Classifying Finite Dimensional Cubulations of Tubular Groups

Daniel J. Woodhouse

ABSTRACT. A tubular group is a group that acts on a tree with $\mathbb{Z}^2$ vertex stabilizers and $\mathbb{Z}$ edge stabilizers. This paper develops further a criterion of Wise and determines when a tubular group acts freely on a finite dimensional CAT(0) cube complex. As a consequence, we offer a unified explanation of the failure of separability by revisiting the nonseparable 3-manifold group of Burns, Karrass, and Solitar and relating it to the work of Rubinstein and Wang. We also prove that if an immersed wall yields an infinite dimensional cubulation, then the corresponding subgroup is quadratically distorted.

1. Introduction

A *tubular group* G is a group that splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and $\mathbb{Z}$ edge groups. A tubular group is the fundamental group of a graph of spaces X with each vertex space homeomorphic to a torus and each edge space homeomorphic to a cylinder. The graph of spaces X is a *tubular space.* In this paper, all tubular groups will be finitely generated, and thus tubular spaces will be compact. Examples of tubular groups were used by Brady and Bridson [1] to provide groups with isoperimetric function n^α for all α in a dense subset of $[2, \infty)$. Cashen [3] provided a method of determining when two tubular groups are quasi-isometric. Wise [9] gave a criterion that determines whether or not a tubular group acts freely on a CAT(0) cube complex and characterized which tubular groups act cocompactly on a CAT(0) cube complex. This paper determines which tubular groups act on finite dimensional CAT(0) cube complexes.

DEFINITION 1.1. An *infinite cube* in a CAT(0) cube complex is the union of an ascending sequence of n-cubes c_n of $\widetilde{X}$ such that c_n is a subcube of c_{n+1} for each n.

Wise reduces the existence of cubulations to a combinatorial criterion called *equitable sets.* Given an equitable set, we can construct a finite set of *immersed walls.* An immersed wall is a graph immersed π_1-injectively in X, such that $\widetilde{\Lambda}$ lifts to a two-sided embedding $\widetilde{\Lambda} \to \widetilde{X}$. By two-sided we mean that the image of $\widetilde{\Lambda}$ in $\widetilde{X}$ is contained in a neighborhood homeomorphic to $\widetilde{\Lambda} \times [-1, 1]$. The set of all such lifts gives a G-invariant set $\mathcal{W}$ of *walls.* The pair $(\widetilde{X}, \mathcal{W})$ is a *wallspace.* In this paper, all references to "immersed walls" will be in reference to immersed

walls obtained from an equitable set. Note that the theorems referring to immersed walls will not apply to any other kind of immersed walls. Section 4 defines the notion of a *dilating wall,* which can be recognized through a combinatorial criterion. As explained in Proposition 4.12, a wall not being dilated means that its G-translates intersecting a vertex space in $\widetilde{X}$ can be partitioned into finitely many sets of pairwise nonintersecting walls.

The following is proven in Theorem 4.10.

THEOREM 1.2. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following are equivalent*:

(1) *The dual cube complex $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional.*
(2) *The dual cube complex $C(\widetilde{X}, \mathcal{W})$ contains an infinite cube.*
(3) *One of the immersed walls is dilated.*

Corollary 4.11 states that it is decidable if a given immersed wall is dilated. However, in contrast to cocompactness, there is no known simple criterion to determine whether or not a given tubular group acts on a finite dimensional CAT(0) cube complex. See Example 5.1 for an example of a direct proof that a specific tubular group does not posses an equitable set that can produce nondilated immersed walls.

A group G is *separable* if every finitely generated subgroup $H \leq G$ is the intersection of all finite index subgroups containing H. Burns, Karrass, and Solitar [2] gave the first example of a nonseparable 3-manifold group. Niblo and Wise [5] reproved this result. Rubinstein and Wang [6] produced an example of a nonembedded immersed surface S in a graph manifold M such that $\widetilde{S} \rightarrow \widetilde{M}$ is injective and any two $\pi_1 M$ translates of $\widetilde{S}$ intersect. It follows that $\pi_1 S$ is not separable in $\pi_1 M$. This is an application of the geometric interpretation for separability given by Scott [8]. Theorem 1.2 carries information about the separability of the associated codimension-1 subgroups of G and provides a new proof of the nonsubgroup separability of certain 3-manifold groups that conceptually unifies the result of Burns, Karrass, and Solitar with the geometric proof of Rubinstein and Wang (see Example 2.1).

Finally, the following theorem is a consequence of Theorem 1.2 together with a study of the geometry of the walls.

THEOREM 1.3. *Let the tubular space X associated to the tubular group G have a finite set of immersed walls. If the corresponding wallspace $(\widetilde{X}, \mathcal{W})$ has quasi-isometrically embedded walls, then the dual CAT(0) cube complex $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Example 6.2 demonstrates that quasi-isometrically embedded walls are not a necessary condition for finite dimensionality.

1.1. Tubular Groups and Their Cubulations

Let G be a tubular group with associated tubular space X. Let X_v and X_e denote vertex and edge spaces in this graph of spaces, where v and e are a vertex and edge in the underlying graph Γ. Each X_v is homeomorphic to the torus $S^1 \times S^1$, and each X_e is homeomorphic to the cylinder $S^1 \times [-1, 1]$. Choosing an orientation for an edge e, let $\overleftarrow{e}$, $\overrightarrow{e}$ denote the initial and terminal vertices, let $\overleftarrow{X_e}$, $\overrightarrow{X_e}$ denote the corresponding boundary circles in X_e, let $\overleftarrow{f_e} : \overleftarrow{X_e} \to X_{\overleftarrow{e}}$, and let $\overrightarrow{f_e} : \overrightarrow{X_e} \to X_{\overrightarrow{e}}$ denote the attaching maps. Let $\widetilde{X}_{\widetilde{v}}$ and $\widetilde{X}_{\widetilde{e}}$ denote vertex and edge spaces in the universal cover where $\widetilde{v}$ and $\widetilde{e}$ are a vertex and edge in the Bass–Serre tree $\widetilde{\Gamma}$. There is a G action on both $\widetilde{X}$ and $\widetilde{\Gamma}$ such that $g\widetilde{X}_{\widetilde{v}} = \widetilde{X}_{g\widetilde{v}}$.

Let $\alpha, \beta : S^1 \to S^1 \times S^1$ be a transverse pair of geodesic immersed circles in a torus. The *geometric intersection number* of α, β is the number of points $(s, t) \in S^1 \times S^1$ such that $\alpha(s) = \beta(t)$. Let $[\alpha]$, $[\beta]$ be a pair of homotopy classes of immersed circles in the torus. The *geometric intersection number* of $[\alpha]$, $[\beta]$ is the geometric intersection number for any choice of transverse geodesic immersed circles representing these homotopy classes. Let $\#[\alpha, \beta]$ denote the geometric intersection number. If $B = \{\beta_i\}$ is a set of geodesic circles, then let $\#[\alpha, B] := \sum_i \#[\alpha, \beta_i]$. Viewing $[\alpha]$, $[\beta]$ as elements of $\pi_1 X_v \cong \mathbb{Z}^2$, we recall from [9] that $\#[\alpha, \beta] = \det[[\alpha], [\beta]]$.

An *equitable set* for G is a collection of sets $\{S_v\}_{v \in V(\Gamma)}$, where each S_v is a finite set of geodesic immersed circles $\alpha : S^1 \to X_v$ such that $\langle [\alpha] \mid \alpha \in S_v \rangle$ generates a finite index subgroup of $\pi_1 X_v$ and $\#[\overleftarrow{f_e}, S_{\overleftarrow{e}}] = \#[\overrightarrow{f_e}, S_{\overrightarrow{e}}]$ for all edges e in Γ. It is often convenient to allow S_v to contain "repeats", that is, distinct elements that are homotopic. It is assumed, however, that distinct parallel circles in a vertex space have disjoint image.

