\widetilde{\Gamma}$ the universal cover of Γ equipped with a distance function $\widetilde{d}$ lifted from Γ. Then $\pi_1(\Gamma) \cong F_r$ acts by deck transformations, hence isometries, on $\widetilde{\Gamma}$. A lift $\widetilde{g}$ of g corresponds to a unique automorphism Φ representing φ. For each $w \in F_r$ and $x \in \widetilde{\Gamma}$, we have $\Phi(w)\widetilde{g}(x) = \widetilde{g}(wx)$. Define the pseudo-distance d_∞, for each $x, y \in \widetilde{\Gamma}$, by $d_\infty(x,y) = \lim_{k\to\infty} \frac{1}{\lambda^k} d(\widetilde{g}^k(x), \widetilde{g}^k(y))$. Then T_+ is the quotient of $\widetilde{\Gamma}$ under $x \sim y$ when $d_\infty(x,y) = 0$.

To describe the attracting lamination we need the following: Let Γ be a marked graph with universal cover $\widetilde{\Gamma}$ and projection map $p\colon \widetilde{\Gamma} \to \Gamma$. By a **line** in $\widetilde{\Gamma}$ we mean a proper embedding of the real line $\widetilde{\lambda}\colon \mathbb{R} \to \widetilde{\Gamma}$, modulo reparametrization. We denote by $\widetilde{\mathcal{B}}(\Gamma)$ the space of lines in $\widetilde{\Gamma}$ with the compact-open topology (generated by the open sets $\widetilde{\mathcal{U}}(\widetilde{\gamma}) := \{L \in \widetilde{\mathcal{B}}(\Gamma) \mid \widetilde{\gamma}$ is a finite subpath of $L\}$). A **line** in Γ is then the projection $p \circ \widetilde{\lambda}\colon \mathbb{R} \to \Gamma$ of a line $\widetilde{\lambda}$ in $\widetilde{\Gamma}$. We denote by $\mathcal{B}(\Gamma)$ the space of lines in Γ with the quotient topology induced by the natural projection map from $\widetilde{\mathcal{B}}(\Gamma)$ to $\mathcal{B}(\Gamma)$. One can then define $\mathcal{U}(\gamma) := \{L \in \mathcal{B}(\Gamma) \mid \gamma$ is a finite subpath of $L\}$, which generate the topology on $\mathcal{B}$. For a marked graph Γ, we say a line γ in Γ is **birecurrent** if each finite subpath of γ occurs infinitely often as an unoriented subpath in each end of γ.

Construction 2.2 (Attracting lamination Λ)**.** Fix a fully irreducible $\varphi \in Out(F_r)$ and consider any train track representative $g\colon \Gamma \to \Gamma$ for φ. Given any edge e in Γ, there exists a $k > 0$ such that the following is a sequence of nested open sets: $\mathcal{U}(e) \supset \mathcal{U}(g^k(e)) \supset \mathcal{U}(g^{2k}(e)) \dots$ The **attracting lamination** Λ is the set of birecurrent lines in the intersection. We often use the same notation for the total lift $\widetilde{\Lambda}$ of Λ to the universal cover. The meaning should be clear from context.

Remark 2.3 (Viewing Λ in trees $T \in \mathrm{CV}_r$)**.** The definition of Λ is well-defined, independent of the choice of train track representative; see [BFH97, Lemma 1.12] for proof. Once a basepoint lift is chosen in $\widetilde{\Gamma}$, one can identify $\partial\widetilde{\Gamma}$ with the hyperbolic boundary ∂F_r of F_r. This allows identification of $\widetilde{\Lambda}$ with a set of unordered pairs of points in ∂F_r, so $\widetilde{\Lambda}$ is also well-defined. Then define the realization of Λ in a general point of CV_r represented by a marked graph Γ' with universal cover $\widetilde{\Gamma}'$ and a chosen basepoint in $\widetilde{\Gamma}'$: Use the identifications $\partial\widetilde{\Gamma} \cong \partial F_r \cong \partial\widetilde{\Gamma}'$, to obtain $\widetilde{\mathcal{B}}(\Gamma) \cong \widetilde{\mathcal{B}}(\Gamma')$, identifying $\widetilde{\Lambda} \subset \widetilde{\mathcal{B}}(\Gamma)$ with a subset of $\widetilde{\mathcal{B}}(\Gamma')$ called the **realization** of $\widetilde{\Lambda}$ in $\widetilde{\Gamma}'$. Via the projection $\widetilde{\mathcal{B}}(\Gamma') \to \mathcal{B}(\Gamma')$, we obtain the **realization** of $\widetilde{\Lambda}$ in Γ'.

2.4. **Folds and splits.** Throughout this subsection, I will denote a (possibly infinite) interval.

A **fold path** in CV_r is a continuous, injective, proper function $\mathcal{F}\colon I \to CV_r$ defined by

1. a continuous 1-parameter family of marked graphs $t \to \Gamma_t$ and
2. a family of homotopy equivalences $h_{ts}\colon \Gamma_s \to \Gamma_t$ defined for $s \le t \in I$, each marking-preserving, satisfying:

- **(Train track property)**: h_{ts} is a local isometry on each edge for all $s \le t \in I$ and
- **(Semiflow property)**: $h_{ut} \circ h_{ts} = h_{us}$ for all $s \le t \le u \in I$ and $h_{ss}\colon \Gamma_s \to \Gamma_s$ is the identity for all $s \in I$.

We call $\mathcal{F}$ a **fold line** when $I = \mathbb{R}$. When $I = [a, b]$ for some $a < b$ and $h_{sa}(\Gamma_a)$ is homeomorphic to $h_{ta}(\Gamma_a)$ for each $a < s < t < b$, we call $\mathcal{F}$ a **fold**. In an abuse of notation, sometimes we use the same terminology for the quotient map h_{ba} as for $\mathcal{F}$. A fold with only 2 edges in the support is **simple**.

For a free, simplicial F_r-tree T, a **Λ-isometry** on T is an F_r-equivariant map $F_t\colon T \to T_+$ such that, for each leaf L of Λ realized in T, the restriction of F_t to L is an isometry onto a bi-infinite geodesic in T_+. Since F_t is continuous, there is a well-defined map of directions DF_t, with a restriction D_pF_t to the set of directions at p for each $p \in T$. A fold $\mathcal{F}$ is **Λ-legal** if h_{ts} is a Λ-isometry for each t, s. Note that a Λ-legal fold cannot identify the directions in a Λ-legal turn.

Viewed as a quotient map, a fold induces a map on directions and a gate structure in which **gates** are defined as equivalence classes of directions identified by the fold. We call this structure **weighted**, as each turn $\{d_1, d_2\}$ in a gate has an associated **length** $\ell(\{d_1, d_2\})$ that is the length of the initial segments of the edges e_1 and e_2, in the directions d_1 and d_2, identified by the fold. In particular $\ell(\{d_1, d_2\}) \le \ell(e_1), \ell(e_2)$.

Lemma 2.4. *Suppose that F_t is a Λ-isometry. Let $d_0, ..., d_m$ be directions in a gate. Then*

$$\min_{i=1,\dots,m}\{\ell(\{d_{i-1}, d_i\})\} \le \ell(\{d_0, d_m\}). \tag{1}$$

Proof. By induction, it suffices to show the lemma when $m = 2$. In this case, the initial segments of d_0 and d_1 of length $\ell(\{d_0, d_1\}) \le \min\{\ell(\{d_0, d_1\}), \ell(\{d_1, d_2\})\}$ have the same F_t-image, and the initial segments of d_1 and d_2 of length $\ell(\{d_1, d_2\}) \le \min\{\ell(\{d_0, d_1\}), \ell(\{d_1, d_2\})\}$ have the same F_t-image. Hence, at least the initial segments of d_0 and d_2 of at least length $\min\{\ell(\{d_0, d_1\}), \ell(\{d_1, d_2\})\}$ have the same F_t-image. $\square$

We call a continuous, injective, proper function $\mathcal{F}'\colon I \to CV_r$ defined by a continuous 1-parameter family of marked graphs $t \to \Gamma_t$ a **split path** if $\mathcal{F}(t) = \mathcal{F}'(a + b - t)$ is a fold path. We call $\mathcal{F}'$ a **split** if $\mathcal{F}$ is a fold, **Λ-legal** if $\mathcal{F}$ is, and **simple** if $\mathcal{F}$ is. Note that, since a Λ-legal fold must be a Λ-isometry, a Λ-legal split must also restrict to an isometry on each leaf L of Λ.

2.5. **Weak train tracks.** A **normalized weak train track** for φ is a $T \in \widehat{CV_r}$ on which a Λ-isometry exists. A **weak train track** is an element of CV_r represented by a normalized weak train track.

As explained in [HM11, Theorem 5.8], the choices of Λ-isometry on a fixed normalized weak train track T can be nonunique, and are parametrized by a closed interval (that is possibly a single point). In [HM11, §6.2], Handel and Mosher define a **right-most isometry** k_T^+ and [HM11, Lemma 6.3] provides that this assignment of k_T^+ is continuous (further explanation is provided in the remark following the lemma in [HM11]). In this paper, we always equip a normalized weak train track with its right-most Λ-isometry.

Having a canonical choice of a Λ-isometry ensures having a unique fold between elements of the axis bundle where one exists.

Lemma 2.5. *Let T, T' be normalized weak train tracks. If there is a fold $f : T \to T'$ such that $k_{T'}^+ f = k_T^+$, then f is unique.*

Proof. Suppose $f' : T \to T'$ is another fold such that $k_{T'}^+ f' = k_T^+$. Since k_T^+ and $k_{T'}^+$ are Λ-isometries, f and f' must be Λ-legal. This implies that for each leaf L of Λ, we have $f(L) = f'(L) = L$, as realized in T'.

Now let x be a point in T and L a leaf of Λ passing through x. Then $f(x)$ and $f'(x)$ lie on the realization of L in T'. But $f(x)$ and $f'(x)$ map to the same point $k_T^+(x)$ under $k_{T'}^+$, so we must have $f(x) = f'(x)$ otherwise $k_{T'}^+$ would not be a Λ-isometry. $\square$

As with folds, there is a **weighted induced gate structure** on T for each Λ-isometry $F_t\colon T \to T_+$.

Definition 2.6 (Fully preprincipal)**.** A (normalized) weak train track is **fully preprincipal** if, in the induced gate structure from its rightmost Λ-isometry, each vertex has ≥ 3 gates. This generalizes the [Pfa24] notion, by leaving out PNPs restrictions.

2.6. **Axis bundle.** The axis bundle $\mathcal{A}_\varphi$ for a nongeometric fully irreducible $\varphi \in Out(F_r)$ was first introduced in [HM11], where 3 equivalent definitions were given, with proof of their equivalence found in [HM11, Theorem 1.1]. Further description is given in [MP16]. We give here the only one of the three definitions we use: Fix a normalization of T_+. Then:

$$\widehat{\mathcal{A}_\varphi} = \{\text{free simplicial } F_r\text{-trees } T \in \widehat{CV_r} \mid \exists\ \Lambda\text{-isometry } f_T\colon T \to T_+\}.$$

In other words, $\widehat{\mathcal{A}_\varphi}$ is the set of normalized weak train tracks in $\widehat{CV_r}$. The **axis bundle** $\mathcal{A}_\varphi$ is the set of weak train tracks in CV_r for φ, i.e. $\mathcal{A}_\varphi$ is the image of $\widehat{\mathcal{A}_\varphi}$ under the projectivization of $\widehat{CV_r}$.

By [HM11, Lemma 5.1], each weak train track in CV_r is represented by a unique normalized weak train track in $\widehat{CV_r}$; equivalently, the projection $\widehat{CV_r} \to CV_r$ restricts to a bijection $\widehat{\mathcal{A}_\varphi} \to \mathcal{A}_\varphi$. As such, we may occasionally blur the distinction between weak train tracks and normalized weak train tracks.

We note that $\mathcal{A}_\varphi$ is also the union of the images of all fold lines $\mathcal{F}\colon \mathbb{R} \to CV_r$ such that $\mathcal{F}(\mathrm{t})$ converges in $\overline{CV_r}$ to T_-^φ as $t \to -\infty$ and to T_+^φ as $t \to +\infty$. An important example of such a fold line is a periodic fold line for a **Stallings fold decomposition** of a train track representative $g\colon \Gamma \to \Gamma$: At an illegal turn for g, fold two maximal initial segments with the same image to obtain a map $\mathfrak{g}_1\colon \Gamma_1 \to \Gamma$ of the quotient graph Γ_1. Repeated the process for $\mathfrak{g}_1$ and recursively. If some $\mathfrak{g}_k\colon \Gamma_{k-1} \to \Gamma$ has no illegal turn, then $\mathfrak{g}_k$ is a homeomorphism and the fold sequence is complete. Taking a rotationless power avoids the homeomorphism.

Several crucial properties of the axis bundle are recorded in [HM11, Theorem 6.1, Lemma 6.2]. We summarize a few here as Proposition 2.7.

Proposition 2.7 ([HM11])**.** *Let $\varphi \in \mathrm{Out}(F_r)$ be nongeometric fully irreducible. Then the map* $\mathrm{vol}\colon \widehat{\mathcal{A}_\varphi} \to (0,\infty)$ *is a surjective, φ-equivariant homotopy equivalence, where φ acts on $(0,+\infty)$ by multiplication by $\frac{1}{\lambda}$.*

The φ-action gives a means to decompose $\mathcal{A}_\varphi$ into compact fundamental domains.

Lemma 2.8. *Suppose that $\varphi \in \mathrm{Out}(F_r)$ is nongeometric fully irreducible. Then each fundamental domain of the φ-action on $\mathcal{A}_\varphi$ contains only finitely many preprincipal train tracks.*

Proof. First note that, since the fundamental domain is compact, it can intersect only finitely many simplices. So it suffices to show that each simplex contains only finitely many preprincipal train tracks of φ.

Since each vertex of each preprincipal train track T contains ≥ 3 gates, each vertex must be mapped by the Λ-isometry F_t to a branch point. Further, F_t, as a Λ-isometry, is an isometry on each edge. Since T_+ has only finitely many orbits of branch points, this gives a finite list of possible edge lengths in T. Together with the domain containing only finitely many simplices, this gives only finitely many possibilities for T. $\square$

We use the following construction, allowing us to connect the axis bundle to train track representatives.

Construction 2.9 (Train tracks)**.** Weak train tracks can be constructed from train track representatives:

Let $g\colon \Gamma \to \Gamma$ be a train track representative of a nongeometric fully irreducible $\varphi \in Out(F_r)$. Recall Construction 2.1 and let T_k denote the simplicial F_r-tree obtained from $\widetilde{\Gamma}$ by assigning the metric $d_k(x,y) = \frac{1}{\lambda^k} d(\widetilde{g}^k(x), \widetilde{g}^k(y))$, identifying each $x, y \in \widetilde{\Gamma}$ with $d_k(x,y) = 0$, and then equipping the quotient graph with the metric induced by d_k. Then, for each i, a basepoint-preserving lift of g induces a basepoint-preserving F_r-equivariant map $\widetilde{g}_{i+1,i}\colon T_i \to T_{i+1}$ restricting to an isometry on each edge. Define a direct system $\widetilde{g}_{j,i}\colon T_i \to T_j$ inductively by $\widetilde{g}_{j,i} = \widetilde{g}_{j,j-1} \circ \widetilde{g}_{j-1,i}$. Then $\widetilde{\Gamma}$ is a weak train track where the Λ-isometry $g_\infty\colon \widetilde{\Gamma} \to T_+$ is the direct limit map. $\widetilde{\Gamma}$ is called a **train track**. $TT(\varphi)$ denotes the set of train tracks for φ.

$$\begin{tikzcd}
\widetilde{\Gamma} = T_0 \arrow[r, "\widetilde{g}_{1,0}"'] \arrow[rrrr, bend left, "g_\infty"] & T_1 \arrow[r, "\widetilde{g}_{2,1}"'] \arrow[rrr, bend left, "g_{1,\infty}"] & T_2 \arrow[r, "\widetilde{g}_{3,2}"'] \arrow[rr, bend left, "g_{2,\infty}"] & \dots & T_+
\end{tikzcd}$$

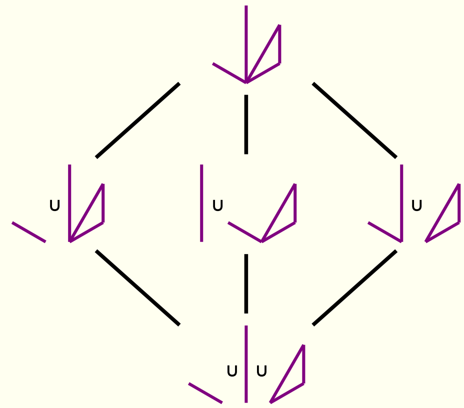

FIGURE 3. The cut decompositions of the top graph, arranged in the partial ordering of fineness.

2.7. **Ideal Whitehead graphs.** The ideal Whitehead graph of a nongeometric fully irreducible $\varphi \in \mathrm{Out}(F_r)$ is first defined in [HM11]. One can reference [Pfa12] and [HM11] for alternative definitions and its outer automorphism invariance. Further explanation yet can be found in [MP16, §2.8, §2.10].

Definition 2.10 (Ideal Whitehead graph $IW(\varphi)$)**.** Let $\varphi \in Out(F_r)$ be nongeometric fully irreducible with lifted attracting lamination $\widetilde{\Lambda}$ realized in T_+. $\widetilde{IW(\varphi)}$ is the union of the components with at least three vertices of the graph that has a vertex for each distinct leaf endpoint and an edge connecting the vertices for the endpoints of each leaf. F_r acts freely, properly discontinuously, and cocompactly in such a way that the restriction to each component of $\widetilde{IW(\varphi)}$ has trivial stabilizer. The **ideal Whitehead graph** $IW(\varphi)$ is the quotient under this action. It has only finitely many components.

