OUTER SPACE FOR RAAGS

COREY BREGMAN, RUTH CHARNEY, and KAREN VOGTMANN

Abstract
For any right-angled Artin group A_Γ, we construct a finite-dimensional space $\mathcal{O}_\Gamma$ on which the group $\mathrm{Out}(A_\Gamma)$ *of outer automorphisms of A_Γ acts with finite point stabilizers. We prove that $\mathcal{O}_\Gamma$ is contractible, so that the quotient is a rational classifying space for* $\mathrm{Out}(A_\Gamma)$. *The space $\mathcal{O}_\Gamma$ blends features of the symmetric space of lattices in $\mathbb{R}^n$ with those of* outer space *for the free group F_n. Points in $\mathcal{O}_\Gamma$ are locally CAT(0) metric spaces that are homeomorphic (but not isometric) to certain locally CAT(0) cube complexes, marked by an isomorphism of their fundamental group with A_Γ.*

Contents

1. Introduction . . . 1033
2. Preliminaries . . . 1038
3. Combinatorial and metric structure of blowups . . . 1048
4. Hyperplanes in Γ-complexes . . . 1061
5. Parallelotope structures on blowups . . . 1069
6. The space of skewed Γ-complexes with untwisted markings . . . 1076
7. Outer space $\mathcal{O}_\Gamma$ . . . 1079
References . . . 1106

1. Introduction

A lattice Λ in a semisimple Lie group G acts discretely on the symmetric space G/K, and a very well-developed theory shows that the algebraic structure of Λ is intimately connected to the geometric structure of G/K. The study of surface mapping class groups by Thurston, Harvey, and Harer among others borrowed ideas from this classical subject, using Teichmüller space as a substitute for the symmetric space, and this point of view proved to be extremely fruitful. An analogue of symmetric spaces and Teichmüller spaces called Culler–Vogtmann's *outer space* was later produced for the purpose of studying the group of outer automorphisms of a free group (see [18]). The

DUKE MATHEMATICAL JOURNAL
Vol. 172, No. 6, © 2023 DOI 10.1215/00127094-2023-0007
Received 18 September 2020. Revision received 19 February 2022.
First published online 3 April 2023.
2020 *Mathematics Subject Classification*. Primary 20F28; Secondary 20F36, 20F65.

study of this group, space, and action has frequently been guided by Thurston's ideas, but there are some respects in which $\mathrm{Out}(F_n)$ more closely resembles a lattice than a mapping class group. For example, mapping class groups are automatic (see [29]), while for $n \geq 3$, $\mathrm{Out}(F_n)$ (see [9]) and $\mathrm{GL}(n, \mathbb{Z})$ (see [23]) are not.

In this article, we study outer automorphism groups of right-angled Artin groups, a class which includes both $\mathrm{Out}(F_n)$ and the most basic lattice, $\mathrm{GL}(n, \mathbb{Z}) = \mathrm{Out}(\mathbb{Z}^n)$. Recall that a *right-angled Artin group (RAAG)* is defined by a presentation with a finite set of generators together with relations specifying that some of the generators commute. A convenient way of expressing this is to draw a graph Γ with one vertex for every generator and one edge connecting each pair of commuting generators; the resulting RAAG is denoted A_Γ. In recent years, RAAGs and their automorphism groups have played a prominent role in geometric group theory and low-dimensional topology. RAAGs are linear groups and they arise naturally as subgroups of many other groups such as mapping class groups, Coxeter groups, and more general Artin groups (see, e.g., [15], [16], [19], [26]). Conversely, while all subgroups of free (or free abelian) groups are themselves free (or free abelian), a surprisingly diverse array of groups can be realized as subgroups of RAAGs, including surface groups and many 3-manifold groups (see [2], [24], [32]). The fact that the fundamental group of every closed hyperbolic 3-manifold virtually embeds in a RAAG was central to Agol's proof of the virtual Haken conjecture in [1], the final step in Thurston's program to classify 3-manifolds. The diversity of subgroups has also made RAAGs a fertile source of counterexamples for a variety of conjectures (see [3], [17]).

To date, outer automorphism groups of RAAGs have primarily been studied from an algebraic point of view (see, e.g., [13], [14], [20], [22]). As the case of mapping class groups and $\mathrm{Out}(F_n)$ clearly demonstrates, geometric approaches to studying such groups can be very effective. In this work, we focus on constructing an analogue of outer space for RAAGs that will allow us to apply similar methods to the study of $\mathrm{Out}(A_\Gamma)$. Some initial steps in this direction appear in previous papers. In [11], Charney and Vogtmann, together with Crisp, constructed a candidate outer space for 2-dimensional RAAGs (those for which Γ contains no triangles), but there is no apparent way to generalize this to higher dimensions. Then in [12], together with Stambaugh, they constructed a contractible space K_Γ with a proper action of a certain subgroup of $\mathrm{Out}(A_\Gamma)$. This subgroup, denoted $U(A_\Gamma)$, is made up of "untwisted" outer automorphisms of A_Γ that behave more like automorphisms of free groups. In particular, it excludes transvections between commuting pairs of generators. Here, we use the space K_Γ as a starting point to build an outer space for the full outer automorphism group.

Outer space for free groups, CV_n, can be described as a space of marked metric graphs with fundamental group F_n, where the marking specifies an isomorphism of

π_1 with F_n. Similarly, the symmetric space $Q_n = \mathrm{SO}(n)\backslash \mathrm{SL}(n,\mathbb{R})$ can be described as the space of marked flat tori with fundamental group $\mathbb{Z}^n$, where the marking gives an isomorphism of π_1 with $\mathbb{Z}^n$. Thus the basic objects in Q_n (tori) are all homeomorphic but have different flat metrics, while the basic objects in CV_n (graphs) have different homeomorphism types as well as different metrics. These different homeomorphism types, however, all have a common quotient, an n-petaled rose, obtained by collapsing any maximal tree. For a general RAAG, there is a canonical construction of a CAT(0) cube complex $\mathbb{S}_\Gamma$ with fundamental group A_Γ, known as the *Salvetti complex*, which has a k-torus for each k-clique in Γ. In the new outer space, this complex plays the role of the n-petaled rose. The basic objects in our outer space $\mathcal{O}_\Gamma$ are locally CAT(0) metric spaces (Y,d) containing contractible subspaces (analogous to maximal trees) that can be collapsed to produce a quotient homeomorphic to $\mathbb{S}_\Gamma$. Each (Y,d) is made up of a collection of (intersecting) flat tori marked by the free abelian groups generated by cliques in Γ. A point in $\mathcal{O}_\Gamma$ consists of one of these metric spaces (Y,d) marked by an isomorphism of $\pi_1(Y)$ with A_Γ.

More precisely, the spaces Y are homeomorphic (but not isometric) to nonpositively curved cube complexes called Γ*-complexes*, which were previously introduced in [12]. Marked Γ-complexes form a partially ordered set whose geometric realization is the simplicial complex K_Γ mentioned above. In K_Γ, Γ-complexes are viewed as combinatorial objects, not as metric objects, and the markings are of restricted type, allowing only an action of the subgroup $U(A_\Gamma)$. In the current paper, Γ-complexes are endowed with locally CAT(0) metrics that make the interior of each "cube" isometric to a Euclidean parallelotope. We call this a *skewed* Γ-complex. The objects (Y,d) in $\mathcal{O}_\Gamma$ are isometric to skewed Γ-complexes. The markings are arbitrary, and objects are equivalent if they are isometric by a map that commutes with the marking, up to free homotopy. As in the special cases of $\mathrm{GL}(n,\mathbb{Z})$ acting on Q_n and $\mathrm{Out}(F_n)$ acting on CV_n, $\mathrm{Out}(A_\Gamma)$ acts on $\mathcal{O}_\Gamma$ by changing the marking. Our main theorem states the following.

THEOREM 1.1
For any right-angled Artin group A_Γ, the space $\mathcal{O}_\Gamma$ is finite-dimensional, contractible, and the action of the group $\mathrm{Out}(A_\Gamma)$ *has finite point stabilizers.*

We now give a brief outline of the proof. The proof begins with the space K_Γ which, as noted above, was shown in [12] to be contractible. The passage from K_Γ to $\mathcal{O}_\Gamma$ involves several intermediate steps. First, we embed K_Γ into a new space Σ_Γ by endowing Γ-complexes with metrics making "cubes" into orthotopes, that is, orthogonal products of intervals of various lengths; these are called *rectilinear* Γ-complexes. In the case of a free group, this corresponds to embedding the spine of outer space into the full outer space CV_n. As in the case of K_Γ, the action on Σ_Γ is restricted to the subgroup $U(A_\Gamma)$. This is a result of allowing only certain types of markings, called *untwisted markings*. It is easy to show that K_Γ is a deformation retract of Σ_Γ, so that Σ_Γ is contractible.

Next, we allow the orthotopes in a Γ-complex to skew so that they become parallelotopes. This is done in a controlled manner, resulting in an *allowable parallelotope structure* which is still locally CAT(0). The collection of skewed Γ-complexes with untwisted markings is denoted $\mathcal{T}_\Gamma$. We show that there is a deformation retraction of $\mathcal{T}_\Gamma$ onto Σ_Γ defined by straightening the parallelotopes, so $\mathcal{T}_\Gamma$ is also contractible.

The action on $\mathcal{T}_\Gamma$ is still restricted to the subgroup $U(A_\Gamma)$. To get a space on which all of $\mathrm{Out}(A_\Gamma)$ acts we must allow for transvections between commuting elements; these are called *twists*. To see how this is done, consider the case of a marked metric torus T^n. One can think of a change of marking as either a change in the isomorphism $\pi_1(T^n) \to \mathbb{Z}^n$, or as a change in the shape of the parallelotope whose quotient is T^n. To reconcile these viewpoints in the case of a skewed Γ-complex, we put an equivalence relation on the points in $\mathcal{T}_\Gamma$. Namely, two skewed Γ-complexes with specified markings are equivalent if they are isometric by a map that takes one marking to the other (up to homotopy), where this map need *not* preserve the combinatorial structure. Then up to equivalence, we can accomplish twists by adjusting the skewing of appropriate tori in the Γ-complex.

The points in the new outer space $\mathcal{O}_\Gamma$ are equivalence classes of points in $\mathcal{T}_\Gamma$, thus there is a natural surjection $\mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$. The proof of Theorem 1.1 consists in showing that this map is a fibration with contractible fibers. The key problem is understanding to what extent the combinatorial structure on a marked skewed Γ-complex is determined by its metric. For this, we divide the hyperplanes into two classes, twist-minimal and twist-dominant, and show that the twist-minimal hyperplanes are completely determined by the metric. The twist-dominant hyperplanes on the other hand, depend on the shapes of the parallelotopes and can vary within a fiber. To show contractibility, we encode the allowable skewings by a vector in a Euclidean space and prove that the set of points in a fiber corresponds to a convex subspace of this Euclidean space.

Theorem 1.1 is a first step toward a more geometric study of $\mathrm{Out}(A_\Gamma)$. It leads to many natural questions, a few of which we now discuss briefly.

The dimension of $\mathcal{O}_\Gamma$ can be computed (with some effort) by looking at the graph Γ. As is the case for symmetric spaces and Teichmüller spaces, the action of $\mathrm{Out}(A_\Gamma)$ on $\mathcal{O}_\Gamma$ is not cocompact, and this dimension is quite a bit larger than the virtual cohomological dimension (VCD) of $\mathrm{Out}(A_\Gamma)$. An algebraic algorithm for computing this VCD has been established by Day, Sale, and Wade [21]. For both $\mathrm{GL}(n, \mathbb{Z})$ and $\mathrm{Out}(F_n)$, there is an equivariant deformation retract (a "spine") of dimension equal to the VCD, and it would be interesting to find an analogous spine for $\mathcal{O}_\Gamma$. The construction of such a spine might be fairly subtle, as it was shown in [28] that the dimension of K_Γ, though often equal to the VCD of $U(A_\Gamma)$, is sometimes strictly larger. As an aside, we remark that no natural spine has yet been constructed for the action of the mapping class group of a closed surface on Teichmüller space.

Much of the work on $\mathrm{Out}(F_n)$ and $\mathrm{GL}(n,\mathbb{Z})$ (as well as surface mapping class groups) depends on understanding the structure of the associated space at or near infinity, for example, by adding a "boundary" that compactifies either the space or its quotient, and studying the action on this boundary. Thurston compactified Teichmüller space by embedding it into the space of projective length functions for the fundamental group of the surface, outer space can be compactified by embedding it into the space of projective length functions on F_n, and the symmetric space Q_n embeds into the space of projective length functions on $\mathbb{Z}^n$. Vijayan [33] initiated a study of length functions on A_Γ, which was further developed by Beyrer and Fioravanti [5], who used length functions to compactify the "untwisted" outer space K_Γ of [12]. A different way of understanding the structure at infinity is by "bordifying" the space, which compactifies not the space but rather the quotient. There are bordifications of Q_n (defined in much more generality by Borel and Serre [6]) and CV_n (defined by Bestvina and Feighn [4]). These were used to prove that the respective groups are virtual duality groups in the sense of Bieri and Eckmann. Is there an analogous bordification of $\mathcal{O}_\Gamma$? The question is subtle, as Brück and Wade [10] showed that $\mathrm{Out}(A_\Gamma)$ is not always a virtual duality group.

A space is a *classifying space for proper G-actions* if fixed point sets of finite subgroups are contractible. Such a space is called an *$\underline{E}G$-space*. These are useful, for example, for studying centralizers of finite-order elements. In addition, we recall that the Baum–Connes conjecture relates the topological K-theory of the reduced C^*-algebra of G to an appropriate equivariant homology theory evaluated at $\underline{E}G$. Both Q_n and CV_n are classifying spaces for proper actions, so it is natural to ask whether $\mathcal{O}_\Gamma$ is likewise for $\mathrm{Out}(A_\Gamma)$.

Finally, both symmetric spaces and outer space for free groups can be equipped with useful metrics (though the most intensively studied metric structure on outer space is an asymmetric metric). A geometric approach often gives a simpler, more natural explanation for algebraic features of the group. Is there a good metric on $\mathcal{O}_\Gamma$? How do geodesics in this metric behave?

The paper is organized as follows. In Section 2, we establish basic terminology, recall the construction of the space K_Γ, and embed it into a space Σ_Γ. In Section 3, we establish some basic properties of Γ-complexes which will be needed later on. In Section 4, we introduce the notion of twist-dominant and twist-minimal hyperplanes and investigate the extent to which these notions depend on the choice of Γ-structure and the marking. In Section 5, we define an allowable parallelotope structure on a Γ-complex and show that the resulting path metric is locally CAT(0). In Section 6, we prove that the space $\mathcal{T}_\Gamma$ of skewed Γ-complexes deformation retracts to Σ_Γ, hence is contractible. Finally, in Section 7, we define our outer space $\mathcal{O}_\Gamma$ and show that the natural map $\mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$ is a fibration with contractible fibers.

2. Preliminaries

We fix a finite simplicial graph $\Gamma = (V, E)$ throughout the paper, and denote by A_Γ the associated right-angled Artin group. In this section we give a brief account of the contents of [12]. We refer the reader to [12] for further details.

2.1. Graph terminology

For $v \in V$, the *link*, $\mathrm{lk}(v)$, is the full subgraph of Γ spanned by vertices adjacent to v, and the *star*, $\mathrm{st}(v)$, is the full subgraph spanned by $\mathrm{lk}(v)$ and v. If $W \subset V$, then $\mathrm{lk}(W) = \bigcap_{w \in W} \mathrm{lk}(w)$ and $\mathrm{st}(W)$ is the full subgraph spanned by $\mathrm{lk}(W) \cup W$.

Define $v \leq w$ to mean $\mathrm{lk}(v) \subseteq \mathrm{st}(w)$. This can happen in one of two ways: either $\mathrm{lk}(v) \subseteq \mathrm{lk}(w)$, in which case we write $v \leq_f w$, or $\mathrm{st}(v) \subseteq \mathrm{st}(w)$, in which case we write $v \leq_t w$. These are mutually exclusive unless $v = w$.

The following elementary lemma puts a restriction on the star- and link-orderings.

LEMMA 2.1
If $u \leq_t v \leq_f w$, then either $v = u$ or $v = w$.

Proof
Suppose that $u \neq v$. Since $u \in \mathrm{lk}(v)$ and $\mathrm{lk}(v) \subseteq \mathrm{lk}(w)$, $u \in \mathrm{lk}(w)$. Since $\mathrm{st}(u) \subseteq \mathrm{st}(v)$, this implies that $w \in \mathrm{lk}(v)$. Hence $v \leq_t w$, which is impossible unless $v = w$. □

If $v \leq_* w$ and $w \leq_* v$, then we say that v and w are *equivalent* and write $w \sim_* v$, where $* = f, t$ or $\emptyset$. The notation $v \leq_* w$ is justified by the fact that the induced relation on equivalence classes $[v]$ is a partial ordering. It will often be important to be more specific, so if $\mathrm{lk}(v) = \mathrm{lk}(w)$, then we say that v and w are *fold-equivalent*, and if $\mathrm{st}(v) = \mathrm{st}(w)$, then we say v and w are *twist-equivalent.*

For each $v \in V$, we divide the elements of $V_{\geq v} = \{u \mid u \geq v\}$ into two groups; namely,

- $\mathrm{lk}^+(v) = \{u \mid u \geq v \text{ and } u \in \mathrm{lk}(v)\} = \{u \in V \mid u \geq_t v, u \neq v\}$,
- $\mathrm{dlk}(v) = \{u \mid u \geq v \text{ and } u \notin \mathrm{lk}(v)\} = \{u \in V \mid u \geq_f v\}$.

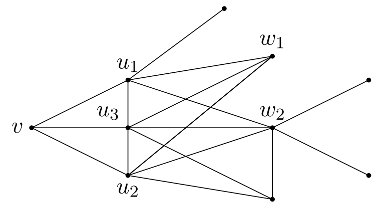

Figure 1. $\mathrm{lk}(v) = \{u_1, u_2, u_3\}$, $\mathrm{lk}^+(v) = \{u_3\}$, $\mathrm{dlk}(v) = \{v, w_1, w_2\}$.

(See Figure 1 for an example.) Observe that $\mathrm{dlk}(v)$ is equal to the "double link" $\mathrm{lk}(\mathrm{lk}(v))$, that is, every vertex in $\mathrm{dlk}(v)$ commutes with every vertex in $\mathrm{lk}(v)$. Also observe that if $u, u' \geq_t v$, then u is connected to u', so $\{v\} \cup \mathrm{lk}^+(v)$ is a clique.

The following distinction will be critical when we define the points in our new outer space.

Definition 2.2
A vertex $v \in \Gamma$ is *twist-dominant* if there is some $u \neq v$ with $v \geq_t u$, and is *twist-minimal* otherwise.

Note that elements of $\mathrm{lk}^+(v)$ are all twist-dominant, while elements of $\mathrm{dlk}(v)$ may be either twist-dominant or twist-minimal.

2.2. *Salvetti complexes*

For a simplicial graph Γ, the *Salvetti complex* $\mathbb{S}_\Gamma$ is a cube complex with one k-cube for each k-clique in Γ; in particular, it has a single 0-cube (for the empty clique) and a 1-cube for each vertex (=1-clique) of Γ. The 2-skeleton of $\mathbb{S}_\Gamma$ is the standard presentation complex for A_Γ, so $\pi_1(\mathbb{S}_\Gamma) \cong A_\Gamma$. The addition of higher-dimensional cubes guarantees that $\mathbb{S}_\Gamma$ satisfies Gromov's link condition; that is, all links are flag. Therefore, if all cubes of $\mathbb{S}_\Gamma$ are identified with standard Euclidean cubes $[0,1]^k$, then the induced path metric on $\mathbb{S}_\Gamma$ is nonpositively curved (locally CAT(0)) and its universal cover $\widetilde{\mathbb{S}}_\Gamma$ is CAT(0). In Figure 2, we show a simple example of a graph Γ and its Salvetti complex. In this example, the Salvetti is made of two tori glued along a circle labeled b plus a loop labeled d at the basepoint. In the right-hand picture we have cut open the tori.

Throughout this paper, we will assume familiarity with the language of locally CAT(0) and CAT(0) cube complexes, including hyperplanes, minsets, and so on, as can be found, for example, in [8].

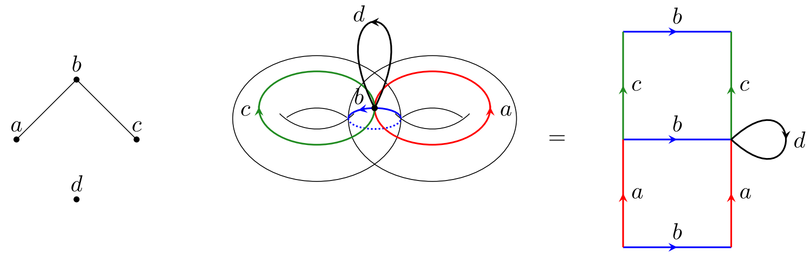

Figure 2. (Color online) A graph Γ and its Salvetti $\mathbb{S}_\Gamma$.

2.3. Γ-*Whitehead partitions*

2.3.1. *Definition and examples*

Let $V \cup V^{-1}$ be the generators of A_Γ and their inverses, and let m be a vertex of Γ. A Γ-*Whitehead partition* $\mathcal{P}$ *based at* m is a partition of $V \cup V^{-1}$ into three parts P^+, P^- (called the *sides* of $\mathcal{P}$) and $\mathrm{lk}(\mathcal{P})$, where

- $\mathrm{lk}(\mathcal{P})$ consists of all generators that commute with m, and their inverses;
- the sides of $\mathcal{P}$ form a thick partition of $V \cup V^{-1} \setminus \mathrm{lk}(\mathcal{P})$ (recall that a partition is *thick* if it has at least two elements on each side);
- m and m^{-1} are in different sides of $\mathcal{P}$;
- if $v \neq w$ are in the same component of $\Gamma \setminus \mathrm{st}(m)$, then v, v^{-1}, w and w^{-1} are all in the same side of $\mathcal{P}$.

A more succinct way to define a Γ-Whitehead partition $\mathcal{P}$ based at m is by forming a graph $\Gamma^\pm$ with one vertex for each element of $V \cup V^{-1}$ and an edge between distinct vertices x and y whenever x and y commute but are not inverses. If we let $\mathrm{lk}^\pm(m)$ be the link of m in $\Gamma^\pm$ and let $\mathcal{C}(m) = \{m, m^{-1}, C_1, \dots, C_k\}$ be the components of $\Gamma^\pm \setminus \mathrm{lk}^\pm(m)$, then

- $\mathrm{lk}(\mathcal{P})$ consists of vertices in $\mathrm{lk}^\pm(m)$, and
- the sides of $\mathcal{P}$ form a thick partition of $\mathcal{C}(m)$ that separates m from m^{-1}.

The components $C_1, \dots, C_k$ are called m-*components.* Thus m together with any proper subset of m-components gives one side of a valid Γ-Whitehead partition based at m.

A Γ-Whitehead partition $\mathcal{P}$ based at m determines an automorphism $\phi(\mathcal{P}, m)$ of A_Γ called a Γ-*Whitehead automorphism*. Examples of Γ-Whitehead automorphisms include partial conjugations and elementary folds (for details see [12]). Different bases for $\mathcal{P}$ give different automorphisms, but the partition $\mathcal{P}$ itself does not depend on the choice of base, and we will often not specify a base. Note that a Γ-Whitehead partition is completely determined by giving one of its sides.

Example 2.3
The following are three examples of Γ-Whitehead partitions for the graph Γ depicted in Figure 2:

- $\mathcal{P} = (P^+|P^-|\operatorname{lk}(\mathcal{P})) = (\{b,d\}|\{b^{-1},d^{-1}\}|\{a,a^{-1},c,c^{-1}\})$,
- $\mathcal{Q} = (Q^+|Q^-|\operatorname{lk}(\mathcal{Q})) = (\{a,d\}|\{a^{-1},d^{-1},c,c^{-1}\}|\{b,b^{-1}\})$,
- $\mathcal{R} = (R^+|R^-|\operatorname{lk}(\mathcal{R})) = (\{a,c,d\}|\{a^{-1},d^{-1},c^{-1}\}|\{b,b^{-1}\})$.

Here $\mathcal{P}$ is based at b, $\mathcal{Q}$ is based at a, and $\mathcal{R}$ can be based at either a or c.

2.3.2. *Properties*

Definition 2.4
If v and v^{-1} are in different sides of $\mathcal{P}$, then we say that $\mathcal{P}$ *splits* v. Define split$(\mathcal{P})$ to be the set of vertices of Γ that are split by $\mathcal{P}$, and

$$\max(\mathcal{P}) = \big\{v \in V \mid v \text{ is a maximal element in } \operatorname{split}(\mathcal{P})\big\},$$

where maximality is with respect to the relation "$\leq$" defined above. For a vertex $v \in V$, it is convenient to also define $\max(v) = \{v\}$.

LEMMA 2.5
If $\mathcal{P}$ is based at m and $\mathcal{P}$ splits v, then $v \leq_f m$.

Proof
Since $\mathcal{P}$ splits v, v is not in the link of m. Suppose that w is in the link of v. Since $\mathcal{P}$ splits v, v and w are not in the same component of $\Gamma - \operatorname{st}(m)$, so w must be in the link of m. This shows that $v \leq_f m$. ☐

It follows that the elements of $\max(\mathcal{P})$ are precisely the bases of $\mathcal{P}$, and they are all fold-equivalent.

LEMMA 2.6
If $\mathcal{P}$ splits a twist-dominant vertex v, then $\max(\mathcal{P}) = \{v\}$.

Proof
Let $m \in \max(\mathcal{P})$. By Lemma 2.5, $v \leq_f m$. Since v is twist-dominant, there is a $w \neq v$ with $v \geq_t w$. But then $v = m$ by Lemma 2.1. ☐

We extend our orderings on vertices of Γ to Γ-Whitehead partitions by declaring $\mathcal{P} \leq_* \mathcal{Q}$ for $* = f$ or t if for some (and therefore any) $v \in \max(\mathcal{P})$ and $w \in \max(\mathcal{Q})$ we have $v \leq_* w$.

2.3.3. Adjacency, compatibility, consistency

Definition 2.7
Let $\mathcal{P}$ and $\mathcal{Q}$ be Γ-Whitehead partitions. We say that $\mathcal{P}$ and $\mathcal{Q}$ are *adjacent* if $\max(\mathcal{P}) \subset \mathrm{lk}(\mathcal{Q})$. A vertex v is *adjacent* to $\mathcal{P}$ if $v \in \mathrm{lk}(\mathcal{P})$, and v and w are *adjacent* if they are adjacent in Γ.

Since all elements of $\max(\mathcal{P})$ have the same link, $\max(\mathcal{P}) \subset \mathrm{lk}(\mathcal{Q})$ if and only if $\max(\mathcal{Q}) \subset \mathrm{lk}(\mathcal{P})$, that is, the definition is symmetric.

Warning. In [12], we said "$\mathcal{P}$ and $\mathcal{Q}$ *commute*" instead of "$\mathcal{P}$ and $\mathcal{Q}$ are *adjacent*." There are two reasons for changing the terminology here. First, two partitions based at the same vertex v do not "commute" in the sense of [12] even though the generator v certainly commutes with itself; this caused confusion for several readers. The second reason is that the definition of "commute" written in [12] is not actually the one used in the proofs of the lemmas: we mistakenly added a condition in the definition requiring that the twist-equivalence classes of $\max(\mathcal{P})$ and $\max(\mathcal{Q})$ be different. The proofs of all lemmas in that paper about commuting partitions are correct, however, if one replaces "commuting" by the definition of "adjacent" given above.

Definition 2.8
Let $\mathcal{P}$ and $\mathcal{Q}$ be distinct Γ-Whitehead partitions.

(1) $\mathcal{P}$ and $\mathcal{Q}$ are *compatible* if either $\mathcal{P}$ and $\mathcal{Q}$ are adjacent or they have sides $P^\times$ and $Q^\times$ with $P^\times \cap Q^\times = \emptyset$.
(2) Sides $P^\times$ of $\mathcal{P}$ and $Q^\times$ of $\mathcal{Q}$ are *consistent* if either $\mathcal{P}$ and $\mathcal{Q}$ are adjacent or $P^\times \cap Q^\times \neq \emptyset$.

If $\mathcal{P}$ and $\mathcal{Q}$ are compatible but are not adjacent, then exactly three of the four possible choices of pairs of sides are consistent, by Lemma 3.6 of [12]. (If they are adjacent, then any choice of sides is consistent.)

Define an involution $P^\times \mapsto \overline{P^\times}$ that switches sides of $\mathcal{P}$, that is, $\overline{P^+} = P^-$ and $\overline{P^-} = P^+$.

LEMMA 2.9
If $\mathcal{P}$ and $\mathcal{Q}$ are compatible but not adjacent and $P^\times \cap Q^\times = \emptyset$, then $P^\times \cap \mathrm{lk}(\mathcal{Q}) = \emptyset$ so $P^\times \subset \overline{Q^\times}$; similarly, $Q^\times \subset \overline{P^\times}$.

Proof
This is Lemma 3.4 of [12]. It is illustrated in Figure 3. □

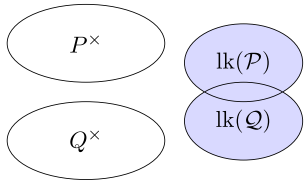

Figure 3. (Color online) Nonadjacent partitions $\mathcal{P}$ and $\mathcal{Q}$ have disjoint sides $P^\times$ and $Q^\times$ that are also disjoint from $\mathrm{lk}(\mathcal{P}) \cup \mathrm{lk}(\mathcal{Q})$ (see Lemma 2.9).

2.4. *Blowups*

In this section, we fix a collection $\Pi = \{\mathcal{P}_1, \dots, \mathcal{P}_k\}$ of pairwise-compatible Γ-Whitehead partitions and construct a locally CAT(0) cube complex $\mathbb{S}^\Pi$ with fundamental group A_Γ, whose edges are labeled either by a partition in Π or by a vertex of Γ.

Definition 2.10

A choice of sides for a set of Γ-Whitehead partitions is *consistent* if each pair is consistent. A consistent choice of sides $P_i^\times$ for all $\mathcal{P}_i \in \Pi$ is a *region*.

LEMMA 2.11

Any consistent choice of sides for a subset of Π *can be extended to a region.*

Proof

This is Lemma 3.9 in [12]. It follows easily by induction on k, the number of partitions. □

Regions will form the vertices of our cube complex. To describe the higher-dimensional cubes, it is convenient to define a graph Γ_Π that realizes our notion of "adjacency" for partitions in Π.

Definition 2.12

Let Π be a collection of pairwise-compatible Γ-Whitehead partitions. Then Γ_Π is the (simplicial) graph with

- one vertex for each element of $V \cup \Pi$, and
- an edge between A and B whenever A and B are adjacent according to Definition 2.7, that is, $\max(A) \subset \mathrm{lk}(B)$.