The main theorem proven in [9] is the following:

THEOREM 1.4 (Wise). *A tubular group G acts freely on a* CAT(0) *cube complex if and only if there is an equitable set for G.*

As we repeatedly use this construction, here is an outline of how an equitable set provides a CAT(0) cube complex with a free G-action.

Let each α in the equitable set be given by a distinct copy of S^1 mapping into X. Taking the disjoint union of these circles yields a map $\bigsqcup_{\alpha \in \bigcup S_v} S^1 \to X$. Since $\#[\overleftarrow{f_e}, S_{\overleftarrow{e}}] = \#[\overrightarrow{f_e}, S_{\overrightarrow{e}}]$, we can choose a bijection between the intersection points of $\overleftarrow{f_e}$ and the intersection points of $\overrightarrow{f_e}$. Extend $\sqcup S^1$ to a space homeomorphic to a graph by attaching the endpoints of an arc to each pair of corresponding intersection points in our equitable set. The connected components of the resulting graph are *horizontal immersed walls*. The immersions of the circles into X can be extended along the arcs by mapping each arc into the associated edge space. For each edge space, there is a *vertical immersed wall,* which is a circle embedding along $S^1 \times \{0\} \subseteq S^1 \times [1, -1]$. The complete collection of immersed walls obtained in this manner are the *immersed walls obtained from the equitable*
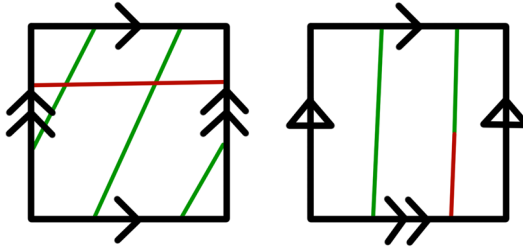
Figure 1 The graph of spaces X, with an immersed wall given by the equitable set

set. All immersed walls mentioned in this paper will be obtained from such a construction. Note further that an equitable set produces a finite set of immersed walls.

Let Λ be an immersed wall. Its universal cover $\widetilde{\Lambda}$ embeds in the universal cover $\widetilde{X}$. A lift of $\widetilde{\Lambda}$ to $\widetilde{X}$ has an image contained in a product neighborhood homeomorphic to $\widetilde{\Lambda} \times [-1, 1]$, so $\widetilde{X} - \widetilde{\Lambda}$ has exactly two connected components. Taking the full set of such lifts, we obtain a wallspace $(\widetilde{X}, \mathcal{W})$, which, via Sageev's construction [7], gives a G-action on a CAT(0) cube complex (see [4] for a comprehensive overview of Sageev's construction).

REMARK 1.5. In the presented construction, the resulting dual CAT(0) cube complex depends on the choice of correspondence in each edge space and on how the interior of each arc embeds in the cylinder. Example 5.3 shows that finite dimensionality of the cubulation is dependent on the choice of correspondence between the intersection points at either end of the edge space. By contrast, Theorem 1.2 and the definition of a dilated immersed wall imply that finite dimensionality is invariant of the subsequent choice of immersion of the interior of the arcs in the edge spaces. In [10], we use the results presented here to show that acting freely on a finite dimensional CAT(0) cube complex is equivalent to virtual specialness for tubular groups.

2. A Motivating Example

The next example motivates the definitions and theorems given later in this paper that identify which immersed walls give infinite cubulations.

EXAMPLE 2.1. Consider the tubular group $G = \langle a, b, t \mid [a, b] = 1, tat^{-1} = b \rangle$. This is a cyclic HNN extension of $\mathbb{Z}^2 = \langle a, b \rangle$. The associated graph of spaces has a torus vertex space X_v and a cylindrical edge space X_e that is attached along geodesic paths associated to a and b. Consider the immersed wall given by the equitable set $\{a, ab^2\}$ and the choice of arcs shown at right in Figure 1.
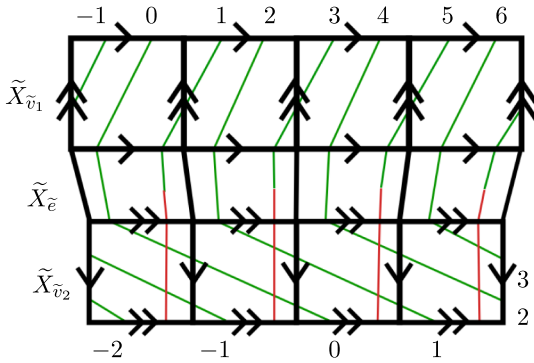
Figure 2 A portion of the universal cover $\widetilde{X}$, consisting of parts of $\widetilde{X}_{\widetilde{v}_1}$, $\widetilde{X}_{\widetilde{e}}$, and $\widetilde{X}_{\widetilde{v}_2}$. The numbers at the top show the indexing of the walls in $\widetilde{X}_{\widetilde{v}_1}$, and the numbers at the bottom show the indexing of the walls in $\widetilde{X}_{\widetilde{v}_2}$

An infinite set of pairwise crossing walls in $(\widetilde{X}, \mathcal{W})$ will be exhibited to demonstrate that $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional. The immersed wall in this example is the prototype for a general dilated immersed wall, which will be discussed in Section 4. Choose a vertex space $\widetilde{X}_{\widetilde{v}_1}$ in the universal cover. The circle ab^2 in X_v is covered by an infinite set of parallel lines in $\widetilde{X}_{\widetilde{v}_1}$. Index these lines consecutively by the elements of $\mathbb{Z}$; see the top of Figure 2.

Choose a vertex space $\widetilde{X}_{\widetilde{v}_2}$ adjacent to $\widetilde{X}_{\widetilde{v}_1}$ via an edge space attached along an a-line in $\widetilde{X}_{\widetilde{v}_1}$. The walls indexed $2n$ in $\widetilde{X}_{\widetilde{v}_1}$ travel through to $\widetilde{X}_{\widetilde{v}_2}$, where they either cover ab^2 or a. By adding $+1$ to each index if necessary, assume that they cover ab^2. Then the walls indexed $2n + 1$ also travel through to $\widetilde{X}_{\widetilde{v}_2}$, where they cover a. Observe that the wall indexed $2n$ in $\widetilde{X}_{\widetilde{v}_1}$ can be indexed n in $\widetilde{X}_{\widetilde{v}_2}$ and the indexing of the lines remains consecutive. This follows from how the immersed wall in the base space determines the walls in the universal cover (see Figure 2).

The odd indexed walls in $\widetilde{X}_{\widetilde{v}_1}$ are now parallel to a in $\widetilde{X}_{\widetilde{v}_2}$. Thus, the even walls in $\widetilde{X}_{\widetilde{v}_1}$ cross the odd walls in $\widetilde{X}_{\widetilde{v}_1}$.