Using Remark 2.3, one can view $\widetilde{IW(\varphi)}$ in any F_r-tree $T \in TT(\varphi^k)$.

Definition 2.11 (Principal points)**.** Given a branch point b of T_+, the lifted ideal Whitehead graph $\widetilde{IW}(\varphi)$ has one component, denoted $\widetilde{IW}_b(\varphi)$, whose edges, realized as lines in T_+, all contain b. This relationship gives a one-to-one correspondence between components of $\widetilde{IW}(\varphi)$ and branch points of T_+. Given a branch point b of T_+, let $\widetilde{IW}_b(\varphi; T)$ denote the realization of $\widetilde{IW}_b(\varphi)$ in T. This makes sense by viewing the ideal Whitehead graph in terms of the lamination leaves, as in Definition 2.10. We call a point v in T **principal** for f if there exists a branch point b of T_+ such that $f(v) = b$ and v lies in come leaf of $\widetilde{IW}_b(\varphi; T)$.

It is shown in [HM11], and can be ascertained from the alternative $IW(\varphi)$ definitions given there (and in [Pfa12]) that a principal point downstairs either has 3 periodic directions or is the endpoint of a PNP.

2.8. **Cut Decompositions.** Suppose $\varphi \in \mathrm{Out}(F_r)$ is nongeometric fully irreducible. In this subsection we describe a methodology developed by Handel and Mosher in [HM11, §4] for using cut vertices in $IW(\varphi)$ to obtain train track representatives with varied numbers of PNPs. We use the methodology in Proposition 2.17 to construct fully preprincipal train track representatives of φ that realize all possible "cut decompositions" of $IW(\varphi)$.

Suppose that G is a graph that can be written as a union of two nontrivial subgraphs of G intersecting in a single vertex $v \in VG$, then we call v a **cut vertex** of G. By a **cut decomposition** of G, we mean a collection of nontrivial connected subgraphs $\{G_1, \dots, G_k\}$ of G satisfying that

(1) $G = \cup G_j$ and
(2) $G_i \cap G_j$ is either empty or a vertex for each $i \neq j$.

As an example, we have listed out the cut decompositions of the graph at the top of Figure 3 in the figure. The cut decompositions are arranged in the partial order of fineness.

The ideal Whitehead graph has another interpretation in terms of singular leaves of $\widetilde{\Lambda}$. Here, a leaf of $\widetilde{\Lambda}$ is **singular** if it shares a ray with another leaf.

Definition 2.12 ($LW(\widetilde{v}, T)$, $SW(\widetilde{v}, T)$)**.** Let T be a weak train track and $\widetilde{v}$ a principal point of T. The **local Whitehead graph** $LW(\widetilde{v}; T)$ has a vertex for each direction at $\widetilde{v}$ and an edge connecting the vertices representing a pair of directions $\{d_1, d_2\}$ if the turn $\{d_1, d_2\}$ is taken by the realization in T of a leaf of $\widetilde{\Lambda}$. The **stable Whitehead graph** $SW(\widetilde{v}; T)$ at $\widetilde{v}$ is the subgraph of $LW(\widetilde{v}; T)$ obtained by restricting to the **principal directions**, i.e. those containing singular rays emanating from $\widetilde{v}$.

Each $SW(\widetilde{v}; T)$ sits inside $\widetilde{IW}(\varphi)$ as follows: A vertex of $SW(\widetilde{v}; T)$ corresponds to a singular leaf ray $\widetilde{R}$ emanating from $\widetilde{v}$. The endpoint of this ray corresponds to a vertex of $\widetilde{IW}(\varphi)$. An edge of $SW(\widetilde{v}; T)$ corresponds to a singular leaf based at $\widetilde{v}$ (as in Definition 2.11). This leaf also gives an edge of $\widetilde{IW}(\varphi)$.

The following is a restatement of [HM11, Lemma 5.2], focused on our purposes.

Lemma 2.13 ([HM11])**.** *Suppose that T is a weak train track and $F_T \colon T \to T_+$ a Λ-isometry. Suppose that b is a branch point of T_+ and $\{\widetilde{w_i}\} \subset T$ is the set of principal vertices mapped by F_T to b. Then*

(1) $\widetilde{IW}_b(\varphi; T) = \cup SW(\widetilde{w_i}; T)$.

(2) For each $i \neq j$, the intersection $SW(\widetilde{w_i}; T) \cap SW(\widetilde{w_j}; T)$ is at most one vertex. In the case where there is a vertex P in the intersection, we have that P is a cut point of $\widetilde{IW}(\varphi)$, separating $SW(\widetilde{w_i}; T)$ from $SW(\widetilde{w_j}; T)$ in $\widetilde{IW}(\varphi)$.

Definition 2.14 (Local decomposition)**.** The cut decomposition in Lemma 2.13 is the **local decomposition** of T.

Let T, T' be weak train tracks. As in [HM11], one says T **is split at least as much as** T' if the local decomposition $\widetilde{IW}(\varphi) = \bigcup SW(v_j; T)$ is at least as fine as the local decomposition $\widetilde{IW}(\varphi) = \bigcup SW(w_i; T')$. That is, for each principal vertex v_j of T, there exists a principal vertex w_i of T' such that $SW(v_j; T) \subset SW(w_i; T)$, where the inclusion takes place in $\widetilde{IW}(\varphi)$, realized as a decomposition, as above.

Suppose $\varphi \in \mathrm{Out}(F_r)$ is a rotationless nongeometric fully irreducible outer automorphism and $\mathcal{G}$ is a cut decomposition of $IW(\varphi)$. [HM11, Lemma 4.3] provides a method to obtain a train track representative g of φ so that $\widetilde{\Gamma}$ has local decomposition $\mathcal{G}$. More specifically, g is obtained from a stable train track representative h of φ via iteratively "splitting open" at cut vertices of the stable Whitehead graphs, as follows.

Construction 2.15 (Splitting open a cut vertex)**.** Suppose w is a cut vertex of a stable Whitehead graph $SW(f, u)$ of a train track representative $f \colon \Gamma \to \Gamma$ of φ and G_1, G_2 are nontrivial subgraphs of $SW(f, u)$ meeting only at w and with $SW(f, u) = G_1 \cup G_2$. Then w is represented by a fixed direction d_0 at u, as well as a collection $d_1, \ldots, d_N$ of directions mapped to d_0 by powers of Df. Let E_0 be the edge in the direction of d_0. See Figure 4 left. In the top row of the figure, we have drawn the graph Γ. In the bottom row, we draw a 'blown-up' view where we insert the local Whitehead graph $LW(v, T)$ at each vertex v, with the purple edges lying in $SW(v, T)$ and the red edges lying in $LW(v, T)$ but not $SW(v, T)$.

We explain how to form an NP via splitting open at E_0. The same procedure should simultaneously be applied to the edge in the direction of each of $d_1, \ldots, d_N$.

G_1 and G_2 correspond to a bipartition of $\mathcal{D}(u) \backslash \{d_0\}$ satisfying that the directions of each f-taken turn are in the same partition element. Create from Γ a new graph Γ' where

- $V\Gamma' = (V\Gamma \backslash \{u\}) \cup \{u_1, u_2\}$, and
- each edge $e_j = [v_j', u]$ is replaced with $[v_j', u_1]$, and
- each edge $e_j' = [v_j'', u]$ is replaced with $[v_j', u_2]$, and
- $E_0 = [u, v]$ is replaced by the 2 edges $[u_1, v]$ and $[u_2, v]$, and
- all other edges remain the same.

See the middle image of Figure 4.

The map f' is the same as f except that E_0 is replaced by E_1 in the image of any edge when u was passed through via G_1 and by E_2 when u was passed through via G_2, and analogous alterations occur for the edges in the directions of $d_1, \ldots, d_N$. The images of E_1 and E_2 are that of E, but that the image of E_1 now starts with E_1 and that of E_2 with E_2. Note that E_0 was a fixed direction, as it was represented by a vertex in $SW(f, u)$, so that this map of the E_j makes sense. In the cases of $d_1, \ldots, d_N$, instead of the image of E_1, E_2 starting with respectively E_1, E_2, the f-image of e would start with $Df(e)$.

We call this procedure **splitting u open along** E.

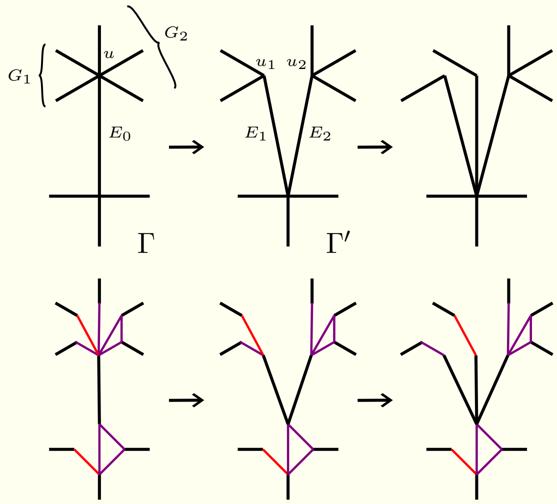

FIGURE 4. Splitting open a cut vertex. In the top row, we have drawn the graph Γ. In the bottom row, we draw a 'blown-up' view where we insert the local Whitehead graph $LW(v,T)$ at each vertex v, with the purple edges lying in the stable Whitehead graph $SW(v,T)$ and the red edges lying in $LW(v,T)$ but not $SW(v,T)$.

Lemma 2.16. *Suppose that g is a fully preprincipal train track representative of a rotationless nongeometric fully irreducible $\varphi \in \mathrm{Out}(F_r)$. Then splitting open at a cut vertex yields a train track representative of φ that can be Λ-legal split to obtain a fully preprincipal train track representative of φ.*

Proof. Suppose g is fully preprincipal and u was split open along $E = [u,v]$. Note that v still has ≥ 3 gates. Our concern is that either u_1 or u_2 has only 2 gates. Without loss of generality we assume u_1 has only 2 gates and that the directed edges at u are $E_1, e_1, \dots, e_n$. Then the only possible Λ-taken turns at u are of the form $\{E_1, e_j\}$. So it is possible to split open u_1 to replace $\{E_1, e_1, \dots, e_n\}$ with $\{\bar{E}_1 e_1, \dots, \bar{E}_1 e_n\}$. See Figure 4 right. The process can be repeated at u_2, if necessary. ☐

Proposition 2.17. *Suppose $\varphi \in \mathrm{Out}(F_r)$ is rotationless nongeometric fully irreducible. Then each cut decomposition of $IW(\varphi)$ is realized by a fully preprincipal train track representative of φ.*

Proof. By [Pfa24, Proposition 7.1], φ has a fully preprincipal train track representative with the minimal number of PNPs (this is explicitly stated in the ageometric case, but follows the same argumentation in the parageometric case). Repeated application of Lemma 2.16 yields a train track representative whose stable Whitehead graphs give the desired cut decomposition. ☐

The stable axis bundle was introduced in [HM11, §6.5] as an object of interest and was used extensively in [MP16]. We expand upon the notion of the stable axis bundle to define an axis bundle for each cut decomposition of the ideal Whitehead graph.

Definition 2.18 ($\mathcal{G}$-axis bundle $\mathcal{GA}_\varphi$)**.** Suppose $\varphi \in \mathrm{Out}(F_r)$ is nongeometric fully irreducible and $\mathcal{G}$ is a cut decomposition of $IW(\varphi)$. The **$\mathcal{G}$-axis bundle $\mathcal{GA}_\varphi$** is the set of all weak train tracks whose local decomposition is at most as coarse as $\mathcal{G}$. Under this terminology, the **stable axis bundle** is the $\mathcal{G}$-axis bundle where $\mathcal{G}$ is the coarsest possible local decomposition $\{IW(\varphi)\}$.

The partial ordering of fineness will be important for us because of the role it plays in [HM11, Proposition 5.4], which we record here as Proposition 2.19. We will use this proposition to ensure that each weak train track is contained in the branched cube determined by a fully preprincipal train track with local decomposition as split as its own (Proposition 3.14).

Proposition 2.19 ([HM11])**.** *Let $\varphi \in Out(F_r)$ be nongeometric fully irreducible. Then for any train track representative $g\colon \Gamma \to \Gamma$ for φ with associated Λ-isometry $g_\infty\colon \widetilde{\Gamma} \to T_+$, there exists an $\varepsilon > 0$ so that, if*

$f: T \to T_+$ is any Λ-isometry, if g_∞ splits at least as much as f, and if $Len(T) \leq \varepsilon$, then there exists a unique equivariant edge-isometry $h: \widetilde{\Gamma} \to T$ such that $g_\infty = f \circ h$. Moreover, h is a Λ-isometry.

Proposition 2.20. *Let $\varphi \in \mathrm{Out}(F_r)$ be nongeometric fully irreducible and let $\mathcal{G}$ be a cut decomposition of $IW(\varphi)$. Then the $\mathcal{G}$ axis bundle $\mathcal{G}\mathcal{A}_\varphi$ is connected.*

Proof. Suppose $S_1, S_2 \in \mathcal{G}\mathcal{A}_\varphi$. By Proposition 2.17, there is a fully preprincipal $T \in \mathcal{A}_\varphi$ with local decomposition $\mathcal{G}$. Using Proposition 2.7, we shift S_1 and S_2 so that the ε requirement in Proposition 2.19 is satisfied and then use Proposition 2.19 to know that both S_1 and S_2 can be obtained from T by folding. In other words, both S_1 and S_2 can be connected to T by a fold path, thus by a path. □

3. Branched cubes in the axis bundle

In this section we start building a cubist complex structure on the axis bundle $\mathcal{A}_\varphi$ by describing the branched cubes. We then show some combinatorial properties of the interaction between these branched cubes. Unless otherwise indicated, we assume throughout this section that $\varphi \in \mathrm{Out}(F_r)$ is nongeometric fully irreducible, $T \in \mathcal{A}_\varphi$ is a fully preprincipal weak train track endowed with the weighted gate structure induced by its rightmost Λ-isometry $F_t: T \to T_+$, and $\mathcal{T}_0$ is the set of illegal turns in T.

3.1. **Description of the branched cubes.** For each $(\ell_\tau) \in \prod_{\tau \in \mathcal{T}_0}[0, \ell(\tau)]$, let $T_{(\ell_\tau)} \in \widehat{\mathrm{CV}}_r$ denote the metric graph obtained from T by folding each turn $\tau \in \mathcal{T}_0$ along initial segments of length ℓ_τ.

Definition 3.1. Let $T \in \mathcal{A}_\varphi$ be fully preprincipal. The **branched cube** at T is the set

$$B_T = \left\{ T_{(\ell_\tau)} \mid (\ell_\tau) \in \prod_{\tau \in \mathcal{T}_0} [0, \ell(\tau)] \right\} \subset \widehat{\mathrm{CV}}_r.$$

For future reference, note that for each $T_{(\ell_\tau)} \in B_T$ there is a canonical fold $h: T \to T_{(\ell_\tau)}$.

Lemma 3.2. *Each branched cube B_T is a subset of $\mathcal{A}_\varphi$. Further, if $T \in \mathcal{G}\mathcal{A}_\varphi$, then B_T is a subset of $\mathcal{G}\mathcal{A}_\varphi$.*

Proof. Suppose $T' \in B_T$. To show the first statement, it suffices to show that there exists a Λ-isometry $F_{T'}: T' \to T_+$. Consider the natural fold $f: T \to T'$. Since f folds along illegal turns, we know both that f restricts to an isometry on each leaf of Λ and that, if $f(x) = f(y)$, then $F_T(x) = F_T(y)$. Thus, we can define $F_{T'}$ as the quotient of F_T induced by the quotient map f.

Now suppose in addition that $T \in \mathcal{G}\mathcal{A}_\varphi$. The second statement follows from the fact that folding cannot increase the number of PNPs and a cut decomposition is determined by PNPs. □

One drawback to the notation $T_{(\ell_\tau)}$ is that we can have $T_{(\ell_\tau)} = T_{(\ell'_\tau)}$ for $(\ell_\tau) \neq (\ell'_\tau)$. To overcome this, we define functions capturing how much folding we actually do, as opposed to how much folding we were instructed to do: Let $\{d_1, d_2\}$ be a turn in T. Let d_i denoted the initial direction of the oriented edge e_i. For each $T' \in B_T$, we define $x_{\{d_1,d_2\}}(T')$ to be the length of the largest initial segments in e_1 and e_2 identified by the fold $T \to T'$. For example, if d_1 and d_2 do not lie in the same gate, then $x_{\{d_1,d_2\}}(T') = 0$.

Lemma 3.3. *Suppose $T \in \mathcal{A}_\varphi$ is fully preprincipal and $d_0, ..., d_m$ is a sequence of directions in T. Then*

$$\min_{i=1,\ldots,m} \{x_{\{d_{i-1},d_i\}}\} \leq x_{\{d_0,d_m\}}. \tag{2}$$

Proof. The proof of this is exactly the same as Lemma 2.4. □

Together with a decomposition of B_T into cubes as will be described in Section 3.3, the functions x_τ will serve as a kind of coordinate system.

3.2. **Examples of branched cubes.** Before analyzing the structure of the branched cubes in general, we take a moment to go through some examples. For simplicity, let us assume that only one gate, G, of T has more than one direction.