The link of a vertex $A \in \Gamma_\Pi$ will be denoted $\mathrm{lk}_\Pi(A)$, the star by $\mathrm{st}_\Pi(A)$, and the double link by $\mathrm{dlk}_\Pi(A)$.

Every $v \in V \cup V^{-1}$ is in P_i^+, P_i^- or $\mathrm{lk}(\mathscr{P}_i)$ for each i. If $v \notin \mathrm{lk}(\mathscr{P}_i)$, define the *$v$-side* of $\mathscr{P}_i$ to be the side containing v. Then the set of v-sides for those $\mathscr{P}_i$ that are not adjacent to v form a consistent set and can be extended to a region. Any such region is called a *terminal region for v*.

Definition 2.13
The *blowup* $\mathbb{S}^\Pi$ is a cube complex with one vertex for each region $\mathbb{r} = \{P_1^\times, \ldots, P_k^\times\}$. The edges of $\mathbb{S}^\Pi$ are constructed as follows.

- If two regions differ only by changing the side of $\mathscr{P}_i$, then we connect them by an (unoriented) edge labeled $\mathscr{P}_i$.
- If $\mathbb{r}$ is a terminal region for v, then the region $\mathbb{r}^{*v}$ obtained by switching sides of all $\mathscr{P}_i$ that split v is a terminal region for v^{-1}, and we connect the two by an oriented edge labeled v that goes from $\mathbb{r}^{*v}$ to $\mathbb{r}$.

Higher-dimensional cubes are attached whenever a set of edges forms the 1-skeleton of a cube whose labels span a clique in Γ_Π.

From the definition, we immediately see the following.

- There is an edge labeled v terminating at the vertex $\mathbb{r} = \{P_1^\times, \ldots, P_k^\times\}$ if and only if for each i, either $v \in P_i^\times$ or $v \in \mathrm{lk}(\mathscr{P}_i)$. If no $\mathscr{P}_i$ splits v, then an edge labeled v in $\mathbb{S}^\Pi$ is a loop at $\mathbb{r}$.
- There is an edge labeled $\mathscr{P}_j$ with one endpoint at $\mathbb{r} = \{P_1^\times, \ldots, P_k^\times\}$ if and only if for each $i \neq j$, either $\mathscr{P}_i$ and $\mathscr{P}_j$ are adjacent or some side of $\mathscr{P}_j$ is contained in $P_i^\times$. In particular, if $\mathscr{P}_i$ and $\mathscr{P}_j$ are not adjacent, then both $P_j^\times \cap P_i^\times$ and $\overline{P}_j^\times \cap P_i^\times$ are nonempty. An edge labeled $\mathscr{P}_i$ is never a loop.

In Figures 4–6, we show three blowups of $\mathbb{S}_\Gamma$ for the graph Γ shown in Figure 2. As before, edges with the same label in the right-hand diagram are identified. In Figures 4 and 5, the blowups are two tori identified along a circle, with an extra edge attached. In Figure 6, the blowup is two tori identified along a cylinder, with an

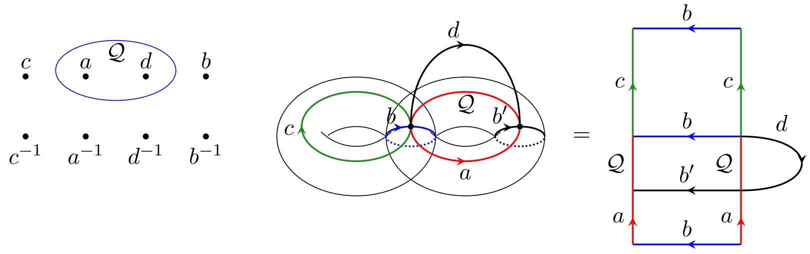

Figure 4. (Color online) The blowup $\mathbb{S}^{\mathcal{Q}}$ for $\mathcal{Q} = (\{a, d\}|\{a^{-1}, d^{-1}, c, c^{-1}\}|\{b, b^{-1}\})$.

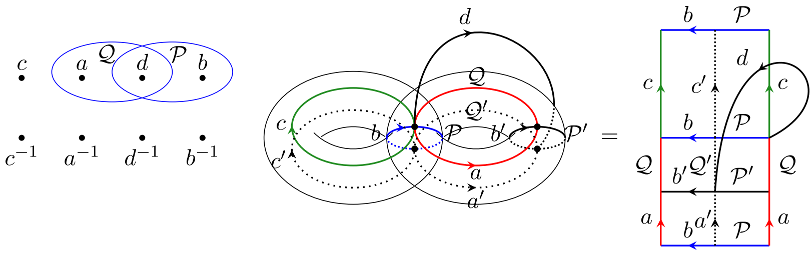

Figure 5. (Color online) The blowup $\mathbb{S}^{\Pi}$ for $\Pi = \{\mathcal{P}, \mathcal{Q}\}$, $\mathcal{P} = (\{b, d\}|\{b^{-1}, d^{-1}\}|\{a, a^{-1}, c, c^{-1}\})$.

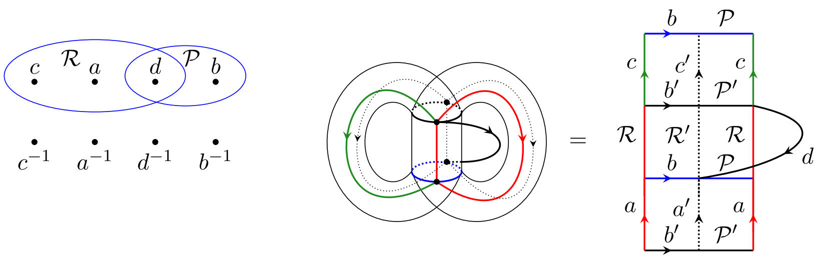

Figure 6. (Color online) The blowup $\mathbb{S}^{\Pi}$ for $\Pi = \{\mathcal{P}, \mathcal{R}\}$, $\mathcal{R} = (\{a, c, d\}|\{a^{-1}, d^{-1}, c^{-1}\}|\{b, b^{-1}\})$.

extra edge attached. The structure of blowups will be explored in much more detail in Section 3.

2.5. *Collapsing hyperplanes*

Definition 2.14

Let H be a hyperplane in a cube complex X. The closure of the set of cubes that intersect H is called the *hyperplane carrier* $\kappa(H)$, and the *hyperplane collapse* associated to H is the map c_H on X that collapses $\kappa(H)$ to H.

Recall from [12] that hyperplanes in $\mathbb{S}^{\Pi}$ are characterized by the fact that the set of edges they intersect is exactly the set of edges with a given label $A \in V \cup \Pi$. We say the hyperplane is *labeled* by A.

PROPOSITION 2.15 ([12, Theorem 4.6])

If $\mathcal{P} \in \Pi$, and $H_{\mathcal{P}}$ is the hyperplane in $\mathbb{S}^{\Pi}$ labeled by $\mathcal{P}$, then the image of $\mathbb{S}^{\Pi}$ under $c_{H_{\mathcal{P}}}$ is isomorphic to $\mathbb{S}^{\Pi - \mathcal{P}}$.

The *standard collapse* $c_\pi \colon \mathbb{S}^\Pi \to \mathbb{S}^\emptyset = \mathbb{S}_\Gamma$ is the map that collapses all hyperplanes whose labels are in Π.

2.6. Untwisted outer space Σ_Γ

Recall that the *untwisted subgroup* $U(A_\Gamma) \leq \mathrm{Out}(A_\Gamma)$ is the subgroup generated by Γ-Whitehead automorphisms, graph automorphisms, and inversions. By work of Laurence [27] and Servatius [31], $U(A_\Gamma)$ together with automorphisms $v \mapsto vw$ for $v \leq_t w$ (called *twists*) generate the full group $\mathrm{Out}(A_\Gamma)$. In this section, we recall the main theorem of [12] and use it to define a contractible space Σ_Γ on which $U(A_\Gamma)$ acts properly. We first recall the space K_Γ studied in [12].

Definition 2.16
A cube complex X is a Γ-*complex* if it is isomorphic to a blowup $\mathbb{S}^\Pi$ for some Π. A Γ-*complex collapse* $c \colon X \to \mathbb{S}_\Gamma$ is the composition of an isomorphism $X \cong \mathbb{S}^\Pi$ with the standard collapse $\mathbb{S}^\Pi \to \mathbb{S}_\Gamma$.

Example 2.17
If Γ has no edges, then a Γ-complex is a connected graph with no univalent or bivalent vertices and no separating edges, and a Γ-complex collapse contracts a maximal tree to a point.

A *marked* Γ-*complex* is an equivalence class of pairs (X, g), where
- X is a Γ-complex,
- $g \colon X \to \mathbb{S}_\Gamma$ is a homotopy equivalence, and
- $(X', g') \sim (X, g)$ if there is a cube complex isomorphism $i \colon X' \to X$ with $g \circ i \simeq g'$.

A marking $h \colon X \to \mathbb{S}_\Gamma$ is *untwisted* if the composition of a homotopy inverse h^{-1} with some (and hence any) Γ-complex collapse induces an element of the untwisted subgroup $U(A_\Gamma)$.

If a hyperplane collapse $c_H \colon X' \to X$ is a homotopy equivalence, then we set

$$(X', h \circ c_H) > (X, h).$$

This induces a partial order on Γ-complexes with untwisted markings. The *spine* K_Γ is the geometric realization of the resulting poset, that is, it is a simplicial complex, where a k-simplex is a Γ-complex with an untwisted marking together with a chain of k hyperplane collapses, each of which is a homotopy equivalence to another Γ-complex with an untwisted marking.

THEOREM 2.18 ([12, Theorems 4.17 and 6.24])
The spine K_Γ is contractible, and $U(A_\Gamma)$ acts properly and cocompactly on K_Γ.

We now define the space Σ_Γ by viewing the cubes of a Γ-complex X as metric objects, each isometric to an orthogonal product of intervals of various lengths, that is, an *orthotope*. The result is a locally CAT(0) complex X which we will call a *rectilinear* Γ-complex. All edges dual to the same hyperplane in X have the same length, called the *width* of the hyperplane. A point in Σ_Γ is then a marked rectilinear Γ-complex (X, h), where h is untwisted and the cube complex isomorphism in the definition of the equivalence relation must be an isometry on each orthotope. In the case where Γ has no edges, the spine K_Γ is the same as the spine of (reduced) outer space, as originally defined in [18], and Σ_Γ is reduced outer space itself.

The spine K_Γ embeds in Σ_Γ as follows: the image of a vertex $[(X, h)]$ of K_Γ is determined by the property that all edges of X have length $1/n$, where n is the number of hyperplanes in X. The image of each higher-dimensional simplex is the linear span of its vertices.

PROPOSITION 2.19
K_Γ is a deformation retract of Σ_Γ.

Proof
Σ_Γ contains the set $P\Sigma_\Gamma$ of marked metric Γ-complexes $[(X, h)]$ for which the sum of the hyperplane widths in X equals 1. Note that the image of our embedding of K_Γ into Σ_Γ is contained in $P\Sigma_\Gamma$. The map $\Sigma_\Gamma \to P\Sigma_\Gamma$ that scales all edge lengths simultaneously is a deformation retraction.

The subspace $P\Sigma_\Gamma$ decomposes into a union of open simplices, one for each marked Γ-complex $[(X, h)]$, of dimension one less than the number of hyperplanes in X. The points in this simplex are obtained by varying the widths of the hyperplanes while keeping the sum equal to 1. For each such simplex, consider the barycentric subdivision of its closure, and let $K[(X, h)]$ be the subcomplex of this barycentric subdivision spanned by the barycenters of faces that are actually contained in $P\Sigma_\Gamma$. It is easy to see that $K[(X, h)]$ is equal to the image of K_Γ under the embedding described above, and is a deformation retract of Σ_Γ. $\square$

COROLLARY 2.20
The space Σ_Γ is contractible.

3. Combinatorial and metric structure of blowups

Throughout this section, we fix a compatible set Π of Γ-Whitehead partitions. To prove our main theorem we will have to understand the structure of the blowup $\mathbb{S}^\Pi$ in some detail. We gather some facts about blowups here.

3.1. Basics

The following basic features of blowups $\mathbb{S}^\Pi$ are either part of Theorem 3.14 of [12] or follow immediately from the definition of $\mathbb{S}^\Pi$.

(1) $\mathbb{S}^\Pi$ is a locally CAT(0) cube complex; that is, the path metric induced by making each k-cube isometric to $[0,1]^k$ is locally CAT(0).

(2) The subcomplex $\mathbb{C}^\Pi \subset \mathbb{S}^\Pi$ consisting of cubes all of whose edge labels are in Π is CAT(0) and locally convex, and it contains all vertices of $\mathbb{S}^\Pi$.

(3) The standard collapse map c_π maps all of $\mathbb{C}^\Pi$ to the single vertex in $\mathbb{S}_\Gamma$.

(4) The set of edges of $\mathbb{S}^\Pi$ with a given label $A \in V \cup \Pi$ is the set of edges that intersect a single hyperplane, which we will call H_A. All hyperplanes in $\mathbb{S}^\Pi$ are of this form.

(5) Each hyperplane H_A inherits a cube complex structure from $\mathbb{S}^\Pi$ whose edges are labeled by the elements of $V \cup \Pi$ that are adjacent to A, that is, by elements in $\mathrm{lk}_\Pi(A)$.

(6) There is at most one edge with a given label at any vertex of $\mathbb{S}^\Pi$.

Another way to define the subcomplex $\mathbb{C}^\Pi$ is to observe that the set of sides of the partitions in Π form a *pocset*, that is, a partially ordered set with an order reversing involution $P \mapsto \overline{P}$ such that pairs P, $\overline{P}$ are unrelated; this follows from Lemma 2.9. Any pocset satisfying suitable finiteness conditions gives rise to a CAT(0) cube complex (see, e.g., [30]), and $\mathbb{C}^\Pi$ is isomorphic to the cube complex associated to the pocset of sides of Π.

3.2. Adjacent labels

In this section, we show that there is a unique cube in $\mathbb{S}^\Pi$ for every maximal clique in the graph Γ_Π, that is, any maximal set of pairwise adjacent elements of $V \cup \Pi$.

We begin with existence, for which the following definition is useful.

Definition 3.1

Let $\mathcal{P} \in \Pi$. For $v \in V \cup V^{-1} \setminus \mathrm{lk}(\mathcal{P})$, the *$v$-side of $\mathcal{P}$* is the side containing v. For $\mathcal{Q} \in \Pi \setminus \{\mathcal{P}\}$ not adjacent to $\mathcal{P}$, the *$\mathcal{Q}$-side of $\mathcal{P}$* is the side containing some side of $\mathcal{Q}$ (there is a unique such side by Lemma 2.9).

Stated in terms of hyperplanes, $H_\mathcal{P}$ splits the subspace $\mathbb{C}^\Pi$ into two components. If $v \notin \mathrm{lk}(\mathcal{P})$, then the v-side of $H_\mathcal{P}$ is the side containing the terminal vertex of some (hence every) edge labeled v. If $\mathcal{Q}$ and $\mathcal{P}$ are distinct and not adjacent, then $H_{\mathcal{Q}}$ does not intersect $H_{\mathcal{P}}$ and the $\mathcal{Q}$-side of $H_{\mathcal{P}}$ is the side containing $H_{\mathcal{Q}}$.

PROPOSITION 3.2

Let Π be a compatible set of k Γ-Whitehead partitions, and let $\mathcal{A} = \{A_1, \dots, A_\ell\}$ be the vertices of a maximal clique in Γ_Π. Then there is a cube in $\mathbb{S}^\Pi$ with edge labels $\{A_1, \dots, A_\ell\}$.

Proof

Let $\mathcal{A} \cap V = \{v_1, \dots v_r\}$ and $\mathcal{A} \cap \Pi = \{\mathcal{Q}_1, \dots, \mathcal{Q}_s\}$, so

- $\Pi = \{\mathcal{Q}_1, \dots, \mathcal{Q}_s, \mathcal{P}_1, \dots, \mathcal{P}_t\}$ with $s + t = k$, and
- $\mathcal{A} = \{v_1, \dots, v_r, \mathcal{Q}_1, \dots, \mathcal{Q}_s\}$ with $r + s = \ell$.

For any choices of sides $Q_i^\times$ of $\mathcal{Q}_i$ for $i = 1, \dots, s$ and exponents $v_j^\times = v_j$ or v_j^{-1} for $j = 1, \dots, r$, we will find a region $\mathbb{r}$ which is a terminal region for each $\mathcal{Q}_i$ and $v_j^\times$. These 2^ℓ regions (some of which may coincide, as we will see) form the vertices of an ℓ-dimensional cube in $\mathbb{S}^\Pi$ with edges labeled by the elements of $\mathcal{A}$.

To define the region associated to $\{v_1^\times, \dots, v_r^\times, Q_1^\times, \dots, Q_s^\times\}$, we will start with the sides $Q_i^\times$. We then need to choose a side of each $\mathcal{P}_i$. Since $\mathcal{A}$ is a maximal clique, for each $\mathcal{P}_i$ there is some A_j not adjacent to $\mathcal{P}_i$. If A_j is a vertex v_j, let $P_i^\times$ be the $v_j^\times$-side of $\mathcal{P}_i$, and if A_j is a Γ-Whitehead partition $\mathcal{Q}_j$, let $P_i^\times$ be the $\mathcal{Q}_j$-side of $\mathcal{P}_i$. To see that $P_i^\times$ does not depend on the choice of A_j, observe that if $\mathcal{P}_i$ is not adjacent to either A_j or A_k, then the fact that $w_j \in \max(A_j)$ and $w_k \in \max(A_k)$ are joined by an edge in Γ implies that all of $\{w_j, w_j^{-1}, w_k, w_k^{-1}\}$ are on the same side of $\mathcal{P}_i$, so the A_j-side of $\mathcal{P}_i$ is the same as the A_k-side of $\mathcal{P}_i$.

Now let $\mathbb{r} = \{Q_1^\times, \dots, Q_s^\times, P_1^\times, \dots, P_t^\times\}$. To see that this is a region, we must show that any two elements either belong to adjacent partitions or intersect nontrivially. Each pair $\mathcal{Q}_i, \mathcal{Q}_j$ is adjacent. If $\mathcal{Q}_j$ is not adjacent to $\mathcal{P}_i$, then we have chosen the $\mathcal{Q}_j$-side $P_i^\times$ of $\mathcal{P}_i$. Since $P_i^\times$ contains an entire side of $\mathcal{Q}_j$, it intersects both sides of $\mathcal{Q}_j$ nontrivially. If $\mathcal{P}_i$ and $\mathcal{P}_j$ are not adjacent, let A_k be an element of $\mathcal{A}$ that is not adjacent to $\mathcal{P}_i$. We argue by contradiction: suppose that $P_i^\times \cap P_j^\times = \emptyset$. If A_k is a vertex v_k, then $v_k^\times \in P_i^\times$, so $v_k^\times \notin P_j^\times$. Since $v_k^\times \notin \mathrm{lk}(\mathcal{P}_j)$ by Lemma 2.9, this contradicts our choice of $P_j^\times$. If A_k is a partition $\mathcal{Q}_k$ and $Q_k^\times \subset P_i^\times$, then $\max(\mathcal{Q}_k) \not\subset \mathrm{lk}(\mathcal{P}_j)$ and neither side of $\mathcal{Q}_k$ is contained in $P_j^\times$, again contradicting our choice of $P_j^\times$.

The region $\mathbb{r}$ is a terminal region for each $v_i^\times$. If we use $(v_i^\times)^{-1}$ instead of $v_i^\times$, then we get another region, terminal for $(v_i^\times)^{-1}$. These two regions may be the same if v_i and v_i^{-1} are on the same side of each $\mathcal{P}_j$, in which case the edge labeled v_i is a loop. Switching sides of any $\mathcal{Q}_i$ gives another region, with an edge labeled $\mathcal{Q}_i$ joining the two (this edge is never a loop). Thus we have the 1-skeleton of an ℓ-dimensional cube in $\mathbb{S}^\Pi$, which is filled in since all of the edge labels are adjacent. □

COROLLARY 3.3
Two hyperplanes H_A and H_B intersect if and only if A and B are adjacent.

Proof
If A and B are adjacent, then it follows from Proposition 3.2 that there is a square with sides labeled A and B, so the hyperplanes H_A and H_B intersect at the midpoint of that square. Conversely, if H_A and H_B intersect, then there is a pair of edges dual to these hyperplanes that bound a square, so A and B must be adjacent since by the construction of $\mathbb{S}^\Pi$, we only fill in squares when labels are adjacent. $\square$

Remark 3.4
Corollary 3.3 says that Γ_Π is the *crossing graph* for $\mathbb{S}^\Pi$ as defined in [30].

PROPOSITION 3.5
Any cubes c, c' in $\mathbb{S}^\Pi$ with the same edge labels are parallel, that is, S^Π contains a subcomplex isomorphic to $c \times [0,n]$ for some $n \in \mathbb{Z}$, with $c = c \times \{0\}$ and $c' = c \times \{n\}$.

Proof
If c and c' share a vertex, then they must be equal, so we may assume that they are disjoint. Recall that $\mathbb{C}^\Pi$ is CAT(0), hence connected, and contains every vertex of $\mathbb{S}^\Pi$. Let p be a minimal-length edge path from c to c' that is contained in $\mathbb{C}^\Pi$. The CAT(0) property implies that p crosses each hyperplane at most once. The first edge of p is labeled by some partition $\mathcal{P}$. Since p has minimal length, $\mathcal{P}$ is distinct from all of the edge labels of c. Let $\mathbb{r} = \{P^\times, \dots\}$ be the initial vertex of p, where $P^\times$ is a side of $\mathcal{P}$, and let $\mathbb{r}'$ be the terminal vertex.

Suppose now that some edge label B of c is not adjacent to $\mathcal{P}$. If $B = \{v\}$, then $P^\times$ contains v, so both $\mathbb{r}$ and $\mathbb{r}'$ use this side. The first edge of the path p switches sides of $\mathcal{P}$, that is, crosses the hyperplane $H_{\mathcal{P}}$, so in order to reach $\mathbb{r}'$ it must cross $H_{\mathcal{P}}$ again, contradicting the assumption that it is the shortest path. If $B = \mathcal{Q}$, then the side of $\mathcal{Q}$ that appears in $\mathbb{r}$ is neither contained in $P^\times$ nor contains $P^\times$ (since there are edges labeled both $\mathcal{P}$ and $\mathcal{Q}$ at $\mathbb{r}$). Since changing sides of $\mathcal{Q}$ is allowed at $\mathbb{r}'$, it follows that the side of $\mathcal{P}$ at $\mathbb{r}'$ must also be equal to $P^\times$. As before, the initial edge of the path p crosses $H_{\mathcal{P}}$, so in order to reach $\mathbb{r}'$ it must cross $H_{\mathcal{P}}$ again, contradicting the assumption that it is shortest.

We conclude that B and $\mathcal{P}$ are adjacent for all edge labels B of c, so there is a cube $c \times e_B \subset \mathbb{S}^\Pi$. The side c'' of this cube opposite from c is closer to c', and we can continue to build a product neighborhood $c \times [0,n]$ until we reach c'. $\square$

COROLLARY 3.6
For every maximal collection $\{A_1, \ldots, A_k\}$ of pairwise adjacent labels, there is a unique maximal cube in $\mathbb{S}^\Pi$ with those edge labels.

Proof
Existence is Proposition 3.2. Uniqueness follows from Proposition 3.5, since the existence of two distinct parallel cubes implies that $\{A_1, \ldots, A_k\}$ is not maximal. □

3.3. Characteristic cycles

Definition 3.7
Let v be a vertex of Γ, and let e_v be an edge of $\mathbb{S}^\Pi$ labeled by v. Choose a minimal-length edge path $p(e_v)$ in $\mathbb{C}^\Pi$ from the terminal vertex $\tau(e_v)$ to the initial vertex $\iota(e_v)$. The loop $\chi_v = p(e_v) \cup e_v$ is called a *characteristic cycle* for v.

Since $\mathbb{C}^\Pi$ is contractible, the standard collapse map takes χ_v to the loop in $\mathbb{S}_\Gamma$ representing v. By the construction of $\mathbb{S}^\Pi$, a characteristic cycle for v has one edge labeled v and one edge labeled $\mathcal{P}$ for each partition $\mathcal{P} \in \Pi$ that splits v. Such a path crosses the same hyperplanes as a locally geodesic loop β_v representing v (see Figure 7). Since v is not adjacent to any other label on an edge crossed by χ_v, e_v must lie in β_v. Similarly, any edge e_A in χ_v for which $v \in \max(A)$ must lie in β_v.

3.3.1. Characteristic cycles and partitions
In this section, we give a more detailed description of characteristic cycles χ_v in terms of partitions that split v. This depends on the following observation.

LEMMA 3.8
Suppose that $\mathcal{P}$ and $\mathcal{Q}$ are compatible and both split v. Let P and Q be the v-sides of $\mathcal{P}$ and $\mathcal{Q}$. If $\mathcal{P}$ and $\mathcal{Q}$ are not adjacent, then either $P \subset Q$ or $Q \subset P$. If $\mathcal{P}$ and $\mathcal{Q}$ are adjacent, then $P \not\subset Q$ and $Q \not\subset P$.

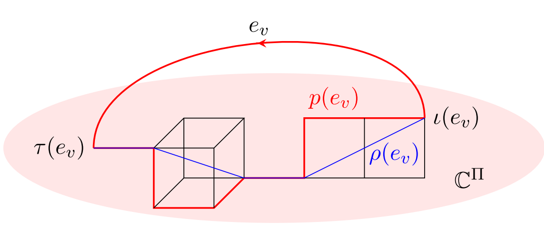

Figure 7. (Color online) The local geodesic $\beta_v = e_v \cup \rho(e_v)$ and a characteristic cycle $p(e_v) \cup e_v$ containing e_v.

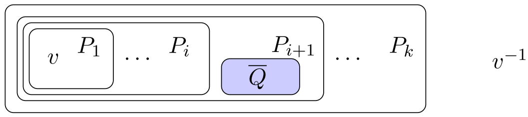

Figure 8. (Color online) Partitions $\mathcal{P} \in \Pi$ with $v \in \max(\mathcal{P})$ are nested. Partitions $\mathcal{Q} \in \Pi$ that do not split v and are not adjacent to v have a side $\overline{Q}$ in the nest. ($\overline{Q}$ is the side that does not contain v.)

Proof
$P \cap Q$ contains v so is not empty, and $\overline{P} \cap \overline{Q}$ contains v^{-1} so is not empty. Since $\mathcal{P}$ and $\mathcal{Q}$ are compatible, either they are adjacent or $P \subset Q$ or $Q \subset P$ by Lemma 2.9. If $\mathcal{P}$ and $\mathcal{Q}$ are adjacent, then P intersects $\mathrm{lk}(Q)$, so $P \not\subset Q$, similarly $Q \not\subset P$. ☐

Now fix a vertex $v \in \Gamma$, and for each $\mathcal{P} \in \Pi$ that is not adjacent to v, let P denote the v-side of $\mathcal{P}$, and let $\overline{P}$ denote the side that does not contain v (note that v^{-1} may be in P or in $\overline{P}$). Let $\mathcal{P}_1, \dots, \mathcal{P}_k$ be the partitions in Π that have v as a maximal element (i.e., are based at v). By Lemma 3.8, the v-sides P_i are nested; that is, after possibly reordering we have $P_1 \subset \cdots \subset P_k$ (see Figure 8). For notational convenience, set $P_0 = \{v\}$ and $\mathcal{P}_0 = \{P_0 | \overline{P}_0 | \mathrm{lk}^{\pm}(v)\}$, and let $P_{k+1} = \overline{P}_0 \setminus \{v^{-1}\}$. The differences $dP_i = P_{i+1} \setminus P_i$ for $i = 1, \dots, k$ are called the *pieces* of the nest.

If $\mathcal{Q} \in \Pi$ is not adjacent to v and does not split v, then Lemma 2.9 implies that some side of $\mathcal{Q}$ is contained in either P_i or $\overline{P}_i$ for each i; since $\mathcal{Q}$ does not split v, this must be the side that does not contain v, which we have called $\overline{Q}$. We conclude that $\overline{Q}$ is contained in some piece dP_i of the nest.

Let Π_v denote the set of partitions Π that split v. Note that in addition to the partitions $\mathcal{P}_i$, Π_v may contain partitions that split v but do not have v as a maximal element; such partitions may be adjacent to each other. A characteristic cycle χ_v has one edge for each element of Π_v, so in particular one edge for each $\mathcal{P}_i$. Let $\mathcal{S}_i$ be the consistent set of sides

$$\begin{cases} Q & \text{if } \mathcal{Q} \in \Pi_v, Q \supseteq P_i, \\ \overline{Q} & \text{if } \mathcal{Q} \in \Pi_v, Q \subsetneq P_i, \\ Q & \text{if } \mathcal{Q} \in \Pi \setminus \Pi_v \text{ is not adjacent to v,} \end{cases}$$

and let $\overline{\mathcal{S}}_i$ be the set obtained from $\mathcal{S}_i$ by replacing P_i by $\overline{P}_i$. Since the P_i are nested, changing sides of $\mathcal{P}_i$ does not change the fact that the relevant intersections are nonempty, so $\overline{\mathcal{S}}_i$ is still consistent. Either set can be completed to a region by any consistent choice of sides of the $\mathcal{R} \in \Pi$ that are adjacent to v. One endpoint of the edge in χ_v labeled $\mathcal{P}_i$ is a region that extends the set $\mathcal{S}_i$; call this endpoint x_i. The

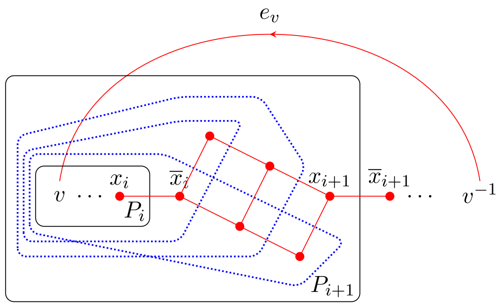

Figure 9. (Color online) A characteristic cycle for v has one edge e_v and one edge for each partition that splits v. The partitions based at v are nested. Partitions that split v but are not based at v are indicated by dotted lines; these have max $>_f v$. Partitions are adjacent if and only if they cross.

other endpoint $\overline{x}_i$ of this edge is obtained by switching P_i to $\overline{P}_i$; this extends $\overline{\mathcal{S}}_i$ (see Figure 9).