We will now show that the walls indexed $1, 2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\widetilde{v}_1}$ are pairwise intersecting. Indeed, as shown before, the wall numbered 1 intersects all the even walls in $\widetilde{X}_{\widetilde{v}_1}$. To repeat the above argument, consecutively index the walls intersecting $\widetilde{X}_{\widetilde{v}_2}$ as ab^2-lines, by the elements of $\mathbb{Z}$, so that the walls indexed $2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\widetilde{v}_1}$ are then indexed $1, 2, 4, \ldots, 2^{n-1}, \ldots$ in $\widetilde{X}_{\widetilde{v}_2}$. Arguing inductively, given the vertex space $\widetilde{X}_{\widetilde{v}_i}$ containing walls parallel to ab^2 indexed $1, 2, 4, 8, \ldots, 2^n, \ldots$, there exists an adjacent vertex space $\widetilde{X}_{\widetilde{v}_{i+1}}$ where the wall indexed 1 will cross all the other walls, which can be reindexed $1, 2, 4, 8, \ldots, 2^n, \ldots$ in $\widetilde{X}_{\widetilde{v}_{i+1}}$. Therefore, there is an infinite family of pairwise crossing walls intersecting $\widetilde{X}_{\widetilde{v}_i}$. It will be shown in Proposition 3.4 that such a family determines an infinite cube.
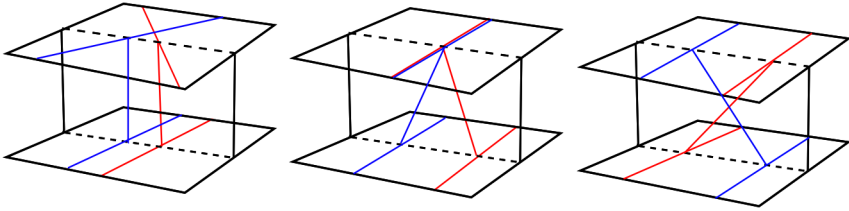
Figure 3 A regular intersection on the left and non-regular intersections on the right and center

REMARK 2.2. The existence of an infinite set of pairwise crossing walls in $\widetilde{X}$ implies the nonseparability of the associated codimension-1 subgroup by Scott's criterion [8]. The tubular group of Example 2.1 is the 3-manifold group that Burns, Karrass, and Solitar [2] proved was not subgroup separable. The group that they used to prove nonsubgroup separability is a subgroup of the stabilizer of this wall. Niblo and Wise [5] reproved this result by finding a family of subgroups that are not contained in any finite index subgroup. The stabilizer of the wall belongs to that family.

3. Infinite Dimensional Cubulations Must Contain Infinite Cubes

The goal of this section is to prove Proposition 3.4, a weaker version of implication (1) $\Rightarrow$ (2) from Theorem 1.2.

REMARK 3.1. A distinction will be made between two types of intersections occurring between a pair of walls in $(\widetilde{X}, \mathcal{W})$ (see Figure 3).

Regular: An intersection in a vertex space as between nonparallel lines in $\mathbb{R}^2$.
Nonregular: An intersection in a vertex space as overlapping parallel lines in $\mathbb{R}^2$ or an intersection occurring in an edge space.

Note that walls may intersect more than once, and a given pair of walls may have multiple regular and nonregular intersections. Note that if $\widetilde{\Lambda}$ and $\widetilde{\Lambda}'$ regularly intersect in $\widetilde{X}_{\widetilde{v}}$, then $\widetilde{\Lambda}$ and $g\widetilde{\Lambda}'$ also regularly intersect in $\widetilde{X}_{\widetilde{v}}$ for all $g \in G_{\widetilde{v}}$. Two walls $\widetilde{\Lambda}_1$, $\widetilde{\Lambda}_2$ that intersect a vertex space $\widetilde{X}_{\widetilde{v}}$ but do not intersect regularly are *parallel in $\widetilde{X}_{\widetilde{v}}$*.

The following short lemma is used extensively throughout the paper.

LEMMA 3.2. *Let $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ be walls in $(\widetilde{X}, \mathcal{W})$ that intersect $\widetilde{X}_{\widetilde{v}}$. Let $\widetilde{Y} \subset \widetilde{X}$ be a subtree of spaces containing $\widetilde{X}_{\widetilde{v}}$ such that $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2$ have no regular intersections in $\widetilde{Y}$. Then every vertex or edge space in $\widetilde{Y}$ that intersects $\widetilde{\Lambda}_1$ also intersects $\widetilde{\Lambda}_2$.*

Proof. Let $\widetilde{X}_{\widetilde{u}}$ be a vertex space in $\widetilde{Y}$. If $\widetilde{\Lambda}_1 \cap \widetilde{X}_{\widetilde{u}}$ and $\widetilde{\Lambda}_2 \cap \widetilde{X}_{\widetilde{u}}$ are a pair of parallel (possibly overlapping) lines in $\widetilde{X}_{\widetilde{u}}$, then $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2$ intersect an identical set of edge spaces adjacent to $\widetilde{X}_{\widetilde{u}}$ and thus an identical set of vertex spaces adjacent

to $\widetilde{X}_{\widetilde{u}}$. Therefore, it can be deduced inductively, starting with $\widetilde{X}_{\widetilde{v}}$, that $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ must intersect an identical set of vertex and edge spaces in $\widetilde{Y}$. □

LEMMA 3.3. *Let S be a finite set of walls in $(\widetilde{X}, \mathcal{W})$ that pairwise intersect. There exists a vertex space $\widetilde{X}_{\widetilde{v}}$ that intersects every wall in S.*

Proof. Each wall in S maps to a connected subset of the Bass–Serre tree. These subsets pairwise intersect; therefore, by application of Helly's theorem for trees, there exists some vertex in the total intersection. Therefore, every wall in S intersects the corresponding vertex space. □

A set of horizontal walls in $(\widetilde{X}, \mathcal{W})$ that pairwise regularly intersect is *crossing*.

PROPOSITION 3.4. *If there are crossing sets of arbitrary finite cardinality in $(\widetilde{X}, \mathcal{W})$, then there exists an infinite cube in $C(\widetilde{X}, \mathcal{W})$.*

Proof of Proposition 3.4. Let $\{K_i\}$ be a sequence of crossing sets such that $|K_i| \to \infty$ as $i \to \infty$. There are only finitely many wall orbits, so assume that each K_i consists exclusively of translates of a single wall $\widetilde{\Lambda}$. By Lemma 3.3, for every i, there exists a vertex $\widetilde{v}_i$ in $\widetilde{\Gamma}$ such that every wall in K_i intersects $\widetilde{X}_{\widetilde{v}_i}$. Since G acts cocompactly on $\widetilde{\Gamma}$, after passing to a subsequence of $\{K_i\}$, we can assume that the $\{\widetilde{v}_i\}$ lie in a single G-orbit $G\widetilde{v}$. Choose g_i such that $g_i\widetilde{v}_i = \widetilde{v}$ and replace K_i with $g_i^{-1}K_i$. Therefore, all walls in K_i intersect the vertex space $\widetilde{X}_{\widetilde{v}}$ for each i.

There are only finitely many parallelism classes of lines in $\widetilde{X}_{\widetilde{v}}$ belonging to walls in $\mathcal{W}$, so there must exist Q_0, a set of walls parallel in $\widetilde{X}_{\widetilde{v}}$, with the following property: *there is no upper bound on the size of crossing subsets of Q_0.* Given Q_0, there must exist a crossing set $\{\widetilde{\Lambda}_1, \widetilde{\Lambda}'_1\} \subset Q_0$. The regular intersection points between $\widetilde{\Lambda}_1, \widetilde{\Lambda}'_1$ must be in vertex spaces other than $\widetilde{X}_{\widetilde{v}}$. Choose $\widetilde{\Lambda}_1, \widetilde{\Lambda}'_1$ such that they intersect in $\widetilde{X}_{\widetilde{v}_1}$, and the distance from $\widetilde{v}$ to $\widetilde{v}_1$ in $\widetilde{\Gamma}$ is minimized over all such choices of intersection points and choices of $\widetilde{\Lambda}_1, \widetilde{\Lambda}'_1$.

By Lemma 3.2, setting $\widetilde{Y} \subset \widetilde{X}$ to be the union of vertex and edge spaces between and including $\widetilde{v}$ and $\widetilde{v}_1$, every other wall in Q_0 also intersects $\widetilde{X}_{\widetilde{v}_1}$. Repeating the argument, there exists a set Q_1 of walls parallel in $\widetilde{X}_{\widetilde{v}_1}$ such that there is no upper bound on the size of crossing subsets. Either $\widetilde{\Lambda}_1$ or $\widetilde{\Lambda}'_1$ will therefore intersect every wall in Q_1, so assume that it is $\widetilde{\Lambda}_1$. Iterate this argument to obtain a sequence of walls $\{\widetilde{\Lambda}_i\}_{i=1}^{\infty}$ that pairwise intersect.