Example 3.4. Suppose G has exactly 2 directions, which we denote as d_1, d_2. The set of train tracks B that can be obtained from T by folding the unique turn $\{d_1, d_2\}$ in G is homeomorphic to an interval. The coordinate $x_{\{d_1,d_2\}}$ on B is the length folded.

Example 3.5. Suppose G has exactly 3 directions, which we denote as d_1, d_2, d_3 (top left in the figure). Then the set of illegal turns is $\mathcal{T}_0 = \{\{d_1, d_2\}, \{d_1, d_3\}, \{d_2, d_3\}\}$.

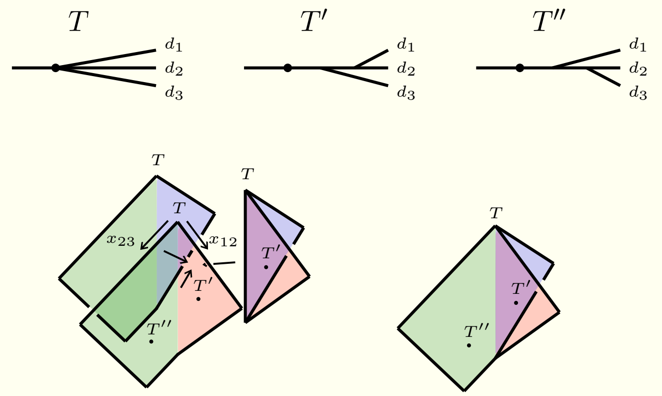

For each subset $\mathcal{T} \subset \mathcal{T}_0$ containing at most two of the three turns, let $C_{\mathcal{T}}$ denote the set of train tracks obtained by folding T along the turns in $\mathcal{T}$. Then $C_{\mathcal{T}}$ is homeomorphic to a $|\mathcal{T}|$-dimensional cube for each $\mathcal{T}$. The coordinates of the cube are the lengths x_τ we folded the turns $\tau \in \mathcal{T}$. If $\mathcal{T}' \subset \mathcal{T}$, then $C_{\mathcal{T}'}$ is a *splitting face* of $C_{\mathcal{T}}$ and consists of all points where the coordinates x_τ for $\tau \in \mathcal{T} \setminus \mathcal{T}'$ are zero.

The interaction between $C_{\mathcal{T}}$ and $C_{\mathcal{T}'}$ is more interesting when neither $\mathcal{T}$ nor $\mathcal{T}'$ is a subset of the other. Let T' be a train track obtained from T by folding $\{d_1, d_2\}$ by x_{12} and $\{d_2, d_3\}$ by x_{23}, where $x_{12} \geq x_{23}$. Then, by definition, $T' \in C_{\{\{d_1,d_2\},\{d_2,d_3\}\}}$. However, T' can also be obtained from T by folding $\{d_1, d_2\}$ by x_{12} and $\{d_1, d_3\}$ by x_{13}. This is because along the interval in which we are identifying d_1 and d_3, we have already identified d_1 and d_2. Thus $T' \in C_{\{\{d_1,d_2\},\{d_1,d_3\}\}}$ as well. Now let T'' be a train track obtained from T by folding $\{d_1, d_2\}$ by x_{12} and $\{d_2, d_3\}$ by x_{23}, where $x_{12} < x_{23}$. Then $T'' \in C_{\{\{d_1,d_2\},\{d_2,d_3\}\}}$ but $T'' \notin C_{\{\{d_1,d_2\},\{d_1,d_3\}\}}$. This argument reveals that $C_{\{\{d_1,d_2\},\{d_2,d_3\}\}}$ meets $C_{\{\{d_1,d_2\},\{d_1,d_3\}\}}$ along the *slice* $\{(x_{12}, x_{23}) \mid x_{12} \geq x_{23}\}$.

A similar argument holds for any pair $C_{\mathcal{T}}$ and $C_{\mathcal{T}'}$ where $\mathcal{T}$ and $\mathcal{T}'$ are 2-element sets. Thus the set of train tracks B that can be obtained by folding T, being $C_{\{\{d_1,d_2\},\{d_1,d_3\}\}} \cup C_{\{\{d_1,d_2\},\{d_2,d_3\}\}} \cup C_{\{\{d_1,d_3\},\{d_2,d_3\}\}}$, is a branched cube as shown in the bottom right of the image.

Example 3.6. In general, a subset $\mathcal{T}$ of $\mathcal{T}_0$ specifies a cube $C_{\mathcal{T}}$ of train tracks obtained by folding T along the turns in $\mathcal{T}$ if and only if the turns in $\mathcal{T}$ can be folded independently of one another. For example, if $\mathcal{T}$ contains a triplet of the form $\{\{d_1, d_2\}, \{d_1, d_3\}, \{d_2, d_3\}\}$, then folding $\{d_1, d_2\}$ and $\{d_1, d_3\}$ will also force folding $\{d_2, d_3\}$. On the other hand, if $\mathcal{T}$ is of the form $\{\{d_1, d_2\}, \{d_2, d_3\}, \{d_3, d_4\}\}$, then independent folding is possible. We formalize this idea later by considering the elements of $\mathcal{T}_0$ as edges of a graph on the elements of G. The set of train tracks B that can be obtained by folding T is the union of the $C_{\mathcal{T}}$ as $\mathcal{T}$ ranges over such *independent* subsets of $\mathcal{T}_0$.

The possible intersections between the cubes $C_{\mathcal{T}}$ also get more complicated as the number of elements in G grows. For example, suppose $\mathcal{T} = \{\{d_1, d_2\}, \{d_2, d_3\}, \{d_3, d_4\}, \{d_4, d_5\}\}$ and $\mathcal{T}' = \{\{d_1, d_2\}, \{d_3, d_4\}, \{d_1, d_4\}\}$. Then, using the coordinates $(x_{12}, x_{23}, x_{34}, x_{45})$ on $C_{\mathcal{T}}$ given by the fold lengths, we claim that $C_{\mathcal{T}}$ meets $C_{\mathcal{T}'}$ along the slice $S_{\mathcal{T},\mathcal{T}'} = \{(x_{12}, x_{23}, x_{34}, x_{45}) \mid x_{23} \leq \min\{x_{12}, x_{34}\} \ \&\ x_{45} = 0\}$.

For the $S_{\mathcal{T},\mathcal{T}'} \subseteq C_{\mathcal{T}} \cap C_{\mathcal{T}'}$ direction, consider a train track $T \in C_{\mathcal{T}}$ satisfying $x_{23} \leq \min\{x_{12}, x_{34}\}$ and $x_{45} = 0$. We can also reach T by first folding $\{d_1, d_2\}$ for x_{12} and $\{d_3, d_4\}$ for x_{34}, then folding $\{d_1, d_4\}$ for x_{23}. See Figure 5, top middle. The point is that since $x_{23} \leq \min\{x_{12}, x_{34}\}$, along the folded segments of d_1 and d_4, we have d_1 is already folded with d_2, as is d_3 with d_4. So folding $\{d_1, d_4\}$ or $\{d_2, d_3\}$ is the same.

For $\supseteq$, consider a train track $T \in C_{\mathcal{T}} \cap C_{\mathcal{T}'}$. Then necessarily $x_{45} = 0$ since the direction d_5 is not folded at all for train tracks in $C_{\mathcal{T}'}$. Meanwhile, for all train tracks in $C_{\mathcal{T}'}$, the turn $\{d_2, d_3\}$ must be folded at most as much as $\{d_1, d_4\}$. Under the coordinates $(x_{12}, x_{23}, x_{34}, x_{45})$, the amount of folding for the former is x_{23} while that of the latter is $\min\{x_{12}, x_{23}, x_{34}\}$. Thus $x_{23} \leq \min\{x_{12}, x_{23}, x_{34}\}$, or equivalently, $x_{23} \leq \min\{x_{12}, x_{34}\}$.

One can similarly verify that using the coordinates (x_{12}, x_{34}, x_{14}) on $C_{\mathcal{T}'}$ given by lengths of folding, $C_{\mathcal{T}'}$ meets $C_{\mathcal{T}}$ along the slice $S_{\mathcal{T}',\mathcal{T}} = \{(x_{12}, x_{34}, x_{14}) \mid x_{14} \leq \min\{x_{12}, x_{34}\}\}$.

In general, the slices of intersection will be given by inequalities determined from a graph on the elements of G, with edges $\mathcal{T}_0$, as described above.

3.3. **The decomposition into cubes.** We view $\mathcal{T}_0$ as a set of possible edges in a (simple, undirected) graph whose vertex set is $\mathcal{D}(T)$. In particular, every subset $\mathcal{T} \subset \mathcal{T}_0$ determines such a graph $\mathfrak{G}_{\mathcal{T}}$ that is the graph with an edge for each turn in $\mathcal{T}$. We say $\mathcal{T}$ is **independent** if $\mathfrak{G}_{\mathcal{T}}$ has no cycles.

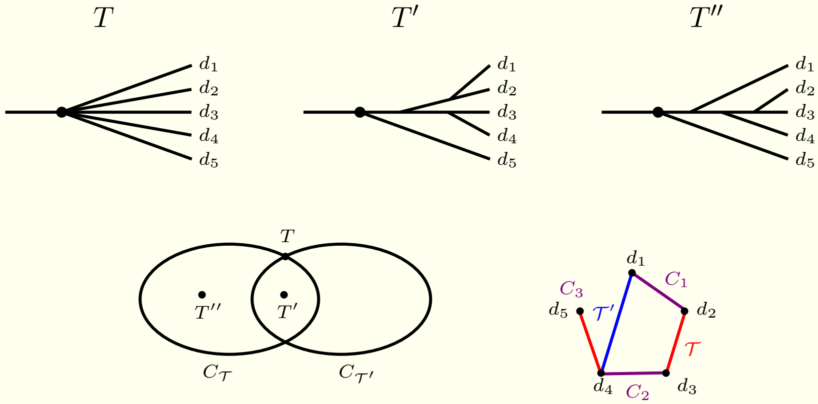

FIGURE 5. An example where T has exactly one gate G of more than one direction and G has exactly 5 directions. The folded train track T' is in $C_{\mathcal{T}} \cap C_{\mathcal{T}'}$. The folded train track T'' is in $C_{\mathcal{T}}$ but not $C_{\mathcal{T}'}$. Bottom right is the graph for determining the relevant inequalities.

Definition 3.7 (Cubes)**.** Let $\mathcal{T} \subset \mathcal{T}_0$ be an independent subset. The **cube** $C_{T,\mathcal{T}}$ is defined as

$$C_{T,\mathcal{T}} = \left\{ T_{(x_\tau,0)} \mid (x_\tau) \in \prod_{\tau\in\mathcal{T}} [0, \ell(\tau)] \right\},$$

where $(x_\tau, 0)$ is the element of $\prod_{\tau\in\mathcal{T}_0}[0, \ell(\tau)]$ with coordinate x_τ for $\tau \in \mathcal{T}$ and 0 for $\tau \in \mathcal{T}_0 \backslash \mathcal{T}$.

By definition, we have $C_{T,\mathcal{T}} \subset B_T$. Further, for each $\tau \in \mathcal{T}$, we have $x_\tau(T_{(x_\tau,0)}) = x_\tau$. The following lemma states that the functions x_τ parametrize $C_{T,\mathcal{T}}$ as a cube.

Lemma 3.8. *The map $\prod_{\tau\in\mathcal{T}}[0, \ell(\tau)] \to C_{T,\mathcal{T}}$ defined by $(x_\tau) \mapsto T_{(x_\tau,0)}$ is a homeomorphism.*

Proof. The map is continuous and surjective by definition. Since $\prod_{\tau\in\mathcal{T}}[0, \ell(\tau)]$ is compact, it suffices to show injectivity. From the proof of Lemma 3.2, the natural fold $f : T \to T'$ satisfies $F_{T'} f = F_T$. By Lemma 2.5, f is the unique fold satisfying this property. Injectivity thus follows since each x_τ is defined as the length of the largest initial segments with the same f-image. □

The functions $x_{\{d_1,d_2\}}$, for $\{d_1, d_2\} \notin \mathcal{T}$, are PL functions on $C_{T,\mathcal{T}}$ in the coordinates x_τ, for the $\tau \in \mathcal{T}$:

Lemma 3.9. *Suppose $x \in C_{T,\mathcal{T}}$ and $\{d_1, d_2\} \notin \mathcal{T}$. Recall the graph $\mathfrak{G}_{\mathcal{T}}$ corresponding to $\mathcal{T}$.*

(a) If there is a path in $\mathfrak{G}_{\mathcal{T}}$ connecting d_1 and d_2, let p be the shortest such path. Then

$$x_{\{d_1,d_2\}} = \min_{\tau\in p} x_\tau. \tag{3}$$

(b) If there is no path in $\mathfrak{G}_{\mathcal{T}}$ connecting d_1 and d_2, then $x_{\{d_1,d_2\}}$ is identically zero.

Proof. We first show (b). If there is no path in $\mathfrak{G}_{\mathcal{T}}$ connecting d_1 and d_2, then either d_1 and d_2 are not directions at the same gate, or they lie in the same gate but are not folded. In both cases, $x_{\{d_1,d_2\}} = 0$.

For (a), we have $x_{\{d_1,d_2\}} \geq \min_{\tau\in p} x_\tau$ since the initial segments of length $\min_{\tau\in p} x_\tau$ of each direction in p, including d_1 and d_2, are all identified together. Conversely, let $\tau_0 \in p \subset \mathcal{T}$ be such that $\min_{\tau\in p} x_\tau = x_{\tau_0}$ and let T' be the element of $C_{T,\mathcal{T}}$ with the same coordinates x_τ, for $\tau \in \mathcal{T}$, as x except $x_{\tau_0} = 0$. Then $T' \in C_{T,\mathcal{T}'}$ for $\mathcal{T}' = \mathcal{T}\backslash\{\tau_0\}$. Since $\mathcal{T}$ was independent, there is no path in $\mathfrak{G}_{\mathcal{T}'}$ connecting d_1 and d_2. Hence, by (2), $x_{\{d_1,d_2\}}(T') = 0$. But x is obtained from T' by folding τ_0 for x_{τ_0}, so $x_{\{d_1,d_2\}} \leq x_{\tau_0} = \min_{\tau\in p} x_\tau$. □

The branched cube B_T is the union of the cubes $C_{T,\mathcal{T}}$ for the independent sets of turns $\mathcal{T} \subset \mathcal{T}_0$:

Proposition 3.10. *Let T be a fully preprincipal train track. Then*

$$B_T = \bigcup_{\textit{independent } \mathcal{T} \subset \mathcal{T}_0} C_{T,\mathcal{T}}.$$

Proof. Again, each $C_{T,\mathcal{T}} \subset B_T$ by the definitions. So we consider $T_{(\ell_t)} \in B_T$ and show $T_{(\ell_t)} \in C_{T,\mathcal{T}}$ for some $\mathcal{T}$. Define an independent $\mathcal{T} \subset \mathcal{T}_0$ inductively as follows: Start with $\mathcal{T} = \varnothing$. At each stage, consider the turns $\tau \in \mathcal{T}_0$ such that $\mathcal{T} \cup \{\tau\}$ is independent. Among such τ, pick one such that ℓ_τ is maximal, and add that turn to $\mathcal{T}$. Since $\mathcal{T}_0$ is finite, this process terminates eventually and we have an independent subset $\mathcal{T}$.

We claim $T_{(\ell_\tau)} = T_{(\ell_\tau,0)}$. First observe that we can obtain $T_{(\ell_\tau)}$ from T by first folding the $\tau \in \mathcal{T}$ by ℓ_τ, to get $T_{(\ell_\tau,0)}$, then folding the remaining $\tau \notin \mathcal{T}$ by ℓ_τ. But for each remaining turn $\{d_1, d_2\}$, we have $x_{\{d_1,d_2\}}(T_{(\ell_\tau,0)})$ is the minimum of the ℓ_τ as τ ranges over edges in $\mathcal{T}$ connecting d_1 to d_2. The value $\ell_{\{d_1,d_2\}}$ cannot exceed this value, or we would have chosen $\{d_1, d_2\}$ to be in $\mathcal{T}$. That is, the amount of folding we have to do on $\{d_1, d_2\}$ was already done in $T_{(\ell_\tau,0)}$, so there is no extra folding to be done, and $T_{(\ell_\tau)} = T_{(\ell_\tau,0)}$. $\square$

Our next task is to compute the intersections between the cubes $C_{T,\mathcal{T}}$ in B_T. Let $\mathcal{T}$ and $\mathcal{T}'$ be two independent subsets of $\mathcal{T}_0$. Let the components of the graph $\mathfrak{G}_{\mathcal{T}\cap\mathcal{T}'}$ be $C_1, ..., C_k$. Note that a single disconnected vertex is also a component. We define the **slice** $S_{\mathcal{T},\mathcal{T}'}$ to be the subset of $C_{\mathcal{T}}$ consisting of all points satisfying the following inequalities for each $i, j = 1, ..., k$:

$$(4) \qquad \begin{cases} x_{\{d_i,d_j\}} \le \min\{x_{\{d_i,d_i'\}}, x_{\{d_j,d_j'\}}\} & \text{if } \exists\ \{d_i,d_j\} \in \mathcal{T} \text{ and } \{d_i',d_j'\} \in \mathcal{T}' \text{ connecting } C_i \ \&\ C_j \\ x_{\{d_i,d_j\}} = 0 & \text{if } \exists\ \{d_i,d_j\} \in \mathcal{T} \text{ connecting } C_i \ \&\ C_j \text{, but no such } \tau' \in \mathcal{T}'. \end{cases}$$

Once can verify that this gives the correct set of inequalities for an example in Example 3.6, using the graph in Figure 5 bottom right.