We can now describe an arbitrary characteristic cycle χ_v in terms of partitions (refer to Figure 9). Start with any consistent choice $\mathcal{S}$ of sides of the $\mathcal{R} \in \Pi$ that are adjacent to v, and let x_0 be the region extending $\mathcal{S}$ that is given by choosing the v-side of every partition that is not adjacent to v. Define a partition $\mathcal{Q} \in \Pi_v$ to be *innermost* if its v-side Q does not contain the v-side of any other element of Π_v. By Lemma 3.8, all innermost partitions in Π_v are adjacent. For the first edge of χ_v, we may choose the edge labeled by any innermost $\mathcal{Q} \in \Pi_v$. For the next edge, we may choose any edge labeled by $\mathcal{Q}' \in \Pi_v$ that is innermost in $\Pi_v \setminus \mathcal{Q}$. The following edge is labeled by any innermost element of $\Pi_v \setminus \{\mathcal{Q}, \mathcal{Q}'\}$, and so on, and the loop is closed by an edge labeled e_v.

If no two partitions that split v are adjacent, the description of characteristic cycles in terms of partitions is particularly simple, since then the v-sides of all elements of Π_v are nested so any characteristic cycle χ_v consists of an edge path dual to the nest plus an edge e_v connecting its endpoints. This characteristic cycle is a local geodesic in $\mathbb{S}^\Pi$. In particular, we record the following.

LEMMA 3.9
If v is twist-dominant, then any characteristic cycle χ_v for v is an edge path in $\mathbb{S}^\Pi$ labeled by v and the partitions that split v. Furthermore, χ_v is a local geodesic in $\mathbb{S}^\Pi$.

Proof

If v is twist-dominant, then all partitions that split v have max $= \{v\}$, so none of them are adjacent. $\square$

3.3.2. Characteristic cycles and minsets

Since $\mathbb{S}^\Pi$ is locally CAT(0), its universal cover $\widetilde{\mathbb{S}}^\Pi$ is CAT(0). We will label edges and hyperplanes in $\widetilde{\mathbb{S}}^\Pi$ with the same label as their images in $\mathbb{S}^\Pi$. The group A_Γ acts on $\widetilde{\mathbb{S}}^\Pi$ via deck transformations (preserving labels), using the identification of $\pi_1(\mathbb{S}^\Pi)$ with A_Γ induced by the standard collapse map $c_\pi : \mathbb{S}^\Pi \to \mathbb{S}_\Gamma$. The following lemma uses standard CAT(0) methods to investigate the relation between characteristic cycles and this action.

LEMMA 3.10

Let $A \in \Pi \cup V$ be a label, and let $v \in \max(A)$.

(1) *The minset of v in $\widetilde{\mathbb{S}}^\Pi$ decomposes as a product $\alpha_v \times \widetilde{H}_v$, where α_v is an axis for v containing an edge $\tilde{e}_v$ and $\widetilde{H}_v$ is the dual hyperplane.*

(2) *For each edge in $\mathbb{S}^\Pi$ labeled A, there is a unique edge e_v such that e_A and e_v are contained in a local geodesic β_v, hence every characteristic cycle for v containing e_v contains e_A, and vice versa.*

(3) *The carrier $\kappa(H_A)$ lies in the image of* $\mathrm{Min}(v)$, *and the induced cubical structures on H_A and H_v are isomorphic.*

(4) *If w is adjacent to A, then the carrier of H_A contains a characteristic cycle for w.*

Proof

Consider the minset of v in the universal cover $\widetilde{\mathbb{S}}^\Pi$. By standard properties of CAT(0) spaces, $\mathrm{Min}(v)$ decomposes as an orthogonal product $\alpha_v \times Y$, where Y is a convex subspace of $\widetilde{\mathbb{S}}^\Pi$ and α_v is an axis for v. The image of α_v under the projection $\widetilde{\mathbb{S}}^\Pi \to \mathbb{S}^\Pi$ is a local geodesic β_v. By the comments after Definition 3.7 we may assume that β_v contains an edge e_v, and thus that α_v contains a lift $\tilde{e}_v$ of e_v. We conclude that $\widetilde{H}_v$ must contain a copy of Y.

Conversely, we claim that every edge dual to $\widetilde{H}_v$ lies on an axis for v, so by convexity this copy of Y contains $\widetilde{H}_v$. Suppose that $\tilde{e}'_v$ is another edge dual to $\widetilde{H}_v$, separated from $\tilde{e}_v$ by a square whose other label is $A \in \mathrm{lk}_\Pi(v)$. Let χ_v be a characteristic cycle for v containing e_v. Since every edge label B on χ_v splits v, we have $\mathrm{lk}_\Pi(B) \supseteq \mathrm{lk}_\Pi(v) \ni A$, so $\mathbb{S}^\Pi$ contains an annulus $\chi_v \times e_A$. The boundary of this annulus is two characteristic cycles, one containing e_v and one containing the image e'_v of $\tilde{e}'_v$, so these two characteristic cycles are homotopic, and correspond to two different axes for v, one containing $\tilde{e}_v$ and one containing $\tilde{e}'_v$. Since any two edges dual to H_v can be connected by a sequence of squares, this proves (1).

As observed above, for any $A \in \Pi$ with $v \in \max(A)$, the local geodesic β_v containing e_v also contains a (unique) edge labeled e_A. It follows that the axis through $\tilde{e}_v$ contains a lift of e_A, hence the dual hyperplane $\widetilde{H}_A$ contains a subspace parallel to $\widetilde{H}_v$. Since every edge that is adjacent to A is also adjacent to v, these two hyperplanes must, in fact, be isomorphic. Thus the carrier of $\widetilde{H}_A$ lies entirely in the minset of v and every edge dual to $\widetilde{H}_A$ lies on an axis containing an e_v edge. This proves (2) and (3).

For (4), note that since w is adjacent to v, e_v and e_w span a cube in the carrier of H_v. Let χ_w be a characteristic cycle containing e_w. The label on every edge of this cycle is also adjacent to v, so the entire characteristic cycle is contained in the carrier of H_v. □

COROLLARY 3.11
If an edge $\widetilde{e}$ of $\widetilde{\mathbb{S}}^\Pi$ is in $\mathrm{Min}(v)$*, then its image in $\mathbb{S}^\Pi$ is labeled either by v, by some partition that splits v, or some label that is adjacent to v.*

Proof
By Lemma 3.10(1), $\mathrm{Min}(v) \cong \alpha_v \times \widetilde{H}_v$, and we may assume that α_v is a lift of the local geodesic β_v described in Section 3.3. An edge e_A of $\mathbb{S}^\Pi$ can only be in β_v if $A = v$ or A splits v. (*Warning*: splitting v does not guarantee that e_A will be contained in β_v unless $\max(A) = v$.) The hyperplane H_v is parallel to a subcomplex with all labels adjacent to v. □

3.4. *Subcomplexes of $\mathbb{S}^\Pi$ associated to a generator*

Fix a compatible set Π of Γ-Whitehead partitions. We will use the graph Γ_Π defined in Definition 2.12, with vertices $V \cup \Pi$, to describe certain subcomplexes of $\mathbb{S}^\Pi$ associated to a generator $v \in V$. We remark that Γ_Π can be used to encode the fold relation: $A \leq_f B$ if and only if $\mathrm{lk}_\Pi(A) \subseteq \mathrm{lk}_\Pi(B)$. However, it does *not* encode the twist relation; this will be explored further in Section 4.

Definition 3.12
Given a set of vertices Λ of Γ_Π, the *span of* Λ, denoted $\mathrm{span}(\Lambda)$, is the subcomplex of $\mathbb{S}^\Pi$ consisting of those cubes with all edge labels in Λ.

Example 3.13
We have $\mathrm{span}(\Pi) = \mathbb{C}^\Pi$.

Now fix $v \in V$, let e_v be an edge labeled by v, and let H_v be the hyperplane in $\mathbb{S}^\Pi$ dual to e_v. The carrier $\kappa(H_v)$ is a product

$$\kappa(H_v) = e_v \times \mathbb{K}_{\mathrm{lk}(v)},$$

where $\mathbb{K}_{\mathrm{lk}(v)}$ is the connected component of span($\mathrm{lk}_\Pi(v)$) that contains the terminal vertex x of e_v.

Since $v \in \mathrm{dlk}_\Pi(v)$, some connected component of span($\mathrm{dlk}_\Pi(v)$) contains x. Denote this component by $\mathbb{K}_{\mathrm{dlk}(v)}$. Since every vertex of $\mathrm{lk}_\Pi(v)$ is linked to every vertex of $\mathrm{dlk}_\Pi(v)$, the product of these two subcomplexes is also a subcomplex of $\mathbb{S}^\Pi$:

$$\mathbb{K}_v = \mathbb{K}_{\mathrm{dlk}(v)} \times \mathbb{K}_{\mathrm{lk}(v)}.$$

Example 3.14
If v is twist-dominant, then $\mathrm{dlk}(v) = \{v\}$, so $\mathrm{dlk}_\Pi(v)$ consists of v and partitions based at v. These are precisely the labels in any characteristic cycle for v (see the discussion at the end of Section 3.3.1), so the characteristic cycle χ_v containing x is one component of span($\mathrm{dlk}_\Pi(v)$) Thus,

$$\mathbb{K}_v = \chi_v \times \mathbb{K}_{\mathrm{lk}(v)} \cong \chi_v \times H_v,$$

and $\mathbb{K}_v$ is equal to the image in $\mathbb{S}^\Pi$ of the minset of v in $\widetilde{\mathbb{S}}^\Pi$.

If v is twist-minimal, then $\mathbb{K}_v$ can be considerably larger and more complicated than the image of Min(v). However, the following lemma holds for any $v \in V$.

PROPOSITION 3.15
The subcomplex $\mathbb{K}_{\mathrm{lk}(v)}$ contains at least one characteristic cycle for every $u \in \mathrm{lk}(v)$, and $\mathbb{K}_{\mathrm{dlk}(v)}$ contains at least one characteristic cycle for every $w \in \mathrm{dlk}(v)$.

Proof
Let e_v be an edge in $\mathbb{S}^\Pi$ labeled v, and let x be the terminal vertex of e_v. Then $\mathbb{K}_v$ contains $\kappa(H_v)$, so the first statement follows from Lemma 3.10(4).

For the second statement, let $w \in \mathrm{dlk}(v)$, and recall that the labels on a characteristic cycle χ_w consist of w and all partitions $\mathcal{P} \in \Pi$ that split w. If $\mathcal{P}$ is based at m and splits w, then $\mathrm{lk}(m) \supseteq \mathrm{lk}(w) \supseteq \mathrm{lk}(v)$, so $m \in \mathrm{dlk}(v)$. This shows that all characteristic cycles χ_w are contained in span($\mathrm{dlk}_\Pi(v)$). It remains to check that the component of span($\mathrm{dlk}_\Pi(v)$) containing x also contains a characteristic cycle for w. For this it suffices to find an edge e_w in the same component as x.

Let e_w be an edge labeled w whose terminal vertex y has minimal distance in $\mathbb{C}^\Pi$ to x. (Recall that $\mathbb{C}^\Pi$ contains all vertices and is CAT(0).) If $y = x$, then we are done; otherwise connect y to x by a minimal-length edge path p in $\mathbb{C}^\Pi$. We claim that this edge path lies entirely in $\operatorname{span}(\operatorname{dlk}(v))$.

To see this, let $\mathcal{P}_1, \dots, \mathcal{P}_k$ be the successive labels on the path p (all of these labels are partitions). Since the path has minimal length, each $\mathcal{P}_i$ occurs only once. The vertex y is a terminal region for w, x is a terminal region for y, and the two regions differ by changing the sides of each $\mathcal{P}_i$ on the path, say, from P_i to $\overline{P}_i$.

If $\mathcal{P}_i$ is not in $\operatorname{lk}_\Pi(w)$, then it is not in $\operatorname{lk}_\Pi(v)$ either, so v and w must be in different sides of $\mathcal{P}_i$, specifically $w \in P_i$ and $v \in \overline{P}_i$. Since each $\mathcal{P}_i$ is a Γ-Whitehead partition, this means v and w are in different components of $\Gamma \setminus \operatorname{st}(\mathcal{P}_i)$ for all i. But $\operatorname{lk}(v) \subset \operatorname{lk}(w)$, so this can only happen if $\mathcal{P}_i \in \operatorname{dlk}_\Pi(v)$. Thus we will be done if we can show that no $\mathcal{P}_i$ is adjacent to w

Suppose to the contrary that some partition along the path is in $\operatorname{lk}_\Pi(w)$; let $\mathcal{P}_i$ be the first such partition. We first claim that $\mathcal{P}_i$ is adjacent to $\mathcal{P}_{i-1}$. If not, then there is a unique pair of sides of $\mathcal{P}_i$ and $\mathcal{P}_{i-1}$ with empty intersection. Since $P_{i-1} \cap P_i$, $\overline{P}_{i-1} \cap P_i$, and $\overline{P}_{i-1} \cap \overline{P}_i$ all correspond to vertices of the path p, the empty intersection must be $P_{i-1} \cap \overline{P}_i$. Since $\mathcal{P}_{i-1}$ is not in $\operatorname{lk}_\Pi(w)$, $w \in P_{i-1}$, as observed in the previous paragraph. But Lemma 2.9 implies that $\operatorname{lk}(\mathcal{P}_i)$, which contains w, does not intersect P_{i-1}, giving a contradiction.

Since $\mathcal{P}_i$ is adjacent to $\mathcal{P}_{i-1}$ we can reroute the path p to obtain a new path with the same edge labels that crosses $\mathcal{P}_i$ before it crosses $\mathcal{P}_{i-1}$. Repeating the argument, we can arrange that $\mathcal{P}_i$ labels the first edge of the path, so this edge has an endpoint at y. Filling in a square with edge labels w and $\mathcal{P}_i$, we obtain an edge labeled w that is closer to x, contradicting our original choice of e_w. □

Now let $\widetilde{\mathbb{K}}_v \subset \widetilde{\mathbb{S}}^\Pi$ be the connected component of the lift of $\mathbb{K}_v$ containing an axis for v. This decomposes as a product $\widetilde{\mathbb{K}}_v = \widetilde{\mathbb{K}}_{\operatorname{dlk}(v)} \times \widetilde{\mathbb{K}}_{\operatorname{lk}(v)}$.

COROLLARY 3.16

$\widetilde{\mathbb{K}}_v$ is preserved by the action of the special subgroup $A_{\operatorname{dlk}(v)} \times A_{\operatorname{lk}(v)}$, and $\widetilde{\mathbb{K}}_{\operatorname{dlk}(v)}$ contains an axis for every element of the group $A_{\operatorname{dlk}(v)}$. If $\alpha_v \subset \widetilde{\mathbb{K}}_{\operatorname{dlk}(v)}$ is the axis for v, then $\alpha_v \times \widetilde{\mathbb{K}}_{\operatorname{lk}(v)}$ is the minset of v.

Proof

First note that the subcomplexes $\mathbb{K}_{\operatorname{dlk}(v)}$ and $\mathbb{K}_{\operatorname{lk}(v)}$ are locally convex in $\mathbb{S}^\Pi$. This follows from the fact that a cube lies in one of these subcomplexes if and only if its edges all lie in that subcomplex. By general properties of CAT(0) spaces, a locally convex embedding of a subspace lifts to a globally convex embedding on universal covers and induces an injective map on fundamental groups.

It follows from Proposition 3.15 that under the standard collapse map, the image of $\pi_1(\mathbb{K}_{\mathrm{dlk}(v)})$ in $\pi_1(\mathbb{S}_\Gamma) = A_\Gamma$ is the subgroup $A_{\mathrm{dlk}(v)}$ and the image of $\pi_1(\mathbb{K}_{\mathrm{lk}(v)})$ is $A_{\mathrm{lk}(v)}$. Hence these subgroups preserve the lifts $\widetilde{\mathbb{K}}_{\mathrm{dlk}(v)}$ and $\widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$. Since these subspaces are convex in $\widetilde{\mathbb{S}}^\Pi$, they contain axes for each element of the corresponding subgroup.

The last statement follows by Lemma 3.10(1) since $\widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$ is parallel to and isomorphic to $\widetilde{H}_v$. □

3.5. Branch loci

In Section 7, we will be given a locally CAT(0) space X with fundamental group A_Γ and will need to construct an isomorphism of X with some blowup $\mathbb{S}^\Pi$. We will do this using the action of A_Γ on the universal cover $\widetilde{X}$. In this section, we discuss features of the action of A_Γ on $\widetilde{\mathbb{S}}^\Pi$ that will help in this task.

Definition 3.17

A point $x \in \mathrm{Min}(v) \subset \widetilde{\mathbb{S}}^\Pi$ is a *branch point for* v if the link of x in $\mathrm{Min}(v)$ is strictly smaller than the link of x in $\widetilde{\mathbb{S}}^\Pi$. Denote the branch locus of v by $\mathrm{br}(v)$.

(Recall that the link of a point x in a CAT(0) metric space X is defined to be the boundary of a small ball centered at x. This is standard terminology; the reader should not confuse this with the graphical links used elsewhere in this paper.)

If v is central, then $\mathrm{Min}(v) = \widetilde{\mathbb{S}}^\Pi$ and hence $\mathrm{br}(v) = \emptyset$. No Γ-Whitehead partition can split a central v, so in every blowup a characteristic cycle for v consists of a single edge which is a loop. For the rest of this section, we assume that v is not central, and show that in this case the location of hyperplanes in $\widetilde{\mathbb{S}}^\Pi$ is determined by branch loci of minsets.

PROPOSITION 3.18

Let H_A be a hyperplane in $\mathbb{S}^\Pi$ with $v \in \max(A)$, and let $\widetilde{H}_A$ be a lift of H_A to $\widetilde{\mathbb{S}}^\Pi$. If v is not central, then each component of the boundary of $\kappa(\widetilde{H}_A)$ contains a branch point of $\mathrm{Min}(v)$.

Proof

Let e_A be an edge dual to H_A. By Lemma 3.10(3), we know that e_A is contained in some characteristic cycle χ_v for v. Let e_B be the edge following e_A in χ_v, so that either $B = v$ or B is a partition that splits v.

If B is a partition based at w and $w >_f v$, choose $u \in \mathrm{lk}(B) \setminus \mathrm{lk}(A)$. Denote the common endpoint of e_A and e_B by x, and let $\partial_x(A)$ and $\partial_x(B)$ be the components of $\partial(\kappa(H_A))$ and $\partial(\kappa(H_B))$, respectively, that contain x. Then $\partial_x(A) \cong H_v$ is a subcomplex of $\partial_x(B)$, but $\partial_x(B)$ is strictly larger, since $\kappa(H_B)$ contains a square with edge labels u and B, and that square is not in $\kappa(H_A)$. Thus there is a point $x' \in \partial_x(A)$ that is adjacent to some edge e_C with C not adjacent to v. This means that no lift of e_C is contained in $\mathrm{Min}(v) \cong \alpha_v \times \widetilde{H}_A \cong \alpha_v \times \widetilde{H}_v$; that is, any lift $\tilde{x}'$ of x' lying on $\partial(\kappa(\widetilde{H}_A))$ is a branch point for v.

If $v \in \max(B)$, then we need to choose our characteristic cycle carefully and look more closely at the vertex x. To this end, we recall the description of characteristic cycles from Section 3.3.1. If $\mathcal{P}_1, \dots, \mathcal{P}_k$ are the partitions in Π that are based at v, then the v-sides P_i of the $\mathcal{P}_i$ are nested and, for notational convenience we set $P_0 = \{v\}$, $\mathcal{P}_0 = \{P_0 | \overline{P}_0 | \mathrm{lk}^{\pm}(v)\}$, and $P_{k+1} = \overline{P}_0 \setminus \{v^{-1}\}$, so (after possibly reordering) we have

$$P_0 \subset P_1 \subset \cdots \subset P_k \subset P_k \subset P_{k+1}.$$

Since A is based at v, we have $A = \mathcal{P}_i$ for some $i = 0, \dots, k$ and the vertex x corresponds to a region that extends the consistent set $\overline{\mathcal{S}}_i$ given by

$$\overline{\mathcal{S}}_i = \begin{cases} Q & \text{if } \mathcal{Q} \text{ splits v and } Q \supsetneq P_i, \\ \overline{Q} & \text{if } \mathcal{Q} \text{ splits v and } Q \subseteq P_i, \\ Q & \text{if } \mathcal{Q} \text{ does not split } v \text{ and is not adjacent to v,} \end{cases}$$

where Q is the v-side of $\mathcal{Q}$.

Since $v \in \max(B)$, there is no $\mathcal{Q}$ in Π whose v-side Q satisfies $P_i \subsetneq Q \subsetneq P_{i+1}$, so for any characteristic cycle the edge labeled $A = \mathcal{P}_i$ is followed by an edge labeled $B = \mathcal{P}_{i+1}$ (this situation is illustrated in Figure 10). We claim that for some characteristic cycle χ_v there is an edge e_C at this vertex whose label C is not adjacent to v and does not split v, so by Lemma 3.11 no lift of this edge is in $\mathrm{Min}(v)$. But $\partial(\widetilde{H}_A) \subset \mathrm{Min}(v)$ does contain a lift of x, so that is a branch point.

Recall that if a partition $\mathcal{Q}$ is not adjacent to v and does not split v, then it has a side $\overline{Q}$ sitting in some piece $dP_j = P_j \setminus P_{j-1}$ of the nest; call this the *nesting side* of $\mathcal{Q}$. We say $\mathcal{Q}$ is *outermost* if $\overline{Q}$ is not properly contained in any other such nesting

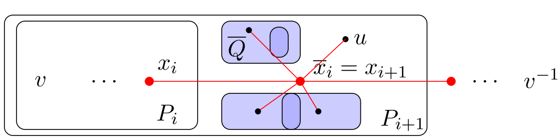

Figure 10. (Color online) If there is no $\mathcal{Q}$ with $P_i \subsetneq Q \subsetneq P_{i+1}$, then the remaining edges at the vertex $\overline{x}_i = x_{i+1}$ are either adjacent to v or correspond to those $\overline{Q}$ and u in $dP_i = P_i \setminus P_{i-1}$ that are outermost.

side. We call a vertex u *outermost* if u is not adjacent to v and is not contained in any nesting side.

Since v is not central and $P_{i+1} \neq P_i$, the piece $dP_{i+1} = P_{i+1} \setminus P_i$ must contain at least one outermost side or vertex; let C be the corresponding label. If C is a partition, then the condition that C is outermost guarantees that both sides of C are consistent with $\overline{\mathcal{S}}_i$. We claim that we can extend $\overline{\mathcal{S}}_i = \mathcal{S}_{i+1}$ to a region that is an endpoint of an edge e_C; that is, we can choose sides of all remaining $\mathcal{Q}$ that are consistent with each other and with both sides of C.

The remaining $\mathcal{Q}$ are those that are adjacent to v. These do not split v and are adjacent to every partition that does split v. Suppose that such a $\mathcal{Q}$ is not adjacent to C. If C is an element of $V \cup V^{-1}$, choose the side of $\mathcal{Q}$ that contains C; the result is a terminal region for C, that is, there is an edge labeled C at the corresponding vertex.

If C is a partition, let $C^{\times}$ denote the nesting side. Both sides of $\mathcal{Q}$ intersect $\mathrm{lk}(v)$, so they cannot be contained in $C^{\times}$. It follows from Lemma 2.9 that some side $Q^{\times}$ must contain $C^{\times}$. Note that $Q^{\times}$ intersects both sides of C, and also intersects the previously chosen side of any partition not adjacent to $\mathcal{Q}$, so the complete set of chosen sides is a region. Since C is outermost, switching sides of C still gives a region, and we have found our vertex x. □

For a generator v, Lemma 3.10(1) gives a decomposition $\mathrm{Min}(v) \cong \alpha_v \times \widetilde{H}_v$, where α_v is an axis containing an edge $\widetilde{e}_v$ and $\widetilde{H}_v$ is the hyperplane dual to $\widetilde{e}_v$. Let

$$\mathrm{pr}_v : \mathrm{Min}(v) \cong \alpha_v \times \widetilde{H}_v \to \alpha_v$$

be the (nearest-point) projection map corresponding to this decomposition.

If $\mathbb{S}^{\Pi}$ is a blowup with the standard collapse marking, then an axis α_v in $\widetilde{\mathbb{S}}^{\Pi}$ is transverse to some lift $\widetilde{H}_A$ of a hyperplane H_A if and only if $A = v$ or A is a partition that splits v. In either case, we say that $\widetilde{H}_A$ *splits* v.

COROLLARY 3.19
If v is not central, then the image $\mathrm{pr}_v(\mathrm{br}(v))$ of the branch set under projection to α_v is a set of discrete points and closed intervals. Each component of the complement of this image crosses exactly one hyperplane, which lifts a hyperplane in $\mathbb{S}^{\Pi}$ labeled either by v or by a partition $\mathcal{P}$ based at v.

Proof
First note that being a branch point is a closed condition and pr_v is a closed map, so $\mathrm{pr}_v(\mathrm{br}(v))$ is closed. By Corollary 3.16, the minset $\mathrm{Min}(v)$ decomposes as $\alpha_v \times \widetilde{\mathbb{K}}_{\mathrm{lk}(v)} \subseteq \widetilde{\mathbb{K}}_{\mathrm{dlk}(v)} \times \widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$, and by Lemma 3.10, we may assume that α_v contains lifts of all edges in $\mathbb{S}^{\Pi}$ labeled by some A with $v \in \max(A)$.

Any segment of α_v lying in the interior of a cube $C \subset \widetilde{\mathbb{K}}_{\mathrm{dlk}(v)}$ of dimension at least 2 lies entirely in the branch set, since C is not contained in the minset. So the only segments of α_v which might not be in the image are contained in edges $\tilde{e}$ of $\widetilde{\mathbb{K}}_{\mathrm{dlk}(v)}$. Let $\tilde{e}$ be such an edge, and let $\widetilde{H}$ be the hyperplane dual to $\tilde{e}$. The hyperplane $\widetilde{H}$ projects to a hyperplane H_A in $\mathbb{S}^\Pi$ for some $A \in V \cup \Pi$.

Since v splits A, either $v \in \max(A)$ or any $w \in \max(A)$ satisfies $w >_f v$. If $v \in \max(A)$, then $\kappa(\widetilde{H}) = \widetilde{e} \times \widetilde{H}$ lies entirely in $\mathrm{Min}(v)$ and hence the interior of $\kappa(\widetilde{H})$ does not contain any branch points. By Proposition 3.18, the two boundary components of $\kappa(\widetilde{H})$ do contain branch points so the two endpoints of $\tilde{e}$ do lie in $\mathrm{pr}_v(\mathrm{br}(v))$.

If $w \in \max(A)$ satisfies $w >_f v$, then there exists $u \in \mathrm{lk}(w)$ with $u \notin \mathrm{lk}(v)$. By Proposition 3.15, $\kappa(H_A)$ contains a characteristic cycle for u. It follows that $\kappa(\widetilde{H})$ contains a square with edges labeled by A and u. This square does not lie in $\mathrm{Min}(v)$ (since u is not in $\mathrm{lk}(v)$). This implies that the closest edge in $\mathrm{Min}(v)$ that is parallel to the edge labeled A is also in a square with one edge outside of $\mathrm{Min}(v)$, so this edge is entirely contained in the branch locus. Since this edge is dual to $\widetilde{H}$, it projects to $\tilde{e}$, so $\tilde{e}$ is contained in $pr_v(\mathrm{br}(v))$. ☐

If v is twist-dominant, then every partition in Π_v has $\max = \{v\}$, so every χ_v is an edge path in $\mathbb{S}^\Pi$ and its lift to $\widetilde{\mathbb{S}}^\Pi$ is an edge path which is an axis α_v for v. By Proposition 3.18, every vertex of α_v is the projection of a branch point and there are no branch points in the interior of edges. We record these observations in the following statement.

COROLLARY 3.20
If v is twist-dominant and not central, then any lift of a characteristic cycle to $\widetilde{\mathbb{S}}^\Pi$ is an axis α_v, and $\mathrm{pr}_v(\mathrm{br}(v))$ *is precisely the set of vertices of α_v.*

4. Hyperplanes in Γ-complexes

Let (X, h) be a point of Σ_Γ, that is, a rectilinear Γ-complex with an untwisted marking. If we choose an isomorphism $X \cong \mathbb{S}^\Pi$ the hyperplanes of X acquire labels, and we can use these labels to define what it means for a hyperplane to be twist-dominant or twist-minimal. In this section, we show that this designation is independent of the isomorphism and can be detected using the action of A_Γ on $\widetilde{X}$ induced by any untwisted marking.

To this end, let $\mathcal{C}(X)$ be the *crossing graph* of X, that is, the graph whose vertices are the hyperplanes of X, and where two vertices are connected by an edge if the corresponding hyperplanes intersect nontrivially. If we give X the structure of a blowup, then $\mathcal{C}(X) \cong \Gamma_\Pi$.

We defined twist and fold orderings on partitions $\mathcal{P}$ by choosing an element $v \in \max(\mathcal{P})$ and using the twist and fold orderings defined in terms of Γ. The defining graph Γ occurs as a subgraph of Γ_Π, but the corresponding subgraph in $\mathcal{C}(X)$ is not well defined since it depends on a choice of isomorphism $\mathcal{C}(X) \cong \Gamma_\Pi$. We do know that both orderings are well defined on fold-equivalence classes in Γ, so it is natural to try to define these orderings on *fold-equivalence classes* of $\mathcal{C}(X)$, that is, equivalence classes of vertices with the same link. This works well for the fold ordering, but must be modified for the twist ordering, as we will see. In the end, our notions of twist-dominant and twist-minimal will be defined using both $\mathcal{C}(X)$ and the combinatorial structure of X itself.

4.1. *Isomorphisms of* Γ*-complexes*

First, we define twist and fold orderings for hyperplanes in a Γ-complex X and show that, for any isomorphism $X \cong \mathbb{S}^\Pi$, these orderings coincide with the orderings of their labels, as previously defined. Note that the ordering of labels is well defined on fold-equivalence classes, so we need the same to be true here.

Definition 4.1

Let H be a hyperplane in a Γ-complex X. The *link* $\mathrm{lk}(H)$ of H is the link of H in $\mathcal{C}(X)$. In other words, $\mathrm{lk}(H)$ is the set of hyperplanes $K \neq H$ that intersect H nontrivially. The *fold-equivalence class* $[\![H]\!]$ is

$$[\![H]\!] = \{K \mid \mathrm{lk}(K) = \mathrm{lk}(H)\}.$$

We then define $[\![H]\!] \leq_f [\![K]\!]$ if $\mathrm{lk}(H) \subseteq \mathrm{lk}(K)$.

By Corollary 3.3, hyperplanes $H_A \neq H_B$ in $\mathbb{S}^\Pi$ intersect nontrivially if and only if their labels A and B are adjacent, so this coincides with the notion previously defined for $A \leq_f B$.

Giving a combinatorial definition of the twist relation is more subtle, and requires us to look beyond the structure of $\mathcal{C}(X)$ to the combinatorial structure of X itself.

Definition 4.2

We call a hyperplane H *cyclic* if

$$\bigcup_{H' \in [\![H]\!]} \kappa(H') \cong H \times C,$$

where C is a graph homeomorphic to S^1. Define $[\![K]\!] \leq_t [\![H]\!]$ to mean that H is cyclic and $\mathrm{lk}(K) \cup [\![K]\!] \subseteq \mathrm{lk}(H) \cup [\![H]\!]$.