Let $\{\widetilde{\Lambda}_i\}_{i=1}^{\infty}$ be the (sub)sequence of walls reindexed so that they are ordered consecutively in $\widetilde{X}_{\widetilde{v}}$. An ascending sequence of n-cubes c_n in $C(\widetilde{X}, \mathcal{W})$ will now be constructed. Orient all walls not intersecting $\widetilde{X}_{\widetilde{v}}$ toward $\widetilde{X}_{\widetilde{v}}$. Let c_0 be the canonical 0-cube given by orienting all walls toward some $x \in \widetilde{\Lambda}_1 \cap \widetilde{X}_{\widetilde{v}}$. By induction, there is an ascending sequence of cubes $c_0 \subset \cdots \subset c_n$ dual to the walls $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_{n+1}$. Suppose that for some orientation of $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_{n+1}$, the orientation of $\widetilde{\Lambda}_{n+2}$ could not be reversed (see Figure 4). This means that there exists a wall that is parallel to $\widetilde{\Lambda}_{n+2}$, that intersects $\widetilde{X}_{\widetilde{v}}$, and that lies between x and $\widetilde{\Lambda}_{n+2}$. There can only exist finitely many such walls, and, by Lemma 3.2, setting $\widetilde{Y} = \widetilde{X}$,
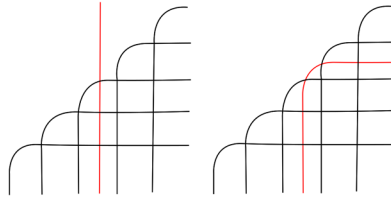
Figure 4   The potential obstruction to an infinite cube on the left, and
the concluding situation on the right. Compare with the picture in [4]

they all intersect $\widetilde{\Lambda}_i$ for $i \neq n + 2$. Replace $\widetilde{\Lambda}_{n+2}$ with the closest such wall to x
to obtain an $(n + 1)$-cube c_{n+1}.   □

4. Dilating Walls and Infinite Dimensional Cubulations

Let $\widetilde{X}_{\widetilde{v}}$ be a vertex space covering X_v. Identify $\widetilde{X}_{\widetilde{v}}$ with $\mathbb{R}^2$ so that a $G_{\widetilde{v}}$-orbit in
$\widetilde{X}_{\widetilde{v}}$ is identified with $\mathbb{Z}^2 \subset \mathbb{R}^2$. Let $\widetilde{\Lambda}$ denote a wall covering the immersed wall Λ,
intersecting $\widetilde{X}_{\widetilde{v}}$. The intersection $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ is a line with rational slope, stabilized by
an infinite cyclic subgroup of $G_{\widetilde{v}}$. Since the slope is rational, there is a maximal
infinite cyclic subgroup $H \leq G_{\widetilde{v}}$ that stabilizes a line perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$.
The elements of H are *perpendicular to* $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$. Note that there are finitely many
$G_{\widetilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\widetilde{v}}$. Given vertices $\widetilde{u}, \widetilde{v}$ in $\widetilde{\Gamma}$, the *carrier of* $\widetilde{u}, \widetilde{v}$,
denoted Carrier$(\widetilde{u}, \widetilde{v})$, is the union of all vertex and edge spaces corresponding
to the vertices and edges on the geodesic in the Bass–Serre tree $\widetilde{\Gamma}$ between and
including $\widetilde{u}$ and $\widetilde{v}$.

The following is elementary.

LEMMA 4.1. *Let T be a torus, and $C \to T$ an immersed geodesic circle with a
choice of lift $\widetilde{C} \to \widetilde{T}$. Given $a, b \in \pi_1 T - \mathrm{Stab}(\widetilde{C})$, there exists $m_a, m_b \in \mathbb{Z} - \{0\}$
such that $a^{m_a} \widetilde{C} = b^{m_b} \widetilde{C}$. Moreover,*

$$\frac{m_a}{m_b} = \frac{\#[C, b]}{\#[C, a]}.$$

LEMMA 4.2. *Let $\widetilde{\Lambda}$ be a wall intersecting the vertex spaces $\widetilde{X}_{\widetilde{v}_0}$ and $\widetilde{X}_{\widetilde{v}_n}$, where
$\widetilde{v}_1, \ldots, \widetilde{v}_{n-1}$ is the sequence of vertices on the geodesic between $\widetilde{v}_0$ and $\widetilde{v}_n$
in $\widetilde{\Gamma}$. Let $\widetilde{e}_i$ be the edge between $\widetilde{v}_{i-1}$ and $\widetilde{v}_i$. Let $G_{\widetilde{e}_i} = \langle g_i \rangle$ for $1 \leq i \leq n$.
Let $\mathrm{Stab}(\widetilde{\Lambda}) \cap G_{\widetilde{v}_i} = \langle \rho_i \rangle$. Let $g_0 \in G_{\widetilde{v}_0}$ and $g_{n+1} \in G_{\widetilde{v}_n}$ be perpendicular to
$\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}_0}$ and $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}_n}$, respectively. Then there exist $\overleftarrow{m}, \overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that
$g_0^{\overleftarrow{m}} \widetilde{\Lambda} = g_{n+1}^{\overrightarrow{m}} \widetilde{\Lambda}$ and*

$$\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{\overleftarrow{m_0} \cdots \overleftarrow{m_n}}{\overrightarrow{m_0} \cdots \overrightarrow{m_n}},$$

*where $\overleftarrow{m_i} = \#[\rho_i, g_{i+1}]$ and $\overrightarrow{m_i} = \#[\rho_i, g_i]$ for $0 \leq i \leq n$. Moreover, the walls in
$\{g_0^{\overleftarrow{m} r} \widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ do not intersect within Carrier$(\widetilde{v}_0, \widetilde{v}_n)$.*

Proof. For each $0 \leq i \leq n$, apply Lemma 4.1 to $g_i, g_{i+1} \in G_{\widetilde{v}_i}$ to deduce that $g_i^{\overleftarrow{m_i}}\widetilde{\Lambda} = g_{i+1}^{\overrightarrow{m_i}}\widetilde{\Lambda}$. Let $\overleftarrow{m} = \overleftarrow{m_0}\cdots\overleftarrow{m_n}$ and $\overrightarrow{m} = \overrightarrow{m_0}\cdots\overrightarrow{m_n}$. Observe that for $0 \leq i \leq n$, we have

$$g_0^{\overleftarrow{m}}\widetilde{\Lambda} = g_0^{\overleftarrow{m_0}\overleftarrow{m_1}\cdots\overleftarrow{m_n}}\widetilde{\Lambda} = g_i^{\overrightarrow{m_0}\cdots\overrightarrow{m_i}\overleftarrow{m_{i+1}}\cdots\overleftarrow{m_n}}\widetilde{\Lambda},$$

so there are no intersections of walls in $\{g_0^{\overleftarrow{m}r}\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ within either $\widetilde{X}_{\widetilde{v}_i}$ or $\widetilde{X}_{\widetilde{e}_i}$. □

DEFINITION 4.3. Let $\widetilde{\Lambda}$ be a wall covering Λ, and let $\widetilde{X}_{\widetilde{v}}$ intersect $\widetilde{\Lambda}$. Observe that $\widetilde{\Lambda}$ also intersects $\widetilde{X}_{g\widetilde{v}}$ for any $g \in \mathrm{Stab}(\widetilde{\Lambda})$. If $h \in G_{\widetilde{v}}$ is perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$, then ghg^{-1} is a perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{g\widetilde{v}}$. By Lemma 4.2 there exist $\overleftarrow{m}, \overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and $\{h^{\overleftarrow{m}r}\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ contains no pair intersecting in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$. The pair $\overleftarrow{m}, \overrightarrow{m}$ are g-*shift exponents*. Note that g-shift exponents are not unique. If $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents, then so are $\overleftarrow{m}n$, $\overrightarrow{m}n$ for $n \in \mathbb{Z} - \{0\}$. The *dilation function of* $\widetilde{\Lambda}$ is the map $R : \mathrm{Stab}(\widetilde{\Lambda}) \to \mathbb{Q}^*$ with $g \mapsto \frac{\overleftarrow{m}}{\overrightarrow{m}}$, where $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents.