Lemma 3.11. *Let $\mathcal{T}$ and $\mathcal{T}'$ be two independent subsets of $\mathcal{T}_0$. The subset of points in $C_{T,\mathcal{T}}$ that also lie in $C_{T,\mathcal{T}'}$ is the slice $S_{\mathcal{T},\mathcal{T}'}$.*

Proof. Suppose $T' = (x_\tau) \in C_{T,\mathcal{T}}$ lies in the slice $S_{\mathcal{T},\mathcal{T}'}$, i.e. the inequalities of Equation (4) are satisfied for each $i, j = 1, ..., k$. Think of T' as obtained from T by (1) folding the $\tau \in \mathcal{T} \cap \mathcal{T}'$ for x_τ, then (2) folding the $\{d_i, d_j\} \in \mathcal{T}\backslash\mathcal{T}'$ for $x_{\{d_i,d_j\}}$. For the folds done in step (2), since $x_{\{d_i,d_j\}} = 0$ if there is no $\tau' \in \mathcal{T}'$ connecting C_i to C_j, we only fold $\{d_i, d_j\}$ for values of i, j for which there exists a $\{d_i', d_j'\} \in \mathcal{T}'$ connecting C_i to C_j.

Since $x_{\{d_i,d_j\}} \le \min\{x_{\{d_i,d_i'\}}, x_{\{d_j,d_j'\}}\}$, by the time we do the step (2) folds, we are folding initial segments of d_i and d_j identified with d_i' and d_j', so we could have equivalently folded $\{d_i', d_j'\}$ by $x_{\{d_i,d_j\}}$. That is, $T' \in C_{T,\mathcal{T}} \cap C_{T,\mathcal{T}'}$.

Conversely, suppose $T' \in C_{T,\mathcal{T}} \cap C_{T,\mathcal{T}'}$. Then for each $\{d_i, d_j\} \in \mathcal{T}\backslash\mathcal{T}'$, considering T' as a point in $C_{T,\mathcal{T}'}$,

$$\begin{aligned} x_{\{d_i,d_j\}} &= \begin{cases} \min\{x_{\{d_i,d_i'\}}, x_{\{d_i',d_j'\}}, x_{\{d_j',d_j\}}\} & \text{if } \exists\ \{d_i',d_j'\} \in \mathcal{T}' \text{ connecting } C_i \ \&\ C_j \\ 0 & \text{otherwise} \end{cases} \\ &\le \begin{cases} \min\{x_{\{d_i,d_i'\}}, x_{\{d_j',d_j\}}\} & \text{if } \exists\ \{d_i',d_j'\} \in \mathcal{T}' \text{ connecting } C_i \ \&\ C_j \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

$\square$

3.4. **Splitting and folding faces.** In this subsection we define the splitting and folding faces of a branched cube. We use this terminology when analyzing the combinatorics of how branched cubes intersect each other.

Let T be a fully preprincipal train track and $\mathcal{T} \subset \mathcal{T}_0$ an independent subset. The **splitting face** of $C_{T,\mathcal{T}}$ associated to $\mathcal{T}' \subset \mathcal{T}$ is the subset of $C_{T,\mathcal{T}}$ defined by $x_\tau = 0$ for all $\tau \in \mathcal{T}'$. Note that this is the same set as $C_{T,\mathcal{T}'}$. The **splitting vertex** of $C_{T,\mathcal{T}}$ is the splitting face associated to $\mathcal{T}$ itself, i.e. the point $\{T\}$. The **folding face** of $C_{T,\mathcal{T}}$ associated to $\mathcal{T}' \subset \mathcal{T}$ is the subset of $C_{T,\mathcal{T}}$ defined by $x_\tau = \ell(\tau)$ for all $\tau \in \mathcal{T}'$.

The splitting/folding faces of B_T will be unions of the splitting/folding faces of the $C_{T,\mathcal{T}}$: The **splitting face** of B_T associated to a choice of partition $G = \sqcup G_i$ for each gate G of T is the subspace of B_T defined by $x_\tau = 0$ whenever $\tau \not\subset G_i$ for all i. The **splitting vertex** of B_T is the splitting face associated to the partition of each gate into one-element sets, so is defined by $x_\tau = 0$ for all τ, and is just the point $\{T\}$. A splitting face of B_T is a proper subset of B_T if and only if at least one of the gate partitions is a nontrivial partition. The **folding face** of B_T associated to a subset $\mathcal{T} \subset \mathcal{T}_0$ is the subspace of B_T defined by $x_\tau = \ell(\tau)$ for each $\tau \in \mathcal{T}$. Such a folding face is a proper subset of B_T if and only if $\mathcal{T}$ is a nonempty subset of $\mathcal{T}_0$.

Lemma 3.12. *A branched cube B_T is the union of its proper folding faces and the interiors of its splitting faces.*

Proof. Suppose $x \in B_T$ does not lie in the interior of B_T, nor of any proper folding face of B_T. Then $x_\tau = 0$ for some $\tau \in \mathcal{T}_0$. Declare two directions d_1, d_2 equivalent if $x_{\{d_1,d_2\}} > 0$. It is clear this relation is reflexive and symmetric. It is transitive by Equation (2). Further, two directions can only be equivalent if they lie in the same gate. Thus, the equivalence classes partition the gates. The partition is nontrivial since $x_\tau = 0$ for some τ. By definition, $x_\tau = 0$ if and only if τ does not lie in an element of this partition. And $x_\tau < \ell(\tau)$ for each τ, or x would lie on a proper folding face of B_T. Thus x lies in the interior of the splitting face associated to this partition. □

3.5. **Union of the branched cubes.** The goal now is to show the branched cubes B_T cover $\mathcal{A}_\varphi$. We first show the following lemma, providing that, as a cube, the vertices of each $C_{T,\mathcal{T}}$ are fully preprincipal points.

Lemma 3.13. *The vertices of $C_{T,\mathcal{T}}$, i.e. the points where each x_τ is either 0 or $\ell(\tau)$, are fully preprincipal.*

Proof. Let T' be a vertex of $C_{T,\mathcal{T}}$. Then T' is obtained from T by folding some collection of illegal turns in T fully. Since each turn folded is folded fully (and by a Λ-sometry), any new vertex created by the fold has at least 3 gates. Further, since each turn folded was illegal, the number of gates at a vertex could not have decreased. Since T was fully preprincipal, each vertex of T has at least 3 gates, and so each vertex of T' has at least 3 gates, and thus T' is also fully preprincipal. □

Proposition 3.14. *Let $S \in \mathcal{A}_\varphi$. Then there is a fully preprincipal $T \in \mathcal{A}_\varphi$ with $S \in B_T$. In fact, we can choose T to have the same local decomposition as S, and so that S is not on a proper folding face of B_T.*

In particular, by Lemma 3.12, for each $\mathcal{G}$, we have that $\mathcal{G}\mathcal{A}_\varphi$ is the union of the interiors of the splitting faces of B_T, as T ranges over all fully preprincipal weak train tracks whose local decomposition is $\mathcal{G}$.

Proof. Let R be the rotationless power of φ and $\mathcal{G}$ the local decomposition of S. By Proposition 2.17, there is a fully preprincipal train track representative $g : \Gamma \to \Gamma$ for φ such that $T = \widetilde{\Gamma}$ has local decomposition $\mathcal{G}$.

Using Proposition 2.7, we can shift T so that the ε requirement in Proposition 2.19 is satisfied and then we can use Proposition 2.19 to know that S can be obtained from T by folding. However, we are not done yet because the fold path α from T to S may pass through multiple branched cubes.

We modify the fold path inductively as follows: If α does not meet a proper folding face of B_T, then α stays inside B_T. Otherwise, if α meets a proper folding face of B_T, then α must meet a folding face of some cube $C_{T,\mathcal{T}}$ at some point S'. We can choose a vertex $T' \neq T$ of $C_{T,\mathcal{T}}$ so that there is a fold path from T' to S'. We then modify α by replacing its initial segment from T to S' by the fold path from T' to S'.

Since there is a fold path from T to T' and from T' to S, the local decomposition of T' is also $\mathcal{G}$. For the same reason, $vol(T') < vol(T)$ and $vol(T') \geq vol(S)$. By Lemma 2.8, this shows the process terminates. □

3.6. **Intersections of the branched cubes.** The goal of this subsection is to analyze the intersection of the branched cubes B_T. The following construction will play a large role.

Construction 3.15 (Peel-off)**.** Let T be an element of $\mathcal{A}_\varphi$ and $v \in VT$. Suppose there is a $d_0 \in \mathcal{D}(v)$ and disjoint nonempty subsets $D_1, D_2 \subset \mathcal{D}(v)$ such that:

a. $\mathcal{D}(v) = \{d_0\} \cup D_1 \cup D_2$.
b. All Λ-leaves that pass through v by entering at d_0 exit through D_1 or D_2.
c. No Λ-leaves pass through v by entering at D_1 and exiting through D_2.
d. All Λ-leaves that pass through v by entering at D_1 and exiting through d_0 travel along the same segment I before meeting a vertex w with ≥ 3 gates.

See Figure 6, where the purple lines are Λ-leaves.

In this case, we say (D_1, D_2, I) is a **possible peel**. Note that, by irreducibility, (c)-(d) imply there are Λ-leaves entering at each of D_1 and D_1 before passing through I.

We can define a fully preprincipal train track T_I by detaching the directions in D_1 at v, attaching them to an endpoint of a copy of I, and attaching the other endpoint of the copy of I to w. We say that T_I is obtained from T by **peeling off** D_1 from D_2 at v along I. Note that there is a folding map $h_I : T_I \to T$ and the preimage of I is the union of two segments I_1 and I_2 that meet at w.

The special case when v is a two-gate vertex will be particularly important in the following. In this case, note that we decrease the valence at a two-gate vertex when splitting from T to T_I.

We say that a fully preprincipal train track T' is obtained by **completely peeling** T if there is a sequence of peel-offs $T = T_0 \to T_1 \to ... \to T_m = T'$ at two-gate vertices.

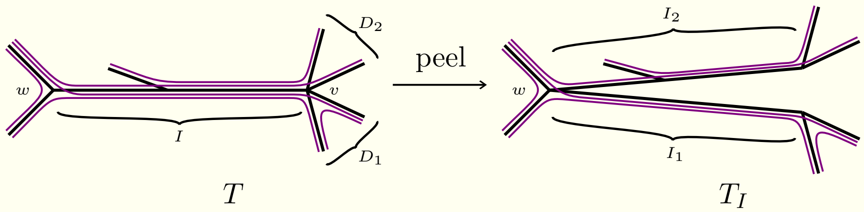

FIGURE 6. Left: A possible peel (D_1, D_2, I). Right: Peeling off D_1 from D_2 at v along I.

Example 3.16. Let T be a point in a branched cube $B_{T'}$. Then T lies in a cube $C_{T',\mathcal{T}}$. Let $\tau = \{d_1, d_2\} \in \mathcal{T}$ be a turn such that $x_\tau(T) > 0$. Let T_1 be the point with the same coordinates as T except that $x_\tau(T_1) = 0$.

Then T can be obtained by first folding T' to T_1 along the turns in $\mathcal{T}\backslash\{\tau\}$ then folding initial segments $I_1 \subset d_1$ and $I_2 \subset d_2$ of length x_τ. Let $h : T' \to T$ be that folding map and v_1, v_2 the terminal points of I_1, I_2 respectively. Let $I \subset T$ be the common image of I_1 and I_2. Let $v = h(v_1) = h(v_2)$ be the endpoint of I other than w, and let d_0 be the direction at v determined by I. Let D_1 be the h-image of the directions at v_1 other than I_1, and let $D_2 = \mathcal{D}(v)\backslash(\{d_0\} \cup D_1)$. Then (D_1, D_2, I) is a possible peel and $T_I = T_1$.

We have a criterion for determining whether we are in the situation of the example.

Lemma 3.17. *Suppose $T \in B_{T'}$. Recall from Definition 3.1 that there is a canonical fold $h : T' \to T$. A split train track T_I lies in $B_{T'}$ if and only if there exists no possible peel (D_1', D_2', I') where $h(I') \supset I$, and $h(D_1') \subset D_1$, and $h(D_2') \subset D_2$. If $T_I \in B_{T'}$, then T_I lies in a splitting face of $B_{T'}$.*

Proof. Suppose T_I lies in $B_{T'}$ and suppose there is a possible peel (D_1', D_2', I') as in the lemma. Let h' be the fold $h' : T' \to T_I$ and h_I the fold $h_I : T_I \to T$. Then the h'-image of the segment I' has to be mapped via h_I over I. But since there are Λ-leaves passing through D_1' and I', and h_I is a Λ-isometry, $h'(I')$ must pass through I_1. Similarly, $h'(I')$ must pass through I_2. This contradicts $h : T' \to T$ from being an isometry on I'. This argument shows that if a possible peel (D_1', D_2', I') exists, then T_I does not lie in $B_{T'}$.

Conversely, suppose $T_I \notin B_{T'}$. Then the preimage $h^{-1}(I)$ must be a union of edge segments, or since T' is fully preprincipal, there would be a ≥ 3-gate vertex in the interior of $h^{-1}(I)$, mapping to a ≥ 3-gate vertex in the interior of I. Consider the edge segments J in $h^{-1}(I)$, having endpoints v_J, w_J mapped to v, w respectively. For each J, consider the leaves that pass through v_J by entering through J. There are three cases (indicated in the upper left image of the figure):

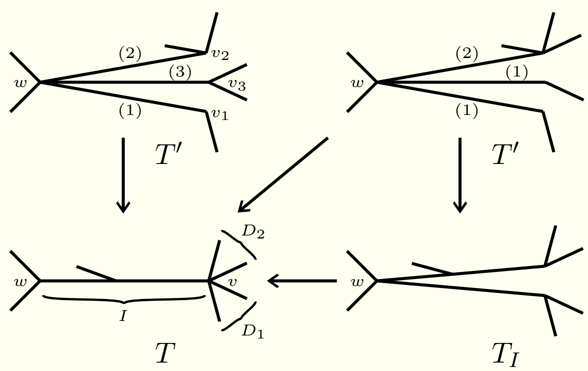

(1) All leaves exit v_J through a direction mapped by h to D_1.

(2) All leaves exit v_J through a direction direction mapped by h to D_2.

(3) Some leaves exit v_J through a direction mapped by h to D_1, while some leaves exit v_J through a direction mapped by h to D_2.

If J is of type (3), one can construct a possible peel by extending J until it meets a vertex. Thus we can assume each J to be of type (1) or (2) as in the upper right image of the figure.

Let $w_1', ..., w_m'$ be the vertices of T' mapped to w by h. Up to reindexing, suppose $w_1', ..., w_p'$ are the vertices that meet *both* segments J of type (1) *and* of type (2). For each $i = 1, ..., p$, we choose a turn τ_i with one direction being a segment of type (1) and the other direction being a segment of type (2).

The goal of the rest of the proof is to locate a cube of the form $C_{T',\mathcal{T}_1\cup\{\tau_1,...,\tau_p\}\cup\mathcal{T}_2}$ containing both T and T_I, hence contradicting the assumption that $T_I \notin B_{T'}$. Toward this end, choose a cube $C_{T',\mathcal{T}}$ containing T and let (x_τ) be the coordinates of T in that cube. Temporarily suppressing the subscript i for notational simplicity, suppose $\tau = \{d_1, d_2\}$, with d_1 determined by the segment of type (1) and d_2 determined by the segment of type (2). Let G be the gate containing d_1 and d_2. Let $G_1 = \{d \in G \mid x_{\{d,d_1\}} > x_{\{d_1,d_2\}}\}$ and $G_2 = \{d \in G \mid x_{\{d,d_1\}} \leq x_{\{d_1,d_2\}}\}$. Thus $G = G_1 \sqcup G_2$, with $d_1 \in G_1$ and $d_2 \in G_2$.

There must be a path in $\mathfrak{G}_\mathcal{T}$ from d_1 to d_2, or we would have $x_{\{d_1,d_2\}} = 0$ which is not the case since τ is folded. Recall from Equation (3) that $x_{\{d_1,d_2\}}$ is the minimum of x_τ as τ ranges over edges of this path. Let $\{d_1', d_2'\}$ be the edge such that $x_{\{d_1',d_2'\}} = x_{\{d_1,d_2\}}$. Then

$$\begin{aligned} x_{\{d_1',d_2'\}} &= x_{\{d_1,d_2\}} \\ &\leq \min\{x_{\{d_1,d_1'\}}, x_{\{d_1',d_2'\}}, x_{\{d_2',d_2\}}\} && \text{by Equation (3)} \\ &\leq \min\{x_{\{d_1,d_1'\}}, x_{\{d_2,d_2'\}}\} \end{aligned}$$

so by Lemma 3.11, T lies in $S_{\mathcal{T},\mathcal{T}'}$ where $\mathcal{T}' = (\mathcal{T} \setminus \{\{d_1', d_2'\}\}) \cup \{\{d_1, d_2\}\}$. (Here $\mathcal{T} \cap \mathcal{T}' = \mathcal{T} \setminus \{\{d_1', d_2'\}\}$, so in the notation of Equation (4), up to relabeling, we can take $d_1, d_1' \in C_1$ and $d_2, d_2' \in C_2$ with $[d_1', d_2'] \in \mathfrak{G}_\mathcal{T}$ and $[d_1, d_2] \in \mathfrak{G}_{\mathcal{T}'}$.) Thus, up to replacing $\mathcal{T}$ by $\mathcal{T}'$, we can assume $\tau = \{d_1, d_2\} \in \mathcal{T}$.