The second condition in the definition of $[\![K]\!] \leq_t [\![H]\!]$ is the analogue of $\mathrm{st}(v) \subseteq \mathrm{st}(w)$ for the twist relation on Γ. However, the second condition alone does not capture the notion of twist-dominance. For instance, if Γ is a 4-cycle with vertices a, b, c, d, then in the Salvetti $\mathbb{S}_\Gamma$ we have $\mathrm{lk}(H_a) \cup [\![H_a]\!] = \{H_a, H_b, H_c, H_d\} = \mathrm{lk}(H_b) \cup [\![H_b]\!]$, but neither a nor b is twist-dominant as generators.

Since not every fold-equivalence class $[\![H]\!]$ is cyclic, we only have $[\![K]\!] \leq_t [\![H]\!]$ when $[\![H]\!]$ is cyclic. Nevertheless, it is transitive: if $[\![K]\!] \leq_t [\![H]\!]$ and $[\![H]\!] \leq_t [\![L]\!]$, then $[\![L]\!]$ must be cyclic so $[\![K]\!] \leq_t [\![L]\!]$. Also note that the analogue of Lemma 2.1 still holds for fold-equivalence classes of hyperplanes: if $[\![K]\!]$, $[\![H]\!]$, and $[\![L]\!]$ are distinct, then $[\![K]\!] \leq_t [\![H]\!] \leq_f [\![L]\!]$ is not possible.

Definition 4.3
A hyperplane H is *twist-dominant* if there is some hyperplane $K \neq H$ with $[\![K]\!] \leq_t [\![H]\!]$; in particular, H must be cyclic. If H is not twist-dominant, then it is *twist-minimal*.

If $X = \mathbb{S}_\Gamma$, then each hyperplane is labeled by a generator and the two notions of twist-dominance coincide. Indeed, a hyperplane H_v of $\mathbb{S}_\Gamma$ is cyclic if and only if v is not fold-equivalent to another generator. Then if there exists $w \neq v$ with $[\![H_w]\!] \leq_t [\![H_v]\!] = \{H_v\}$, this means that $w \leq_t v$, and conversely.

LEMMA 4.4
Let X be a Γ-complex, and choose an isomorphism $X \cong \mathbb{S}^\Pi$. For any hyperplane $H_B \subset \mathbb{S}^\Pi$, $[\![H_B]\!]$ is twist-dominant if and only if there exists H_A such that $\max(A) \leq_t \max(B)$.

Proof
First, suppose that $[\![H_B]\!]$ is twist-dominant. Then there exists H_A such that $[\![H_A]\!] \leq_t [\![H_B]\!]$. As noted above, the fold-equivalence class of the hyperplane H_A in $\mathbb{S}^\Pi$ consists of all hyperplanes $H_{A'}$ with $\max(A) \sim_f \max(A')$. Since $[\![H_B]\!]$ is cyclic, under the collapse map $c : \mathbb{S}^\Pi \to \mathbb{S}_\Gamma$, $[\![H_B]\!]$ maps to a single hyperplane labeled by a generator v. Hence, the fold-equivalence class of v is just $\{v\}$, and all the hyperplanes in $[\![H_B]\!]$ have v as the unique maximal element. Since $[\![H_A]\!] \leq_t [\![H_B]\!]$, this means that $\max(A) \leq_t v = \max(B)$.

Conversely, if there exists H_A such that $\max(A) \leq_t \max(B)$, then any $v \in \max(B)$ is twist-dominant and $[\![H_A]\!] \cup \mathrm{lk}(H_A) \subseteq [\![H_B]\!] \cup \mathrm{lk}(H_B)$. By Lemma 2.6, $\max(B) = \{v\}$ and the hyperplanes with $\max(H) = \{v\}$ coincide with the hyperplanes that split v, which are exactly those occurring along any characteristic cycle χ_v for v. It follows from Example 3.14 that $\mathbb{K}_v \cong \chi_v \times H_v$, but on the other hand this implies that

$$\mathbb{K}_v = \bigcup_{H \in [\![H_v]\!]} \kappa(H).$$

This proves that $[\![H_v]\!] = [\![H_B]\!]$ is cyclic as well, and therefore $[\![H_B]\!]$ is twist-dominant. $\square$

Since the definitions of cyclic hyperplane and the link of a hyperplane depend only the combinatorial structure of X, the following is an immediate corollary.

COROLLARY 4.5
Let $i: \mathbb{S}^\Pi \to \mathbb{S}^\Omega$ be an isomorphism of cube complexes. Then i preserves the twist and fold ordering on edge labels.

4.2. *Untwisted markings*

In this section, we show that we can detect twist-minimal hyperplanes in a Γ-complex X using only the action of A_Γ that an untwisted marking $h\colon X \to \mathbb{S}_\Gamma$ induces on the CAT(0) space $\widetilde{X}$.

We begin by recalling some basic facts about untwisted markings. Define $U^0(A_\Gamma)$ to be the subgroup of $U(A_\Gamma)$ generated by inversions, folds, and partial conjugations.

LEMMA 4.6
For any $v \in V$, both $A_{\mathrm{lk}(v)}$ and $A_{\mathrm{dlk}(v)}$ are invariant up to conjugacy under the action of $U^0(A_\Gamma)$.

Proof
Let $\Delta \subseteq \Gamma$ be any subgraph. We claim that the special subgroup $A_{\mathrm{lk}(\Delta)}$ is invariant up to conjugacy under $U^0(A_\Gamma)$, where $\mathrm{lk}(\Delta) = \bigcap_{v \in \Delta} \mathrm{lk}(v)$. The lemma will follow by taking $\Delta = \{v\}$ and $\Delta = \mathrm{lk}(v)$, respectively. If $\mathrm{lk}(\Delta) = \emptyset$, then we set $A_\emptyset = \{1\}$ which is trivially invariant. Otherwise, assume that $\mathrm{lk}(\Delta) \neq \emptyset$. We consider each type of generator of $U^0(A_\Gamma)$. Clearly, inversions preserve $A_{\mathrm{lk}(\Delta)}$. If $v \in \mathrm{lk}(\Delta)$ and $v <_f w$, then $\Delta \subseteq \mathrm{lk}(v) \subseteq \mathrm{lk}(w)$, hence $w \in \mathrm{lk}(\Delta)$ as well. It follows that $A_{\mathrm{lk}(\Delta)}$ is also invariant under the fold sending v to vw. Finally, consider a partial conjugation by $w \in V$. If Δ is not contained in $\mathrm{st}(w)$, then there exists $v \in \Delta \setminus \mathrm{st}(w)$, hence $\mathrm{lk}(\Delta) \setminus \mathrm{st}(w)$ is contained in the same component of $\Gamma \setminus \mathrm{st}(w)$ as v. Hence a partial conjugation by w preserves $A_{\mathrm{lk}(\Delta)}$ up to conjugacy. On the other hand, if $\Delta \subseteq \mathrm{st}(w)$, then either $w \in \Delta$, whence $\mathrm{lk}(\Delta) \subset \mathrm{st}(w)$, or $w \in \mathrm{lk}(\Delta)$. Either way, any partial conjugation by w preserves $A_{\mathrm{lk}(\Delta)}$, and the claim is proved. $\square$

The next lemma implies that after changing the collapse map, we can always assume that the marking lies in $U^0(A_\Gamma)$.

LEMMA 4.7
Let X be a Γ-complex. Every untwisted marking $h\colon X \to \mathbb{S}_\Gamma$ is a Γ-collapse map c followed by an element of $U^0(A_\Gamma)$,

Proof
By definition, h is untwisted if there is an isomorphism $i : X \cong \mathbb{S}^\Pi$ for some Π so that the composition $h \circ c^{-1}$ (where $c = c_\pi i$ and c^{-1} is a homotopy inverse for c) induces an element $\varphi \in U(A_\Gamma)$ on $\pi_1(\mathbb{S}_\Gamma) = A_\Gamma$. This condition is independent of the choice of i. The subgroup $U(A_\Gamma)$ is generated by inversions, partial conjugations, elementary (right and left) folds, and graph automorphisms. Any product of these is equal to a product with a single graph automorphism σ as the initial element. The automorphism σ permutes V, sends a Γ-Whitehead partition $\mathcal{P}$ to $\sigma\mathcal{P}$, and induces an isomorphism $\mathbb{S}^\Pi$ to $\mathbb{S}^{\sigma\Pi}$, so the composition of the initial Γ-collapse map $c = c_\pi i$ with σ is itself a Γ-collapse map, and the rest of the factors are inversions, partial conjugations, and elementary folds. □

Now let H be a hyperplane in a Γ-complex X, fix an untwisted marking h, and let $\bar{h}\colon \mathbb{S}_\Gamma \to X$ be a homotopy inverse for h, so that $g \in A_\Gamma = \pi_1(\mathbb{S}_\Gamma)$ acts on $\widetilde{X}$ by the deck transformation $\bar{h}_*(g) \in \pi_1(X)$. Define

$$\mathrm{split}_h(H) = \big\{v \in V \mid \text{an axis for } \bar{h}_*(v) \text{ crosses a lift of } H\big\},$$

and let $\max_h(H)$ denote the set of maximal elements in $\mathrm{split}_h(H)$.

A special case is when $h = c$ is just a collapse map. In this setting, the set of elements of $\max_c(H)$ all belong to the same fold-equivalence class of Γ. In the next lemma, we will see that the elements of $\max_h(H)$ also all belong to the same fold-equivalence class and moreover, that this equivalence class is actually independent of the marking h up to graph automorphisms.

LEMMA 4.8
Let X be a Γ-complex, let $h\colon X \to \mathbb{S}_\Gamma$ be an untwisted marking, and let m_h be any element of $\max_h(H)$.

(1) *If $v \in \mathrm{split}_h(H)$, then $v \leq_f m_h$.*

(2) *There is a Γ-collapse map c such that $m_c \sim_f m_h$ for any $m_c \in \max_c(H)$.*

Thus the maximal elements in $\mathrm{split}_h(H)$ *all lie in the fold-equivalence class of* $\max_c(H)$.

Proof
By Lemma 4.7, we can write

$$h = \varphi_n \circ \cdots \circ \varphi_1 \circ c,$$

where $c = c_\pi i : X \cong \mathbb{S}^\Pi \to \mathbb{S}$ is a Γ-collapse map and each φ_j induces an inversion, partial conjugation, or elementary fold. Let $\bar{h}$ denote the homotopy inverse of h. By Lemma 4.6, $U^0(A_\Gamma)$ preserves the special subgroup $A_{\mathrm{dlk}(v)}$ up to conjugacy for every $v \in V$, and by Corollary 3.16, the subcomplex $\widetilde{X}_{\mathrm{dlk}(v)} = i^{-1}(\widetilde{\mathbb{K}}_{\mathrm{dlk}(v)})$ contains an axis for every element of this subgroup. Thus, some translate of $\widetilde{X}_{\mathrm{dlk}(v)}$ contains an axis for $\bar{h}_*(v)$. It follows that every hyperplane H that splits $\bar{h}_*(v)$ has a lift that is dual to an edge in $\widetilde{X}_{\mathrm{dlk}(v)}$, or in other words, if $v \in \mathrm{split}_h(H)$, then $v \leq_f m_c$. Thus (1) will follow immediately from (2).

We will prove (2) by induction on n. By definition, if $v \in \mathrm{split}_h(H)$, then the image of an $\bar{h}_*(v)$-axis crosses H at least once. For the purpose of this proof, we will need to keep track of more information about how many times it crosses H. Begin by choosing an orientation for H, or equivalently, for a dual edge to H. (H is orientable since X is a special cube complex.) If p is an edge path in X, then we define the *net crossing number* of p with H to be

$$n(p, H) = \#\text{ positive crossings of } H - \#\text{ negative crossings of } H.$$

Note that two paths that are homotopic rel endpoints have the same net crossing number with respect to any hyperplane. For a generator $v \in V$, set $n_h(v, H) = n(p_v, H)$, where p_v is some (hence any) loop in X representing $\bar{h}_*(v)$. This is independent of basepoint since changing basepoints conjugates p_v by a path connecting the two basepoints and hence leaves the net crossing number unchanged. In particular, $n_h(vw, H) = n_h(v, H) + n_h(w, H)$.

Note that if $n_h(v, H) \neq 0$, then v necessarily lies in $\mathrm{split}_h(H)$, but the converse need not be true. In addition to property (2), we will prove by induction that the following property holds:

(3) For some $v \in \max_h(H)$, $n_h(v, H) \neq 0$.

For $n = 0$, $h = c$, so (2) is true trivially and for any $v \in \max_h(H)$, a characteristic cycle for v crosses H exactly once, so $n_h(v, H) = \pm 1$.

Now set $h' = \varphi_{n-1} \circ \cdots \circ \varphi_1 \circ c$, with homotopy inverse $\bar{h}'$, and assume by induction that (2) and (3) hold for h'. If φ_n is an inversion $v \mapsto v^{-1}$, then $c \circ \bar{h}$ and $c \circ \bar{h}'$ agree on every generator except v. Furthermore, $\bar{h}'_*(v)$ and $\bar{h}_*(v) = \bar{h}'_*(v^{-1}) = \bar{h}'_*(v)^{-1}$ have the same axis in $\widetilde{X}$, so there is no change in which hyperplanes this axis crosses hence no change in the splitting set. Only the sign of the net crossing numbers with these hyperplanes change.

If φ_n is a partial conjugation, then $c \circ \bar{h}_*(v)$ is conjugate to $c \circ \bar{h}'_*(v)$ for every generator v, so an axis for $\bar{h}_*(v)$ is just a translate of an axis for $\bar{h}'_*(v)$. Thus, the former crosses some lift of H if and only if the latter crosses some lift of H, and again there is no change is the splitting set or the net crossing numbers for H.

It remains to consider the case that φ_n is a right fold $v \mapsto vw^{-1}$ for some $w \geq_f v$ (the case of left folds is symmetric). Again $c \circ \bar{h}_*$ and $c \circ \bar{h}'_*$ agree on every generator except v, and $c \circ \bar{h}_*(v) = c \circ \bar{h}'_*(vw)$. So the only possible change is that after composing with φ_n, v may be added to or removed from $\operatorname{split}_{h'}(H)$ and the net crossing of v with H may change.

Suppose that v is in $\operatorname{split}_h(H)$, but not in $\operatorname{split}_{h'}(H)$. By induction, we know that $m_{h'} \sim_f m_c$, and as observed above, $v \leq_f m_c$. Thus, $v \leq_f m_{h'}$, so adding v to $\operatorname{split}_{h'}(H)$ does not change its maximal equivalence class and (2) and (3) remain valid.

Next, suppose that $v \in \operatorname{split}_{h'}(H)$. If $\operatorname{split}_{h'}(H)$ contains more than one maximal element with nonzero net crossing number, then removing v from $\operatorname{split}_{h'}(H)$ or changing its net crossing number will again preserve properties (2) and (3).

Thus, we need only consider the case where v is the unique element of $\max_{h'}(H)$ with $n_{h'}(v, H) \neq 0$. Since $n_{h'}(vw, H) = n_{h'}(v, H) + n_{h'}(w, H)$, either $n_{h'}(w, H)$ or $n_{h'}(vw, H)$ must also be nonzero. In the former case, w lies in $\operatorname{split}_{h'}(H)$ and since $v \leq_f w$, this contradicts our assumption that v is the unique maximal element with nonzero net crossing number. In the latter case, since $n_h(v, H) = n_{h'}(vw, H) \neq 0$, we conclude that v is also in $\operatorname{split}_h(H)$ and its net crossing number remains nonzero, so (2) and (3) still hold for h. This completes the induction. □

Remark 4.9
Suppose that v is twist-dominant. Then there are no elements $w \neq v$ with $v \leq_f w$. It follows that any element of $U^0(A_\Gamma)$ takes v to a conjugate of itself or its inverse, so the image in X of an axis for $\bar{h}_*(v)$ is the same for every marking h as in the lemma. Moreover, any hyperplane H crossed by this axis has $\max_h(H) = \{v\}$.

More generally, if we drop the assumption that $h^{-1}c \in U^0(A_\Gamma)$, then we have the following corollary.

COROLLARY 4.10
Let X be a Γ-complex, and let $h, h' \colon X \to \mathbb{S}_\Gamma$ be untwisted markings. Then $[\max_{h'}(H)] = \sigma[\max_h(H)]$ *for some graph automorphism σ.*

Proof
We have $h' = \psi h$ for some $\psi \in U(A_\Gamma)$. Write $\psi = \varphi \circ \sigma$, where σ is a graph

automorphism and $\varphi \in U^0(A_\Gamma)$. Then $\max_{\sigma h}(H) = \sigma \max_h(H)$, so by Lemma 4.8, $[\max_{h'}(H)] = [\max_{\sigma h}(H)] = \sigma[\max_h(H)]$. □

COROLLARY 4.11

Let H be a hyperplane in a Γ-complex X, and let $h\colon X \to \mathbb{S}_\Gamma$ be an untwisted marking. Then H is twist-minimal (resp., twist-dominant) if and only if $[\max_h(H)]$ is twist-minimal (resp., twist-dominant).

Proof

Choose an isomorphism $X \cong \mathbb{S}^\Pi$, and let $h = c_\pi$. Then $H = H_A$ for some $A \in V \cup \Pi$ and $[\max_h(H_A)] = [\max(A)]$. □

Definition 4.12

Given an untwisted marking $h\colon X \to \mathbb{S}_\Gamma$ and a generator $v \in V$, we define $\mathrm{Min}_h(v)$ to be $\mathrm{Min}(\bar{h}_*(v)) \subset \widetilde{X}$. Similarly, we define the branch locus $\mathrm{br}_h(v)$ to be the set of points in $\mathrm{Min}_h(v)$ whose link in $\widetilde{X}$ strictly larger than the link in $\mathrm{Min}_h(v)$.

If we choose an identification of X with $\mathbb{S}^\Pi$, then in terms of this definition $\mathrm{Min}(v) = \mathrm{Min}_{c_\pi}(v)$ and $\mathrm{br}(v) = \mathrm{br}_{c_\pi}(v)$, where $c_\pi\colon \mathbb{S}^\Pi \to \mathbb{S}_\Gamma$ is the standard collapse map. Using Lemma 4.8, we can identify when a hyperplane is contained in the minset for a general untwisted marking.

PROPOSITION 4.13

Let $h\colon X \to \mathbb{S}_\Gamma$ be an untwisted marking, let H be a hyperplane of X, and let $v \in$ $\mathrm{split}_h(H)$. Then $v \in \max_h(H)$ if and only if there is a lift $\widetilde{H}$ contained in $\mathrm{Min}_h(v)$, and in this case, both components of $\partial\kappa(\widetilde{H})$ contain points in $\mathrm{br}_h(v)$.

Proof

It is easily seen that this property is preserved by graph automorphisms, so it suffices to consider the case where $hc^{-1} \in U^0(A_\Gamma)$ for some collapse map c. Fix an identification of X with $\mathbb{S}^\Pi$, and let $c = c_\pi$. Then by Lemma 4.8, for any hyperplane H we have $[\max_h(H)] = [\max_c(H)]$. Consider the subspace $\widetilde{\mathbb{K}}_v = \widetilde{\mathbb{K}}_{\mathrm{dlk}(v)} \times \widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$. By Lemma 4.6, since $U^0(A_\Gamma)$ preserves the subgroups $A_{\mathrm{dlk}(v)}$ and $A_{\mathrm{lk}(v)}$ up to conjugacy, taking a translate if necessary, we may assume that $\widetilde{\mathbb{K}}_{\mathrm{dlk}(v)}$ contains an axis for $\bar{h}_*(v)$ (see Corollary 3.16). Call this axis α_v^h. Then $\widetilde{\mathbb{K}}_v$ contains $\mathrm{Min}_h(v) = \alpha_v^h \times \widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$.

By assumption, the axis α_v^h crosses some lift $\widetilde{H}$ of H. Let $m \in \max_c(H)$. If v is not maximal in $\mathrm{split}_h(H)$, then $v <_f m$. In this case, $\widetilde{H}$ is isomorphic to $\widetilde{\mathbb{K}}_{\mathrm{lk}(m)}$ which is strictly bigger than $\widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$, hence $\widetilde{H}$ is not contained in $\mathrm{Min}_h(v)$.

Conversely, if v is maximal, then $v \sim_f m$ and we can identify $\widetilde{H} = \widetilde{\mathbb{K}}_{\mathrm{lk}(v)} = \widetilde{\mathbb{K}}_{\mathrm{lk}(m)}$. It follows that the entire carrier $\kappa(\widetilde{H})$ is contained in $\mathrm{Min}_h(v)$. In addition, $\widetilde{\mathbb{K}}_{\mathrm{dlk}(v)}$ also contains an axis α_m for m with respect to the marking c and $\mathrm{Min}_c(m) = \alpha_m \times \widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$. Since $m \in \max_c(H)$, this minset contains $\kappa(\widetilde{H})$ and by Proposition 3.18, both components of $\partial\kappa(\widetilde{H})$ contain points in the branch locus of $\mathrm{Min}_c(m)$. Since $\mathrm{Min}_c(m)$ and $\mathrm{Min}_h(v)$ are both metrically the product of a real line with $\widetilde{\mathbb{K}}_{\mathrm{lk}(v)}$, any point in $\partial\kappa(\widetilde{H})$ that is branch for one of these minsets is branch for the other. Thus both components of $\partial\kappa(\widetilde{H})$ also contain points in $\mathrm{br}_h(v)$. □

5. Parallelotope structures on blowups

In this section, we consider blowups $\mathbb{S}^\Pi$ as metric objects, where we now allow some of the cubes in $\mathbb{S}^\Pi$ to be skewed in certain directions, so that edges spanning a "cube" are no longer necessarily orthogonal. We call these *skewed blowups*.

An n-dimensional *Euclidean parallelotope* F is a metric space isometric to the image of the unit cube $[0,1]^n \subset \mathbb{R}^n$ under some element of $\mathrm{GL}(n,\mathbb{R})$. If e is an edge of F, then the *midplane* H_e is the convex hull of the midpoints of the edges parallel to e. A parallelotope F is an *orthotope* if any two edges at a vertex are orthogonal or, equivalently, the dihedral angle between any two midplanes is a right angle.

5.1. *Allowable parallelotope structures*

In a blowup $\mathbb{S}^\Pi$ every edge e has a label, that is, $e = e_A$, where $A \in V \cup \Pi$. By Corollary 3.3, there is a square in $\mathbb{S}^\Pi$ spanned by e_A and e_B if and only if A and B are adjacent. We will say that A, B are *twist-related* if $\max(A) \leq_t \max(B)$ or vice versa.

Definition 5.1

Let $\mathbf{c}$ be a maximal cube of $\mathbb{S}^\Pi$ with outgoing edges $e_1,\dots,e_n$ at a vertex $p \in \mathbf{c}$. Let A_i be the label of e_i, choose $v_i \in \max(A_i)$, and let $\mathrm{st}^+(v_i) = \{v_i\} \cup \mathrm{lk}^+(v_i)$. Given $d_\mathbf{c}$ a parallelotope metric on $\mathbf{c}$, we realize $d_\mathbf{c}$ via an embedding $\rho\colon \mathbf{c} \hookrightarrow \mathbb{R}^n$ which sends p to 0. Regarding $\rho(e_i)$ as vectors in $\mathbb{R}^n$, set

$$K_i = \text{ the subspace of } \mathbb{R}^n \text{ spanned by } \rho(e_k) \text{ with } v_k \in \mathrm{st}^+(v_i) \text{ and}$$
$$L_{ij} = K_i \cap K_j.$$

The metric $d_\mathbf{c}$ on $\mathbf{c}$ is *allowable* if whenever v_i and v_j are not twist-related, then

$$L_{ij}^\perp \cap K_i \text{ is orthogonal to } L_{ij}^\perp \cap K_j.$$

Note that if v_i, v_j are not twist-related, then $\mathrm{st}^+(v_i) \cap \mathrm{st}^+(v_j) = \mathrm{lk}^+(v_i) \cap \mathrm{lk}^+(v_j)$ so in this case, L_{ij} is the subspace spanned by the $\rho(e_k)$ with $v_k \in \mathrm{lk}^+(v_i) \cap \mathrm{lk}^+(v_j)$.

LEMMA 5.2
An allowable parallelotope metric is determined by edge lengths and the angles between twist-related edges.

Proof
Any parallelotope metric is determined by edge lengths and the angles between edges, so we must show that in an allowable metric, the angles between non-twist-related edges are determined by those between twist-related edges.

Suppose that e_i and e_j have labels that are not twist-related, so L_{ij} is the span of edges e_k with $v_k \in \mathrm{lk}^+(v_i) \cap \mathrm{lk}^+(v_j)$. The proof is by induction on $\dim(L_{ij})$. If $\dim(L_{ij}) = 0$, then the angle between e_i and e_j must be $\frac{\pi}{2}$. If $\dim(L_{ij}) > 0$, write $e_i = e_i' + \ell_i$, where ℓ_i is the orthogonal projection of e_i onto L_{ij} and e_i' is the projection onto $L_{ij}^{\perp} \cap K_i$. Similarly, write $e_j = e_j' + \ell_j$. Then $e_i \cdot e_j = (e_i' + \ell_i) \cdot (e_j' + \ell_j) = \ell_i \cdot \ell_j$. Since ℓ_i and ℓ_j are linear combinations of the $e_k \in L_{ij}$, this dot product is determined by the dot products of these e_k. The dot products $e_k \cdot e_k$ are the squares of the lengths of the e_k, which are given. If $v_k \in \mathrm{lk}^+(v_i) \cap \mathrm{lk}^+(v_j)$, then $\mathrm{lk}^+(v_k)$ is strictly contained in $\mathrm{lk}^+(v_i)$, so for two edges $e_k \neq e_l \in L_{ij}$, the subspace L_{kl} has dimension strictly smaller than $\dim(L_{ij})$. Thus, by induction, the angle between edges in L_{ij} is determined by edge lengths and the angles between twist-related edges, so the same holds for e_i, e_j. □

Definition 5.3
An *allowable parallelotope structure* $\mathcal{F}$ on $\mathbb{S}^{\Pi}$ is an assignment of a parallelotope metric to each cube $\mathbf{c}$ of $\mathbb{S}^{\Pi}$ such that
(1) the metric on each maximal cube is allowable,
(2) if $\mathbf{c}'$ is a face of $\mathbf{c}$, then the metric on $\mathbf{c}'$ is the restriction of the metric on $\mathbf{c}$, and
(3) if $\max(A) = \{v\}$ is twist-dominant, then for any B adjacent to A, the angle between e_A and e_B is equal to the angle between e_v and e_B.

The parallelotope structure in which every k-cube is isometric to the Euclidean cube $[0,1]^k$ is clearly allowable; it will be called the *standard structure* and denoted $\mathcal{E}$. If all parallelotopes in $\mathcal{F}$ are orthotopes, then the structure $\mathcal{F}$ will be called *rectilinear*. These too are clearly allowable.

If A and B are adjacent labels in $\Pi \cup V$, then there is at least one parallelogram $F \in \mathcal{F}$ with edges labeled e_A and e_B. If $\mathcal{F}$ is allowable, then condition (2) guarantees that the angle between these edges is the same for any such F, and we will denote this angle by $\alpha_{A,B}$. By Lemma 5.2, the entire structure $\mathcal{F}$ is determined by the lengths of the edges e_A and the angles $\alpha_{A,B}$ for twist-related A, B.

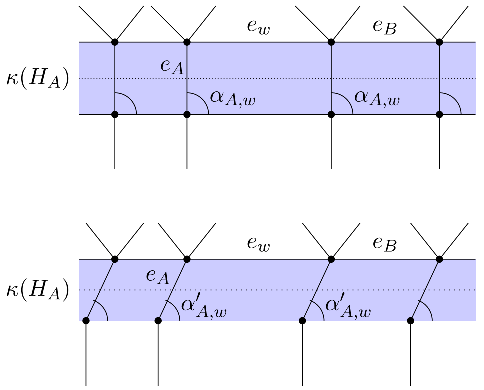

Figure 11. (Color online) Rotating H_A in the direction $w \in \mathrm{lk}^+(A)$. Here B splits w.

An allowable parallelotope structure $\mathcal{F}$ induces a (path) metric $d_{\mathcal{F}}$ on $\mathbb{S}^\Pi$. Different parallelotope structures may induce the same metric on $\mathbb{S}^\Pi$; for example, if $\mathbb{S}^\Pi = \mathbb{S}_\Gamma$ is an n-torus consisting of a single parallelotope F with sides identified, then changing F by any element of $\mathrm{GL}(n,\mathbb{Z})$ results in the same metric $d_{\mathcal{F}}$.

Note that an edge path which was convex in the standard cube complex structure $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{E}})$ is no longer necessarily convex in the metric space $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$. We define a *hyperplane* H_A in $(\mathbb{S}^\Pi, \mathcal{F})$ to be the set of midplanes dual to edges with label A. This is the usual notion of hyperplane if $\mathcal{F} = \mathcal{E}$, but for arbitrary $\mathcal{F}$ lifts of hyperplanes are no longer necessarily convex in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$.

5.2. *Rotating a hyperplane in* $\mathbb{S}^\Pi$

Definition 5.4
Suppose that $(\mathbb{S}^\Pi, \mathcal{F})$ is an allowable parallelotope structure. Let $A \in \Pi \cup V$, let H_A be the hyperplane in $\mathbb{S}^\Pi$ labeled A, and let $v \in \max(A)$ and $w \in \mathrm{lk}^+(v)$. Then *rotating* H_A *in the direction of* w means changing the angle $\alpha_{A,w}$ to $\alpha'_{A,w}$, so that for every B that splits w, the angle between the edges e_A and e_B is $\alpha'_{A,w}$. More generally, *rotating* H_A means rotating it in the direction of one or more $w \in \mathrm{lk}^+(v)$. The length of the edge e_A remains unchanged under rotation (see Figure 11).

Rotating a hyperplane H_A in an allowable parallelotope structure $\mathcal{F}$ gives rise to a new parallelotope structure $\mathcal{F}'$ which still satisfies the first two conditions for allowability. This is because the subspaces K_i in the definition of an allowable parallelotope are unchanged by the rotation. However, if A is twist-dominant, then to achieve the third allowability condition, one needs to do comparable rotations to every hyperplane $H_{A'}$ with $\max(A') = \max(A) = \{v\}$.

Recall that we have a partial ordering on equivalence classes in V given by $[v] \leq [w]$ if $\mathrm{lk}(v) \subseteq \mathrm{st}(w)$. Choose a total ordering $\prec$ on V consistent with this partial ordering. Given a compatible collection of partitions Π, we can extend this to a total order on $\Pi \cup V$ satisfying $[\max(A)] < [\max(B)] \Rightarrow A \prec B$.