LEMMA 4.4. *The map* $R : \mathrm{Stab}(\widetilde{\Lambda}) \to \mathbb{Q}^*$ *is a homomorphism and does not depend on the choice of vertex space* $\widetilde{X}_{\widetilde{v}}$ *or perpendicular element* $h \in G_{\widetilde{v}}$*. Moreover, if* R' *is the dilation function of* $\widetilde{\Lambda}'$*, where* $\widetilde{\Lambda}' = g'\widetilde{\Lambda}$*, then* $R(g'^{-1}gg') = R'(g)$ *for* $g \in \mathrm{Stab}(\widetilde{\Lambda}')$*.*

Proof. Fix a choice of $\widetilde{X}_{\widetilde{v}}$ and let $h \in G_{\widetilde{v}}$ be a primitive element perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$. To see that R is well defined, take any two choices of g-shift exponents $\overleftarrow{m}, \overrightarrow{m}$ and $\overleftarrow{m}', \overrightarrow{m}'$. Since $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and $h^{\overleftarrow{m}'}\widetilde{\Lambda} = gh^{\overrightarrow{m}'}g^{-1}\widetilde{\Lambda}$, we have $gh^{\overleftarrow{m}\,\overrightarrow{m}'}g^{-1}\widetilde{\Lambda} = h^{\overleftarrow{m}\,\overleftarrow{m}'}\widetilde{\Lambda} = gh^{\overrightarrow{m}\,\overleftarrow{m}'}g^{-1}\widetilde{\Lambda}$, which implies that $\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{\overleftarrow{m}'}{\overrightarrow{m}'}$. Now suppose that $h' \in G_{\widetilde{v}}$ is any other choice element perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$; then $h' = h^n$ with $n \in \mathbb{Z} - \{0\}$. Therefore, $(h')^{\overleftarrow{m}}\widetilde{\Lambda} = h^{n\overleftarrow{m}}\widetilde{\Lambda} = gh^{n\overrightarrow{m}}g^{-1}\widetilde{\Lambda} = g(h')^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$, and $\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{n\overleftarrow{m}}{n\overrightarrow{m}}$, so $\overleftarrow{m}, \overrightarrow{m}$ are also g-shift exponents with respect to h'.

To see that R is independent of the choice of $\widetilde{X}_{\widetilde{v}}$, suppose that $\widetilde{\Lambda}$ also intersects $\widetilde{X}_{\widetilde{w}}$. Let $h_{\widetilde{u}} \in G_{\widetilde{u}}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{u}}$, and $h_{\widetilde{w}} \in G_{\widetilde{w}}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}}$. Then by Lemma 4.2 there exist $p, q \in \mathbb{Z} - \{0\}$ such that $h_{\widetilde{w}}^p\widetilde{\Lambda} = h_{\widetilde{u}}^q\widetilde{\Lambda}$. Let $g \in \mathrm{Stab}(\widetilde{\Lambda})$, and let $\overleftarrow{m}, \overrightarrow{m}$ be g-shift exponents with respect to $\widetilde{X}_{\widetilde{u}}$; then deduce that $h_{\widetilde{w}}^{\overleftarrow{m}p}\widetilde{\Lambda} = h_{\widetilde{u}}^{\overleftarrow{m}q}\widetilde{\Lambda} = gh_{\widetilde{u}}^{\overrightarrow{m}q}g^{-1}\widetilde{\Lambda} = gh_{\widetilde{w}}^{\overrightarrow{m}p}g^{-1}\widetilde{\Lambda}$. This means that $\overleftarrow{m}p$, $\overrightarrow{m}p$ are g-shift exponents with respect to $\widetilde{X}_{\widetilde{w}}$

Let $\widetilde{\Lambda}' = g'\widetilde{\Lambda}$ with $g' \in G$. Let $g \in \mathrm{Stab}(g'\widetilde{\Lambda})$. Then we know that $g'^{-1}gg' \in \mathrm{Stab}(\widetilde{\Lambda})$, so there exist $g'^{-1}gg'$-shift exponents $\overleftarrow{m}, \overrightarrow{m}$ for $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{\widetilde{v}}$ so that $h^{\overleftarrow{m}}\widetilde{\Lambda} = (g'^{-1}gg')h^{\overrightarrow{m}}(g'^{-1}gg')^{-1}\widetilde{\Lambda}$, where h is perpendicular to $\widetilde{\Lambda}$ in $\widetilde{X}_{\widetilde{v}}$. This can be rewritten as $(g'hg'^{-1})^{\overleftarrow{m}}(g'\widetilde{\Lambda}) = g(g'hg'^{-1})^{\overrightarrow{m}}g^{-1}(g'\widetilde{\Lambda})$, where $g'hg'^{-1}$ is perpendicular to $\widetilde{\Lambda}'$ in $\widetilde{X}_{g'\widetilde{v}}$. This means that $\overleftarrow{m}, \overrightarrow{m}$ are g-shift exponents for $g'\widetilde{\Lambda}$ with respect to $\widetilde{X}_{g'\widetilde{v}}$, and $R'(g) = R(g'^{-1}gg')$.

To see that R is a homomorphism, let $g_1, g_2 \in \mathrm{Stab}(\widetilde{\Lambda})$, let $\overleftarrow{m_1}, \overrightarrow{m_1}$ be g_1-shift exponents of $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{\widetilde{v}}$ so $h^{\overleftarrow{m_1}}\widetilde{\Lambda} = g_1 h^{\overrightarrow{m_1}} g_1^{-1}\widetilde{\Lambda}$, and let $\overleftarrow{m_2}, \overrightarrow{m_2}$ be g_2-shift exponents of $\widetilde{\Lambda}$ with respect to $\widetilde{X}_{g_1\widetilde{v}}$ so $g_1 h^{\overleftarrow{m_2}} g_1^{-1}\widetilde{\Lambda} = g_2 g_1 h^{\overrightarrow{m_2}} g_1^{-1} g_2^{-1}\widetilde{\Lambda}$. Then $h^{\overleftarrow{m_1}\overleftarrow{m_2}}\widetilde{\Lambda} = g_2 g_1 h^{\overrightarrow{m_1}\overrightarrow{m_2}} g_1^{-1} g_2^{-1}\widetilde{\Lambda}$. Therefore, $R(g_1 g_2) = \overleftarrow{m_1}\overleftarrow{m_2}/(\overrightarrow{m_1}\overrightarrow{m_2}) = R(g_1)R(g_2)$. $\qquad\square$

DEFINITION 4.5. If the dilation function R has infinite image, then $\widetilde{\Lambda}$ is *dilated*. Otherwise, $\widetilde{\Lambda}$ is *nondilated* (and the image of R is either trivial or $\{-1, 1\}$). An element $g \in \mathrm{Stab}\,\widetilde{\Lambda}$ is *dilated* if $R(g) \neq \pm 1$. Lemma 4.4 implies that if $\widetilde{\Lambda}$ is (non)dilated, then all its G-translates are (non)dilated. Therefore, it makes sense to say that an immersed wall Λ is (non)dilated if its associated walls are (non)dilated walls.

A wall $\widetilde{\Lambda}$ being dilated means there is an infinite family of walls $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}}$, intersecting $\widetilde{X}_{\widetilde{v}}$, that become closer together while traveling from $\widetilde{X}_{\widetilde{v}}$ to $\widetilde{X}_{g\widetilde{v}}$. Considering higher powers of g, some infinite subset of these walls must intersect each other. The proof of the following proposition makes this idea precise.