If $\mathfrak{G}_\mathcal{T}$ contains an edge $\{d, d_1\}$ where $d \in G_2$ then, since $x_{\{d,d_1\}} \leq x_{\{d_1,d_2\}}$, we have $T \in S_{\mathcal{T},\mathcal{T}'}$ for $\mathcal{T}' = (\mathcal{T} \setminus \{\{d, d_1\}\}) \cup \{\{d, d_2\}\}$ by Lemma 3.11. (Here $\mathcal{T} \cap \mathcal{T}' = \mathcal{T} \setminus \{\{d, d_1\}\}$, so in the notation of Equation (4), up to relabeling, we can take $d \in C_1$ and $d_1, d_2 \in C_2$ with $(d, d_1) \in \mathcal{T}$ and $(d, d_2) \in \mathcal{T}'$.) Thus, up to replacing $\mathcal{T}$ by $\mathcal{T}'$, we can assume no edges connect d_1 and G_2.

Similarly, if $\mathfrak{G}_\mathcal{T}$ contains an edge $\{d, d_2\}$ where $d \in G_1$, then since $x_{\{d,d_1\}} > x_{\{d_1,d_2\}}$, we have $x_{\{d,d_2\}} = x_{\{d_1,d_2\}}$ by Equation (2). Thus, by the symmetric argument, we can assume no edges connect d_2 and G_1.

The conclusion is that $T \in C_{T',\mathcal{T}_1 \cup \{\tau_1,\dots,\tau_p\} \cup \mathcal{T}_2}$ where $\mathcal{T}_j$ is an independent subset whose elements lie in G_j (for some i), for $j = 1, 2$. Within this cube, the split train track T_I is the point with the same coordinates as T except $x_{\tau_i} = 0$ for each $i = 1, \dots, p$. In particular T_I lies in a splitting face of $B_{T'}$. Contradiction. ☐

This criterion allows us to show that each T' has a unique complete peeling.

Lemma 3.18. *Each $T' \in \mathcal{A}_\varphi$ has a unique complete peel. More specifically, if T' lies in the interior of a splitting face of B_T then the complete peel of T' is T.*

Proof. By Proposition 3.14, we know that T' lies in the interior of a splitting face of some B_T. Then T' lies in the interior of some cube $C_{T,\mathcal{T}}$. If $x_\tau(T') > 0$, then we can run Example 3.16 to peel to the point T_1' with the same coordinates as T', except that $x_\tau(T_1') = 0$. Repeating this argument inductively, we get to a complete peel, which must then be the splitting vertex of B_T, namely T.

Conversely, by Lemma 3.17, each peel of T' at a 2-gate vertex lies in a splitting face of B_T, for otherwise there is a possible peel at a 2-gate vertex of T, but since T' is fully preprincipal these cannot exist. ☐

Lemma 3.19. *Suppose T' lies in the branched cube B_T. Then the complete peel of T' lies in B_T as well.*

Proof. We apply Lemma 3.17 repeatedly. At every peeling at a 2-gate vertex, we stay in $B_{T'}$, for otherwise there is a possible peel at a 2-gate vertex of T', but since T' is fully preprincipal these cannot exist. ☐

Proposition 3.20. *Distinct branched cubes B_{T_1} and B_{T_2} cannot intersect away from their folding faces.*

Proof. Assume otherwise, then there exists some T lying in the interior of a splitting face of B_{T_1} and that of B_{T_2}. Taking their complete peels, we get $T_1 = T_2$ by Lemma 3.18. ☐

Proposition 3.21. *Let T' be a fully preprincipal weak train track on a folding face F of B_T. Then $B_T \cap B_{T'}$ is a splitting face of $B_{T'}$ contained in F. See Figure 7.*

Proof. For each point $T'' \in B_T \cap B_{T'}$, there is a fold path from $T' \in F$ to T'', hence $B_T \cap B_{T'}$ is contained in F. It remains to show that this intersection is a splitting face of $B_{T'}$.

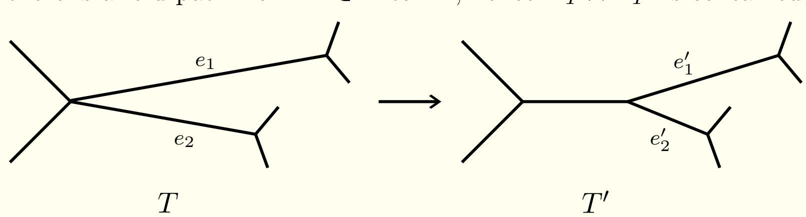

Let (x_τ') be the coordinates of T' in B_T and $\mathcal{T}$ the subset of $\mathcal{T}_0$ consisting of elements τ for which $x_\tau' < \ell(\tau)$. Each $\tau \in \mathcal{T}$ specifies an element τ° of $\mathcal{T}_0'$, namely the unfolded portion of the directions in τ. More precisely, suppose $\tau = \{d_1, d_2\}$, and suppose d_i is the germ of the edge e_i. Then since $x_\tau < \ell(\tau)$, there are some terminal segments $e_1' \subset e_1$ and $e_2' \subset e_2$ so that the images of e_1' and e_2' in T' determine two directions d_1' and d_2' at a vertex. Moreover, d_1' and d_2' lie in the same gate, since

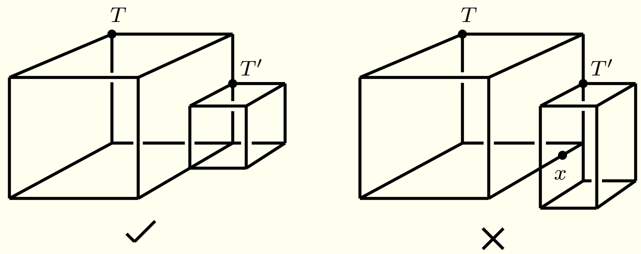

FIGURE 7. If T' be a fully preprincipal train track on a folding face of B_T, then $B_T \cap B_{T'}$ is a splitting face of $B_{T'}$, as in the left, and *not* on the right.

e'_1 and e'_2 would be folded if one folds τ completely. The turn τ° determined by τ is $\{d'_1, d'_2\}$. Furthermore, $\ell(\tau^\circ) \leq \ell(\tau) - x'_\tau$.

Let $\mathcal{T}^\circ \subset \mathcal{T}'_0$ denote the subset consisting of the τ_0 arising as such. We claim $\mathcal{T}^\circ$, when considered as a graph with vertex set G', is a disjoint union of complete subgraphs of $\mathcal{T}'_0$.

For this, it suffices to show that if $\{d'_0, d'_1\} \in \mathcal{T}^\circ$ and $\{d'_1, d'_2\} \in \mathcal{T}^\circ$, then $\{d'_0, d'_2\} \in \mathcal{T}^\circ$. Let $\{d_0, d_1\} \in \mathcal{T}_0$ be a turn determining $\{d'_0, d'_1\}$ and $\{d_2, d_3\} \in \mathcal{T}_0$ a turn determining $\{d'_1, d'_2\}$. The directions d_0, d_1, d_2, d_3 must lie in the same gate of T, since otherwise d_1 and d_2 cannot pass through the same direction d'_1. Thus $\{d_0, d_3\} \in \mathcal{T}_0$ and determines $\{d'_0, d'_2\} \in \mathcal{T}^\circ$, as desired. So the claim is proved.

By the claim, $\mathcal{T}^\circ$ determines a partition of G', namely where two directions d'_1, d'_2 lie in the same subset if and only if $\{d'_1, d'_2\} \in \mathcal{T}^\circ$. The intersection $B_T \cap B_{T'}$ is the splitting face associated to this partition. ☐

Proposition 3.22. *Each folding face of a branched cube B_T is a union of splitting faces of branched cubes.*

Proof. Fix a folding face F. By Proposition 3.14, each point $x \in F$ lies in the interior of a splitting face S of a branched cube $B_{T'}$. Lemma 3.19 implies that T', being the complete peel of x, must lie in B_T. Suppose T' lies in a folding face F' of B_T, and suppose F' is minimal with respect to this property. Since there is a fold path from T' to x, we have that F is a subfolding face of F'.

By Proposition 3.21, $B_T \cap B_{T'}$ is a splitting face S' of $B_{T'}$ contained in F'. Since $x \in B_T \cap B_{T'} = S'$, we have that S is a subsplitting face of S'.

Suppose F is a proper subset of F'. Then by the minimality of F', we have $T' \in F' \backslash F$, which implies that the interiors of the subsplitting faces of $S' = B_T \cap B_{T'}$ cannot meet F, contradicting the choice of S as a splitting face that contains x in its interior. Thus $F = F'$. In particular T' lies in F, which implies that $S' = B_T \cap B_{T'}$ lies in F as well. This argument shows that F is the union of such splitting faces S'. ☐

3.7. **The unique successor property.** In this subsection, we prove one final combinatorial property of the branched cubes.

Proposition 3.23. *For each point T in the axis bundle, the set of branched cubes $B_{T'}$ for which T lies in the interior of a folding face of $B_{T'}$, once partially ordered by inclusion, has a unique maximal element.*

Proof. Suppose otherwise that there are two maximal branched cubes $B_{T'_1}, B_{T'_2}$ for which T lies in the interior of a folding face of $B_{T'_i}$. Suppose T lies in a cube $C_{T'_2,\mathcal{T}} \subset B_{T'_2}$. Without loss of generality suppose $\mathcal{T}$ is a maximal independent set. As explained in Example 3.16, each coordinate x_τ gives a possible peel, and the act of peeling determines a split path. If all such split paths lie in $B_{T'_1}$, then the interior of $B_{T'_2}$ would meet $B_{T'_1}$. But then by Proposition 3.20, $B_{T'_2}$ lies on a folding face of $B_{T'_1}$, contradicting maximality of $B_{T'_2}$. Thus some split path, determined by some possible peel (D_1, D_2, I), does not lie in $B_{T'_1}$.

We apply Lemma 3.17 to $B_{T'_1}$ and the possible peel (D_1, D_2, I) to obtain a possible peel (D'_1, D'_2, I') of T'_1 which maps to (D_1, D_2, I) in the sense of the lemma. Let T'' be the train track obtained by splitting this possible peel. Since T'_1 lies in a folding face of $B_{T''}$, Proposition 3.21 states that $B_{T'_1} \cap B_{T''}$ is a splitting face of $B_{T'_1}$, so T cannot lie in the interior of a folding face of $B_{T'_1}$ unless $B_{T'_1} \cap B_{T''} = B_{T'_1}$, see Figure 8. This is equivalent to $B_{T'_1} \subset B_{T''}$, but this contradicts the maximality of $B_{T'_1}$. ☐

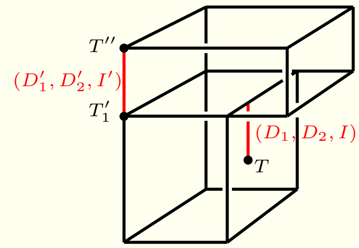

FIGURE 8. $B_{T_1'} \cap B_{T''}$ is a splitting face of $B_{T_1'}$, so T cannot lie in the interior of a folding face of $B_{T_1'}$ unless $B_{T_1'} \cap B_{T''} = B_{T_1'}$.

4. Cubist complexes

In this section we introduce the definition of a cubist complex. We choose to present the definition in a general abstract setting in anticipation of future applications. Because of this, we have to first introduce an abstract definition of a branched cube in §4.1. The definition of a cubist complex will then appear in §4.2.

Using the properties established in the previous section, we then show the axis bundle is a cubist complex in §4.3. Finally, we define the cardiovascular system of a cubist complex and study its properties in §4.4.

4.1. **Abstract definition of branched cubes.** We first present an abstract definition of branched cubes. The properties and terminology here are all motivated by the axis bundle setting of Section 3.

Let $\mathcal{G} = \{G_1, ..., G_m\}$ be a finite collection of finite sets. We denote by $\Lambda^2 G_i$ the set of unordered pairs of distinct elements of G_i. (This notation is motivated from exterior products of vector spaces.) Let $\mathcal{T}_0 = \Lambda^2 G_1 \sqcup ... \sqcup \Lambda^2 G_m$. As in Section 3, we can view each subset $\mathcal{T} \subset \mathcal{T}_0$ as a graph $\mathfrak{G}_{\mathcal{T}}$ with vertex set $\sqcup_{i=1}^m G_i$. We say that $\mathcal{T}$ is **independent** if $\mathfrak{G}_{\mathcal{T}}$ has no cycles.

Suppose we have a $\ell(\tau) \in \mathbb{R}_{>0}$ associated to each $\tau \in \mathcal{T}_0$ satisfying Equation (1). For each independent subset $\mathcal{T} \subset \mathcal{T}_0$, we define the **cube** $C_{\mathcal{T}}$ to be the metric space $\prod_{\tau \in \mathcal{T}} [0, \ell(\tau)]$. The coordinates of $C_{\mathcal{T}}$ are denoted $x_\tau^{(\mathcal{T})}$, for each $\tau \in \mathcal{T}$.

We extend the notation $x_{\{d_1,d_2\}}^{(\mathcal{T})}$ to make sense for each $\{d_1, d_2\} \in \mathcal{T}_0$ as follows: If there is a path in $\mathfrak{G}_{\mathcal{T}}$ connecting d_1 to d_2, consider the shortest such path (which is unique since $\mathcal{T}$ has no cycles) and define $x_{\{d_1,d_2\}}^{(\mathcal{T})}$ as the minimum over $x_\tau^{(\mathcal{T})}$ as τ ranges over the edges of this path. If $d_1 = d_2$, then $x_{\{d_1,d_2\}}^{(\mathcal{T})} = \infty$. Also, if $\tau \in \mathcal{T}$, then $x_\tau^{(\mathcal{T})}$ retains its original definition as a coordinate on $C_{\mathcal{T}}$. If there is no path in $\mathcal{T}$ connecting d_1 to d_2, define $x_{\{d_1,d_2\}}^{(\mathcal{T})} = 0$.

A **fold path** in $C_{\mathcal{T}}$ is an oriented path of the form $\alpha(s) = (\alpha_t(s))_{t \in \mathcal{T}}$ where each α_t is an nondecreasing function. A **subcube** of $C_{\mathcal{T}}$ is a subset of the form $\prod_{\tau \in \mathcal{T}} [x_\tau', x_\tau'']$ where $x_\tau' \leq x_\tau''$ for each τ. Each subcube can be viewed as an isometrically embedded copy of some other cube such that the image of each folding path is a folding path. As another subcube example, let $p_0 \in C_{\mathcal{T}}(\ell_t)$ and then the set of $p \in C_{\mathcal{T}}$ for which there is a fold path from p_0 to p is a subcube. We refer to this set as the **subcube determined by** p_0.

Finally, we define the slice $S_{\mathcal{T},\mathcal{T}'} \subset C_{\mathcal{T}}$ for each ordered pair of independent subsets $\mathcal{T}, \mathcal{T}' \subset \mathcal{T}_0$: Let $C_1, \dots, C_k$ be the components of $\mathfrak{G}_{\mathcal{T} \cap \mathcal{T}'}$. Define $S_{\mathcal{T},\mathcal{T}'}$ to be the subset of $C_{\mathcal{T}}$ consisting of all points satisfying the following inequalities for each $i, j = 1, \dots, k$:

$$\begin{cases} x_{\{d_i,d_j\}}^{(\mathcal{T})} \leq \min\{x_{\{d_i,d_i'\}}^{(\mathcal{T})}, x_{\{d_j,d_j'\}}^{(\mathcal{T})}\} & \text{if } \exists \ \{d_i, d_j\} \in \mathcal{T} \text{ and } \{d_i', d_j'\} \in \mathcal{T}' \text{ connecting } C_i \ \& \ C_j \\ x_{\{d_i,d_j\}}^{(\mathcal{T})} = 0 & \text{if } \exists \ \{d_i, d_j\} \in \mathcal{T} \text{ but no } \tau' \in \mathcal{T}' \text{ connecting } C_i \ \& \ C_j. \end{cases}$$

Definition 4.1. A **branched cube** $B_{\mathcal{G}}$ associated to $\mathcal{G}$ is a metric space that is the union of cubes $C_{\mathcal{T}}$, as $\mathcal{T}$ ranges over all independent subsets of $\mathcal{T}_0$, for some choice of $\ell(\tau) \in \mathbb{R}_{>0}$ associated to the elements $\tau \in \mathcal{T}_0$ satisfying Equation (1), so that for each pair of independent subsets $\mathcal{T}_1, \mathcal{T}_2 \subset \mathcal{T}_0$, we have

$$C_{\mathcal{T}_1} \cap C_{\mathcal{T}_2} = S_{\mathcal{T}_1,\mathcal{T}_2} = S_{\mathcal{T}_2,\mathcal{T}_1},$$

with the function $x^{(\mathcal{T}_1)}_{\{d_i,d_j\}}$ identified with the function $x^{(\mathcal{T}_2)}_{\{d'_i,d'_j\}}$ for each i,j for which there is a $\{d_i,d_j\}\in\mathcal{T}$ and a $\{d'_i,d'_j\}\in\mathcal{T}'$ connecting C_i and C_j. This ensures that the functions $x^{(\mathcal{T})}_\tau$, as $\mathcal{T}$ ranges over all independent subsets of $\mathcal{T}_0$, can be patched together into a function x_τ on $B_\mathcal{G}$.

The **splitting face** of $B_\mathcal{G}$ associated to partitions $G_i = G_{i,1}\sqcup\ldots\sqcup G_{i,m_i}$ is the subspace defined by $x_\tau = 0$ whenever $\tau\not\subset G_{i,m_i}$ for all i. In particular, the **splitting vertex** of B_G is the splitting face associated to the partition of each G_i into one-element sets, defined by $x_\tau = 0$ for all τ. The **folding face** of $B_\mathcal{G}$ associated to an independent subset $\mathcal{T}\subset\mathcal{T}_0$ is the subspace defined by $x_\tau = \ell(\tau)$ for all $\tau\in\mathcal{T}$.