PROPOSITION 5.5
Every allowable parallelotope structure on $\mathbb{S}^\Pi$ can be obtained from an orthotope structure on $\mathbb{S}^\Pi$ by a sequence of rotations.

Proof
Suppose that $\mathcal{F}$ is an allowable parallelotope structure on $\mathbb{S}^\Pi$, and let $\alpha_{A,B}$ denote the angle between edges e_A, e_B for any adjacent pair A, B. Let $\mathcal{F}_0$ denote the rectilinear structure with the same edge lengths as $\mathcal{F}$. Using the total order $\prec$, we will rotate the hyperplanes in $\mathcal{F}_0$ in descending order and show inductively that after rotating H_A, we get a parallelotope structure on $\mathbb{S}^\Pi$ satisfying the following.

(i) The metric on each parallelotope is allowable and agrees on common faces.

(ii) For all $B, C \succeq A$ such that $\max(B) \leq_t \max(C)$, the angle between e_B and e_C equals $\alpha_{B,C}$.

Say by induction that we have rotated all the hyperplanes $H_{A'}$ with $A \prec A'$. Rotating H_A only changes the angles between e_A and other edges. By induction, condition (ii) is already satisfied whenever $B, C \succ A$. We now rotate H_A so that condition (ii) also holds when $A = B$, that is, when $A \prec C$ and $\max(A) \leq_t \max(C)$. As observed above, rotating preserves allowability of individual parallelotopes, and by definition it agrees on common faces, so condition (i) continues to hold.

At the end of this process, when we have rotated all the hyperplanes as needed, we arrive at a parallelotope structure in which the angles between any two edges e_A, e_B with $\max(A) \leq_t \max(B)$, agree with those in $\mathcal{F}$. This implies that this structure also satisfies the third condition for allowability. So by Lemma 5.2, it must in fact be equal to $\mathcal{F}$. ☐

PROPOSITION 5.6
Let $\mathcal{F}$ be an allowable parallelotope structure on $\mathbb{S}^\Pi$, and suppose that the induced path metric $d_{\mathcal{F}}$ is locally CAT(0). Suppose that $\mathcal{F}'$ is obtained from $\mathcal{F}$ by a hyperplane rotation. Then

(1) *$d_{\mathcal{F}'}$ is also locally CAT(0), and*

(2) *any twist-minimal hyperplane which is locally convex with respect to $d_{\mathcal{F}}$ remains locally convex with respect to $d_{\mathcal{F}'}$.*

Proof

The local geometry at a point p is determined by the geometry of its link. Thus, it suffices to show that rotating a single hyperplane H does not change the isometry type of links in $\mathbb{S}^{\Pi}$. The carrier $\kappa(H)$ either has two boundary components, each isometric to H, or (if the dual edge is a loop in $\mathbb{S}^{\Pi}$) these boundary components may be identified to each other. In either case, we will denote (the image of) this boundary by $\partial\kappa(H)$ and the interior by $\kappa^{\circ}(H) = \kappa(H) - \partial\kappa(H)$. Setting $Y = \mathbb{S}^{\Pi} - \kappa^{\circ}(H)$, we have

$$\mathbb{S}^{\Pi} = \kappa(H) \cup_{\partial\kappa(H)} Y.$$

Rotating H changes the parallelotope structure only on cubes meeting the interior of $\kappa(H)$, leaving those in $\partial\kappa(H)$ and Y unchanged. Hence, it suffices to show that if x is a vertex lying in $\partial\kappa(H)$, then the rotation does not change the induced metric on the link of x in $\kappa(H)$ (though it may change the metric on individual simplices in that link).

To see this, note that if H is dual to e_A with $v \in \max(A)$, then the carrier of H decomposes as $\kappa(H) = e_A \times \mathbb{K}_{\mathrm{lk}(v)}$, and by Proposition 3.15, $\mathbb{K}_{\mathrm{lk}(v)}$ contains a characteristic cycle for every $w \in \mathrm{lk}(v)$. Consider the subspace $\mathbb{K}_{\mathrm{lk}^+(v)} \subset \mathbb{K}_{\mathrm{lk}(v)}$ spanned by the characteristic cycles for $w \in \mathrm{lk}^+(v)$. Elements of $\mathrm{lk}^+(v)$ commute and are twist-dominant, so this subspace is a torus with a flat metric. Moreover, as elements of $\mathrm{lk}^+(v)$ commute with every element of $\mathrm{lk}(v)$ we have a further (combinatorial) decomposition $\kappa(H_A) = e_A \times \mathbb{K}_{\mathrm{lk}^+(v)} \times \mathbb{K}_{\mathrm{lk}(v)\setminus \mathrm{lk}^+(v)}$. The edge e_A can only rotate in the direction of $\mathbb{K}_{\mathrm{lk}^+(v)}$. Thus, viewing $e_A \times \mathbb{K}_{\mathrm{lk}^+(v)}$ geometrically as the product of an interval and a torus, this rotation changes only the width of the interval. In particular, the rotation does not change the local geometry of $\kappa(H)$. This proves (1).

For (2), let L be any twist-minimal hyperplane of $\mathbb{S}^{\Pi}$, and let p be a point of L. Then as was just shown, the local metrics at p with respect to $d_{\mathcal{F}}$ and $d_{\mathcal{F}'}$ are the same. Since L is twist-minimal, it is preserved setwise by rotation. Thus if L was locally convex before rotation, it remains locally convex afterward. □

COROLLARY 5.7

If $\mathcal{F}$ is an allowable parallelotope structure on $\mathbb{S}^{\Pi}$, then the induced path metric $d_{\mathcal{F}}$ is locally CAT(0).

Proof

This follows from Propositions 5.6(1) and 5.5, since any orthotope structure is CAT(0) by Gromov's link condition. □

As noted above, subcomplexes which are locally convex in $(\mathbb{S}^\Pi, \mathcal{E})$ may no longer be convex in a general allowable parallelotope structure $(\mathbb{S}^\Pi, \mathcal{F})$. The following lemma specifies two exceptions that will be important in the sequel.

LEMMA 5.8
Suppose that $\mathcal{F}$ is an allowable parallelotope structure on $\mathbb{S}^\Pi$.

- *Let $v \in V$ be twist-dominant. Then any lift of a characteristic cycle for v in $\mathbb{S}^\Pi$ is convex in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$.*
- *Let $A \in V \cup \Pi$ be a label with $v \in \max(A)$. If v is twist-minimal, then any lift of the hyperplane H_A is convex in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$.*

Proof
If v is twist-dominant, then Definition 5.3(3) guarantees that consecutive edges in a characteristic cycle for v have angle π in $(\mathbb{S}^\Pi, d_{\mathcal{F}})$. Since $(\mathbb{S}^\Pi, d_{\mathcal{F}})$ is locally CAT(0), the lift of the characteristic cycle to $\widetilde{\mathbb{S}}^\Pi$ is geodesic and convex. The second statement follows from Propositions 5.5 and 5.6(2). □

5.3. Straightening an allowable parallelotope structure
In this section, we show how to straighten an allowable parallelotope structure $\mathcal{F}$ on $\mathbb{S}^\Pi$ to obtain an orthotope structure, while maintaining allowability throughout the process.

Remark 5.9
It will be convenient to describe the straightening process in terms of what it does to the edges of $\mathbb{S}^\Pi$, rather than its dual hyperplanes. In particular, if an edge e_A is dual to a hyperplane H_A and $m \in \max(A)$, then we say that m is a *maximal element of* e_A.

We begin by straightening a single parallelotope $F \in \mathcal{F}$. The straightening procedure for F will depend only on the equivalence classes $[\max(A)]$ of the edges e_A in F. Therefore, it suffices to describe the straightening process in the case where all edges are labeled e_v for some $v \in V$.

Fix a vertex x in F with all angles acute or right. Let E be the set of edges emanating from x. We can view E as a set of n linearly independent vectors in the positive orthant of $\mathbb{R}^n$. Let $\prec$ be a total ordering on V as described in Section 5.2. For each edge e_v, define subspaces

$$K_v = \text{span of } \{e_w \in E \mid w \in \mathrm{st}^+(v)\},$$
$$K_v^\prec = \text{span of } \{e_w \in E \mid w \in \mathrm{st}^+(v), v \prec w\}.$$

For $e_v, e_w \in E$, set $L_{v,w} = K_v \cap K_w$. Recall that F is allowable if whenever v, w are not twist-related, $K_v \cap L_{v,w}^{\perp}$ is orthogonal to $K_w \cap L_{v,w}^{\perp}$.

Now define a new basis $\{b_v\}$ for $\mathbb{R}^n$ as follows. For each v, let b_v be the unit normal vector to $K_v^{\prec}$ in $K_v^{\prec} \oplus \langle e_v \rangle$. In the case where $K_v^{\prec}$ is empty, b_v is just the unit vector in the direction of e_v. With respect to this basis, we have

$$e_v = r_v b_v + \sum_w r_{v,w} b_w$$

for some $r_v > 0$, $r_{v,w} \geq 0$, where the sum is taken over all w with $e_w \in K_v^{\prec}$. In particular, this set of vectors $\{b_w\}$ is also a basis for $K_v^{\prec}$.

LEMMA 5.10
F is allowable if and only if for any two edges $e_v, e_w \in E$, b_v is orthogonal to b_w. That is, the vectors $\{b_v\}$ span an orthotope.

Proof
Assume that F is allowable. Suppose that v and w are twist-related, and say that $v \prec w$. In this case, $K_w^{\prec} \oplus \langle e_w \rangle \subset K_v^{\prec}$, so by definition, b_w lies in $K_v^{\prec}$ and hence it is orthogonal to b_v. (This is always true, even without assuming allowability.)

So now suppose that v and w are not twist-related. Then

$$L_{v,w} = K_v \cap K_w = K_v^{\prec} \cap K_w^{\prec}$$

since any $u \in \mathrm{st}^+(v) \cap \mathrm{st}^+(w)$ must be strictly greater than either v or w with respect to the ordering $\leq_t$, and hence also with respect to $\prec$. Since $b_v \in (K_v^{\prec})^{\perp} \subset L_{v,w}^{\perp}$, and $b_w \in (K_w^{\prec})^{\perp} \subset L_{v,w}^{\perp}$, the allowability condition implies that b_v and b_w are orthogonal.

Conversely, assume that all of the b_* vectors are orthogonal to each other. For v, w not twist-related, a basis for $K_v \cap L_{v,w}^{\perp}$ is given by the set of b_u with $v \leq_t u$ and $w \not\leq_t u$, and similarly, a basis for $K_w \cap L_{v,w}^{\perp}$ is given by the set of b_z with $w \leq_t z$ and $v \not\leq_t z$. These sets are disjoint, and any two such b_u and b_z are orthogonal, so $K_v \cap L_{v,w}^{\perp}$ is orthogonal to $K_w \cap L_{v,w}^{\perp}$ as required. □

Next we describe a process for straightening F. For $t \in [0, 1]$, set

$$e_v^t = s_t\Big(r_v b_v + t \sum_w r_{v,w} b_w\Big),$$

where $s_t \in \mathbb{R}^+$ is chosen so that $\|e_v^t\| = \|e_v\|$. Then $e_v^1 = e_v$ and $e_v^0 = \|e_v\| b_v$. Let F^t be the parallelotope spanned by $\{e_v^t\}$.

LEMMA 5.11

If F is allowable, then F^t is allowable for all $t \in [0,1]$, and F^0 is an orthotope.

Proof

At all times t, the subspaces $K_v^{\prec}(F^t)$ remain unchanged, that is, $K_v^{\prec}(F^t) = K_v^{\prec}(F)$ for all t, and likewise for $K_v^{\prec}(F^t) \oplus \langle e_v^t \rangle$. Hence the normal vectors b_v remain fixed throughout the process. By the previous lemma, F is allowable if and only if all of the b_v vectors are orthogonal, or equivalently, F^0 is an orthotope. $\square$

We now want to apply the straightening procedure simultaneously to all parallelotopes in $\mathcal{F}$. Suppose that two maximal parallelotopes F and F' share a face F_0 in $\mathbb{S}^{\Pi}$. If e_A is an edge lying in F_0, with $v \in \max(A)$, then for any w with $v <_t w$, F and F' must each contain an edge with maximal element w. Since w is twist-dominant, the allowability condition implies that these edges both lie along an axis for w, hence they are parallel. Since the straightening procedure on e_A depends only on these edges, it follows that the straightening in F and F' agree on this face. Moreover, the same argument applied to the edges with maximal element w shows that these edges remain parallel throughout the straightening process. Thus, we obtain a consistent straightening, $\mathcal{F}^t$, of the entire complex which remains allowable at all times t. We call $(\mathbb{S}^{\Pi}, \mathcal{F}^t)$ the *straightening path* for $(\mathbb{S}^{\Pi}, \mathcal{F})$.

6. The space of skewed Γ-complexes with untwisted markings

We are now ready to define a space $\mathcal{T}_\Gamma$ of skewed Γ-complexes with untwisted markings, that serves as an intermediary between Σ_Γ and the full outer space $\mathcal{O}_\Gamma$.

6.1. *Skewed Γ-complexes*

Let X be a Γ-complex, and let $\mathcal{F}$ be a parallelotope structure on X. Define $\mathcal{F}$ to be *allowable* if there is some isomorphism $\mathbb{S}^{\Pi} \cong X$ such that the pullback of $\mathcal{F}$ is an allowable parallelotope structure on $\mathbb{S}^{\Pi}$.

LEMMA 6.1

Allowability of a parallelotope structure on X is independent of the choice of isomorphism $\mathbb{S}^{\Pi} \cong X$.

Proof

By Corollary 4.5, the twist relation is independent of the isomorphism $\mathbb{S}^{\Pi} \cong X$. Hence if X is isomorphic to both $\mathbb{S}^{\Pi}$ and $\mathbb{S}^{\Pi'}$, then conditions (1), (2), and (3) of Definition 5.3 are satisfied by the pullback structure on $\mathbb{S}^{\Pi}$, if and only if they are satisfied for the pullback structure on $\mathbb{S}^{\Pi'}$, showing that allowability is also independent of the isomorphism. □

Definition 6.2
A *skewed* Γ*-complex* is a Γ-complex X together with an allowable parallelotope structure $\mathcal{F}$. If all of the parallelotopes $F \in \mathcal{F}$ are orthotopes, we will call $(X, \mathcal{F})$ a *rectilinear* Γ*-complex*, and if all parallelotopes are isometric to $[0,1]^k$, we will write $\mathcal{F} = \mathcal{E}$ and call $(X, \mathcal{E})$ a *standard* Γ*-complex*.

6.2. Definition of $\mathcal{T}_\Gamma$

We now add untwisted markings to skewed Γ-complexes to form a space $\mathcal{T}_\Gamma$.

Definition 6.3
A *marked*, skewed Γ-complex is a triple $(X, \mathcal{F}, h)$, where $(X, \mathcal{F})$ is a skewed Γ-complex and $h\colon X \to \mathbb{S}_\Gamma$ is an untwisted homotopy equivalence; that is, for any Γ-collapse map $c : X \to \mathbb{S}_\Gamma$, the composition $c \circ h^{-1} : \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ induces an element of $U(A_\Gamma)$ (where h^{-1} is a homotopy inverse to h). Two marked, skewed Γ-complexes $(X, \mathcal{F}, h)$ and $(X', \mathcal{F}', h')$ are equivalent if there is a combinatorial isometry $i : (X, \mathcal{F}) \to (X', \mathcal{F}')$ (i.e., a map which preserves both the combinatorial structure and the metric on each parallelotope) that commutes with the markings up to homotopy, that is, $h \simeq h' \circ i$.

The space $\mathcal{T}_\Gamma$ is the space of equivalence classes of marked skewed Γ-complexes with untwisted markings:

$$\mathcal{T}_\Gamma = \big\{\text{marked, skewed } \Gamma\text{-complexes } (X, \mathcal{F}, h) \mid h \text{ is untwisted}\big\}/\sim.$$

We will denote the equivalence class of $(X, \mathcal{F}, h)$ by $[X, \mathcal{F}, h]$.

Given a Γ-complex X and untwisted marking $h : X \to \mathbb{S}_\Gamma$, let $U_{X,h}$ denote the subset of $\mathcal{T}_\Gamma$ obtained by equipping X with all possible allowable parallelotope structures, that is,

$$U_{X,h} = \big\{[X, \mathcal{F}, h] \in \mathcal{T}_\Gamma\big\}.$$

We will call this a *cell* in $\mathcal{T}_\Gamma$. It comes equipped with a natural topology as a subspace of a Euclidean space determined by the parallelotopes in $\mathcal{F}$ and subject to the allowability conditions in Definition 5.3. Metrically, collapsing a hyperplane in X corresponds to letting the length of the dual edges go to zero. The closure of $U_{X,h}$ thus consists of the cells $U_{X',h'}$ such that there exists a hyperplane collapse map $c : X \to X'$ with h homotopic to $h' \circ c$. The topology on $\mathcal{T}_\Gamma$ is therefore determined as a complex of spaces comprised of the cells $U_{X,h}$, where cells are identified by collapse maps as just described. (For a more detailed description of complexes of spaces and their properties, see [25, Chapter 4.G].)

6.3. Contractibility of $\mathcal{T}_\Gamma$

We will show that $\mathcal{T}_\Gamma$ is contractible by finding a deformation retraction of $\mathcal{T}_\Gamma$ onto the subspace of rectilinear marked Γ-complexes; this is the space Σ_Γ defined in Section 2.6, which we know is contractible. In other words, we want to find a way to straighten marked, skewed Γ-complexes in a way that maintains allowability and extends to a continuous straightening of the whole of $\mathcal{T}_\Gamma$.

In order to straighten a skewed Γ-complex $(X, \mathcal{F})$, we choose an identification of X with $\mathbb{S}^\Pi$ for some Π and apply the straightening process described in Section 5.3. We need to show that this is independent of the isomorphism $X \cong \mathbb{S}^\Pi$. We note that the labeling on $\mathbb{S}^\Pi$ was used in the straightening process only to order the edges in K_v. By Corollary 4.5, any combinatorial isomorphism $i : \mathbb{S}^\Pi \to \mathbb{S}^{\Pi'}$ preserves the twist ordering $\leq_t$ on edge labels, so in fact we need only be concerned about what it does to the ordering $\prec$ within each twist equivalence class.

To address this problem, we will need to choose preferred representatives for points in $\mathcal{T}_\Gamma$. Let X be a Γ-complex, and let $h\colon X \to \mathbb{S}_\Gamma$ be an untwisted marking. By Lemma 4.7, there exists a blowup $\mathbb{S}^\Pi$ and an isomorphism of cube complexes $i : X \to \mathbb{S}^\Pi$ such that $c_\pi \circ i \circ h^{-1} \in U^0(A_\Gamma)$. Suppose that $j\colon X \to \mathbb{S}^\Omega$ is another such isomorphism.

LEMMA 6.4

Let $i\colon X \to \mathbb{S}^\Pi$ and $j\colon X \to \mathbb{S}^\Omega$ be as above. For any twist-dominant v, $j \circ i^{-1}$ takes edges with maximal element v to edges with maximal element v (cf. Remark 5.9).

Proof

For $i\colon X \to \mathbb{S}^\Pi$ and $j\colon X \to \mathbb{S}^\Omega$ as above, the composition $c_\omega \circ j \circ i^{-1} \circ c_\pi^{-1}$ induces an element of $U^0(A_\Gamma)$. Since any element of $U^0(A_\Gamma)$ takes every twist-dominant generator v to a conjugate of itself, the map $j \circ i^{-1} : \mathbb{S}^\Pi \to \mathbb{S}^\Omega$ takes an axis for v in $\widetilde{\mathbb{S}}^\Pi$ (with respect to the standard metric) to an axis for v in $\widetilde{\mathbb{S}}^\Omega$. Edges with maximal element v lie on such an axis, thus they map to edges with the same maximal element. □

We can now define the deformation retraction $R_t : \mathcal{T}_\Gamma \to \Sigma_\Gamma$ as follows. Let $(X, \mathcal{F}, h)$ represent a point in $\mathcal{T}_\Gamma$, and choose a cubical isomorphism $i : X \to \mathbb{S}^\Pi$ as in Lemma 6.4. Using this isomorphism, we can identify parallelotope structures on

$\mathbb{S}^\Pi$ with parallelotope structures on X. Thus the straightening path for $(\mathbb{S}^\Pi, \mathcal{F})$ gives a path in $\mathcal{T}_\Gamma$ defined by $R_t[X, \mathcal{F}, h] = [X, \mathcal{F}^t, h]$.

LEMMA 6.5
The deformation retraction $R_t : \mathcal{T}_\Gamma \to \Sigma_\Gamma$ is well defined and continuous.

Proof
The straightening path depends only on which edges in the parallelotope structure are twist-dominant. If $i: X \to \mathbb{S}^\Pi$ and $j: X \to \mathbb{S}^\Omega$ are two identifications of X with blowups, then by Lemma 6.4, $j \circ i^{-1}$ takes twist-dominant edges to twist-dominant edges, so for each t the straightening path induced by i is isometric to the straightening path induced by j.

It is clear from the definition of the straightening path that R_t is continuous on each cell $U_{X,h}$ of $\mathcal{T}_\Gamma$. It suffices to show that R_t is also continuous on the closure of each cell. The closure of $U_{X,h}$ consists of all the cells $U_{X',h'}$ such that there exists a collapse map $c : X \to X'$ with h homotopic to $h' \circ c$. Since straightening paths preserve edge lengths, a path $[X, \mathcal{F}^t, h]$ in $U_{X,h}$ will collapse to a path in $U_{X',h'}$ when the appropriate edge lengths go to zero. Moreover, since the straightening paths in every cell are defined using the same ordering $\prec$ on V, this path will agree with R_t on $U_{X',h'}$. □

In light of Corollary 2.20, we conclude the following.

COROLLARY 6.6
The space $\mathcal{T}_\Gamma$ is contractible.

7. Outer space $\mathcal{O}_\Gamma$

7.1. Definition of $\mathcal{O}_\Gamma$ and the map $\Theta: \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$

We now define a new space $\mathcal{O}_\Gamma$ by forgetting the combinatorial structure on skewed Γ-complexes and allowing arbitrary markings. Thus a point in $\mathcal{O}_\Gamma$ is an equivalence class of triples (Y, d, f) such that

- (Y, d) is a locally CAT(0) metric space that is isometric to $(\mathbb{S}^\Pi, d_{\mathcal{F}})$ for some skewed blowup $(\mathbb{S}^\Pi, \mathcal{F})$,
- $f : Y \to \mathbb{S}_\Gamma$ is a homotopy equivalence, and
- $(Y, d, f) \sim (Y', d', f')$ if there is an isometry $i : (Y, d) \to (Y', d')$ with $f' \circ i \simeq f$.

The full group $\mathrm{Out}(A_\Gamma)$ acts on the left on $\mathcal{O}_\Gamma$ by changing the marking f.

PROPOSITION 7.1
The action of $\mathrm{Out}(A_\Gamma)$ *on* $\mathcal{O}_\Gamma$ *has finite stabilizers.*

Proof
The element of $\mathrm{Out}(A_\Gamma)$ induced by a homotopy equivalence $g : \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ fixes the point $[Y, d, f]$ if and only if $f^{-1} \circ g \circ f$ is homotopic to an isometry of (Y, d). Thus, the stabilizer of a point $[Y, d, f]$ can be identified with the group of isometries of Y up to homotopy.

Since $(\mathbb{S}^\Pi, \mathcal{F})$ has no free faces, each (Y, d) has the geodesic extension property. It follows from [8, Lemma II.6.16] that the minset of the center of A_Γ is all of Y, so by the flat torus theorem (Theorem II.7.1 there) Y splits as a product $Y = Y_0 \times T_{\mathcal{Z}(A_\Gamma)}$, where $T_{\mathcal{Z}(A_\Gamma)}$ is a torus of dimension equal to the rank of the center $\mathcal{Z}(A_\Gamma)$. Moreover, by [8, Theorem II.6.17], $\mathrm{Isom}(Y)$ is a topological group with finitely many components, and the connected component of the identity is generated by translations of $T_{\mathcal{Z}(A_\Gamma)}$. As every such translation is homotopic to the identity, the group of isometries of Y up to homotopy is a quotient of the group of path components of $\mathrm{Isom}(Y)$, hence finite as claimed. □

In fact, as shown by Bregman in [7], the group of path components of $\mathrm{Isom}(Y)$ injects into $\mathrm{Out}(A_\Gamma)$.

To finish the proof of Theorem 1.1 we need to show that $\mathcal{O}_\Gamma$ is contractible. To do this, we define a map $\Theta \colon \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$ by forgetting the parallelotope structure on $X \in \mathcal{T}_\Gamma$ and just viewing it as a CAT(0) metric space. The remainder of this section is devoted to proving the following theorem.

THEOREM 7.2
The map $\Theta \colon \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$ *is a fibration with contractible fibers. Hence* $\mathcal{O}_\Gamma$ *is contractible.*

Since the inclusion map $\Sigma_\Gamma \hookrightarrow \mathcal{T}_\Gamma$ is a homotopy equivalence by Lemma 6.5, the map

$$\Sigma_\Gamma \hookrightarrow \mathcal{T}_\Gamma \xrightarrow{\Theta} \mathcal{O}_\Gamma$$

that forgets the orthotope structure on X is also a homotopy equivalence. We will show in Corollary 7.17 below that this map is an embedding.

COROLLARY 7.3
The restriction of Θ *to* Σ_Γ *is a homotopy equivalence* $\Sigma_\Gamma \simeq \mathcal{O}_\Gamma$.

The proof of Theorem 7.2 has two major components. The first is to show that the map Θ is surjective. This is by no means obvious since the markings in $\mathcal{T}_\Gamma$ must be untwisted, whereas the markings in $\mathcal{O}_\Gamma$ are unrestricted. Finding a point in the fiber over some $(Y, d, f) \in \mathcal{O}_\Gamma$ means finding a skewed blowup structure $(\mathbb{S}^\Pi, \mathcal{F})$ on Y such that f^{-1} followed by the standard collapse map is untwisted. To do this, we first decompose Y into parallelotopes, then identify the Γ-Whitehead partitions in the blowup structure, and finally calculate the composition $c_\pi \circ f^{-1}$.

The second component of the proof is to show that the fibers are contractible. To do this, we fix a point in the fiber and describe a process of "shearing" edges dual to a hyperplane in this Γ-complex. We then prove that every point in the fiber can be obtained by a series of "zero-sum shearings" of the initial point. This set of shearings spans a linear subspace of a Euclidean space, hence is contractible.

7.2. *Surjectivity of* Θ

The first step in proving Theorem 7.2 is to show that the inverse image of an arbitrary point in $\mathcal{O}_\Gamma$ is nonempty.

PROPOSITION 7.4

$\Theta\colon \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$ *is* $U(A_\Gamma)$*-equivariant and surjective.*

Equivariance under the action of $U(A_\Gamma)$ is clear from the definition of Θ whereas surjectivity is not, since markings in $\mathcal{O}_\Gamma$ can differ by any element of $\mathrm{Out}(A_\Gamma)$. The key is to show that an appropriate change of skewed blowup structure on a point of $\mathcal{T}_\Gamma$ will have the effect of composing the collapse marking with a twist. The proof of Proposition 7.4 will occupy the remainder of this subsection.

For skewed blowups, the end result of the retraction R_t defined in Section 6.3 followed by scaling the edge lengths linearly gives a continuous "straightening map" $s_{\mathcal{F}}\colon (\mathbb{S}^\Pi, \mathcal{F}) \to (\mathbb{S}^\Pi, \mathcal{E})$ that sends each parallelotope to a unit cube. The standard collapse map $c_\pi\colon \mathbb{S}^\Pi \to \mathbb{S}_\Gamma$ induces a collapse map $c_\pi^{\mathcal{F}} = c_\pi \circ s_{\mathcal{F}}$ on $(\mathbb{S}^\Pi, d_{\mathcal{F}})$, called a *straighten-collapse* map.

Definition 7.5

An automorphism $\phi \in \mathrm{Out}(A_\Gamma)$ is *realized* by an isometry $i\colon (\mathbb{S}^\Pi, d_{\mathcal{F}}) \to (\mathbb{S}^\Omega, d_{\mathcal{G}})$ if $c_\omega^{\mathcal{G}} \circ i \circ (c_\pi^{\mathcal{F}})^{-1}$ induces ϕ on $\pi_1(\mathbb{S}_\Gamma) = A_\Gamma$:

$$\begin{array}{ccc} (\mathbb{S}^\Pi, d_{\mathcal{F}}) & \xrightarrow{\;i\;} & (\mathbb{S}^\Omega, d_{\mathcal{G}}) \\ \downarrow{\scriptstyle c_\pi^{\mathcal{F}}} & & \downarrow{\scriptstyle c_\omega^{\mathcal{G}}} \\ \mathbb{S}_\Gamma & \xrightarrow{\;\phi\;} & \mathbb{S}_\Gamma \end{array}$$

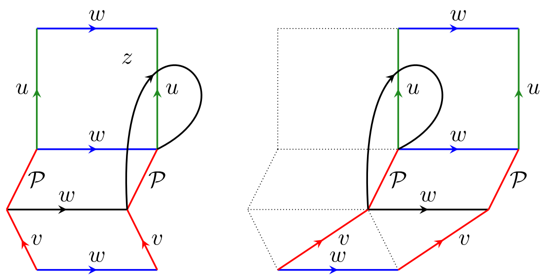

Figure 12. (Color online) Parallelotope structures $\mathcal{F}$, $\mathcal{G}$ on $\mathbb{S}^{\mathcal{P}}$ such that $(\mathbb{S}^{\mathcal{P}}, d_{\mathcal{F}})$ is isometric to $(\mathbb{S}^{\mathcal{P}}, d_{\mathcal{G}})$ and $c^{\mathcal{G}}_{\omega} \circ i \circ (c^{\mathcal{F}}_{\pi})^{-1}$ induces a twist $v \mapsto vw$.

Note that we are not requiring i to be a *combinatorial* isometry, just an isometry. The realization of a combinatorial isometry is always untwisted. In Figure 12, we illustrate an isometry between two skewed blowups that realizes an elementary twist $v \mapsto vw$; one should think of these blowups as giving two different parallelotope decompositions of the same space, and the isometry as the identity. The following lemma explains in general how to realize a twist $v \mapsto vw$ in the case that v is twist-minimal.

LEMMA 7.6
Let $\mathcal{F}$ be an allowable parallelotope structure on $\mathbb{S}^{\Pi}$, and let $\tau \colon v \mapsto vw$ be an elementary twist. If v is twist-minimal, then τ can be realized by an isometry i : $(\mathbb{S}^{\Pi}, d_{\mathcal{F}}) \to (\mathbb{S}^{\Pi}, d_{\mathcal{G}})$ for some allowable parallelotope structure $\mathcal{G}$ on $\mathbb{S}^{\Pi}$.