PROPOSITION 4.6. *If $(\widetilde{X}, \mathcal{W})$ contains a dilated wall, then there are crossing sets of arbitrary finite cardinality.*

Proof. Let $\widetilde{\Lambda}_0$ be a dilated wall intersecting $\widetilde{X}_{\widetilde{v}}$, and let $h \in G_{\widetilde{v}}$ be perpendicular to $\widetilde{\Lambda}_0 \cap \widetilde{X}_{\widetilde{v}}$. Choose $g \in \mathrm{Stab}(\widetilde{\Lambda}_0)$ such that $R(g) > 1$. By Lemma 4.7 the set of all G-translates of $\widetilde{\Lambda}_0$ intersecting $\widetilde{X}_{\widetilde{v}}$ that do not regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$ can be partitioned:

$$\mathcal{P} = \{\{h^{\overleftarrow{n}r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}}, \ldots, \{h^{\overleftarrow{n}r}\widetilde{\Lambda}_p\}_{r\in\mathbb{Z}}\},$$

where $h^{\overleftarrow{n}}\widetilde{\Lambda}_i = g h^{\overrightarrow{n}} g^{-1}\widetilde{\Lambda}_i$ and $\frac{\overleftarrow{n}}{\overrightarrow{n}} = R(g)$.

Let $s \geq 1$ and $\widetilde{M} = h^{\overleftarrow{n}s}\widetilde{\Lambda}_0$. We will show that $Q_s^\infty = \{\widetilde{M}, g^{-1}\widetilde{M}, g^{-2}\widetilde{M}, \ldots\}$ contains a cardinality s crossing set. All $g^{-i}\widetilde{M}$ are distinct because if $g^{-p}\widetilde{M} = g^{-q}\widetilde{M}$ for $p \neq q$, then $h^{\overleftarrow{n}s}\widetilde{\Lambda}_0 = g^{p-q}h^{\overleftarrow{n}s}\widetilde{\Lambda}_0 = g^{p-q}h^{\overleftarrow{n}s}g^{q-p}\widetilde{\Lambda}_0$, which would imply $R(g^{p-q}) = 1$, but R is a homomorphism, so this would contradict $R(g) > 1$. This implies that $Q_s^\infty = g(Q_s^\infty - \{\widetilde{M}\})$. Note that none of the walls $\widetilde{M}, \ldots, g^{1-s}\widetilde{M}$ regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$ since

$$g^{-i}\widetilde{M} = g^{-i}h^{\overleftarrow{n}s}\widetilde{\Lambda}_0 = g^{-i}h^{\overleftarrow{n}i}\overleftarrow{n}^{s-i}g^i\widetilde{\Lambda}_0 = h^{\overrightarrow{n}i}\overleftarrow{n}^{s-i}\widetilde{\Lambda}_0,$$

and the right-hand side belongs to the set $\{h^{\overleftarrow{n}r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}} \in \mathcal{P}$ for $0 \leq i \leq s - 1$.

Suppose there exists k such that $g^{-k}\widetilde{M}$ regularly intersects $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$. Let k be the minimal such value and note that $k \geq s$. Lemma 3.2 implies that $g^{-i}\widetilde{M}$ does not regularly intersect $\widetilde{\Lambda}_0$ in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$, which implies that $g^{-k}\widetilde{M}$ regularly intersects $g^{-i}\widetilde{M}$ for $0 \leq i < k$. Then $Q_s^k = \{\widetilde{M}, \ldots, g^{-k}\widetilde{M}\}$ is a crossing set of cardinality at least $s + 1$ since $g^{-k}\widetilde{M}$ regularly intersects all the other elements of Q_s^k and $g^{-1}(Q_s^k - \{g^{-k}\widetilde{M}\}) = Q_s^k - \{\widetilde{M}\}$.

Suppose that $g^{-i}\widetilde{M}$ does not regularly intersect $\widetilde{\Lambda}_0$ in Carrier$(\widetilde{v}, g\widetilde{v})$ for all $i \geq 0$. Then each $g^{-i}\widetilde{M}$ belongs to some subset in the partition $\mathcal{P}$. Therefore, there exist $r_i \in \mathbb{Z}$ and $\sigma(i) \in \{0, \ldots, p\}$ such that $g^{-i}\widetilde{M} = h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)}$. We will obtain a contradiction by showing that $\{r_i\}$ is a bounded sequence, contradicting that Q_s^∞ is an infinite set. From construction of the sequence

$$g^{-1}h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)} = g^{-1}(g^{-i}\widetilde{M}) = g^{-(i+1)}\widetilde{M} = h^{\overleftarrow{n} r_{i+1}}\widetilde{\Lambda}_{\sigma(i+1)}$$

and from the properties of $\mathcal{P}$ we can infer that

$$h^{\overleftarrow{n} r_{i+1}}\widetilde{\Lambda}_{\sigma(i+1)} = g^{-1}h^{\overleftarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i)} = h^{\overrightarrow{n} r_i}g^{-1}\widetilde{\Lambda}_{\sigma(i)}.$$

Hence,

$$h^{\overleftarrow{n} r_{i+1} - \overrightarrow{n} r_i}\widetilde{\Lambda}_{\sigma(i+1)} = g^{-1}\widetilde{\Lambda}_{\sigma(i)},$$

so $a_i = \overleftarrow{n} r_{i+1} - \overrightarrow{n} r_i$ is a bounded integer sequence because the right-hand side can only be one of a finite number of walls. Therefore, $r_{i+1} = \frac{\overrightarrow{n}}{\overleftarrow{n}}r_i + a_i$ is bounded since $\frac{1}{R(g)} = \frac{\overrightarrow{n}}{\overleftarrow{n}} < 1$. $\qquad\square$

The following lemma is used in the proof of Proposition 4.6.

LEMMA 4.7. *Fix $g \in \mathrm{Stab}(\widetilde{\Lambda}_0)$, and $h \in G_{\widetilde{v}}$ perpendicular to $\widetilde{\Lambda}_0 \cap X_{\widetilde{v}}$. Let U be the set of G-translates of $\widetilde{\Lambda}_0$ that intersect $X_{\widetilde{v}}$ and do not regularly intersect $\widetilde{\Lambda}_0$ in Carrier$(\widetilde{v}, g\widetilde{v})$. Then there exist $\overleftarrow{n}, \overrightarrow{n} \in \mathbb{Z} - \{0\}$ such that $R(g) = \frac{\overleftarrow{n}}{\overrightarrow{n}}$ and a subset $\{\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_p\} \subset U$ such that there is a partition of P,*

$$\mathcal{P} = \{\{h^{\overleftarrow{n} r}\widetilde{\Lambda}_0\}_{r\in\mathbb{Z}}, \ldots, \{h^{\overleftarrow{n} r}\widetilde{\Lambda}_p\}_{r\in\mathbb{Z}}\},$$

and such that $h^{\overleftarrow{n}}\widetilde{\Lambda}_i = gh^{\overrightarrow{n}}g^{-1}\widetilde{\Lambda}_i$ for each i.

Proof. Let Carrier$(\widetilde{v}, g\widetilde{v})$ be the union of the consecutive vertex and edge spaces $X_{\widetilde{v}_0}, X_{\widetilde{e}_1}, \ldots, X_{\widetilde{e}_\ell}, X_{\widetilde{v}_\ell}$. For each i, there exists primitive $\rho_i \in G_{\widetilde{v}_i}$ such that $\mathrm{Stab}(\widetilde{\Lambda}_0) \cap G_{\widetilde{v}_i} = \langle \rho_i^{k_i} \rangle$ for some $k_i \in \mathbb{Z} - \{0\}$. Let $\widetilde{\Lambda}$ be any wall in U. Since $\widetilde{\Lambda}$ does not regularly intersect $\widetilde{\Lambda}_0$ in $X_{\widetilde{v}_i}$, observe that $\mathrm{Stab}(\widetilde{\Lambda}) \cap G_{\widetilde{v}_i} = \langle \rho_i^{j_i} \rangle$ for some $j_i > 0$.