A **fold path** in $B_\mathcal{G}$ is an oriented path that is locally a fold path in each cube that it lies in. Let p_0 be a point in $B_\mathcal{G}$. Consider the set of points p in $B_\mathcal{G}$ for which there is a folding path from p_0 to p. This set is also the union of subcubes determined by p_0 in each cube that contains p_0. Note that this set is in general not a branched cube. We refer to it as the **generalized branched cube** determined by p_0.

Lemma 4.2. *We have the following properties of branched cubes.*

(1) A finite product of branched cubes is a branched cube.
(2) Each splitting face of a branched cube is a branched cube.
(3) Each folding face of a branched cube is a branched cube.

Proof. For (1), we have $B_\mathcal{G}\times B_{\mathcal{G}'}\cong B_{\mathcal{G}\sqcup\mathcal{G}'}$. Because of this, it suffices to show (2) and (3) in the case when $\mathcal{G}$ just contains one finite set G.

For (2), consider the splitting face S associated to a partition $G = G_1\sqcup\ldots\sqcup G_m$. Then for every independent subset $\mathcal{T}\subset\mathcal{T}_0$, we have $C_\mathcal{T}\cap S = C_{\mathcal{T}_1\sqcup\ldots\sqcup\mathcal{T}_m} = C_{\mathcal{T}_1}\times\ldots\times C_{\mathcal{T}_m}$ where $\mathcal{T}_i$ is the subset of $\mathcal{T}$ consisting of elements lying within G_i, which is independent as a subset of $\Lambda^2 G_i$. Conversely, given independent subsets $\mathcal{T}_i\subset\Lambda^2 G_i$, we have that $\mathcal{T}_1\sqcup\ldots\sqcup\mathcal{T}_m$ is an independent subset of $\mathcal{T}_0$. Thus $S = \bigcup_{\mathcal{T}_1,\ldots,\mathcal{T}_m} C_{\mathcal{T}_1}\times\ldots\times C_{\mathcal{T}_m}\cong B_{G_1}\times\ldots\times B_{G_m}$.

For (3), we first show this for a folding face F associated to a single element $\{d_1,d_2\}$. Let $G_1 = \{d\in G\mid \ell(\{d,d_1\}) > \ell(\{d_1,d_2\})\}$ and $G_2 = \{d\in G\mid \ell(\{d,d_1\})\le\ell(\{d_1,d_2\})\}$. Thus $G = G_1\sqcup G_2$, with $d_1\in G_1$ and $d_2\in G_2$. The same argument as in Lemma 3.17 shows that $F\subset\bigcup_{\mathcal{T}_1,\mathcal{T}_2} C_{\mathcal{T}_1\cup\{d_1,d_2\}\cup\mathcal{T}_2}$ where the union ranges over all independent subsets $\mathcal{T}_1\subset\Lambda^2 G_1$ and $\mathcal{T}_2\subset\Lambda^2 G_2$.

Now $C_{\mathcal{T}_1\cup\{d_1,d_2\}\cup\mathcal{T}_2}\cap F\cong C_{\mathcal{T}_1}\times C_{\mathcal{T}_2}$. Thus $F = \bigcup_{\mathcal{T}_1,\mathcal{T}_2}(C_{\mathcal{T}_1\cup\{d_1,d_2\}\cup\mathcal{T}_2}\cap F)\cong\bigcup_{\mathcal{T}_1,\mathcal{T}_2}(C_{\mathcal{T}_1}\times C_{\mathcal{T}_2}) = B_{G_1}\times B_{G_2}$.

Moreover, under this isomorphism, the maximum value of x_τ for $\tau\in\Lambda^2 G_i$ equals $\ell(\tau)$. Hence, for folding faces of B_G associated to larger subsets, we can run this argument inductively. $\square$

4.2. **The definition of a cubist complex.**

Definition 4.3. A **cubist complex** is an ordered pair $(X,\mathcal{B})$ where X is a topological space and $\mathcal{B}$ is a collection of subspaces, each homeomorphic to a branched cube, and satisfying each of (a)-(d):

a. The space X is the disjoint union of the interiors of the elements of $\mathcal{B}$.
b. For each $B\in\mathcal{B}$, each splitting face of B is an element of $\mathcal{B}$, while each folding face of B is a union of elements of $\mathcal{B}$.
c. For each $B_1,B_2\in\mathcal{B}$, either $B_1\cap B_2=\varnothing$, or $B_1\cap B_2\in\mathcal{B}$ and is a sub-branched cube of B_1 and B_2.
d. For each 0-dimensional branched cube $\{v\}\in\mathcal{B}$, the set of all $B\in\mathcal{B}$ for which v lies in the interior of a folding face of B, once partially ordered by inclusion, has a unique maximal element $B(v)$.

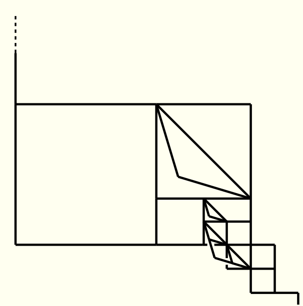

For a cubist complex $(X,\mathcal{B})$, X carries a natural piecewise-linear structure by declaring each branched cube in $\mathcal{B}$ to be PL-isomorphic to a branched cube. In the following, we will always implicitly endow X with this PL-structure. In particular, an **isomorphism** of cubist complexes is a PL isomorphism sending branched cubes to branched cubes. Also, in the following, whenever the collection $\mathcal{B}$ of branched cubes is clear, we simply refer to X as a cubist complex.

An example of a cubist complex is shown to the right. Note that the 'local dimension' of a cubist complex is allowed to jump. In the image, the local dimension at some parts is 2, while at other parts it is 1.

Recall from pg. 2 of the introduction the differences between cubist complexes and cube complexes. We make now item (3) from that comparison precise via the following lemma.

Lemma 4.4. *Let $(X, \mathcal{B})$ be a cubist complex. Let Y be a subspace of X that is a union of elements in $\mathcal{B}$ and which is PL isomorphic to a generalized branched cube itself. Then for each 0-dimensional branched cube $\{v\} \in \mathcal{B}$, the generalized sub-branched cube of Y determined by v is a union of elements of $\mathcal{B}$.*

Proof. Let N be the number of elements of $\mathcal{B}$ that intersect the interior of the sub-branched cube Z of Y determined by v. We proceed by induction on $(\dim Z, N)$ with the lexicographic order. If $\dim Z = 0$, then $Z = \{v\}$ is a 0-dimensional element of $\mathcal{B}$. We move on to the induction step.

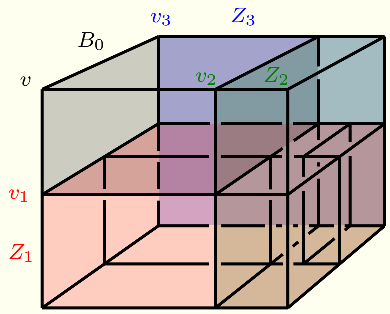

Consider the elements of $\mathcal{B}$ contained in Y that have splitting vertex at v, and let B_0 be the element among these with the highest dimension. Here we use the fact that Y is PL isomorphic to a generalized branched cube to make sure that B_0 is uniquely defined. Let $e_1, ..., e_n$ be the collection of 1-dimensional splitting edges of B_0. Let v_i be the endpoint of e_i that is not v, and let Z_i be the sub-branched cube of Y determined by v_i. For each i, if e_i is an entire splitting edge of Z, then Z_i has lower dimension than Z, otherwise Z_i intersects strictly less elements of $\mathcal{B}$ in their interior than Z does, since B_0 intersects the interior of Z but not that of Z_i. By our induction hypothesis, each Z_i is a union of elements of $\mathcal{B}$. This implies that $Z = B_0 \cup Z_1 \cup ... \cup Z_m$ is a union of elements of $\mathcal{B}$.

□

We refer to the union of the k-dimensional branched cubes in a cubist complex X as the **k-skeleton** of X. We caution that the k-skeleton of X as a cubist complex is different from the k-skeleton of X when considered as a cell complex. For example, if B is a branched cube as in Example 3.5, then the 1-cell that the branching happens along does not lie in the 1-skeleton of X as a cubist complex.

The 1-skeleton of X can be given the structure of a directed graph $\mathfrak{g}_X$ by orienting each 1-dimensional branched cube so that it is a folding path in each branched cube that contains it.

The next two lemmas concern the behaviour of $\mathfrak{g}_X$ in each branched cube.

Lemma 4.5. *Let $(X, \mathcal{B})$ be a cubist complex. Let Y be a subspace of X that is a union of elements in $\mathcal{B}$ and which is PL-isomorphic to a generalized branched cube itself. Let v be the splitting vertex of Y. Then for each vertex w of $\mathfrak{g}_X$ lying in Y, there is a directed edge path in $\mathfrak{g}_X$ from v to w.*

Proof. Let N be the number of elements of $\mathcal{B}$ that intersect the interior of Y. We prove the lemma by induction on $(\dim Y, N)$ with the lexicographic order. If $\dim Y = 0$ or 1, the lemma is clear.

We move on to the induction step. Define the sub-branched cube B_0, the vertices v_i, and the generalized sub-branched cubes Z_i as in Lemma 4.4. If $w = v$, then the lemma holds trivially. Otherwise, since $Y = B_0 \cup Z_1 \cup ... \cup Z_m$, w lies in some Z_i. By induction, there is a directed edge path from the splitting vertex v_i to w. Concatenating this with the edge from v to v_i, we are done. □

Furthermore, we claim that a directed edge path as in Lemma 4.5 is unique up to sweeping the path across 2-dimensional branched cubes.

More precisely, suppose B is a 2-dimensional branched cube. Suppose β and β' are directed edge paths in $\mathfrak{g}_X$ with a common initial vertex v_1 and a common terminal vertex v_2, and each of which is the concatenation of one 1-dimensional splitting face of B and one 1-dimensional folding face of B. Then for each edge path α_1 with terminal vertex v_1 and each edge path α_2 with initial vertex v_2, we say that the directed edge paths $\alpha_1 * \beta * \alpha_2$ and $\alpha_1 * \beta' * \alpha_2$ are related by **sweeping across** B.

Lemma 4.6. *Let $(X, \mathcal{B})$ be a cubist complex. Let Y be a subspace of X that is a union of elements in $\mathcal{B}$ and which is PL isomorphic to a generalized branched cube itself. Let v be the splitting vertex of Y, and let w be some vertex lying in Y. Suppose γ and γ' are directed edge paths in $\mathfrak{g}_X$ from v to w. Then γ and γ' are related by sweeping across finitely many 2-dimensional branched cubes.*

Proof. Define N as in Lemma 4.5. We prove the lemma by induction on $(\dim Y, N)$ with the lexicographic order. If $\dim Y = 0$ or 1, the lemma is clear.

We move on to the induction step. Define the sub-branched cube B_0, the edges e_i, the vertices v_i, and the generalized sub-branched cubes Z_i as in Lemma 4.4. If $w = v$, then the lemma holds trivially. Otherwise the initial edges of γ and γ' must each be one of the e_i.

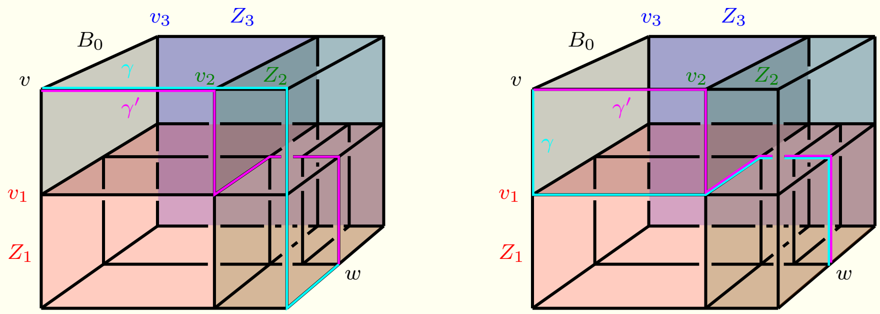

FIGURE 9. For Lemma 4.6, we first argue the case when the initial edges of γ and γ' agree (left), we then argue one case when the initial edges of γ and γ' differ (right).

We first assume the initial edges of γ and γ' are the same e_i (see Figure 9 left). In this case, $\gamma = e_i * \alpha$ and $\gamma' = e_i * \alpha'$ for some directed edge paths α and α' in Z_i from v_i to w. By induction, α and α' are related by sweeping across finitely many 2-dimensional branched cubes, so the same holds for γ and γ'.

It remains to prove the lemma for one choice of γ whose initial edge is e_i and for one choice of γ' whose initial edge is $e_{i'}$, for each pair (i, i'). Note that there is a vertex u on B_0 such that the sub-branched cube Q determined by u equals $Z_i \cap Z_{i'}$. More concretely, there is a 2-dimensional splitting face F of B_0 for which e_i and $e_{i'}$ are among the 1-dimensional splitting faces of F. The vertex u can be characterized as the intersection between the subcubes of F determined by v_i and $v_{i'}$.

By Lemma 4.5, there is a directed edge path α from u to w. Meanwhile, there is a directed edge path β from e_i to u and a directed edge path β' from $e_{i'}$ to u. See Figure 9 right. The edge paths $e_i * \beta * \alpha$ and $e_{i'} * \beta' * \alpha$ are related by sweeping across F. $\square$

Finally, we introduce the following condition, which will come into play in Section 4.4.

Definition 4.7. A **periodic cubist complex** is a connected cubist complex $(X, \mathcal{B})$ together with an isomorphism $\varphi : X \to X$ such that the $\mathbb{Z}$-action generated by φ is free and cocompact.

4.3. **The axis bundle is a cubist complex.** Let $\varphi \in \mathrm{Out}(F_r)$ be nongeometric fully irreducible. We briefly return to the axis bundle setting to show that it has a cubist complex structure.

Recall from Section 3 that for each fully preprincipal element $T \in \mathcal{A}_\varphi$, there is a branched cube B_T. Lemma 4.2 implies that each splitting face of B_T is a branched cube. We define $\mathcal{B}_\varphi$ to be the collection of all the B_T and their splitting faces, as T ranges over all fully preprincipal elements in $\mathcal{A}_\varphi$.

More generally, given any local decomposition $\mathcal{G}$ of $IW(\varphi)$, we define $\mathcal{G}\mathcal{B}_\varphi$ to be the collection of the B_T and their splitting faces, as T ranges over all fully preprincipal elements split at least as much as $\mathcal{G}$.

To distinguish the branched cubes B_T from a general element of $(\mathcal{G})\mathcal{B}_\varphi$, which can be a proper splitting face of B_T, we refer to the former as the **primary branched cubes** in the following proof.

Theorem 4.8. *Suppose that $r \geq 3$ and $\varphi \in \mathrm{Out}(F_r)$ is nongeometric fully irreducible. Then $(\mathcal{A}_\varphi, \mathcal{B}_\varphi)$ is a cubist complex. Furthermore, the action of φ on $\mathcal{A}_\varphi$ makes it into a periodic cubist complex.*

More generally, for each local decomposition $\mathcal{G}$, we have that $(\mathcal{G}\mathcal{A}_\varphi, \mathcal{G}\mathcal{B}_\varphi)$ is a cubist complex, and the action of φ on $\mathcal{G}\mathcal{A}_\varphi$ makes it into a periodic cubist complex.

Proof. Definition 4.3(a) follows from Proposition 3.14 and Proposition 3.20. Definition 4.3(b) follows from the definition of the branched cubes and Proposition 3.22. Definition 4.3(d) is a special case of Proposition 3.23.

We now show Definition 4.3(c). Since branched cubes have disjoint interiors, two branched cubes B_1 and B_2 can only possibly intersect along a union of branched cubes, each of them being (1) a splitting face of B_1 and a branched cube in a folding face of B_2, or (2) a branched cube in the interior of a folding face of B_1 and of B_2. If a branched cube of type (1) arises, Proposition 3.21 implies that $B_1 \cap B_2$ is a sub-branched cube. If a branched cube of type (2) arises, then by Proposition 3.23, the splitting vertex of B_1 and B_2 coincide, so B_1 and B_2 are splitting faces of some primary branched cube B_T, in which case they intersect in a splitting face of B_T. (In fact, this argument shows that there can never be branched cubes of type (2).)

It remains to show that $(\mathcal{G}\mathcal{A}_\varphi, \varphi)$ is a periodic cubist complex. Connectedness follows from Proposition 2.20. Freeness follows from Proposition 2.7. Finally, cocompactness follows from Lemma 2.8. $\square$

4.4. **Cardiovascular system.** Let $(X, \mathcal{B})$ be a cubist complex. We define a directed graph $\mathfrak{c}_X$ as follows:

- The vertex set of $\mathfrak{c}_X$ is the 0-skeleton of X.
- For each vertex v, let $B(v)$ be the branched cube that is maximal with respect to the property that v is contained in the interior of a folding face of $B(v)$, and let $S(v)$ be the splitting vertex of $B(v)$. By Definition 4.3(4), $S(v)$ is well-defined. If $S(v) \neq v$, then we add a directed edge from v to $S(v)$. (If $S(V) = v$ then we do not add any edges.)

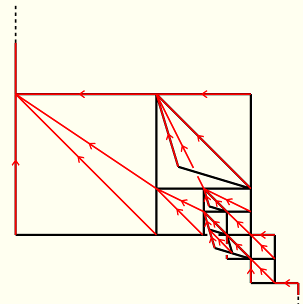

The graph $\mathfrak{c}_X$ is naturally seen as a subspace of X by placing the directed edges along straight lines. We refer to $\mathfrak{c}_X$ as the **cardiovascular system** of X. See the right-hand image for an example of a cardiovascular system.