Proof
Let χ_w be a characteristic cycle for w. Note that w is twist dominant, so χ_w is a local geodesic. The carrier $\kappa(H_v)$ of the hyperplane H_v decomposes combinatorially as a product

$$e_v \times \chi_w \times Z,$$

where Z is the subcomplex of $\mathbb{S}^{\Pi}$ spanned by edges that are adjacent to v and do not split w. The orientation on e_w induces an orientation on all edges of χ_w. We define a new decomposition of $e_v \times \chi_w$ by replacing each edge e_v by the geodesic from its initial vertex to its terminal vertex which cuts diagonally across all the parallelograms in $e_v \times \chi_v$. In a lift of $e_v \times \chi_v$ to $\widetilde{\mathbb{S}}^{\Pi}$, the new edge is a geodesic from the initial point of $\widetilde{e}_v$ to the terminal point of $w\widetilde{e}_v$ (this is what happened in Figure 12, where χ_w consisted of a single edge e_w). Since the structure of Z is unchanged, the new decomposition of $e_v \times \chi_w$ extends to a new parallelotope decomposition of $\kappa(H_v)$

which is combinatorially isomorphic to the old one. It does not change the metric on any parallelotope outside $\kappa(H_v)$, so it extends to a new parallelotope structure $\mathcal{G}$ on $Y = \mathbb{S}^\Pi$. Since v is twist-minimal, skewing a single edge of a characteristic cycle is allowed, so this new parallelotope structure is allowable (see Definition 5.3). Note that the identity on Y is an isometry $(\mathbb{S}^\Pi, d_\mathcal{F}) \to (\mathbb{S}^\Pi, d_\mathcal{G})$ but is not a combinatorial isometry $(\mathbb{S}^\Pi, \mathcal{F}) \to (\mathbb{S}^\Pi, \mathcal{G})$.

The new collapse map $c_\pi^\mathcal{G}$ gives a new action of $\pi_1(\mathbb{S}_\Gamma) = A_\Gamma$ on $\widetilde{Y}$. The only generator whose action has changed is v, whose new axis is the axis that was formerly the axis for vw. □

Notice that in the proof of Lemma 7.6 we skewed a single edge of a characteristic cycle for v. If v is twist-dominant, then we cannot use that trick to realize $\tau \colon v \mapsto vw$, since a characteristic cycle for v must lift to a (straight!) axis for v in $\widetilde{\mathbb{S}^\Pi}$. Instead, we will have to construct a new blowup structure $(S^\Omega, \mathcal{G})$ on Y to realize τ. The idea is to locate branch points and twist-minimal hyperplanes using our identification of Y with $\mathbb{S}^\Pi$, then show that these are metric invariants and use them to construct a new skewed blowup structure $(S^\Omega, \mathcal{G})$ on Y. To make this work, we first need to relate the geometry of $(Y, d) = (\mathbb{S}^\Pi, d_\mathcal{F})$ to the combinatorial structure of $\mathbb{S}^\Pi$. The following proposition is the key.

PROPOSITION 7.7
Let v be a twist-dominant generator of A_Γ. The straightening map $s_\mathcal{F} \colon (\mathbb{S}^\Pi, \mathcal{F}) \to (\mathbb{S}^\Pi, \mathcal{E})$ takes axes for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{F})$ to axes for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{E})$ and the minset of v to the minset of v, where the actions are given by the collapse maps $c_\pi^\mathcal{F}$ and $c_\pi = c_\pi^\mathcal{E}$, respectively. Moreover, $s_\mathcal{F}$ maps branch points for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{F})$ to branch points for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{E})$. The same holds if we replace $c_\pi^\mathcal{E}$ and $c_\pi^\mathcal{F}$ by any untwisted markings h on $(\mathbb{S}^\Pi, \mathcal{E})$ and $h' = h \circ s_F$ on $(\mathbb{S}^\Pi, \mathcal{F})$.

Proof
First assume that the markings are standard collapse maps. Since v is twist-dominant, each characteristic cycle for v in both $(\mathbb{S}^\Pi, \mathcal{F})$ and $(\mathbb{S}^\Pi, \mathcal{E})$ is a geodesic that is the image of an axis by Lemma 5.8. The full minset $\mathrm{Min}(v) \subset \widetilde{\mathbb{S}^\Pi}$ is the convex hull of the lifts of these characteristic cycles, and since $s_\mathcal{F}$ identifies these, it also takes the minset for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{F})$ to the minset for v in $(\widetilde{\mathbb{S}^\Pi}, d_\mathcal{E})$. The last statement about branch points follows from the fact that the straightening map induces a homeomorphism on links.

For a more general untwisted marking h, factor h as $\sigma \circ h_0$ where $h_0 \circ c_\pi^{-1} \in U^0(A_\Gamma)$ and σ is a graph automorphism. Since $U^0(A_\Gamma)$ preserves twist-dominant generators up to conjugacy, the axes and minset of v with respect to h_0 are just trans-

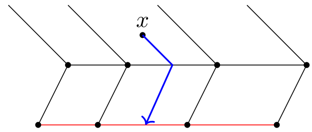

Figure 13. (Color online) The projection map $\mathrm{pr}_w^{\mathcal{F}}$ for w twist-dominant.

lates of the axes and minset with respect to c_π, so the argument above still applies. For the graph automorphism, $\sigma(v) = w$ for some other twist-dominant generator w, so applying the proposition to w gives the same result. □

For the standard metric $d_{\mathcal{E}}$, Lemma 3.10(1) gives a decomposition of the minset of a generator v with respect to the marking c_π as $\mathrm{Min}(v) \cong \alpha_v \times \widetilde{H}_v$, and hence a projection $\mathrm{pr}_v \colon \mathrm{Min}(v) \to \alpha_v$. This projection can be viewed either as the nearest-point (orthogonal) projection, or as collapsing hyperplanes whose labels are adjacent to v. If v is twist-dominant, then by the proposition above, the straightening map takes axes of v in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$ to axes of v in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{E}})$, and likewise minsets to minsets. Thus, we can define an analogous projection in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}})$ by "straightening-projecting-unstraightening" (see Figure 13), that is,

$$\mathrm{pr}_v^{\mathcal{F}} = s_{\mathcal{F}}^{-1} \circ \mathrm{pr}_v \circ s_{\mathcal{F}}.$$

While this is no longer a nearest-point projection, it is again obtained by collapsing all hyperplanes whose labels are adjacent to v. That is, for any parallelotope in the minset, $\mathrm{pr}_v^{\mathcal{F}}$ collapses every edge e_A with $\max(A) \neq \{v\}$ to a point.

PROPOSITION 7.8
Let $\mathcal{F}$ be an allowable parallelotope structure on $\mathbb{S}^\Pi$. Let v be twist-dominant, and suppose that $\tau \colon v \mapsto vw$ is an elementary twist. Then there is an isometry $i : (\mathbb{S}^\Pi, d_{\mathcal{F}}) \to (\mathbb{S}^\Omega, d_{\mathcal{G}})$ that realizes $\tau \circ \varphi$ for some $\varphi \in U^0(A_\Gamma)$ satisfying $\tau \circ \varphi = \varphi \circ \tau$.

Proof
Since v and w commute, there is a vertex $x \in \mathbb{S}^\Pi$ which is a terminal vertex for edges e_v and e_w. Let χ_v and χ_w be characteristic cycles for v and w containing e_v and e_w, $\tilde{x}$ a lift of x to $\widetilde{\mathbb{S}}^\Pi$ and $\tilde{\chi}_v$, $\tilde{\chi}_w$ lifts starting at $\tilde{x}$ of these characteristic cycles. Since both v and w are twist-dominant, $\tilde{\chi}_v$ and $\tilde{\chi}_w$ are contained in axes α_v and α_w through $\tilde{x}$, and the product of these axes is a subcomplex of $\widetilde{\mathbb{S}}^\Pi$ isometric to $\mathbb{E}^2$, with stabilizer the subgroup $\langle v, w\rangle \cong \mathbb{Z}^2$ of A_Γ. The parallelogram in $\widetilde{\mathbb{S}}^\Pi$ spanned by $\tilde{x}$, $v\tilde{x}$, $w\tilde{x}$,

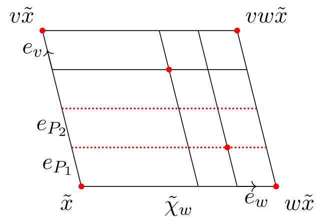

Figure 14. (Color online) Fundamental domain D for $\langle v, w\rangle \cong \mathbb{Z}^2$ on $\alpha_v \times \alpha_w \cong \mathbb{E}^2 \subset \widetilde{\mathbb{S}}^\Pi$. The dots and dotted lines are the projections of all branch points for v and w onto D.

and $vw\tilde{x}$ is a fundamental domain D for this action (see Figure 14). Define a map $p\colon (\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{F}}) \to \alpha_v \times \alpha_w$ by $p = (s_{\mathcal{F}})^{-1} \circ p^\perp \circ s_{\mathcal{F}}$, where $p^\perp$ is the nearest-point projection in $(\widetilde{\mathbb{S}}^\Pi, d_{\mathcal{E}})$. We will be most interested in the restriction of p to $\mathrm{Min}(w)$, projecting $\mathrm{Min}(w)$ onto $\alpha_v \times \alpha_w$.

Claim
Let $\mathrm{br}(v)$ be the set of branch points for v, and let $\mathrm{br}(w)$ be the set of branch points for w. Then $p(\mathrm{br}(v))$ consists of lines parallel to α_w and isolated points, and $p(\mathrm{br}(w))$ consists of isolated points. The isolated points are vertices of $\mathbb{S}^\Pi$.

Proof of claim
There is a branch point for w at a vertex $x \in \chi_w \times H_w \subset \mathbb{S}^\Pi$ if and only if there is an edge e_A at x with $[A, w] \neq 1$. If x is a branch point for w and $x \in \chi_v \times H_v$, then x is also a branch point for v. If x is a branch point for v but not for w, then all edges e_A at x that are not adjacent to v must be adjacent to w. In this case, every point of $x \times \chi_w$ is a branch point for v. □

Let $x \in \mathbb{S}^\Pi$ be a terminal vertex for edges e_v and e_w as above, and let $\tilde{x}$ be a lift to $\tilde{\chi}_v \times \tilde{\chi}_w$. If w is central, then $\mathrm{br}(w)$ is empty. In this case, the characteristic cycle for w consists of the single edge e_w and the only vertices on $\tilde{\chi}_w$ are the w-translates of $\tilde{x}$, but these are not branch points. The same is true for $\tilde{\chi}_v$ if v and w are both central.

Let $B = \mathrm{br}(v) \cup \mathrm{br}(w)$. Note that the decomposition of $\tilde{\chi}_v \times \tilde{\chi}_w$ into parallelograms is completely determined by $p(B) \cup \{\tilde{x}\}$. This is because each edge of this decomposition is on a lift of a characteristic cycle χ_w or χ_v, and each endpoint of this edge corresponds to a branch point in some (parallel) axis for v or w or to a translate of $\tilde{x}$.

We are now ready to replace the action of v by the action of $\tau(v) = vw$. Since $v \leq_t w$, the centralizer of v is equal to the centralizer of vw, so $\mathrm{Min}(v) = \mathrm{Min}(vw)$

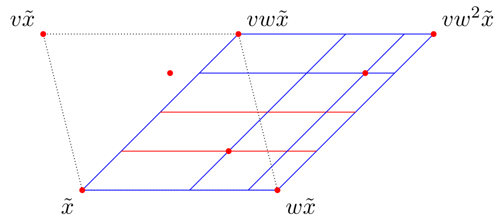

Figure 15. (Color online) Skewing D. The branch points for v and w are the same as the branch points for vw and w.

and $\mathrm{br}(v) = \mathrm{br}(vw)$. Thus replacing v by vw does not change B or the projections of branch points onto the plane $\alpha_v \times \alpha_w$. Replacing the fundamental domain D of $\alpha_v \times \alpha_w$ by a new fundamental domain D' with vertices $\tilde{x}$, $vw\tilde{x}$, $vw^2\tilde{x}$, and $w\tilde{x}$, these projections determine a decomposition of D' into parallelograms (see Figure 15). The decomposition of $\alpha'_v = \alpha_{vw}$, the axis for $\tau(v)$, is in one-to-one correspondence with the decomposition of α_v, since in both cases, the vertices are projections of points in $p(B)$ parallel to an axis α_w. But the decomposition of α_w will change since vertices are now projections of $p(B)$ parallel to the new axis α'_v, instead of the old axis α_v. So for example, two points in $p(B)$ could project to the same point under one of these projections and to distinct points under the other.

We claim that D' together with its decomposition is part of a skewed Γ-complex structure $(\mathbb{S}^\Omega, \mathcal{G})$ on (Y, d). To prove this, we need to do two things. The first is to complete the new parallelogram decomposition of $\alpha_v \times \alpha_w$ to a parallelotope decomposition of all of Y. The second is to find a compatible set Ω of partitions corresponding to this decomposition, that is, a parallelotope structure $\mathcal{G}$ on $\mathbb{S}^\Omega$ making $(\mathbb{S}^\Omega, d_{\mathcal{G}})$ isometric to $(Y, d) = (\mathbb{S}^\Pi, d_{\mathcal{F}})$.

Parallelotope decomposition. We have changed the decomposition of the axis α_w into edges. As collateral damage, we have also changed the decomposition of any characteristic cycle with a lift that intersects α_w. However, the endpoints of the intersection interval are images of branch points for w, so are still vertices in the new decomposition; that is, this segment of the characteristic cycle is the only thing we have changed. (In particular, if the intersection is a single point, then we have not changed this characteristic cycle at all.)

If w commutes with $u \in \max(A)$ for some label A, then the decomposition of every product subcomplex $\chi_w \times \chi_u$ of $\mathbb{S}^\Pi$ is affected by changing the decomposition of α_w. If u also commutes with v, then this is not a problem because then the new decomposition of $\alpha_v \times \alpha_w$ extends to a decomposition of $\alpha_v \times \alpha_w \times e_A \subset \widetilde{\mathbb{S}}^\Pi$.

If $[u, w] = 1$ but $[u, v] \neq 1$, then it may happen that some partition $\mathcal{P}$ that splits v also splits u, so that $\chi_u \times \chi_w$ overlaps $\chi_v \times \chi_w$ in the band $e_{\mathcal{P}} \times \chi_w$. We have

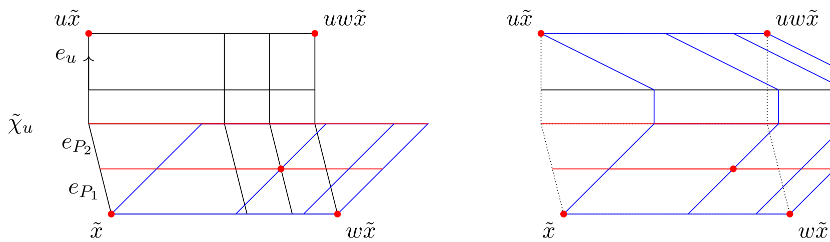

Figure 16. (Color online) Skewing e_u back in $\tilde{\chi}_u \times \alpha_w$ when $e(\mathcal{P}_1)$, $e(\mathcal{P}_2)$ are in χ_u.

changed the decomposition of this band. However, notice that $u \leq_f v$ so u cannot be twist-dominant and moreover $u \leq_t w$. Since u is twist-minimal, we can compensate for what we have done by using the band $e_u \times \alpha_w \subset \tilde{\chi}_u \times \alpha_w$ to skew the characteristic cycle for u back to the original endpoint of $\tilde{\chi}_u$ (see Figure 16). We do not change the angle with α_w in any other band, so preserve the condition of allowability for the new parallelotope structure.

Blowup structure. We need to find a set of partitions Ω corresponding to our new parallelotope decomposition. In particular, we need to show that the new decomposition of α_w comes from a set of Γ-Whitehead partitions that split w. Recall from Section 3.3 that the partitions $\{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$ that split w are nested, that is, their w-sides P_i satisfy $P_0 = \{w\} \subset P_1 \subset P_2 \subset \cdots \subset P_k \subset P_{k+1} = \overline{P}_0 \setminus \{w^{-1}\}$, and if $\mathcal{R}$ is any other partition in Π that is not adjacent to w, then its non-w side $\overline{R}$ is contained in some piece $dP_i = P_i \setminus P_{i-1}$ of the nest.

Since v is twist-dominant, the partitions splitting v are also nested, say, $Q_0 = \{v\} \subset Q_1 \subset Q_2 \subset \cdots \subset Q_\ell \subset Q_{\ell+1} = \overline{Q}_0 \setminus \{v^{-1}\}$, and the pieces $dQ_j = Q_j \setminus Q_{j-1}$ are unions of v-components of $\Gamma^{\pm}$. (Recall from Section 2.3 that a v-component is a connected component of $\Gamma^{\pm} \setminus \mathrm{lk}^{\pm}(v) \setminus \{v, v^{-1}\}$ and that each side of a partition based at v is a union of v-components plus v or v^{-1}.) Since $\mathrm{st}(v) \subseteq \mathrm{st}(w)$, these v-components are unions of w-components plus possibly some elements of $\mathrm{lk}(w)$. Thus the intersection of a set of v-components with a set of w-components is a set of w-components. In particular, each intersection $I_{ij} = dP_i \cap dQ_j$ is a union of w-components.

Each vertex $\mathbb{r}_{ij}$ of $\chi_v \times \chi_w$ is a region that contains the consistent set

$$\{\overline{P}_1, \ldots \overline{P}_{i-1}, P_i, \ldots, P_k, \overline{Q}_1, \ldots, \overline{Q}_{j-1}, Q_j, \ldots, Q_\ell\}.$$

Partitions that are not adjacent to w also are not adjacent to v, so have sides $\overline{R}_i$ that fit into both nests (the sides that do not contain v or w), and $\mathbb{r}_{ij}$ must also contain the consistent set

$$\mathcal{S}_{ij} = \{\overline{P}_1, \ldots \overline{P}_{i-1}, P_i, \ldots, P_k, \overline{Q}_1, \ldots, \overline{Q}_{j-1}, Q_j, \ldots, Q_\ell, R_1, \ldots, R_m\}.$$

The remaining partitions in Π are all adjacent to w.

If $I_{ij} = dP_i \cap dQ_j$ contains some outermost $\overline{R}_s$ or a vertex u outside all of the $\overline{R}_s$, then we can use this to extend $\mathcal{S}_{ij}$ to a region incident to an edge labeled $\mathcal{R}_s$ or u as we did in Section 3.5. This region is a branch point for w in some parallel copy of D, and projects to $\mathbb{r}_{ij}$.

On the other hand, suppose that I_{ij} contains no $\overline{R}_s$ or outermost vertex u. Then no extension of $\mathcal{S}_{ij}$ produces a region incident to an edge labeled $\mathcal{R}_s$ or u. Since every edge that branches off $\mathrm{Min}(w)$ has such a label, no such region gives a branch point for w, that is, $\mathbb{r}_{ij}$ is not in the image of $\mathrm{br}(w)$.

Identifying $(\alpha_v \times \alpha_w) = (\alpha_{vw} \times \alpha_w)$, we get a new fundamental domain D' and a new map $\mathrm{pr}'_w\colon D' \to \alpha_w \cap D'$ which projects along vw-axes. Using pr'_w, project those $\mathbb{r}_{ij}$ that are images of branch points for w to an ordered set of points $(\mathbb{x}_1, \dots, \mathbb{x}_n)$ on $\alpha_w \cap D'$.

Let $I(\mathbb{x}_k)$ be the union of the I_{ij} such that $\mathrm{pr}'_w(\mathbb{r}_{ij}) = \mathbb{x}_k$. Let $P'_1 = \{w\} \cup I(\mathbb{x}_1)$, $P'_2 = P'_1 \cup I(\mathbb{x}_2)$, and so on. Each P'_i is a side of a valid Γ-Whitehead partition $\mathcal{P}'_i$ based at w, since each I_{ij} is a union of w-components.

Let Ω be the collection of Γ-Whitehead partitions obtained from Π by replacing $\mathcal{P}_1, \dots \mathcal{P}_k$ by $\mathcal{P}'_1, \dots \mathcal{P}'_n$. To see that the Ω partitions are pairwise compatible, we need only check that $\mathcal{P}'_i$ is compatible with $\mathcal{R}_j$ for all i, j. We know that the side $\overline{R}_j$ lies in some I_{st}, and hence in some $I(\mathbb{x}_k)$. So by definition, $\overline{R}_j \subset P'_k \setminus P'_{k-1}$ and it follows that $\mathcal{R}_j$ is compatible with $\mathcal{P}'_i$ for all i.

Marking change. Finally, we calculate the effect of replacing the structure $(\mathbb{S}^\Pi, \mathcal{F})$ on Y, with its marking $c_\pi^{\mathcal{F}}$, by the new structure $(\mathbb{S}^\Omega, \mathcal{G})$ and marking $c_\omega^{\mathcal{G}}$.

LEMMA 7.9

Suppose that v is twist-dominant, and let $\tau\colon v \mapsto vw$ be an elementary twist. The composite map $c_\omega^{\mathcal{G}} \circ (c_\pi^{\mathcal{F}})^{-1} : \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ is of the form $\tau \circ \varphi$, where $\varphi \in U^0(A_\Gamma)$ and $\tau \circ \varphi = \varphi \circ \tau$.

Proof

Let $\mu = c_\omega^{\mathcal{G}} \circ (c_\pi^{\mathcal{F}})^{-1}\colon \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$. The corner point $\widetilde{x}$ of the fundamental domains D and D' described above is a terminal vertex of edges $\widetilde{e}_v$ and $\widetilde{e}_w$ in $\widetilde{\mathbb{S}}^\Omega$ as well as in $\widetilde{\mathbb{S}}^\Pi$. Let x be its image in $\mathbb{S}^\Pi$, and for each $u \in V$, let ξ_u be an edge path which goes from x to an e_u edge in $\mathbb{C}^\Pi$, across e_u, and then back to x in $\mathbb{C}^\Pi$. Note that ξ_u crosses a single e_u edge and all other edges are labeled by partitions. We choose ξ_u to have minimal length among all such paths. ξ_u represents the homotopy class $(c_\pi^{\mathcal{F}})^{-1}(u) \in \pi_1(\mathbb{S}^\Pi, x)$.

Lift ξ_u to a path $\widetilde{\xi}_u$ based at $\widetilde{x}$. The endpoint $\widetilde{y}$ of $\widetilde{\xi}_u$ is then $u \cdot \widetilde{x}$ with respect to the $c_\pi^{\mathcal{F}}$-marking. Since ξ_u was taken to be minimal, $\widetilde{\xi}_u$ is a combinatorial geodesic (i.e., it crosses each hyperplane in $\widetilde{\mathbb{S}}^{\Pi}$ at most once), and our choice of x means that $\widetilde{x}$ and $\widetilde{y}$ are vertices in $(\widetilde{\mathbb{S}}^{\Omega}, \mathcal{G})$. Any minimal-length edge path $\widetilde{\eta}_u$ in $(\widetilde{\mathbb{S}}^{\Omega}, \mathcal{G})$ between $\widetilde{x}$ and $\widetilde{y}$ consists of edges that cross a hyperplane which separates $\widetilde{x}$ and $\widetilde{y}$. To calculate $\mu(u)$, it is enough to know the Ω-labels of hyperplanes that are crossed by $\widetilde{\eta}_u$. The only hyperplanes and labels that change as we go from $(\mathbb{S}^{\Pi}, \mathcal{F})$ to $(\mathbb{S}^{\Omega}, \mathcal{G})$ are those with max $= w$. Thus, $\widetilde{\eta}_u$ crosses one hyperplane labeled u, and all other hyperplanes are either labeled by partitions or by w.

It follows that $\mu(u) = w^{n_u} u w^{m_u}$ for some $n_u, m_u \in \mathbb{Z}$. In particular, the twist-component of μ is a product of elementary twists by w. By construction, a $c_\pi^{\mathcal{F}}$-axis for v maps to a $c_\omega^{\mathcal{G}}$-axis for vw, so we know that $\mu(v) = vw$. If $u \neq v$ but $u \leq_t w$, then either u is twist-dominant, so the axis for u has not changed, or we have sheared the e_u edge so that a $c_\pi^{\mathcal{F}}$-axis for u maps to a $c_\omega^{\mathcal{G}}$-axis for u. Thus, $\mu(u) = u$. This proves that the twist component τ of μ is just $\tau\colon v \mapsto vw$. Therefore, we can write $\mu = \tau \circ \varphi$, where φ is a product of folds and partial conjugations by w. Thus, $\varphi \in U^0(A_\Gamma)$ and since τ is a twist by w, $\tau \circ \varphi = \varphi \circ \tau$, as desired. □

This completes the proof of Proposition 7.8. □

We next make some observations about changing the order of elementary twists, folds, and partial conjugations.

Definition 7.10
Let $\tau\colon v \mapsto vw$ be an elementary twist. If v is twist-dominant, then we say that τ is a *TD twist*, and if v is twist-minimal, then we say that τ is a *TM twist*.

LEMMA 7.11
Let $\tau\colon v \mapsto vw$ be an elementary twist.

(1) *Let φ be a partial conjugation or an elementary fold. Then either φ commutes with τ or $\tau\varphi = \alpha\varphi\tau$, where α is a partial conjugation, an elementary fold, or an elementary TM twist by w that commutes with both φ and α.*

(2) *If τ is a TD twist and t is a TM twist, then either t commutes with τ or $\tau t = \alpha t \tau$, where α is an elementary TM twist by w that commutes with both τ and t.*

Proof
(1) First suppose that φ is conjugates a component C of $\Gamma \setminus \mathrm{st}(u)$ by u. Since v and w are connected by an edge in Γ, φ commutes with τ unless $u = v$, in which case $\varphi\tau$ agrees with $\tau\varphi$ except that $\varphi\tau$ conjugates C by vw instead of v. Since $\mathrm{st}(v) \subset \mathrm{st}(w)$,

C is a union of components C_i of $\Gamma \setminus \mathrm{st}(w)$ plus some elements of $\mathrm{st}(w)$, so we can correct this by partially conjugating the C_i by w^{-1}.

Next suppose that φ is a right fold $\rho_{xy}\colon x \mapsto xy$ or left fold $\lambda_{xy}\colon x \mapsto yx$. It cannot be that $w = x$ since that would mean $v \leq_t w \leq_f y$. Therefore, τ commutes with φ unless $v = y$, in which case

$$\varphi\tau\alpha = \tau\varphi,$$

where $\alpha : x \mapsto xw$ if φ is a right fold, or $\alpha : x \mapsto wx$ if φ is a left fold. Note that α may be either a fold if $[x, w] \neq 1$ or a twist if $[x, w] = 1$. Since $x \leq_f y$, this implies that x cannot be twist-dominant, so if α is a twist, then it is a TM twist. In either case, since v commutes with w, α commutes with both τ and φ.

(2) Let $t\colon x \mapsto xy$ be a TM twist. Then $x \neq w$ since w is twist-dominant, so t commutes with τ unless $v = y$. If $v = y$, then x must commute with v and hence also with w. In this case, $\tau t = \alpha t \tau$ where $\alpha : x \mapsto xw$, which is a TM twist commuting with both τ and t. □

Recall that $\mathrm{Out}^0(A_\Gamma)$ is the subgroup of $\mathrm{Out}(A_\Gamma)$ generated by folds, twists, partial conjugations, and inversions. By checking the generators, it is not hard to see that graph automorphisms normalize $\mathrm{Out}^0(A_\Gamma)$, hence it is a normal subgroup.

COROLLARY 7.12

Let $\langle TM \rangle$ denote the subgroup of $\mathrm{Out}^0(A_\Gamma)$ *generated by TM twists, and let G be the subgroup generated by $U^0(A_\Gamma)$ and $\langle TM \rangle$.*

(1) *Any element $g \in G$ can be factored as $g = t_1 \circ \phi_1 = \phi_2 \circ t_2$, where $\phi_i \in U^0(A_\Gamma)$ and $t_i \in \langle TM \rangle$.*

(2) *TD twists normalize G, hence any element of* $\mathrm{Out}^0(A_\Gamma)$ *can be factored as a product of an element of $\langle TD \rangle$, an element of $\langle TM \rangle$, and an element of $U^0(A_\Gamma)$ in any order. The $U^0(A_\Gamma)$ and $\langle TM \rangle$ factors may depend on the choice of order, but the $\langle TD \rangle$ factor remains unchanged.*

Proof

First note that inversions normalize the subgroup of $\mathrm{Out}^0(A_\Gamma)$ generated by folds, twists, and partial conjugations. Thus any inversion can be moved past any twist. For (1), it remains to consider the case where $t_1 = \tau$ is a single TM-twist and $\phi_1 = \varphi_1 \cdots \varphi_n$ is a product of folds and partial conjugations. Applying Lemma 7.11(1) repeatedly gives

$$t_1 \circ \phi_1 = \tau\varphi_1 \cdots \varphi_n = (\varphi_1\alpha_1 \cdots \varphi_n\alpha_n)\tau,$$

where each α_i is either the identity, a partial conjugation, an elementary fold, or a TM twist by the same element w. In particular, all of the α_i's commute with each other.

If all α_i lie in $U^0(A_\Gamma)$, then we are done, but if one or more α_i is a twist, then we must apply the lemma again to move these twists to the right. Since α_i commutes with the other α_j's, only moving it past the φ_j terms can introduce new factors and these, too, will commute with each other and with the α_i. Repeating this process, we can move all of the newly introduced TM twists to the right to obtain a new factorization $\tau\phi_1 = \phi_2 t_2$ as desired.

For (2), the fact that TD twists normalize G follows immediately from Lemma 7.11 since α always lies in G. So for any $h \in \mathrm{Out}^0(A_\Gamma)$, we can write $h = g_1 \circ t = t \circ g_2$, where $t \in \langle TD \rangle$ and $g_i \in G$. By part (1), we can factor g_i into an element of $U^0(A_\Gamma)$ and an element of $\langle TD \rangle$ in either order. By Lemma 7.11(2), we can also switch the order of the TM and TD twists if desired. $\square$

We can now complete the proof that Θ is surjective.

Proof of Proposition 7.4

By definition, a point in $\mathcal{O}_\Gamma$ is a space (Y, d) isometric to a skewed Γ-complex $(\mathbb{S}^\Pi, d_{\mathcal{F}})$, together with a homotopy equivalence $f : Y \to \mathbb{S}_\Gamma$. For the purpose of this proof, we will identify (Y, d) with $(\mathbb{S}^\Pi, d_{\mathcal{F}})$. Then a point in the fiber $\Theta^{-1}(Y, d, f)$ is given by a skewed Γ-complex $(X, \mathcal{G})$, an untwisted homotopy equivalence $h\colon X \to \mathbb{S}_\Gamma$, and an isometry $i : (\mathbb{S}^\Pi, d_{\mathcal{F}}) \to (X, d_{\mathcal{G}})$ such that $h \simeq f \circ i$. If we also choose a combinatorial isometry of X with some blowup $\mathbb{S}^\Omega$, then the picture is

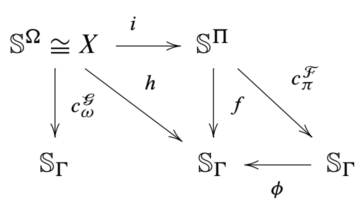

where $h \circ (c^{\mathcal{G}}_\omega)^{-1} \in U(A_\Gamma)$. To prove the proposition, we must find such an $(\mathbb{S}^\Omega, \mathcal{G}, h)$.