Let $G_{\widetilde{e}_i} = \langle g_i \rangle$ for $0 < i < \ell + 1$, and let $g_0 = h$ and $g_{\ell+1} = ghg^{-1}$. By Lemma 4.2 there exist $\overleftarrow{m}, \overrightarrow{m} \in \mathbb{Z} - \{0\}$ such that $h^{\overleftarrow{m}}\widetilde{\Lambda} = gh^{\overrightarrow{m}}g^{-1}\widetilde{\Lambda}$ and

$$\frac{\overleftarrow{m}}{\overrightarrow{m}} = \frac{\overleftarrow{m_0} \cdots \overleftarrow{m_\ell}}{\overrightarrow{m_0} \cdots \overrightarrow{m_\ell}} = \frac{\#[\rho_0, g_1] \cdots \#[\rho_\ell, g_{\ell+1}]}{\#[\rho_0, g_0] \cdots \#[\rho_\ell, g_\ell]},$$

where the j_is have canceled on the right-hand side. Hence, $\frac{\overleftarrow{m}}{\overrightarrow{m}} = R(g)$ because the final ratio is independent of the values of j_i. There are only finitely many values that j_i can take, so we can obtain $\overleftarrow{n}, \overrightarrow{n}$ such that for any $\widetilde{\Lambda}$, we have $h^{\overleftarrow{n}}\widetilde{\Lambda} = gh^{\overrightarrow{n}}g^{-1}\widetilde{\Lambda}$. Since there are only finitely many $G_{\widetilde{v}}$-orbits of walls intersecting $X_{\widetilde{v}}$, the partition is obtained by taking a finite number of representatives $\widetilde{\Lambda}_1, \ldots, \widetilde{\Lambda}_p$, such that the sets $\{h^{\overleftarrow{n} r}\widetilde{\Lambda}_i\}_{r\in\mathbb{Z}}$ are a maximal collection of disjoint sets. $\qquad\square$

We now consider the nondilated walls.

LEMMA 4.8. *Let $\widetilde{\Lambda}$ be a nondilated wall. For any $\widetilde{X}_{\widetilde{v}}$ intersected by $\widetilde{\Lambda}$, there exists $h \in G_{\widetilde{v}}$ perpendicular to $\widetilde{X}_{\widetilde{v}} \cap \widetilde{\Lambda}$ such that if $g \in \mathrm{Stab}(\widetilde{\Lambda})$, then $ghg^{-1}\widetilde{\Lambda} = h\widetilde{\Lambda}$, and the walls in $\{h^r\widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ are pairwise disjoint.*

Proof. Choose a generating set $\{f_1, \ldots, f_n\}$ for $\mathrm{Stab}(\widetilde{\Lambda})$, and fix $h_0 \in G_{\widetilde{v}}$ perpendicular to $\widetilde{X}_{\widetilde{v}} \cap \widetilde{\Lambda}$. By Lemma 4.2 and the assumption that $\widetilde{\Lambda}$ is nondilated, each f_i has an s_i such that $h_0^{s_i}\widetilde{\Lambda} = f_i h_0^{\pm s_i} f_i^{-1}\widetilde{\Lambda}$ and the walls $\{h_0^{s_i r}\widetilde{\Lambda}\}$ pairwise do not intersect in $\mathrm{Carrier}(\widetilde{v}, f_i\widetilde{v})$. Let $s = \mathrm{lcm}\, s_i$ and $h = h_0^s$. Then $h\widetilde{\Lambda} = gh^{\pm 1}g^{-1}\widetilde{\Lambda}$ and $\{h^r\widetilde{\Lambda}\}$ pairwise do not intersect in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$ for all $g \in \mathrm{Stab}(\widetilde{\Lambda})$ since any g can be expressed as the product of the generating set.

For any other $\widetilde{X}_{\widetilde{w}}$ intersected by $\widetilde{\Lambda}$, there is a geodesic $\gamma : I \to \widetilde{\Lambda}$ with initial point $\widetilde{a} \in \widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ and endpoint in $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}}$. Extend γ to a geodesic γ' such that the endpoint of γ' is $g\widetilde{a}$ for some $g \in \mathrm{Stab}(\widetilde{\Lambda})$. Then $h\widetilde{\Lambda} = gh^{\pm 1}g^{-1}\widetilde{\Lambda}$, and for all r, $h^r\widetilde{\Lambda} \cap \widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}} = \emptyset$ since $\widetilde{X}_{\widetilde{w}}$ lies in $\mathrm{Carrier}(\widetilde{v}, g\widetilde{v})$. This implies that the walls in $\{h^r\widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ are pairwise disjoint. $\square$

PROPOSITION 4.9. *Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of nondilated immersed walls. Then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Proof. Let $\{K_i\}$ be a sequence of finite sets of pairwise intersecting walls such that $\lim |K_i| = \infty$. By Lemma 3.3 there exist vertex spaces $\widetilde{X}_{\widetilde{v}_i}$ such that each wall in K_i intersects $\widetilde{X}_{\widetilde{v}_i}$. Since there are finitely many G-orbits of vertices in the Bass–Serre tree, choose a subsequence of $\{K_i\}$ and find a sequence $\{g_i\} \subset G$ such that all $\{g_i K_i\}$ intersect a fixed $\widetilde{X}_{\widetilde{v}}$. Since there are finitely many $G_{\widetilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\widetilde{v}}$, we may restrict to subsets, still of unbounded cardinality, such that all the walls lie in the same $G_{\widetilde{v}}$-orbit.

Fix some $\widetilde{\Lambda}$ in $g_1 K_1$. By Lemma 4.8 there exists $h \in G_{\widetilde{v}}$ such that $\{h^r\widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ consists of pairwise nonintersecting walls. There are only finitely many $G_{\widetilde{v}}$-orbits of $\{h^r\widetilde{\Lambda}\}_{r \in \mathbb{Z}}$, so by the pigeonhole principle there must be some $g_i K_i$ that has more than one wall in one of these orbits. This contradicts that the walls in $g_i K_i$ pairwise intersect. $\square$

THEOREM 4.10. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following are equivalent*:

(1) $C(\widetilde{X}, \mathcal{W})$ *is infinite dimensional.*
(2) $C(\widetilde{X}, \mathcal{W})$ *contains an infinite cube.*
(3) *Some immersed wall is dilated.*

Proof. (3) $\Rightarrow$ (2) follows from Proposition 4.6 and Proposition 3.4.
    (2) $\Rightarrow$ (1) is immediate from the definition of an infinite cube.
    (1) $\Rightarrow$ (3) is the contrapositive of 4.9. $\square$

COROLLARY 4.11. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a set of immersed walls in X. It is decidable whether or not $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Proof. For each immersed wall Λ, the corresponding dilation function can be computed by using Lemmas 4.1 and 4.2 to find their values on a finite generating set of $\pi_1\Lambda$. $\qquad\square$

The following characterization of nondilated immersed walls will be used in [10].