A crucial property of the cardiovascular system is that each vertex has at most one outgoing edge, but possibly multiple incoming edges. In particular, each vertex v has a unique 'successor' $S(v)$. Iteratively determining edges as such, we obtain a directed edge path $\gamma_v = (S^i(v))_{i \geq 0}$, which is a ray if $S^{i+1}(v) \neq S^i(v)$ for each $i \geq 0$, and is a finite path otherwise. Conversely, each maximal directed edge path arises as such.

For the rest of this section, we restrict to the setting where (X, φ) is periodic (recall Definition 4.7). The aim is to deduce the cardiovascular system properties in this setting.

The cocompactness of the $\mathbb{Z}$-action implies each directed edge path $\gamma_v = (S^i(v))_{i \geq 0}$ is eventually periodic:

Proposition 4.9. *Let (X, φ) be a periodic cubist complex. For each vertex v, there exists $P, N \in \mathbb{Z}$, and an $i_0 \geq 0$, such that $S^{i+P}(v) = \varphi^N(S^i(v))$ for all $i \geq i_0$.*

Proof. Since $\langle \varphi \rangle$ acts cocompactly on X, the set V of $\langle \varphi \rangle$-orbits of vertices is finite. We define a map $\mathbb{Z}_{\geq 0} \to V$ by sending i to the orbit of $S^i(v)$. By the pigeonhole principle, there are integers $i_1 > i_2$ such that $S^{i_1}(v)$ and $S^{i_2}(v)$ lie in the same orbit, i.e. there is some $N \in \mathbb{Z}$ such that $S^{i_1}(v) = \varphi^N(S^{i_2}(v))$.

Meanwhile, since φ preserves the cubist complex structure of X, we have $\varphi(S(v)) = S(\varphi(v))$ for any vertex v. Applying this fact repeatedly, we deduce that, for each $i \geq i_2$,

$$S^{i+(i_1-i_2)}(v) = S^{i-i_2}(\varphi^N(S^{i_2}(v))) = \varphi^N(S^i(v)).$$

□

Next, we show that the connectedness of X implies that each directed path $\gamma_v = (S^i(v))_{i \geq 0}$ is actually a ray. To this end, we introduce a measure of distance between vertices of $\mathfrak{c}_X$.

Definition 4.10. Let α be an edge path in Γ. Note that α may not be a directed edge path, i.e. it can traverse some edges of Γ in the opposite direction of their prescribed orientations.

We define the **combinatorial length** of α to be

$$\min\{n \mid \alpha = \alpha_1 * ... * \alpha_n, \text{ for some monotone edge paths } \alpha_j \text{ each lying in a single branched cube }\}$$

where by a monotone edge path, we mean an edge path β where β or $-\beta$ is a directed edge path.

The **combinatorial distance** between two vertices v_0 and v_1 is the minimum combinatorial length of paths between them. It is straightforward to verify that the combinatorial distance is a metric.

Lemma 4.11. *Let v_0 and v_1 be two vertices of X. Suppose there is a directed edge path α from v_0 to v_1 that lies in one branched cube B. Then either $S(v_0) = v_1$, or there is a directed edge path β from v_1 to $S(v_0)$ that lies in one branched cube, so that the $\mathfrak{c}_X$-edge from v_0 to $S(v_0)$ is homotopic to $\alpha * \beta$.*

Proof. Without loss of generality, we can assume that B is the branched cube that is maximal with respect to the property that v_0 is contained in the interior of a folding face of B, so that $S(v_0)$ is the splitting vertex of B. Since α lies on B, we have that v_1 lies on B. By Lemma 4.5, there is a directed edge path β from v_1 to $S(v_0)$ that lies in B. Since $\alpha * \beta$ and the $\mathfrak{c}_X$-edge from v_0 to $S(v_0)$ are paths in B with the same initial and terminal vertices, and B is contractible, they are homotopic. □

Corollary 4.12. *Suppose v_0, $v_1 \in VX$. If the combinatorial distance between v_0 and v_1 is 1, then either $v_0 = S(v_1)$ or $v_1 = S(v_0)$ or $S(v_0) = S(v_1)$, or the combinatorial distance between $S(v_0)$ and $S(v_1)$ is ≤ 1.*

Proof. Up to switching v_0 and v_1, there is a directed edge path α_0 from v_0 to v_1 that lies in one branched cube. By Lemma 4.11, either $v_1 = S(v_0)$, or there is a directed edge path β_0 from v_1 to $S(v_0)$ that lies in one branched cube. By Lemma 4.11 again, either $S(v_0) = S(v_1)$, or there a directed edge path α_1 from $S(v_0)$ to $S(v_1)$ that lies in one branched cube. $\square$

Proposition 4.13. *Let (X, φ) be a periodic cubist complex. For each vertex v, the directed edge path $\gamma_v = (S^i(v))_{i\geq 0}$ is a ray, i.e. $S^{i+1}(v) \neq S^i(v)$ for each $i \geq 0$.*

Proof. Suppose otherwise that there is a vertex v_0 for which $S^{i+1}(v_0) = S^i(v_0)$ for some $i \geq i_0$. Let v_1 be a vertex that is of combinatorial distance 1 away from v_0. We claim that the directed edge path from v_1 stabilizes at the same point as that from v_0, i.e. $S^{i+1}(v_1) = S^i(v_1) = S^i(v_0)$ for $i \geq i_1$.

By Corollary 4.12, $S^{i_0}(v_0)$ and $S^{i_0}(v_1)$ have combinatorial distance at most 1. If $S^{i_0}(v_0) = S^{i_0}(v_1)$ then our claim is clear. Otherwise there is a directed edge path α either from $S^{i_0}(v_0)$ to $S^{i_0}(v_1)$ or from $S^{i_0}(v_1)$ to $S^{i_0}(v_0)$ that lies in one branched cube. The former cannot be true since $S^{i_0}(v_0) = S^{i_0+1}(v_0)$. But then, applying Lemma 4.11, there is a directed edge path β from $S^{i_0}(v_0)$ to $S^{i_0+1}(v_1)$, thus $S^{i_0+1}(v_0) = S^{i_0}(v_0) = S^{i_0+1}(v_1)$ since again $S^{i_0}(v_0) = S^{i_0+1}(v_0)$.

Now let α be an edge path between v_0 and $\varphi(v_0)$. Let v_0, v_1, ..., $v_m = \varphi(v_0)$ be the sequence of vertices on α. For each i, the combinatorial distance between v_{i-1} and v_i is 1, hence by applying our claim in the first paragraph repeatedly, we have $S^i(\varphi(v_0)) = S^i(v_0)$ for all large i. Thus $S^{i_0}(v_0) = S^i(v_0) = S^i(\varphi(v_0)) = \varphi(S^i(v_0)) = \varphi(S^{i_0}(v_0))$, that is, $S^{i_0}(v_0)$ is a fixed point of φ. This contradicts freeness of φ. $\square$

An **artery** of the cardiovascular system is a periodic directed edge path, i.e. directed edge path $A = (v_i)_{i\in\mathbb{Z}}$ for which there exist P, N such that $v_{i+P} = \varphi^N(v_i)$ for all i. Then P is the **period** of A and N the **order**.

Proposition 4.13 implies that $N \neq 0$, since if $N = 0$, we have that $S^P(v_i) = v_{i+P} = v_i$, and the directed edge path starting at v_i is finite. Thus we can define the **average period** of A to be $\frac{P}{N}$.

Proposition 4.14. *Let (X, φ) be a periodic cubist complex. There is at least one, and at most finitely many arteries in the cardiovascular system of X.*

Proof. To show that there is at least one artery, take some vertex v and apply Proposition 4.9 to the directed edge ray $r = (S^i(v))_{i\geq 0}$ to get values of P, N, i_0 so that $S^{i+P}(v) = \varphi^N(S^i(v))$ for all $i \geq i_0$. For each i, we define a vertex v_i by picking k large enough so that $i + kP \geq i_0$, and setting $v_i = \varphi^{-kN}(S^{i+kP}(v))$. Note that if $k_1 < k_2$ are integers such that $i + k_1P, i + k_2P \geq i_0$, then

$$\begin{aligned}\varphi^{-k_2N}(S^{i+k_2P}(v)) = \varphi^{-k_2N}(S^{i+(k_2-1)P+P}(v)) &= \varphi^{-k_2N}\varphi^N(S^{i+(k_2-1)P}(v)) \\ &= \varphi^{-(k_2-1)N}(S^{i+(k_2-1)P}(v)) = ... = \varphi^{-k_1N}(S^{i+k_1P}(v))\end{aligned}$$

so v_i is well-defined. The following similar computation shows that $(v_i)_{i\in\mathbb{Z}}$ is an artery:

$$v_{i+P} = \varphi^{-kN}(S^{i+P+kP}(v)) = \varphi^{-(k-1)N}(S^{i+kP}(v)) = \varphi^N(v_i).$$

To show finiteness, we first claim that there is a uniform bound on the order of the arteries in X: Suppose two arteries A and A' pass through the same $\langle\varphi\rangle$-orbit of vertices, say $v \in A$ and $v' \in A'$ where $v' = \varphi^q(v)$. Then A and A' share the same order. Indeed, if $S^{i+P}(v) = \varphi^N(S^i(v))$ then

$$S^{i+P}(v') = S^{i+P}(\varphi^q(v)) = \varphi^q(S^{i+P}(v)) = \varphi^{N+q}(S^i(v)) = \varphi^N(S^i(\varphi^q(v))) = \varphi^N(S^i(v')).$$

Thus the order of A' divides that of A. Symmetrically, the order of A divides that of A', so they must coincide. Our claim now follows from the fact that there are finitely many $\langle\varphi\rangle$-orbits of vertices.

Let N_0 be the lowest common multiple of the orders of all arteries. Let V_0 be the finite set of $\langle\varphi^{N_0}\rangle$-orbits of vertices. If there are more than $|V_0|$ arteries, then two of them must pass through the same $\langle\varphi^{N_0}\rangle$-orbit of vertices, but then they would actually share some vertex, which would imply the two arteries coincide. $\square$

By Lemma 4.5, each artery A can be homotoped into the reverse of a directed edge line L of the 1-skeleton Γ. Indeed, one can replace each edge $v \to S(v)$ by the reverse of a homotopic directed edge path within the same cube. In this context we say that L is a **simple factorization** of A.

Lemma 4.6 implies that any two simple factorizations of a common artery A are related by sweeping across 2-dimensional branched cubes. The last goal of this section is show that any two simple factorizations of any two arteries are also related in this way.

Recall the notion of the combinatorial distance between two vertices. We define the **combinatorial distance** between two arteries to be the minimum combinatorial distance between their vertices.

Lemma 4.15. *Let (X, φ) be a periodic cubist complex. Suppose A and A' are arteries combinatorial distance 1 apart. Then A and A' have a common average period and they admit a common simple factorization.*

Proof. By definition, up to switching A and A', there exist vertices $v \in A$ and $v' \in A'$, and a directed edge path α_0 from v to v' that lies in one branched cube. By Lemma 4.11, either $S(v) = v'$, or there is a directed edge path β_0 from v' to $S(v)$ so that the $\mathfrak{c}_X$-edge from v to $S(v)$ is homotopic to $\alpha_0 * \beta_0$. The former case cannot happen here or we would have $A = A'$.

We then apply Lemma 4.11 to β_0 to obtain a directed edge path α_1 from $S(v)$ to $S(v')$ so that the $\mathfrak{c}_X$-edge from v' to $S(v')$ is homotopic to $\beta_0 * \alpha_1$.

Repeating this argument, we have edge paths α_k from $S^k(v)$ to $S^k(v')$ and β_k from $S^k(v')$ to $S^{k+1}(v)$ such that the $\mathfrak{c}_X$-edge from $S^k(v)$ to $S^{k+1}(v)$ is homotopic to $\alpha_k * \beta_k$ and the $\mathfrak{c}_X$-edge from $S^k(v')$ to $S^{k+1}(v')$ is homotopic to $\beta_k * \alpha_{k+1}$.

Let P and N be the period and order of A, and let P' and N' be the period and order of A'. For each q, we have $\varphi^{-qNP'}(\alpha_{qPP'})$ is an edge path from $\varphi^{-qNP'}(S^{qPP'}(v)) = v$ to $\varphi^{-qNP'}(S^{qPP'}(v')) = \varphi^{q(PN'-NP')}(v')$ that lies on a branched cube. Since there are only finitely many such edge paths, for some $0 < q_1 < q_2$, we have $\varphi^{-q_1NP'}(\alpha_{q_1PP'}) = \varphi^{-q_2NP'}(\alpha_{q_2PP'})$. In particular, $\varphi^{q_1(PN'-NP')}(v') = \varphi^{q_2(PN'-NP')}(v')$, thus $\varphi^{(q_1-q_2)(PN'-NP')}(v') = v'$. Since φ cannot have fixed points, we have $PN' - NP' = 0$, thus $\frac{P}{N} = \frac{P'}{N'}$.

Finally, the directed edge path $*_{j=-\infty}^{\infty}\varphi^{j(q_2-q_1)NP'}(*_{k=q_1PP'}^{q_2PP'}(\alpha_k * \beta_k))$ is a common simple factorization for A and A'. $\square$

Proposition 4.16. *Let (X, φ) be a periodic cubist complex. Suppose A and A' are arteries. Then A and A' have a common average period and their simple factorizations are related by sweeping across 2-dimensional branched cubes.*

Proof. We first claim that there exist arteries $A_0 = A, A_1, ..., A_n = A'$ such that A_{i-1} and A_i are of combinatorial distance 1 apart for each $1 \leq i \leq n$.

To see this, let α be an edge path between A and A'. Let v_0,...,v_m be the sequence of vertices on α, where $v_0 \in A$ and $v_m \in A'$. For each $i = 0, ..., m$, let r_i be the $\mathfrak{c}_X$-ray starting at v_i. By Proposition 4.9, each r_i eventually converges into an artery A_i.

For each i, the combinatorial distance between v_{i-1} and v_i is 1, hence by Corollary 4.12, either $A_{i-1} = A_i$ or the combinatorial distance between A_{i-1} and A_i is 1. Hence it suffices to discard any repeated arteries.

Now the proposition follows from Lemma 4.6 and Lemma 4.15. $\square$

5. Examples

In this section, we go through some examples of cubist decompositions of axis bundles.

Example 5.1 (Lone axes)**.** In [MP16], it is shown that the axis bundle of an ageometric, fully irreducible $\varphi \in \mathrm{Out}(F_r)$ is a line if and only if both the index satisfies $i(\varphi) = \frac{3}{2} - r$ and no component of the ideal Whitehead graph $IW(\varphi)$ has a cut vertex. See [GP23] and [Pfa24] for concrete examples.

In this case, the cubist decomposition of the axis bundle $\mathcal{A}_\varphi$ can only be a union of 1-cubes and 0-cubes. The cardiovascular system coincides with the axis bundle (with the additional data of being oriented toward the splitting direction). Thus there is exactly one artery and it coincides with the axis bundle.

Example 5.2 (Multiple arteries)**.** Let $\varphi \in \mathrm{Out}(F_3)$ be defined by $\varphi(a) = cbca$, and $\varphi(b) = cbc$, and $\varphi(c) = ac$. This is the same outer automorphism considered in [Pfa24, Example 9.1].

This outer automorphism φ is represented by the fully preprincipal train track map g on the three-petaled rose, as depicted to the right. One can check via a straightforward computation that it has no PNPs. Using the criterion in [Kap14], one can verify that φ is ageometric fully irreducible.

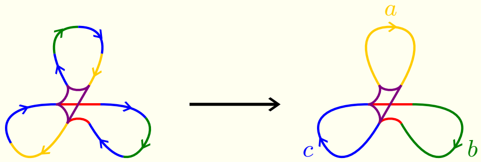

We compute the cubist decomposition of $\mathcal{A}_\varphi$ by the following algorithm:

1. Take $T \in \mathcal{A}_\varphi$ to be the fully preprincipal element that is the universal cover of the domain of the train

track map above. Compute the branched cube $B_T \subset \mathcal{A}_\varphi$. Let $\mathcal{A}$ be the union of the φ-translates of B_T.
2. For each fully preprincipal $T \in \mathcal{A} \subset \mathcal{A}_\varphi$, compute $B_T \subset \mathcal{A}_\varphi$. (There are infinitely many such T but only finitely many φ-orbits, so this is a finite time process.)

a. If any such B_T do not lie in $\mathcal{A}$, we add it to $\mathcal{A}$ and repeat this step.
b. If all such B_T lie in $\mathcal{A}$, the algorithm terminates.

The output $\mathcal{A}$ of the algorithm is the axis bundle $\mathcal{A}_\varphi$: Suppose, for the sake of contradiction, there was a $T' \in \mathcal{A}_\varphi$ with $T' \notin \mathcal{A}$. Then, as in the proof of Proposition 3.14, there is a fold path α from a vertex T of $\mathcal{A}$ to T'. Every time α exits a branched cube in $\mathcal{A}$ it enters another one (for otherwise we could have added that branched cube to $\mathcal{A}$) so at the end of the path, T' lies in $\mathcal{A}$.

Figure 10 shows the result of our computation of $\mathcal{A}_\varphi$. We have drawn one fundamental domain of $\mathcal{A}_\varphi$ under the φ-action; the 1-cubes on the right are sent to the 1-cubes on the left by φ as indicated.