Let $\phi = f \circ (c^{\mathcal{F}}_\pi)^{-1}$. Since graph automorphisms normalize $\mathrm{Out}^0(A_\Gamma)$, we can write $\phi = \phi' \circ \gamma$, where $\phi' \in \mathrm{Out}^0(A_\Gamma)$ and γ is a graph automorphism. Then replacing $\mathbb{S}^\Pi$ by $\mathbb{S}^{\gamma(\Pi)}$ as in the proof of Lemma 4.7, we may assume that $\phi \in \mathrm{Out}^0(A_\Gamma)$. By composing $c^{\mathcal{F}}_\pi$ with an isometry of $\mathbb{S}_\Gamma$, we can change the collapse map as in the proof of Lemma 4.7, thereby removing γ. Without loss of generality, we therefore assume that $\phi \in \mathrm{Out}^0(A_\Gamma)$. By Corollary 7.12, we can factor ϕ as $\phi = \eta \circ t_1 \circ t_2$, where $\eta \in U^0(A_\Gamma)$, t_1 is a product of TD twists, and t_2 is a product of TM twists. Elements of $U^0(A_\Gamma)$ act on the left on both $\mathcal{T}_\Gamma$ and $\mathcal{O}_\Gamma$ and the action commutes with Θ, so the fiber over $(\mathbb{S}^\Pi, d_{\mathcal{F}}, f)$ is isomorphic to the fiber over $\eta^{-1}(\mathbb{S}^\Pi, d_{\mathcal{F}}, f) = (\mathbb{S}^\Pi, d_{\mathcal{F}}, \eta^{-1} f)$. Thus we may assume that $\phi = t_1 \circ t_2$. Moreover, by Lemma 7.6 we can realize t_2 by a change of parallelotope structure on $\mathbb{S}^\Pi$ (which changes the collapse map, but not the metric on $\mathbb{S}^\Pi$), so we may assume that $t_2 = \mathrm{id}$ and write $\phi = \tau_1 \circ \tau_2 \circ \cdots \circ \tau_k$, a product of an elementary TD twist.

By Proposition 7.8, we can find elements $\varphi_i \in U^0(A_\Gamma)$ and a sequence of skewed blowup structures on Y realizing the compositions $\tau_i \circ \varphi_i$. Composing these gives

$$\begin{array}{ccccc} \mathbb{S}^\Omega & \xleftarrow{\;i\;} & \mathbb{S}^\Pi & & \\ \big\downarrow c_\omega^{\mathcal{G}} & & \big\downarrow c_\pi^{\mathcal{F}} & \searrow^{f} & \\ \mathbb{S}_\Gamma & \xleftarrow[\tau_1\varphi_1\cdots\tau_k\varphi_k]{} & \mathbb{S}_\Gamma & \xrightarrow{\;\phi\;} & \mathbb{S}_\Gamma \end{array}$$

By Corollary 7.12, we can rewrite

$$\tau_1\varphi_1 \cdots \tau_k\varphi_k = t' \circ \varphi \circ (\tau_1 \cdots \tau_k) = t' \circ \varphi \circ \phi,$$

where t' is a product of TM twists and $\varphi \in U^0(A_\Gamma)$. By changing the parallelotope structure on $\mathbb{S}^\Omega$, we may again arrange that $t' = \mathrm{id}$, so the diagram above becomes

$$\begin{array}{ccccc} \mathbb{S}^\Omega & \xleftarrow{\;i\;} & \mathbb{S}^\Pi & & \\ \big\downarrow c_\omega^{\mathcal{G}} & & \big\downarrow c_\pi^{\mathcal{F}} & \searrow^{f} & \\ \mathbb{S}_\Gamma & \xleftarrow[\varphi\circ\phi]{} & \mathbb{S}_\Gamma & \xrightarrow{\;\phi\;} & \mathbb{S}_\Gamma \end{array}$$

Setting $h = \varphi^{-1} \circ c_\omega^{\mathcal{G}} = f \circ i^{-1}$, we have $h \circ (c_\omega^{\mathcal{G}})^{-1} = \varphi^{-1} \in U^0(A_\Gamma)$, so $(\mathbb{S}^\Omega, \mathcal{G}, h)$ is the desired point in the fiber. □

7.3. Structure of fibers

7.3.1. Finding twist-minimal hyperplanes

In this section, we show that the set of twist-minimal hyperplanes in a marked twisted Γ-complex depends only on the underlying metric and the marking, that is, on the projection to $\mathcal{O}_\Gamma$.

LEMMA 7.13

Let $[X, \mathcal{F}, h]$ and $[X', \mathcal{F}', h']$ be two points in the fiber over $[Y, d, f]$. The images in Y of twist-minimal hyperplanes in $(X, \mathcal{F})$ and $(X', \mathcal{F}')$ are the same (both set-theoretically and pointwise) and their carriers have the same width.

Proof

Since $(X, \mathcal{F}, h)$ and $(X', \mathcal{F}', h')$ both project to (Y, d, f), we can identify $(X, d_{\mathcal{F}}) \cong$

$(Y,d) \cong (X', d_{\mathcal{F}'})$, that is, we consider $(X,\mathcal{F})$ and $(X',\mathcal{F}')$ to be two different skewed Γ-complex structures on the same underlying space Y. Using this identification, we have $h = f = h'$, so f is untwisted in both of these structures.

Recall that in Section 4.2 we defined the sets $\mathrm{split}_h(H)$ and $\max_h(H)$ for a hyperplane in a rectilinear Γ-complex with an untwisted marking h. The same definitions can be used for a hyperplane H in a skewed Γ-complex $(X,\mathcal{F})$ provided that H is convex, that is, $v \in \mathrm{split}_h(H)$ if an axis α_v crosses some lift of H in $(\widetilde{X},\mathcal{F})$, where the action is given by the isomorphism $h_*\colon \pi_1(X) \cong \pi_1(\mathbb{S}_\Gamma) = A_\Gamma$. In addition, if $s_{\mathcal{F}}$ is the straightening map and $h_{\mathcal{F}} = hs_{\mathcal{F}}^{-1}$, then the induced map $\tilde{s}_{\mathcal{F}}$ on the universal cover is equivariant with respect to the markings determined by h and $h_{\mathcal{F}}$. Thus if some lift of H separates x from vx in $(X,\mathcal{F})$, then the same holds after straightening. In other words, $\mathrm{split}_h(H) = \mathrm{split}_{h_{\mathcal{F}}}(s_{\mathcal{F}}(H))$. Thus, by Corollary 4.11 and Lemma 5.8, a hyperplane H in $(X,\mathcal{F})$ is twist-minimal if and only if any lift $\widetilde{H}$ is convex and $[\max_h(H)]$ is twist-minimal.

Assume that H in $(X,\mathcal{F})$ is twist-minimal, and let $v \in \max_h(H)$. Then some lift $\widetilde{H}$ lies in $\mathrm{Min}_h(v)$ and we can decompose $\mathrm{Min}_h(v)$ as a (not necessarily orthogonal) product $\mathrm{Min}_h(v) = \alpha_v \times \widetilde{H}$. We would like to apply Proposition 4.13, but that proposition was proved only in the context of rectilinear Γ-structures, so we first must straighten $(X,\mathcal{F})$. For twist-minimal elements, the straightening map $s_{\mathcal{F}}$ need not take axes to axes or minsets to minsets, but as observed above, it does take the carrier of $\widetilde{H}$ to the carrier of a hyperplane $\widetilde{H}' = s_{\mathcal{F}}(\widetilde{H})$ that also has maximal element v. Hence the minset of v in the straightened structure $(X,\mathcal{E},h_{\mathcal{F}})$ decomposes as $\alpha'_v \times \widetilde{H}'$, where α'_v is an axis for v with respect to the marking $h_{\mathcal{F}}$. In particular, the straightening map between carriers extends to a homeomorphism between these two minsets. It follows that $s_{\mathcal{F}}^{-1}$ maps branch points in $\kappa(\widetilde{H})$ to branch points in $\kappa(\widetilde{H}')$. By Proposition 4.13, $\kappa(\widetilde{H}')$ contains branch points on both components of its boundary, so the same holds for $\kappa(\widetilde{H})$. The position of $\widetilde{H}$ is determined by the projection of these branch points on α_v via the projection map $\mathrm{pr}_v^{\mathcal{F}} = s_{\mathcal{F}}^{-1} \circ \mathrm{pr}_v \circ s_{\mathcal{F}}$. Moreover, since $\widetilde{H}$ is the convex hull of the $A_{\mathrm{lk}(v)}$-orbit of a point on α_v, the projection map is determined by the CAT(0) metric and the marking h, independent of the choice of point on α_v.

Since $\mathrm{Min}_h(v)$, the projection map to α_v, and the branch locus of v depend only on the CAT(0) metric and the marking, they are the same for $(\widetilde{X},\mathcal{F})$ and $(\widetilde{X}',\mathcal{F}')$. □

7.3.2. Shearing $(\mathbb{S}^\Pi,\mathcal{F})$

Now let (Y,d,f) be an arbitrary point of $\mathcal{O}_\Gamma$. By Proposition 7.4, the fiber $\Theta^{-1}(Y,d,f)$ is nonempty, so we may fix a point $(\mathbb{S}^\Pi,\mathcal{F},h_0)$ in this fiber and use an isometry

$(Y,d) \cong (\mathbb{S}^\Pi, d_{\mathcal{F}})$ to identify Y with $\mathbb{S}^\Pi$ and f with the untwisted marking h_0. After acting by the untwisted subgroup $U(A_\Gamma)$, we may further assume that $f = c_\pi^{\mathcal{F}}$.

If $(X,\mathcal{G},h)$ is any other point in the same fiber, then there is an isometry $i\colon (X,d_{\mathcal{G}}) \to (Y,d)$ with $f \circ i \simeq h$. Using this isometry to identify $(X,d_{\mathcal{G}})$ with (Y,d) and h with f, we can view $(X,\mathcal{G})$ as a different decomposition of the same underlying metric space Y into (unlabeled) parallelotopes. We say that $(X,\mathcal{G})$ is a *Γ-complex structure on Y*. To understand the topology of the fiber, we will compare an arbitrary Γ-complex structure $(X,\mathcal{G})$ with our given structure $(\mathbb{S}^\Pi,\mathcal{F})$.

The action of A_Γ on universal covers is given by f in both cases, so the axes, minsets, and branch points are the same. However, while f is untwisted with respect to both structures, it is a Γ-collapse map only for $(\mathbb{S}^\Pi,\mathcal{F})$ where it is in fact the standard collapse map.

By Lemma 7.13, the set of twist-minimal hyperplanes and their carriers are the same in both structures. Let H be a twist-minimal hyperplane, and let $\widetilde{H}$ be a lift of H to $\widetilde{\mathbb{S}}^\Pi$. The carrier $\kappa(\widetilde{H})$ has two boundary components, ∂_0 and ∂_1. Let x_0 be a branch point in ∂_0; then each of $(\widetilde{\mathbb{S}}^\Pi,\mathcal{F})$ and $(\widetilde{X},\mathcal{G})$ must have an edge dual to $\widetilde{H}$ with one endpoint at x_0. In $(\mathbb{S}^\Pi,\mathcal{F})$ hyperplanes are labeled, so we have $\widetilde{H} = \widetilde{H}_A$ for some $A \in \Pi \cup V$ with $\max(A)$ twist-minimal, and we label this edge e_A. In the skewed Γ-complex structure $(X,\mathcal{G})$ the edge does not have a label, so we will just call it e_H.

By Lemma 5.8, $\widetilde{H}$ is convex. The elements of $\mathrm{lk}^+(A)$ are twist-dominant and commute with each other, so $\widetilde{H}$ contains a subspace of the form $e_A \times \mathbb{E}_A^+$, where $\mathbb{E}_A^+$ is an affine space generated by axes of elements in $\mathrm{lk}^+(A)$. By definition of an allowable parallelotope structure, the edge e_A was obtained from an orthogonal edge by rotating in the direction of $\mathbb{E}_A^+$. The same applies to e_H, since $\mathrm{lk}^+(H) = \mathrm{lk}^+(A)$. Letting $t(e_A)$ and $t(e_H)$ be the endpoints of e_A and e_H in ∂_1, it follows that the subspaces $t(e_A) \times \mathbb{E}_A^+$ and $t(e_H) \times \mathbb{E}_A^+$ agree. So the difference $s_A = t(e_H) - t(e_A)$ is a vector in the vector space U_A^+ spanned by the axes of $\mathrm{lk}^+(A)$. (See Figure 17.)

Finally, note that in defining s_A, we began by choosing an isometry $i\colon (X,d_{\mathcal{G}}) \to (Y,d)$. While this isometry need not be unique, for any other such isometry j, we have

$$j^{-1} \circ i = (j^{-1} \circ f^{-1}) \circ (f \circ i) \simeq h^{-1} \circ h = id.$$

Recall that Y decomposes as an orthogonal product $Y = Y_0 \times T_{\mathcal{Z}(A_\Gamma)}$, where $T_{\mathcal{Z}(A_\Gamma)}$ is a torus of dimension equal to the rank of the center $\mathcal{Z}(A_\Gamma)$. It follows from the work of Bregman [7] that the only isometries of Y that are homotopic to the identity are translations of the central torus $T_{\mathcal{Z}(A_\Gamma)}$. Such a translation has no effect on the relative position of e_H and e_A, so s_A is independent of the choice of i.

Definition 7.14
The vector $s_A = t(e_H) - t(e_A) \in U_A^+$ is the *shear* of e_H relative to e_A.

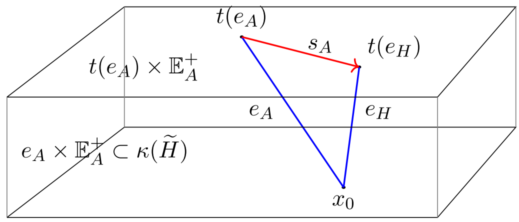

Figure 17. (Color online) Shear of e_H with respect to e_A.

Definition 7.15
A *shearing* of $(\mathbb{S}^\Pi, \mathcal{F})$ is a choice of vector $s_A \in U_A^+$ for every hyperplane H_A, subject to the condition that if $\max(A)$ is twist-dominant, then $s_A = 0$.

We now observe that two Γ-complex structures $(X, \mathcal{G})$ and $(X', \mathcal{G}')$ in the fiber that define the same shearing are the same.

PROPOSITION 7.16
Two points $[X, \mathcal{G}, h]$ and $[X', \mathcal{G}', h']$ in the fiber over $(\mathbb{S}^\Pi, d_{\mathcal{F}}, c_\pi^{\mathcal{F}}) = (Y, d, f)$ are the same if and only if they define the same shearings $\{s_A\}$ and $\{s'_A\}$ of $(\mathbb{S}^\Pi, \mathcal{F})$.

Proof
If $[X, \mathcal{G}, h] = [X', \mathcal{G}', h']$, then there is a combinatorial isometry $i\colon (X, \mathcal{G}) \to (X', \mathcal{G}')$ with $h' \simeq h \circ i$ (i.e., an isomorphism of cube complexes $X \cong X'$ which restricts to an isometry on each parallelotope), so the fact that corresponding edges have the same shearing is clear.

For the converse, suppose that $i\colon (X, d_{\mathcal{G}}) \to (X', d_{\mathcal{G}'})$ is an isometry of underlying metric spaces such that $h' \simeq i \circ h$. Lift i to an equivariant isometry $\tilde{i}\colon (\widetilde{X}, d_{\mathcal{G}}) \to (\widetilde{X}', d_{\mathcal{G}'})$. By Lemma 7.13, the CAT(0) metric and the marking completely determine the twist-minimal hyperplanes, as well as the width of their carriers. Hence $\tilde{i}$ maps each twist-minimal hyperplane $\widetilde{H}$ to a twist-minimal hyperplane $\tilde{i}(\widetilde{H})$. The assumption on shearings now implies that the image of an edge dual to $\widetilde{H}$ is parallel to any edge dual to $\tilde{i}(\widetilde{H})$ in $(\widetilde{X}', \mathcal{G}')$. To show that $\tilde{i}$ is a combinatorial isometry, we will show that it also maps twist-dominant hyperplanes in $(\widetilde{X}, \mathcal{G})$ bijectively to twist-dominant hyperplanes in $(\widetilde{X}', \mathcal{G}')$.

Suppose that v is twist-dominant and does not lie in the center of A_Γ. Then by Propositions 4.13 and 7.7, the hyperplanes split by v are completely determined by the projection maps $\mathrm{pr}_v^{\mathcal{G}}$, $\mathrm{pr}_v^{\mathcal{G}'}$. Thus, to show that i preserves these hyperplanes, it suffices to show that these two projection maps agree. In both cases, the projection map may be thought of as performing a hyperplane collapse along all hyperplanes

$\widetilde{H} \subseteq \mathrm{Min}(v)$ whose maximal equivalence class commutes with v, where the collapse map takes the dual edges e_H to a point. Since i takes twist-minimal hyperplanes to twist-minimal hyperplanes preserving the shearing and length of their dual edges, $\mathrm{pr}_v^{\mathcal{G}}$ and $\mathrm{pr}_v^{\mathcal{G}'}$ agree on twist-minimal hyperplanes $\widetilde{H} \subseteq \mathrm{Min}(v)$. If the maximal element $w \in \mathrm{lk}(v)$ is twist-dominant, then the dual edge to $\widetilde{H}$ lies along an axis α_w by Lemma 5.8. The entire axis is collapsed to a point under either of these projections. Since i takes axes of twist-dominant generators to axes of twist-dominant generators, $\mathrm{pr}_v^{\mathcal{G}}$ and $\mathrm{pr}_v^{\mathcal{G}'}$ also agree along twist-dominant $\widetilde{H} \subseteq \mathrm{Min}(v)$. We conclude that the two projection maps are the same and hence determine the same twist-dominant hyperplanes.

When A_Γ has nontrivial center, X and X' decompose as (nonorthogonal) products with a locally convex torus endowed with a flat metric. In each case, the parallelotope structure on the torus consists of a single parallelotope with opposite faces identified. In particular, any edge in the 1-skeleton of this torus is the image of an axis of some central element. As i is an isometry and $h' \simeq i \circ h$, the torus factors in X and X' agree as marked, metric tori. Thus we may write $X = Z \times T$, $X' = Z' \times T$, where Z, Z' are subcomplexes, and i maps every edge of the T-factor in X parallel to an edge of the T-factor of X'. The above argument now shows that the combinatorial structure on Z and Z' must also agree, and that for every edge e in the 1-skeleton X, $i(e)$ differs by translation in T from an edge in the 1-skeleton of X'. Since the 1-skeleton of X is connected, i differs from a combinatorial isometry by some fixed translation in T. Since any translation is isotopic to the identity, post-composing i with the inverse of this translation gives a combinatorial isometry $i'\colon (X,\mathcal{G}) \to (X',\mathcal{G}')$ which still satisfies $h' \simeq h \circ i'$. □

COROLLARY 7.17
The composite map $\Sigma_\Gamma \hookrightarrow \mathcal{T}_\Gamma \xrightarrow{\Theta} \mathcal{O}_\Gamma$ that forgets the cube complex structure on $[X,\mathcal{G},h]$, is an embedding.

Proof
Suppose that $[X,\mathcal{G},h]$, $[X',\mathcal{G}',h']$ are two rectilinear Γ-complexes in the fiber over $[Y,d,h] \in \mathcal{O}_\Gamma$. Then there is an isometry $i : (X', d_{\mathcal{G}'}) \to (X, d_{\mathcal{G}})$ such that $h' \simeq i \circ h$. Since $\mathcal{G}$, $\mathcal{G}'$ are rectilinear, no shearing of edges dual to twist-minimal hyperplanes is allowed, so by Proposition 7.16, $[X,\mathcal{G},h] = [X',\mathcal{G}',h']$ in $\mathcal{T}_\Gamma$, and hence also in Σ_Γ. □

7.3.3. Zero-sum shearings
We now want to show that given any shearing of $(\mathbb{S}^\Pi, \mathcal{F})$ satisfying a certain *zero-sum condition*, there is a skewed Γ-complex structure $(X,\mathcal{G})$ on Y with that shearing.

Together with Proposition 7.16 this gives us a characterization of all points in the fiber, which we can then use to prove that the fiber is contractible.

Let v be a twist-minimal vertex of Γ. For any $A \subset \Pi \cup V$ with $\max(A) \geq_f v$, we have $\mathrm{lk}^+(A) \subseteq \mathrm{lk}^+(v) \cup \mathrm{dlk}(v)$, so the shearing vector $s_A \in U_A^+$ decomposes as

$$s_A = \ell_A^v + f_A^v,$$

where the first factor lies in the subspace spanned by axes of $\mathrm{lk}^+(A) \cap \mathrm{lk}^+(v)$ and the second by the axes of $\mathrm{lk}^+(A) \cap \mathrm{dlk}(v)$. Note that if $v \in \max(A)$, then $f_A^v = 0$.

Now let χ_v be a characteristic cycle for v in $\mathbb{S}^\Pi$. Let $H_{A_1}, \dots, H_{A_k}$ be the hyperplanes crossed by χ_v, and orient the dual edges to be consistent with the orientation of e_v. For all i, we have $\max(A_i) \geq_f v$, so $s_{A_i} = \ell_{A_i}^v + f_{A_i}^v$. Viewing all of the $\ell_{A_i}^v$ as vectors in the subspace of U_A^+ spanned by axes of $\mathrm{lk}^+(v)$, we can define ℓ_v to be the sum

$$\ell_v = \sum_i \ell_{A_i}^v.$$

Definition 7.18
A shearing $\{s_A\}$ of $(\mathbb{S}^\Pi, \mathcal{F})$ is a *zero-sum shearing* if $\ell_v = 0$ for all twist-minimal v.

PROPOSITION 7.19
If the images of $[X, \mathcal{G}, h]$ *and* $[\mathbb{S}^\Pi, d_\mathcal{F}, c_\pi^\mathcal{F}]$ *in* $\mathcal{O}_\Gamma$ *are equal, then* $(X, \mathcal{G})$ *differs from* $(\mathbb{S}^\Pi, \mathcal{F})$ *by a zero-sum shearing.*

Proof
If $[X, \mathcal{G}, h]$ and $[\mathbb{S}^\Pi, d_\mathcal{F}, c_\pi^\mathcal{F}]$ have the same image in $\mathcal{O}_\Gamma$, then there is an isometry $i\colon (X, d_\mathcal{G}) \to (\mathbb{S}^\Pi, d_\mathcal{G})$ such that $h \simeq c_\pi^\mathcal{F} \circ i$. Any such isometry lifts to an equivariant isometry on universal covers that takes minsets to minsets, axes to axes, and twist-minimal hyperplanes to twist-minimal hyperplanes. Let u be twist-minimal, and let $\chi_u \subseteq \mathbb{S}^\Pi$ be a characteristic cycle for u beginning at a vertex in the image of the branch locus $\mathrm{br}(u)$. Let η_u be a minimal-length edge path in $(X, \mathcal{G})$ homotopic to $i^{-1}(\chi_u)$. Then η_u and χ_u cross the same twist-minimal hyperplanes and lift to homotopic paths in $\widetilde{X} \cong \widetilde{\mathbb{S}}^\Pi$ with endpoints on some axis for u. Thus η_u is a characteristic cycle for u in $(X, \mathcal{G})$. Since only twist-minimal hyperplanes contribute to the total shearing along η_u, we conclude that $\ell_u = 0$. □

Conversely, we claim that any zero-sum shearing corresponds to a point in the fiber.

PROPOSITION 7.20
Let (Y,d,h) be the image of $(\mathbb{S}^\Pi,\mathcal{F},c_\pi^\mathcal{F})$ in $\mathcal{O}_\Gamma$. Any zero-sum shearing of $(\mathbb{S}^\Pi,\mathcal{F})$ is realized by some skewed Γ-complex structure $(X,\mathcal{G})$ on Y such that h is untwisted with respect to this structure and hence $(X,\mathcal{G},h)$ represents a point in the fiber over (Y,d,h).

Before proving this proposition we deduce the following important corollary, which characterizes the fiber in terms of zero-sum shearings.

COROLLARY 7.21
$\Theta(X,\mathcal{G},h)=\Theta(\mathbb{S}^\Pi,\mathcal{F},c_\pi^\mathcal{F})$ if and only if $(X,\mathcal{G})$ differs by a zero-sum shearing from $(\mathbb{S}^\Pi,\mathcal{F})$ and $h\simeq c_\pi^\mathcal{F}\circ i$ for some isometry $i:X\to\mathbb{S}^\Pi$.

Proof
If $\Theta(X,\mathcal{G},h)=\Theta(\mathbb{S}^\Pi,\mathcal{F},c_\pi^\mathcal{F})$, then there exists an isometry $i:(X,d_\mathcal{G})\to(\mathbb{S}^\Pi,d_\mathcal{F})$ such that $h\simeq c_\pi^\mathcal{F}\circ i$, and $(X,\mathcal{G},h)$ differs by a zero-sum shearing from $(\mathbb{S}^\Pi,\mathcal{F},c_\pi^\mathcal{F})$ by Proposition 7.19. Conversely, by Proposition 7.20, if $(X,\mathcal{G})$ differs by a zero-sum shearing from $(\mathbb{S}^\Pi,\mathcal{F})$, then $(X,\mathcal{G})$ is a skewed Γ-complex with an isometry $i:(X,d_\mathcal{G})\to(\mathbb{S}^\Pi,d_\mathcal{F})$ such that $c_\pi^\mathcal{F}\circ i$ is untwisted. Since $h\simeq c_\pi^\mathcal{F}\circ i$, we conclude that $\Theta(X,\mathcal{G},h)=\Theta(\mathbb{S}^\Pi,\mathcal{F},c_\pi^\mathcal{F})$. □

The proof of Proposition 7.20 will occupy the rest of this subsection. As in the proof of surjectivity, we need to find a new decomposition of Y into parallelotopes, a corresponding skewed blowup structure $(\mathbb{S}^\Omega,\mathcal{G})$, and then determine the change of marking $c_\omega^\mathcal{G}\circ(c_\pi^\mathcal{F})^{-1}$.

For each A_i appearing in the characteristic cycle for v, $\ell^v_{A_i}$ decomposes into a sum of components lying along axes for $w\in\mathrm{lk}^+(v)$. The zero-sum condition, $\ell_v=0$, implies that the components of $\ell^v_{A_i}$ along the axis for each w also sum to zero. This means that we can achieve any zero-sum shearing by ordering the twist-dominant elements w_i, then first performing all shears towards w_1, then w_2, and so on. At each stage, we will verify that the resulting parallelotope structure is a skewed Γ-complex with an untwisted marking. At the final stage, we arrive at a skewed Γ-complex $(X,\mathcal{G})$ that differs from $(\mathbb{S}^\Pi,\mathcal{F})$ by the original zero-sum shearing. This will prove the proposition.

Parallelotope decomposition. Assume that we are shearing toward a single twist-dominant element w. We will define the new parallelotope decomposition by determining the hyperplanes dual to the parallelotopes. The twist-minimal hyperplanes in the structure $(\mathbb{S}^\Pi,\mathcal{F})$ will remain hyperplanes in the new decomposition, and we will eventually identify these with the twist-minimal hyperplanes in a new skewed

Γ-complex structure $(X, \mathcal{G})$. The twist-dominant hyperplanes in $(X, \mathcal{G})$ with maximal element w will be defined using a projection of $\mathrm{Min}(w)$ to an axes for w. The remaining twist-dominant hyperplanes will remain unchanged.

Choose a basepoint x_0 in $\mathrm{Min}(w)$ which is the terminal vertex of an edge labeled w and a branch point for w in the structure $(\widetilde{\mathbb{S}}^{\Pi}, \mathcal{F})$. Let α_w be the axis through x_0, viewed as a copy of the real line, based at x_0. We already have one projection $\mathrm{pr}_w^{\mathcal{F}} = s_{\mathcal{F}}^{-1} \mathrm{pr}_w s_{\mathcal{F}}$ from $\mathrm{Min}(w)$ to α_w defined using the skewed blowup structure $(\mathbb{S}^{\Pi}, \mathcal{F})$ on Y. The image of the branch locus $\mathrm{br}(w)$ under this projection is a set of isolated points dividing α_w into edges, and the inverse image of the midpoints of these edges are the hyperplanes H_A with $\max(A) = w$. The image of any edge $e_A \in \mathrm{Min}(w)$ with $\max(A) \neq w$ is a vertex of α_w, while every axis for w is sent isomorphically to α_w.

Since w is twist-dominant the subspace $\mathrm{Min}(w)$ is a subcomplex of $(\widetilde{\mathbb{S}}^{\Pi}, \mathcal{E})$, and therefore also of $(\widetilde{\mathbb{S}}^{\Pi}, \mathcal{F})$ by Proposition 7.7. We define a new projection map on edges of this subcomplex as follows. Every (oriented) axis for w can be identified with the real line $\mathbb{R}$ and this identification is unique up to translation. Thus, segments of an axis can be viewed as vectors in $\mathbb{R}$ (up to translation). We first associate such a vector r_A to each oriented edge e_A in $\mathrm{Min}(w)$. If $\max(A) = w$, then e_A lies in an axis for w and we let r_A be the corresponding vector in $\mathbb{R}$. If $\max(A) \neq w$, then the shearing of e_A is given by a vector $s_A \in U_A^+$. Since we are only shearing toward w, s_A lies along an axis for w and we let $r_A = -s_A$. Note that if $w \notin \mathrm{lk}^+(A)$, then by definition of an allowable shearing, $r_A = 0$.

Now define the new projection map $\mathrm{pr}'_w : \mathrm{Min}(w) \to \alpha_w$ as follows. For any vertex y in $\mathrm{Min}(w)$, choose a minimal-length edge path $e_{A_1} \cdots e_{A_k}$ from x_0 to y, and set $\mathrm{pr}'_w(y) = x_0 + \sum r_{A_i}$. Since the vectors r_A depend only on the label A and the orientation of e_A, this is independent of the choice of path and two vertices connected by an edge e_A will project to points that differ by the vector r_A. Extending this map linearly on each parallelotope gives the desired projection.

We remark that pr'_w can also be viewed as the map which collapses every hyperplane $\widetilde{H}_A$ in $\mathrm{Min}(w)$ that does not split w. The collapse is performed by identifying the hyperplane carrier with the product $e'_A \times \widetilde{H}_A$, where e'_A is the sheared version of e_A (i.e., an interval parallel to $e_A + s_A$) and collapsing every copy of e'_A to a point.

Now let v be any generator that commutes with w, and let χ_v be a characteristic cycle for v in the structure $(\mathbb{S}^{\Pi}, \mathcal{F})$. Then χ_v lifts to a path $p = e_{A_1} \cdots e_{A_k}$ in $\mathrm{Min}(w)$. Since we are only allowing shearing in the direction of w, for each edge e_{A_i} in p, we have $r_{A_i} = s_{A_i} = \ell_{A_i}^v$, so the zero-sum shearing condition says that $\sum r_{A_i} = 0$, or in other words, the two endpoints y and vy of p project to the same point under pr'_w. It follows that pr'_w is equivariant under the action of $A_{\mathrm{st}(w)}$.

Next observe that if an edge e_A in $\mathrm{Min}(w)$ is contained in the branch locus $\mathrm{br}(w)$, then $\max(A)$ must commute with some $u \notin \mathrm{lk}(w)$. Thus $w \notin \mathrm{lk}^+(A)$ and hence pr'_w maps e_A to a single point. It follows that pr'_w takes each connected component of the branch locus to a single point. We declare these projection points to be the new vertices of α_w; note that x_0 is one of these vertices. This subdivides α_w into a new set of edges. The inverse image under pr'_w of these edges form the carriers of the new hyperplanes that split w.

Now consider the hyperplane structure on $\widetilde{Y}$ consisting of the original hyperplanes which do not split w, together with the new hyperplanes that split w. These determine a new (equivariant) parallelotope structure $(\widetilde{X}, \mathscr{G})$: the maximal parallelotopes in $(\widetilde{X}, \mathscr{G})$ are maximal intersections of carriers of these hyperplanes.

More explicitly, parallelotopes in $(\widetilde{\mathbb{S}}^\Pi, \mathscr{F})$ containing no edges e_A with $\max(A) \leq_t w$, remain unchanged in $(\widetilde{X}, \mathscr{G})$. In particular, this is true for all parallelotopes not contained in $\mathrm{Min}(w)$. The $(\widetilde{X}, \mathscr{G})$-structure on $\mathrm{Min}(w)$ consists of parallelotopes whose edges either lie in an axis for w and project under pr'_w to a single edge in α_w, or are parallel to $e_A + s_A$ in some $\kappa(\widetilde{H}_A)$ and project to a single point in α_w. By the equivariance of pr'_w, this descends to a parallelotope structure on the image of $\mathrm{Min}(w)$ in $\mathbb{S}^\Pi$.

It remains to check that this new parallelotope structure is allowable in the sense of Definition 5.3. To see this, note that an allowable metric on a single parallelotope $\mathbf{c}$, as defined in Definition 5.1, depends on the intersections of linear subspaces K_i associated to edges e_i emanating from a fixed vertex. In our current terminology, if $e_i = e_A$, then K_i is the subspace spanned by e_A together with U_A^+. Since this subspace remains unchanged after shearing, the resulting metric on $\mathbf{c}$ is still allowable, so condition (1) of the definition is satisfied. Condition (2), that if $\mathbf{c}'$ is a face of $\mathbf{c}$, then the metric on $\mathbf{c}'$ is the restriction of the metric on $\mathbf{c}$, is obvious. For condition (3), note that if $\max(A) = \{v\}$ is twist-dominant, then both e_A and e_v lie in the image $\overline{\alpha}_v$ of an axis for v and neither of these edges are allowed to shear. Thus if B is adjacent to A, then any change in angle between e_B and e_A or e_v must result from a shearing of the edge e_B. This can only occur if B is twist-minimal, in which case H_B is locally convex and contains $\overline{\alpha}_v$. It follows that any shearing of e_B will change the angles between e_B and any edge in $\overline{\alpha}_v$ by the same amount.

Blowup structure. We have found a new decomposition $(X, \mathscr{G})$ of Y into parallelotopes. The next thing to show is that $(X, \mathscr{G})$ is a Γ-complex; that is, we need to find a new set of partitions Ω such that $(X, \mathscr{G}) = (\mathbb{S}^\Omega, \mathscr{G})$. The only difference between Π and Ω will be the partitions that split our twist-dominant generator w.

Since w is twist-dominant, $\mathrm{Min}(w)$ is a (convex) subcomplex of $(\widetilde{\mathbb{S}}^\Pi, \mathscr{F})$ by Proposition 7.7. If x and y are vertices of $\mathrm{Min}(w)$ which are branch points for w, then there are edges e_A adjacent to x and e_B adjacent to y with $[A, w] \neq 1$ and

$[B, w] \neq 1$. Choose a lift of $\mathbb{C}^{\Pi}$ to $\widetilde{\mathbb{S}}^{\Pi}$; we will abuse notation by calling this $\mathbb{C}^{\Pi}$ as well. Let M_w be the intersection of $\mathrm{Min}(w)$ with $\mathbb{C}^{\Pi}$. Since $\mathbb{C}^{\Pi}$ and $\mathrm{Min}(w)$ are both convex in the straightened version $(\widetilde{\mathbb{S}}^{\Pi}, \mathcal{E})$, their intersection is connected. Since $\mathbb{C}^{\Pi}$ contains all vertices of $\mathbb{S}^{\Pi}$, every branch vertex in $\mathrm{Min}(w)$ has a unique translate in $M_w = \mathbb{C}^{\Pi} \cap \mathrm{Min}(w)$.

LEMMA 7.22

Suppose that e_A and e_B are edges branching off of M_w at vertices x and y, respectively. If A is a partition, let $A^{\times}$ denote the side of A that does not contain w, and if A is a vertex v, let $A^{\times} = v$ if e_v terminates at x, and let $A^{\times} = v^{-1}$ if x is the initial vertex of e_v; define $B^{\times}$ similarly. If $a \in A^{\times}$ is maximal in A and $b \in B^{\times}$ is maximal in B, and a, b lie in the same w-component of $\Gamma^{\pm}$, then x and y project to the same point of α_w under pr'_w.

Proof

Since M_w is a connected subcomplex, we may connect x and y by a minimal-length edge path $e_{A_1}, \ldots, e_{A_r}$ lying in this intersection. We claim that $\max(A_i) \not\leq_t w$ for all i, so each e_{A_i} collapses to a point under pr'_w. Thus pr'_w maps the entire path to a point, showing that $\mathrm{pr}'_w(x) = \mathrm{pr}'_w(y)$.

We argue by contradiction, so let $a_i \in \max(A_i)$, and suppose that $a_i \leq_t w$ for some i. Since $e_{A_i} \subset \mathrm{Min}(w)$, we have $[a_i, w] = 1$. If $[a, a_i] = 1$, then $a_i \leq_t w$ implies that $[a, w] = 1$, so $e_A \subset \mathrm{Min}(w)$, contradicting our hypothesis. Thus we have $[a, a_i] \neq 1$ for all i, and similarly $[b, a_i] \neq 1$.

The lift of the hyperplane H_{A_i} containing e_{A_i} separates M_w into two components. Since a_i does not commute with either a or b, the endpoints of edges labeled e_A and those labeled e_B lie in different components (where orientation matters if A or B is a generator). In terms of partitions, the sides of A and B that do not contain a_i sit in different sides of the partition A_i. Since $a_i \in \mathrm{lk}(w)$ and $A^{\times}$ does not contain w it does not contain a_i either, and similarly $B^{\times}$ does not contain a_i. Thus $A^{\times}$ and $B^{\times}$ are in different sides of A_i. But each side of A_i is a union of a_i-components plus a_i or a_i^{-1}, and, since $a_i \leq_t w$, each a_i-component is a union of w-components plus possibly some elements of $\mathrm{lk}(w)$. Since $a \in A^{\times}$, $b \in B^{\times}$ and neither is in $\mathrm{lk}(w)$, this contradicts the hypothesis that a and b are in the same w-component. □

We now form the new partitions splitting w in the same way we did in Section 7.2. To each branch point $\mathbb{r} \in \mathrm{br}(w) \cap \mathbb{C}^{\Pi}$, associate the union $I(\mathbb{r})$ of the sets $A^{\times}$ for edges e_A incident to $\mathbb{r}$ but not in $\mathrm{Min}(w)$. The new projection pr'_w sends $\mathrm{br}(w) \cap M_w$ to an ordered set of points $(\mathbb{x}_1, \ldots, \mathbb{x}_n)$ on $\alpha_w \cap \mathbb{C}^{\Pi}$, and to each $\mathbb{x}_i$ we associate the union $I(\mathbb{x}_i)$ of the $I(\mathbb{r})$ with $\mathrm{pr}'(\mathbb{r}) = x_i$. Let $P'_1 = \{w\} \cup I(\mathbb{x}_1)$, $P'_2 = P'_1 \cup I(\mathbb{x}_2)$, and so on. By Lemma 7.22, each $I(\mathbb{x}_i)$ is a union of w-components, so each P_i' is a side of a valid Γ-Whitehead partition $\mathcal{P}_i'$ based at w.

Let Ω be the collection of Γ-Whitehead partitions obtained from Π by replacing $\mathcal{P}_1, \dots \mathcal{P}_k$ in Π by $\mathcal{P}_1', \dots \mathcal{P}_n'$. To see that the Ω partitions are pairwise compatible, we need only check that each $\mathcal{P}_i'$ is compatible with those $\mathcal{R} \in \Pi$ that are not adjacent to w. The side $R^\times$ is contained in some outermost $Q^\times$ in some piece dP_i. The partition $\mathcal{Q}$ cannot be adjacent to w, so there is an edge $e_{\mathcal{Q}}$ at a branch point $\mathbb{r} \in M_w$, and $Q^\times \subset I(\mathbb{r})$. Since $I(\mathbb{r}) \subset I(\mathbb{x}_i)$ for some i. it follows that $\mathcal{R}$ is compatible with $\mathcal{P}_i'$.

Marking change. The blowup structure $(X = \mathbb{S}^\Omega, \mathcal{G})$ defined above comes with a collapse map $c_\omega^\mathcal{G} \colon \mathbb{S}^\Omega \to \mathbb{S}_\Gamma$. We now analyze the change in marking induced by the difference between $c_\omega^\mathcal{G}$ and the original collapse map $c_\pi^\mathcal{F}$ from $(\mathbb{S}^\Pi, \mathcal{F})$.

LEMMA 7.23
Suppose that $(X = \mathbb{S}^\Omega, \mathcal{G}, c_\omega^\mathcal{G})$ is a zero-sum shearing of $(\mathbb{S}^\Pi, \mathcal{F}, c_\pi^\mathcal{F}) \in \mathcal{T}_\Gamma$ which differs only in the direction of a twist-dominant generator w. Then the composite map $c_\omega^\mathcal{G} \circ (c_\pi^\mathcal{F})^{-1} : \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ is untwisted.

Proof
Let $\mu = c_\omega^\mathcal{G} \circ (c_\pi^\mathcal{F})^{-1} : \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$. Observe that the only hyperplanes which change from $\mathbb{S}^\Pi$ to $\mathbb{S}^\Omega$ are those with max $= \{w\}$. Following exactly the same argument as in the proof of Lemma 7.9, for each $v \in V$ we have that $\mu(v) = w^{n_v} v w^{m_v}$ for some $n_v, m_v \in \mathbb{Z}$. Thus, nontrivial twists can occur only for $v \leq_t w$. If such a v is twist-dominant, then the characteristic cycle for v is the same in both $\mathbb{S}^\Pi$ and $\mathbb{S}^\Omega$, so $\mu(v) = v$. If v is twist-minimal, then the fact that $(\mathbb{S}^\Omega, \mathcal{G})$ is a zero-sum shearing implies that a characteristic cycle for v with respect to $(\mathbb{S}^\Omega, \mathcal{G})$ has the same endpoints as a characteristic cycle for v with respect to $(\mathbb{S}^\Pi, \mathcal{F})$, hence in this case $\mu(v) = v$ as well. Therefore, μ is untwisted. □

The proof of Lemma 7.23 shows that μ acts trivially on vertices $v \leq_t w$, and one might be tempted to conclude that it shows that the action of μ is entirely trivial. However, if $z \leq_f v \leq_t w$, then a characteristic cycle for z may have an edge e_A that also lies in a characteristic cycle for v. In this case, a zero-sum shearing of the characteristic cycle for v in the direction of w may result in μ acting as a nontrivial fold of w onto z.

Remark 7.24
As observed in the proof of Lemma 7.23, the only hyperplanes that change in a zero-sum shearing are twist-dominant hyperplanes. In particular, the set of twist-minimal hyperplanes which split a particular v does not vary among all zero-sum shearings.

We now finish the proof of Proposition 7.20.

Proof of Proposition 7.20
Order the twist-dominant generators $w_1, \dots, w_n$. We perform an arbitrary zero-sum shearing as a sequence of single generator zero-sum shearings. By the discussion above, we obtain a sequence of skewed blowups

$$(\mathbb{S}^{\Pi}, \mathcal{F}) = (\mathbb{S}^{\Omega_0}, \mathcal{G}_0), (\mathbb{S}^{\Omega_1}, \mathcal{G}_1), \dots, (\mathbb{S}^{\Omega_n}, \mathcal{G}_n) = (\mathbb{S}^{\Omega}, \mathcal{G}),$$

where for $1 \le i \le n$, $(\mathbb{S}^{\Omega_i}, \mathcal{G}_i)$ is obtained from $(\mathbb{S}^{\Omega_{i-1}}, \mathcal{G}_{i-1})$ by a zero-sum shearing in the direction of w_i, and the change in marking μ_i is untwisted by Lemma 7.23. The change in marking from $(\mathbb{S}^{\Pi}, \mathcal{F})$ to $(\mathbb{S}^{\Omega}, \mathcal{G})$ is then a composition of untwisted automorphisms

$$c_{\omega}^{\mathcal{G}} \circ (c_{\pi}^{\mathcal{F}})^{-1} \simeq \mu_n \circ \cdots \circ \mu_1,$$

hence untwisted as well. □

7.4. Contractibility of $\mathcal{O}_\Gamma$

7.4.1. Contractibility of fibers
Let $\mathcal{H}_{\min}$ denote the set of twist-minimal hyperplanes in $(\mathbb{S}^{\Pi}, \mathcal{F}, c_{\pi}^{\mathcal{F}})$. Since these depend only on the metric $d_{\mathcal{F}}$ by Lemma 7.13, the set $\mathcal{H}_{\min}$ is well defined over the whole Θ-fiber containing $(\mathbb{S}^{\Pi}, \mathcal{F}, c_{\pi}^{\mathcal{F}})$. Likewise, the twist-dominant axes remain the same throughout the fiber. The dual edge to $H \in \mathcal{H}_{\min}$ is allowed to shear in the direction of $\mathrm{lk}^+(H)$, and as above, we regard a given shearing s_H as a vector in the vector space U_H^+. (Here to emphasize the independence from Π, we use the notation s_H and U_H^+ rather than specifying a label A and writing s_A and U_A^+.) We now describe the fiber containing $(\mathbb{S}^{\Pi}, \mathcal{F}, c_{\pi}^{\mathcal{F}})$ as a linear subspace of

$$\bigoplus_{H \in \mathcal{H}_{\min}} U_H^+.$$

Let $V_{\min}$ denote the set of twist-minimal vertices. For $v \in V_{\min}$, the only edges which contribute to the shearing of v are those dual to $H \in \mathcal{H}_{\min}$ which split v. By Remark 7.24 and Corollary 7.21, the set of twist-minimal hyperplanes that split v does not change within the fiber. If $v \in \mathrm{split}(H)$, $H \in \mathcal{H}_{\min}$, then the contribution of s_H to l_v is l_H^v, where l_H^v lies in the subspace of U_H^+ corresponding to $\mathrm{lk}^+(H) \cap \mathrm{lk}^+(v)$. If $v \notin \mathrm{split}(H)$, then define $l_H^v \equiv 0$. We then identify l_H^v with a vector in U_v^+ since l_H^v lies in the span of axes in $\mathrm{lk}^+(H) \cap \mathrm{lk}^+(v)$. Thus, for each $v \in V_{\min}$ we can think of ℓ_v as a linear map:

$$l_v : \bigoplus_{H \in \mathcal{H}_{\min}} U_H^+ \to U_v^+,$$
$$\oplus s_H \longmapsto \sum l_H^v.$$

Call the equations $\{l_v = 0 \mid v \in V_{\min}\}$ the *structure equations* for shearings of $(\mathbb{S}^\Pi, \mathcal{F}, c_\pi^{\mathcal{F}})$. We now easily deduce the contractibility of the fibers from Corollary 7.21.

THEOREM 7.25
The fibers of the map $\Theta : \mathcal{T}_\Gamma \to \mathcal{O}_\Gamma$ are contractible.

Proof
The space of solutions to the structure equations is the intersection $\bigcap_{v \in V_{\min}} \ker l_v$, which is a linear subspace of $\bigoplus_{H \in \mathcal{H}_{\min}} U_H^+$ and hence contractible. The preceding discussion shows that this subspace is in one-to-one correspondence with the set of zero-sum shearings of $(\mathbb{S}^\Pi, \mathcal{F}, c_\pi^{\mathcal{F}})$. Thus, by Corollary 7.21, there is a bijection between the space of solutions and points in $\Theta^{-1}([\mathbb{S}^\Pi, d_{\mathcal{F}}, c_\pi^{\mathcal{F}}])$. It is easy to see that this correspondence is a homeomorphism. By Proposition 7.4, every fiber of Θ is a $U(A_\Gamma)$-translate of one containing some $[\mathbb{S}^\Pi, \mathcal{F}, c_\pi^{\mathcal{F}}]$, so every fiber is contractible. □

7.4.2. Contractibility of $\mathcal{O}_\Gamma$
We now finish the proof of Theorem 7.2. By Theorem 7.25, the fibers of Θ are contractible, but since they are not compact, Θ is not a proper map. To conclude that $\mathcal{O}_\Gamma$ is contractible, we will show that Θ is in fact a fibration.

Proof of Theorem 7.2
Since $\mathcal{O}_\Gamma$ is paracompact (the equivariant Gromov–Hausdorff topology is metrizable), it suffices to show that Θ is a fibration when restricted to sufficiently small neighborhoods $U \subset \mathcal{O}_\Gamma$.

We begin by showing that for any point y_0 in $\mathcal{O}_\Gamma$, and any lift $x_0 \in \mathcal{T}_\Gamma$ of y_0, there exist a neighborhood U and a section $s : U \to \Theta^{-1}(U)$ with $s(y_0) = x_0$. Say $x_0 = [X_0, \mathcal{F}_0, h_0]$ and $y_0 = [Y_0, d_0, h_0]$, so that $(X_0, d_{\mathcal{F}_0})$ is isometric to (Y_0, d_0). By Proposition 7.4, it suffices to consider the case when $h_0 = c_0$ is a collapse map.

Consider the fiber over a point $y = [Y, d, c]$ in a small neighborhood U of y_0. To define $s(y)$, we must choose a Γ-complex structure $(X, \mathcal{F})$ on (Y, d). For any such $\mathcal{F}$, the twist-minimal hyperplanes with v as a maximal element are determined by the projection of the branch locus $\mathrm{br}(v)$ on an axis for v. If (Y, d, c) is close to (Y_0, d_0, c_0) in the equivariant Gromov–Hausdorff topology, then these branch loci must also be close, and hence likewise their projections on an axis for v. However, these projections can change in three ways as we move from $[Y_0, d_0, c_0]$ to $[Y, d, c]$.

- The distance between a pair of projection points may expand or contract. This will affect the width of the carrier of the hyperplane separating these projection points.
- One projection point can split into multiple points. This will require introducing new twist-minimal hyperplanes.
- Two or more projection points may coalesce, causing the corresponding hyperplanes to merge.

Shrinking U if necessary, we may avoid the coalescing of projection points and allow only changes of the first two types. Moreover, for U sufficiently small, the new twist-minimal hyperplanes will have carriers of width less than half that of the old twist-minimal hyperplanes, and thus (by abuse of notation) we may consider the set of twist-minimal hyperplanes in $\mathcal{F}_0$ to be a subset of those in $\mathcal{F}$. Then the marking $c : X \to \mathbb{S}_\Gamma$ will correspond to collapsing the newly added hyperplanes, composed with the straighten-collapse map corresponding to c_0.

The collection of twist-minimal hyperplanes at $[Y, d, c]$ is completely determined by the metric d. The axes of twist-dominant generators are determined by the marking c. As seen in Proposition 7.16, once we have determined the twist-minimal hyperplanes, the shearing of their dual edges together with the branch locus completely determines $\mathcal{F}$. Suppose that H is a new hyperplane, not coming from a hyperplane in $\mathcal{F}_0$. We are free to choose the shearing on the dual edge by any vector in U_H^+. Different choices will only affect the determination of twist-dominant hyperplanes. Therefore, we choose the dual edge to be orthogonal to H, of length equal to the width of $\kappa(H)$. If H corresponds to a twist-minimal hyperplane in $\mathcal{F}_0$ which is collapsed by c_0, then we leave the shearing unchanged (i.e., the angle between the dual edge and the axes of $\mathrm{lk}^+(H)$), but adjust the length of the dual edge to take account of the change in the width of the hyperplane carrier. Finally, for twist-minimal hyperplanes not collapsed by c_0, namely, those labeled H_v, we adjust the shearing so that the new characteristic cycle lifts to a path whose endpoints lie on an axis for v. This determines a parallelotope structure $\mathcal{F}$ on (Y, d) with the property that the shearing along twist-minimal hyperplanes satisfies the zero-sum condition relative to any skewed Γ-complex in the fiber over y. Hence by Corollary 7.21, $[X, \mathcal{F}, c]$ also lies in this fiber. Set $s(y) = [X, \mathcal{F}, c]$. Since the construction of $\mathcal{F}$ depends only on the metric d and lift $[Y, \mathcal{F}_0, c_0]$, the map s is well defined and continuous.

Now let Z be any space. Suppose that $f_t : Z \to U$ is a homotopy, and let $\hat{f}_0 : Z \to \Theta^{-1}(U)$ be a lift of f_0. We can lift f_t to a homotopy $g_t = s \circ f_t : Y \to \Theta^{-1}(U)$, but g_0 need not agree with the given lift $\hat{f}_0$. We can correct this by concatenating g_t with a homotopy h_t from $\hat{f}_0$ to g_0 which projects at all times t to the map f_0. To do this, use the fact that the fibers of Θ are convex subspaces of some Euclidean space, so the straight-line homotopy in each fiber from $\hat{f}_0(y)$ to $g_0(y)$ gives such a homotopy h_t. Then, up to reparameterizing the interval, h_t followed by g_t is a lift of f_t.

This shows that Θ is a fibration. Since we have already proved that the fibers are contractible (Theorem 7.25), we conclude that Θ is a homotopy equivalence. By Corollary 6.6, $\mathcal{T}_\Gamma$ is contractible, so the same holds for $\mathcal{O}_\Gamma$. □

Acknowledgments. The authors would like to thank the referees for their careful reading of the paper and many helpful comments.

Bregman's work was partially supported by National Science Foundation (NSF) grant DMS-1906269. Charney's work was partially supported by NSF grant DMS-1607616.

References

[1] I. AGOL, *The virtual Haken conjecture*, with appendix "Filling virtually special subgroups" by I. Agol, D. Groves, and J. Manning, Doc. Math. **18** (2013), 1045–1087. MR 3104553. DOI IncorrectDOIlink. *(1034)*

[2] N. BERGERON and D. T. WISE, *A boundary criterion for cubulation*, Amer. J. Math. **134** (2012), no. 3, 843–859. MR 2931226. DOI 10.1353/ajm.2012.0020. *(1034)*

[3] M. BESTVINA and N. BRADY, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), no. 3, 445–470. MR 1465330. DOI 10.1007/s002220050168. *(1034)*

[4] M. BESTVINA and M. FEIGHN, *The topology at infinity of* $\mathrm{Out}(F_n)$, Invent. Math. **140** (2000), no. 3, 651–692. MR 1760754. DOI 10.1007/s002220000068. *(1037)*

[5] J. BEYRER and E. FIORAVANTI, *Cross ratios on* CAT(0) *cube complexes and marked length-spectrum rigidity*, J. Lond. Math. Soc. (2) **104** (2021), no. 5, 1973–2015. MR 4368668. DOI 10.1112/jlms.12489. *(1037)*

[6] A. BOREL and J.-P. SERRE, *Corners and arithmetic groups*, Comment. Math. Helv. **48** (1973), 436–491. MR 0387495. DOI 10.1007/BF02566134. *(1037)*

[7] C. BREGMAN, *Isometry groups of skewed* Γ*-complexes*, preprint, arXiv:2202.09860v1 [math.GT]. *(1080, 1094)*

[8] M. R. BRIDSON and A. HAEFLIGER, *Metric Spaces of Non-Positive Curvature*, Grundlehren Math. Wiss. **319**, Springer, Berlin, 1999. MR 1744486. DOI 10.1007/978-3-662-12494-9. *(1039, 1080)*

[9] M. R. BRIDSON and K. VOGTMANN, *On the geometry of the automorphism group of a free group*, Bull. London Math. Soc. **27** (1995), no. 6, 544–552. MR 1348708. DOI 10.1112/blms/27.6.544. *(1034)*

[10] B. BRÜCK and R. D. WADE, *A note on virtual duality and automorphism groups of right-angled Artin groups*, preprint, arXiv:2101.08225v1 [math.GR]. *(1037)*

[11] R. CHARNEY, J. CRISP, and K. VOGTMANN, *Automorphisms of 2-dimensional right-angled Artin groups*, Geom. Topol. **11** (2007), 2227–2264. MR 2372847. DOI 10.2140/gt.2007.11.2227. *(1034)*

[12] R. CHARNEY, N. STAMBAUGH, and K. VOGTMANN, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geom. Topol. **21** (2017), no. 2, 1131–1178. MR 3626599. DOI 10.2140/gt.2017.21.1131. *(1034, 1035, 1037, 1038, 1040, 1042, 1043, 1045, 1046, 1047, 1048)*

[13] R. CHARNEY and K. VOGTMANN, *Finiteness properties of automorphism groups of right-angled Artin groups*, Bull. Lond. Math. Soc. **41** (2009), no. 1, 94–102. MR 2481994. DOI 10.1112/blms/bdn108. *(1034)*

[14] ———, "Subgroups and quotients of automorphism groups of RAAGs" in *Low-Dimensional and Symplectic Topology*, Proc. Sympos. Pure Math. **82**, Amer. Math. Soc., Providence, 2011, 9–27. MR 2768650. DOI 10.1090/pspum/082/2768650. *(1034)*

[15] M. T. CLAY, C. J. LEININGER, and J. MANGAHAS, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. **6** (2012), no. 2, 249–278. MR 2914860. DOI 10.4171/GGD/157. *(1034)*

[16] J. CRISP and L. PARIS, *The solution to a conjecture of Tits on the subgroup generated by the squares of the generators of an Artin group*, Invent. Math. **145** (2001), no. 1, 19–36. MR 1839284. DOI 10.1007/s002220100138. *(1034)*

[17] C. B. CROKE and B. KLEINER, *Spaces with nonpositive curvature and their ideal boundaries*, Topology **39** (2000), no. 3, 549–556. MR 1746908. DOI 10.1016/S0040-9383(99)00016-6. *(1034)*

[18] M. CULLER and K. VOGTMANN, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119. MR 0830040. DOI 10.1007/BF01388734. *(1033, 1047)*

[19] M. W. DAVIS and T. JANUSZKIEWICZ, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235. MR 1783167. DOI 10.1016/S0022-4049(99)00175-9. *(1034)*

[20] M. B. DAY, *Peak reduction and finite presentations for automorphism groups of right-angled Artin groups*, Geom. Topol. **13** (2009), no. 2, 817–855. MR 2470964. DOI 10.2140/gt.2009.13.817. *(1034)*

[21] M. B. DAY, A. W. SALE, and R. D. WADE, *Calculating the virtual cohomological dimension of the automorphism group of a RAAG*, Bull. Lond. Math. Soc. **53** (2021), no. 1, 259–273. MR 4224528. DOI 10.1112/blms.12418. *(1036)*

[22] M. B. DAY and R. D. WADE, *Relative automorphism groups of right-angled Artin groups*, J. Topol. **12** (2019), no. 3, 759–798. MR 4072157. DOI 10.1112/topo.12101. *(1034)*

[23] D. B. A. EPSTEIN, J. W. CANNON, D. F. HOLT, S. V. F. LEVY, M. S. PATERSON, and W. P. THURSTON, *Word Processing in Groups*, Jones and Bartlett, Boston, 1992. MR 1161694. *(1034)*

[24] F. HAGLUND and D. T. WISE, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620. MR 2377497. DOI 10.1007/s00039-007-0629-4. *(1034)*

[25] A. HATCHER, *Algebraic Topology*, Cambridge Univ. Press, Cambridge, 2002. MR 1867354. *(1078)*

[26] T. KOBERDA, *Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups*, Geom. Funct. Anal. **22**

(2012), no. 6, 1541–1590. MR 3000498. DOI 10.1007/s00039-012-0198-z. *(1034)*

[27] M. R. LAURENCE, *A generating set for the automorphism group of a graph group*, J. Lond. Math. Soc. (2) **52** (1995), no. 2, 318–334. MR 1356145. DOI 10.1112/jlms/52.2.318. *(1046)*

[28] B. MILLARD and K. VOGTMANN, *Cube complexes and abelian subgroups of automorphism groups of RAAGs*, Math. Proc. Cambridge Philos. Soc. **170** (2021), no. 3, 1–25. MR 4243766. DOI 10.1017/S0305004119000501. *(1037)*

[29] L. MOSHER, *Mapping class groups are automatic*, Ann. of Math. (2) **142** (1995), no. 2, 303–384. MR 1343324. DOI 10.2307/2118637. *(1034)*

[30] M. SAGEEV, "CAT(0) cube complexes and groups" in *Geometric Group Theory*, IAS/Park City Math. Ser. **21**, Amer. Math. Soc., Providence, 2014, 7–54. MR 3329724. DOI 10.1090/pcms/021/02. *(1048, 1050)*

[31] H. SERVATIUS, *Automorphisms of graph groups*, J. Algebra **126** (1989), no. 1, 34–60. MR 1023285. DOI 10.1016/0021-8693(89)90319-0. *(1046)*

[32] H. SERVATIUS, C. DROMS, and B. SERVATIUS, *Surface subgroups of graph groups*, Proc. Amer. Math. Soc. **106** (1989), no. 3, 573–578. MR 0952322. DOI 10.2307/2047406. *(1034)*

[33] A. VIJAYAN, *Compactifying the space of length functions of a right-angled Artin group*, Ph.D. dissertation, Brandeis University, Waltham, 2013. MR 3153211. *(1037)*

Bregman
Department of Mathematics and Statistics, University of Southern Maine, Portland, Maine, USA; corey.bregman@maine.edu

Charney
Department of Mathematics, Brandeis University, Waltham, Massachusetts, USA; charney@brandeis.edu

Vogtmann
Mathematics Institute, University of Warwick, Coventry, United Kingdom; k.vogtmann@warwick.ac.uk