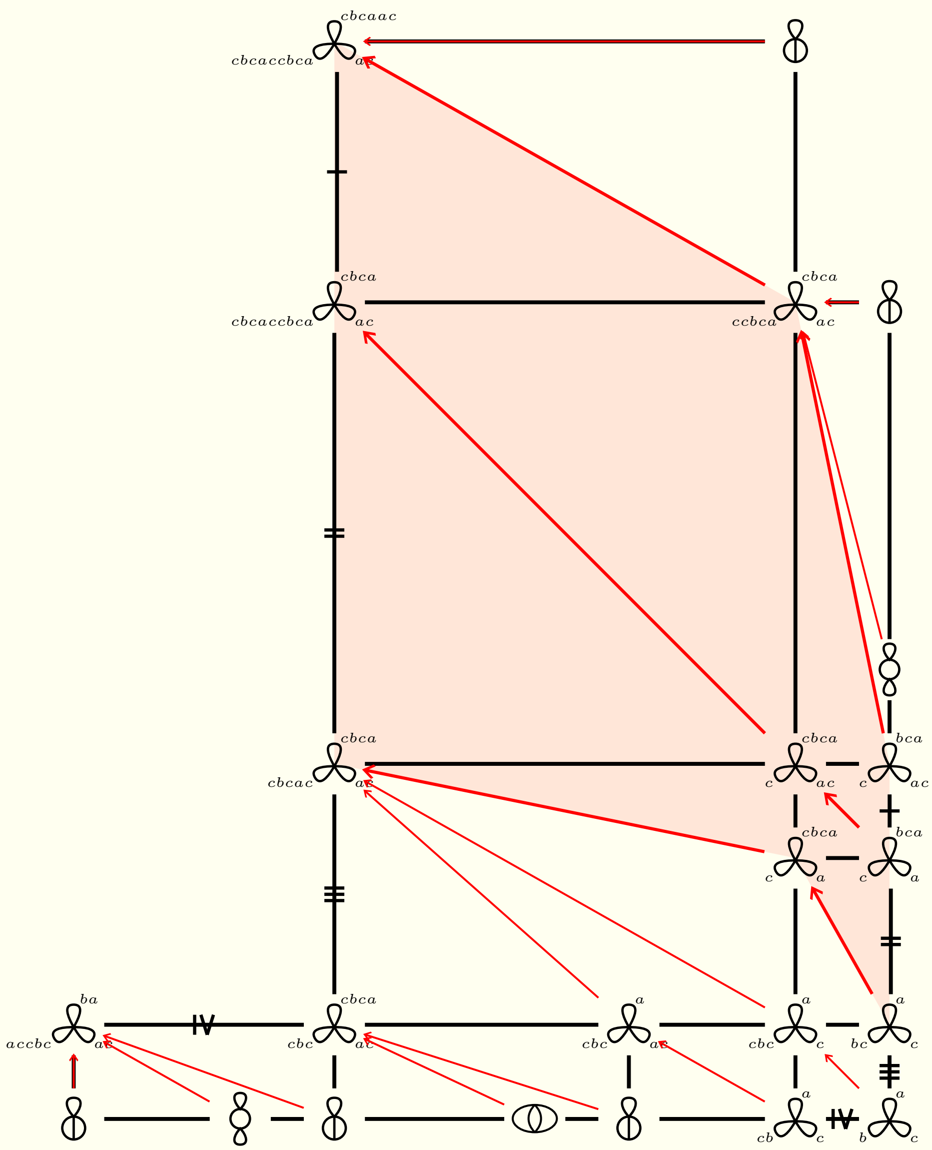

FIGURE 10. The axis bundle of the outer automorphism in Example 5.2.

The length of each edge in the figure is the fold length. These are computed as follows: The normalized eigenvector is $[0.45, 0.29, 0.26]$, so we will start with $\ell(a) \approx 0.45$, and $\ell(b) \approx 0.29$, and $\ell(c) \approx 0.26$. Consider the vertical fold path from the node meeting the edges labeled III and IV on the left. This folds the turn $\{\bar{b}, \bar{c}\}$, so the length folded can be computed to be $\approx 0.29(\frac{0.26}{0.26+0.29+0.26})$, or equivalently $\approx 0.26(\frac{0.26}{0.45+0.26})$. This is ≈ 0.095. This says the lengths on the graph below are $\ell(a') \approx 0.45$, and $\ell(b') \approx 0.185$, and $\ell(c') \approx 0.155$, and $\ell(d') \approx 0.095$, where d' is the edge created by the fold. One then performs similar computations iteratively.

The cardiovascular system is drawn in red, with the arteries bold. In this example there are three arteries, demonstrating that axis bundles do not necessarily have a unique artery.

Example 5.3 (Unique artery, branching)**.** Let $\varphi \in \mathrm{Out}(F_3)$ be defined by $\varphi(a) = a\overline{ca}$, and $\varphi(b) = ac$, and $\varphi(c) = \bar{b}\bar{c}$. This outer automorphism φ is represented by the train track map on the three-petaled rose depicted to the right. A straightforward computation shows this train track map has no PNPs and then the criterion in [Kap14] provides that φ is ageometric fully irreducible. We show the cubist decomposition of $\mathcal{A}_\varphi$ in Figure 11. Here, the 1-cube on the right is sent to the 1-cube on the *top* by φ as indicated.

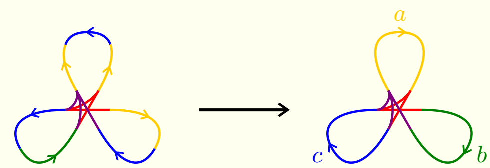

This example exhibits some phenomena that are different from Example 5.2:

(1) There is branching in the axis bundle $\mathcal{A}_\varphi$. More precisely, there are two branched 2-cubes in $\mathcal{A}_\varphi$, modulo the action of φ.
(2) There is a (branched) 2-cube with 4 vertices on one of its folding faces.
(3) There is a unique artery.
(4) The artery meets the boundary of $\mathcal{A}_\varphi$.

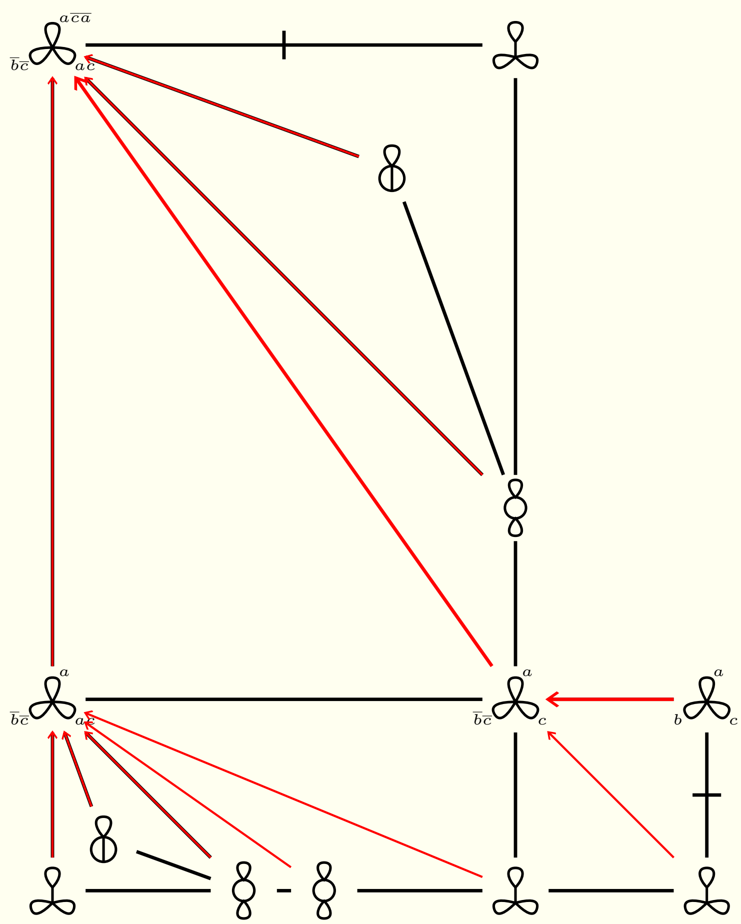

Figure 11. The axis bundle of the outer automorphism in Example 5.3.

Example 5.4 (Full & stable axis bundles)**.** Let $\varphi \in \mathrm{Out}(F_3)$ be defined by $\varphi(a) = ac$, and $\varphi(b) = cbc$, and $\varphi(c) = cbca$. This outer automorphism φ is represented by the train track map on the three-petaled rose in the upper row of the image to the right. A straightforward computation shows that this train track map has one PNP, namely $a * \overline{b_2}$, where b_2 is a suitable suffix of b.

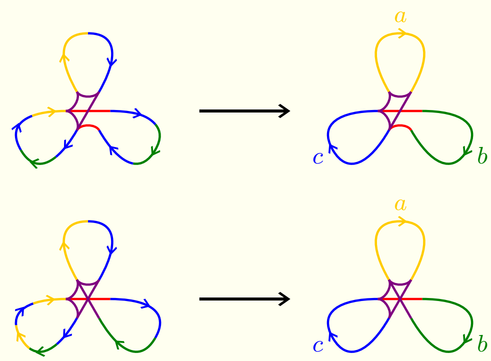

Collapsing this PNP, one obtains the train track map in the bottom row of the image. Note, the stable Whitehead graph of this map is that of the previous one with an additional edge, as one should expect. A straightforward computation shows this train track map has no PNPs and then the criterion in [Kap14] provides that φ is ageometric fully irreducible.

We show the cubist decomposition of $\mathcal{A}_\varphi$ in Figure 12. The 1-cubes on the right are sent to the 1-cubes on the left by φ as indicated. The cardiovascular system of $\mathcal{A}_\varphi$ is drawn in red, with the arteries bold.

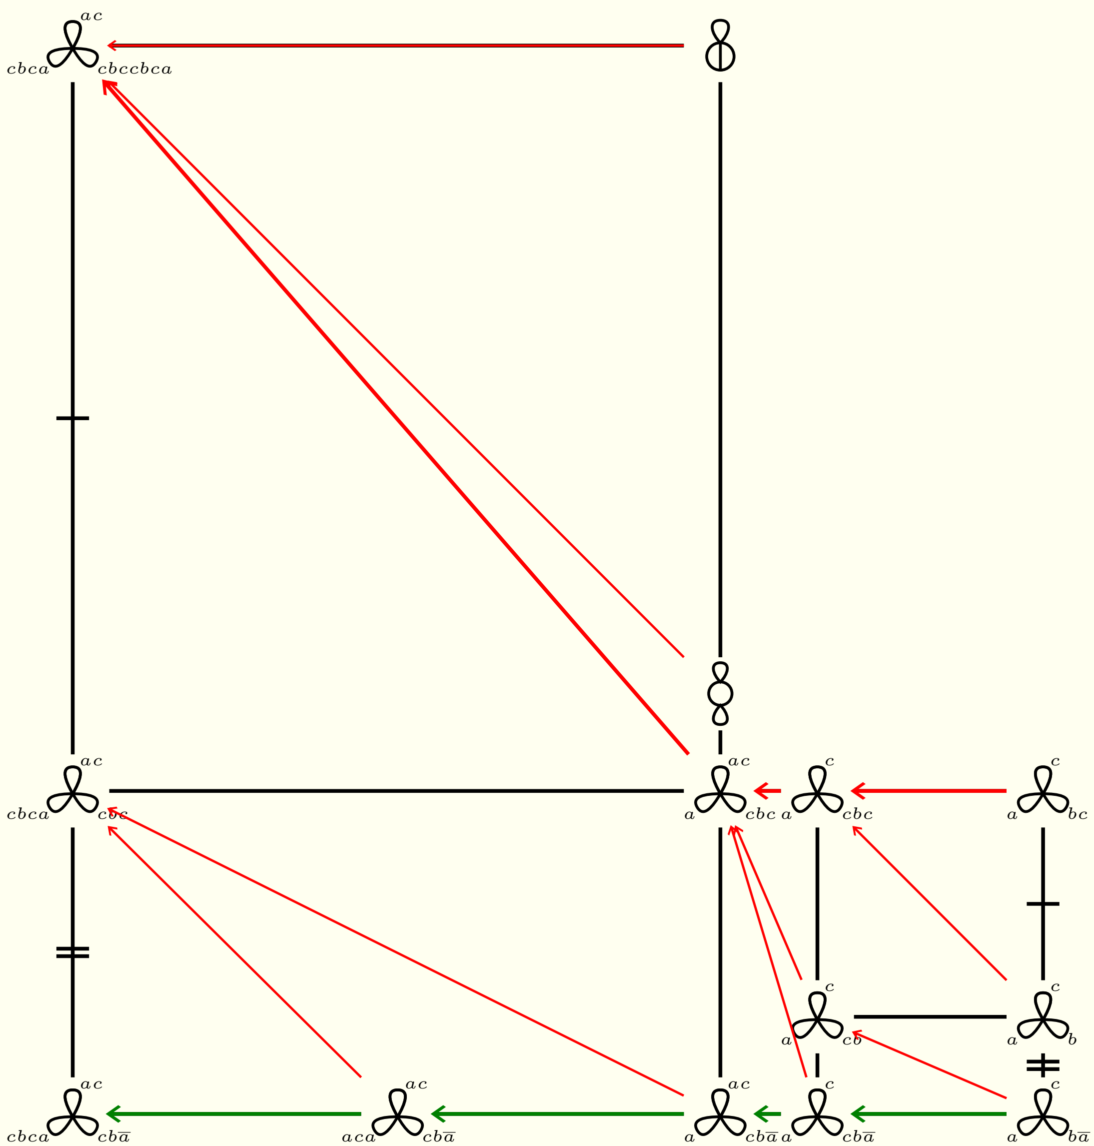

Figure 12. The axis bundle and stable axis bundle in Example 5.4.

As reasoned above, the ideal Whitehead graph $IW(\varphi)$ has a cut vertex. Hence, the stable axis bundle $\mathcal{SA}_\varphi$ is a proper subset of $\mathcal{A}_\varphi$. In the figure, $\mathcal{SA}_\varphi$ is the line on the bottom. The cardiovascular system of $\mathcal{SA}_\varphi$ is drawn in green, with the arteries highlighted. This example demonstrates that the arteries of the stable axis bundle need not agree with the arteries of the full axis bundle.

References

[Ago11] I. Agol. Ideal triangulations of pseudo-Anosov mapping tori. In *Topology and geometry in dimension three*, volume 560 of *Contemp. Math.*, pages 1–17. Amer. Math. Soc., Providence, RI, 2011.

[AKKP19] Y. Algom-Kfir, I. Kapovich, and C. Pfaff. Stable strata of geodesics in outer space. *International Mathematics Research Notices*, 2019(14):4549–4578, 2019.

[BF94] M. Bestvina and M. Feighn. Outer limits. *preprint*, pages 1–19, 1994.

[BFH97] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geometric and Functional Analysis*, 7(2):215–244, 1997.

[BFH00] M. Bestvina, M. Feighn, and M. Handel. The Tits Alternative for Out (F_n) I: Dynamics of exponentially-growing automorphisms. *Annals of Mathematics-Second Series*, 151(2):517–624, 2000.

[BH92] M. Bestvina and M. Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, 1992.

[BV06] M. Bridson and K. Vogtmann. Automorphism groups of free groups, surface groups and free abelian groups. In *Proceedings of Symposia in Pure Mathematics*, volume 74, page 301. Providence, RI; American Mathematical Society; 1998, 2006.

[CL95] M. Cohen and M. Lustig. Very small group actions on r-trees and dehn twist automorphisms. *Topology*, 34(3):575–617, 1995.

[CV86] M. Culler and K. Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones mathematicae*, 84(1):91–119, 1986.

[FH11] M. Feighn and M. Handel. The recognition theorem for $\mathrm{Out}(F_n)$. *Groups Geom. Dyn.*, 5(1):39–106, 2011.

[GP23] D. Gagnier and C. Pfaff. Taking the high-edge route of rank-3 outer space. *International Journal of Algebra and Computation*, 33(08):1659–1685, 2023.

[Ham09] U. Hamenstädt. Geometry of the mapping class groups. I. Boundary amenability. *Invent. Math.*, 175(3):545–609, 2009.

[HM11] M. Handel and L. Mosher. Axes in outer space. *Mem. Amer. Math. Soc.*, 213(1004):vi+104, 2011.

[Kap14] I. Kapovich. Algorithmic detectability of iwip automorphisms. *Bulletin of the London Mathematical Society*, 46(2):279–290, 2014.

[LL03] G. Levitt and M. Lustig. Irreducible automorphisms of F_n have north–south dynamics on compactified outer space. *Journal of the Institute of Mathematics of Jussieu*, 2(01):59–72, 2003.

[Mos03] L. Mosher. Train track expansions of measured foliations. *Preprint*, 2003.

[MP16] L. Mosher and C. Pfaff. Lone axes in outer space. *Algebr. Geom. Topol.*, 16(6):3385–3418, 2016.

[Pfa12] C. Pfaff. *Constructing and Classifying Fully Irreducible Outer Automorphisms of Free Groups*. PhD thesis, Rutgers University, 2012. PhD Thesis,

[Pfa24] C. Pfaff. $\mathrm{Out}(F_r)$ train track automata I: Proper full fold decompositions. *arXiv preprint arXiv:2409.05599*, 2024.

Institute for Advanced Study, Queen's University
Email address: catherine.pfaff@gmail.com
URL: https://mast.queensu.ca/~cpfaff/

Université du Québec à Montréal
Email address: tsang.chi_cheuk@uqam.ca
URL: https://sites.google.com/view/chicheuktsang