PROPOSITION 4.12. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The horizontal walls in $\mathcal{W}$ can be partitioned into a collection $\mathcal{A}$ of subsets such that*:

(1) *The partition is preserved by G.*
(2) *The walls in A are pairwise nonintersecting for each $A \in \mathcal{A}$.*
(3) *For each $\widetilde{X}_{\widetilde{v}}$, only finitely many $A \in \mathcal{A}$ contain walls intersecting $\widetilde{X}_{\widetilde{v}}$.*
(4) *Let $\widetilde{\Lambda} \in A \in \mathcal{A}$ be a wall intersecting $\widetilde{X}_{\widetilde{v}}$. There exists $h \in G_{\widetilde{v}}$ perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ such that $A = \{h^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}}$.*

Proof. Let $\widetilde{\Lambda}$ be a horizontal wall in $\mathcal{W}$. There are finitely many $\mathrm{Stab}(\widetilde{\Lambda})$-orbits of lines of the form $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}}$. Let $\{\ell_i\}_{i=1}^n$ be a set of representatives for these orbits, and let $\{\widetilde{v}_i\}_{i=1}^n$ be vertices in $\widetilde{\Gamma}$ such that $\ell_i = \widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}_i}$. By Proposition 4.8, for each i, there exist $h_i \in G_{\widetilde{v}_i}$ perpendicular to ℓ_i such that $h_i \widetilde{\Lambda} = g h_i g^{-1} \widetilde{\Lambda}$ for all $g \in \mathrm{Stab}(\widetilde{\Lambda})$ and each $A_i = \{h_i^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ is a collection of pairwise disjoint walls. Let $\mathcal{A}_i = \{g A_i \mid g \in G\}$. If $A_i \cap g A_i \neq \emptyset$, then there exist r, s such that $h^s \widetilde{\Lambda} = g h^r \widetilde{\Lambda}$, which implies $h^{-s} g h^r \in \mathrm{Stab}(\widetilde{\Lambda})$. Therefore, for all $t \in \mathbb{Z}$,

$$h^{-s} g h^{t+r} \widetilde{\Lambda} = (h^{-s} g h^r) h^t (h^{-s} g h^r)^{-1} \widetilde{\Lambda} = h^{\pm t} \widetilde{\Lambda},$$

so $g h^{t+r} \widetilde{\Lambda} = h^{\pm t+s} \widetilde{\Lambda}$, and so $g A_i = A_i$. Therefore, each $\mathcal{A}_i$ is a partition of the G-orbit of $\widetilde{\Lambda}$ satisfying (1) and (2).

By Lemma 4.2 there exist p_i, q_i such that $h_1^{p_i} \widetilde{\Lambda} = h_i^{q_i} \widetilde{\Lambda}$. Let $p = p_1 \cdots p_n$, and $\hat{q}_i = p_1 \cdots p_{i-1} q_i p_{i+1} \cdots p_n$. Let $h = h_1^p$. Let $A = \{h^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}}$ and $\mathcal{A} = \{g A \mid g \in G\}$. Note that $\mathcal{A}$ is a common refinement of the partitions $\mathcal{A}_i$, still satisfying (1) and (2) since $\{h_i^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}} \supseteq \{h_i^{\hat{q}_i r} \widetilde{\Lambda}\}_{r\in\mathbb{Z}} = \{h^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}}$. Condition (4) holds since if $\widetilde{\Lambda}$ intersects $\widetilde{X}_{\widetilde{w}}$, then there exist i and $f \in \mathrm{Stab}(\widetilde{\Lambda})$ such that $\widetilde{w} = f\widetilde{v}_i$, so $A = \{f h_i^{\hat{q}_i r} f^{-1} \widetilde{\Lambda}\}$.

Suppose that $\widetilde{X}_{\widetilde{w}}$ is intersected by some infinite set $\{g_j A\}_{j=1}^\infty$. Then there exists an i such that $g_j \widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{w}} = g_j \ell_i$ for infinitely many j. Fix one such g_k; then there is an infinite set $\{g_k^{-1} g_j A\}$ such that all the walls in $g_k^{-1} g_j A$ intersect $\widetilde{X}_{g_k^{-1}\widetilde{w}}$. By our choice of g_k there exists $f \in \mathrm{Stab}(\widetilde{\Lambda})$ such that $\widetilde{\Lambda} \cap \widetilde{X}_{g_k^{-1}\widetilde{w}} = f\ell_i$. But since $A = \{h^r \widetilde{\Lambda}\}_{r\in\mathbb{Z}} = \{h_i^{\hat{q}_i r} \widetilde{\Lambda}\}_{r\in\mathbb{Z}} = \{f h_i^{\hat{q}_i r} f^{-1} \widetilde{\Lambda}\}_{r\in\mathbb{Z}}$, the $G_{g_k^{-1}\widetilde{w}}$-orbit of $f\ell_i$ is contained in finitely many G-translates of A. This contradicts that there is an infinite set $\{g_k^{-1} g_j A\}$ such that all the walls in $g_k^{-1} g_j A$ intersect $\widetilde{X}_{g_k^{-1}\widetilde{w}}$, so (3) holds. $\qquad\square$

5. Computing the Dilation Function

In this section, the results of Section 4 are applied to concrete examples to determine whether or not they are finite dimensional. In this section, we assume that all immersed walls are nontrivial, in the sense that they do not consist of a single immersed circle. This allows us to identify $\mathrm{Stab}(\widetilde{\Lambda})$ with $\pi_1\Lambda$. The main focus will be on computing the dilation function $R : \pi_1\Lambda \to \mathbb{Q}^*$ of an immersed wall. Let $q : \Lambda \to \Omega$ be the quotient map obtained by quotienting each circle in the equitable set to a vertex. This simplifies the computation since R factors through $q_* : \pi_1\Lambda \to \pi_1\Omega$:

$$
\begin{array}{ccc}
\pi_1\Lambda & \xrightarrow{R} & \mathbb{Q}^* \\
q_* \downarrow & \nearrow & \\
\pi_1\Omega & &
\end{array}
$$

Regard Ω as a directed graph by fixing orientations of its edges $E(\Omega)$, choosing the orientation of each edge σ to be consistent with all other edges mapping into the same vertex space of X. Define the *weighting* $\omega : E(\Omega) \to \mathbb{Q}^*$ as follows: for each directed edge $\sigma \in E(\Omega)$, let $\omega(\sigma) = \#[C_\iota : \overleftarrow{f_e}]/\#[C_\tau : \overrightarrow{f_e}]$, where X_e is the edge space σ maps into, and C_ι, C_τ are the elements in the equitable set attached to the initial and terminal ends of σ. Let γ be a combinatorial path in Λ representing an element $[\gamma] \in \pi_1\Lambda$, and let $\overline{\gamma}$ be the combinatorial path in Ω obtained by quotienting the circle-edges of γ to vertices; then $q \circ \gamma$ is homotopic to $\overline{\gamma}$. If $\overline{\gamma} = \sigma_1^{\epsilon_1} \cdots \sigma_n^{\epsilon_n}$ with $\epsilon \in \{\pm 1\}$, then $R([\gamma]) = \omega(\sigma_1)^{\epsilon_1} \cdots \omega(\sigma_n)^{\epsilon_n}$. Note that this does not depend on the choice of representative γ. See Figure 6 for examples of this quotient and edge weightings.

EXAMPLE 5.1. Let $G = \langle a, b, s, t \mid [a, b], s^{-1}abs = a^2, t^{-1}abt = b^2\rangle$. The group G is a free-by-cyclic tubular group that acts freely on a CAT(0) cube complex but cannot act on a finite dimensional CAT(0) cube complex. The proof uses the dilation function.

The equitable set $\{a, b\}$ demonstrates that G acts freely on a CAT(0) cube complex. Any collection of circles parallel to ab^{-1} extends to an immersed wall $\widetilde{\Lambda}$, which must be nondilated since $\widetilde{\Lambda}$ cannot regularly intersect its translates. We claim that any immersed wall containing a circle not parallel to ab^{-1} is dilated. Such an immersed wall must exist since an equitable set for G must generate a finite index subgroup of the vertex group. Suppose there is an immersed wall Λ with equitable set $S = \{v_1, \ldots, v_n\}$, where $v_i = a^{x_i} b^{y_i}$ and we assume that $v_1 \neq a^n b^{-n}$.

The equitable set S must satisfy the equations

$$
2\sum_{i=1}^{n} |x_i| = 2\sum_{i=1}^{n} |y_i| = \sum_{i=1}^{n} |x_i - y_i|. \tag{$*$}
$$

The claim follows by finding a closed path γ in Ω such that $R([\gamma]) \neq \pm 1$. Direct the edges of Ω so that edges exiting the vertex space via the attaching

maps ab are the initial ends. Therefore, the number of edges leaving the vertex corresponding to v_i is $2|x_i - y_i|$, whereas the number of vertices arriving at that vertex is $2|x_i| + 2|y_i|$. If some $|x_i| + |y_i| > |x_i - y_i|$, then since each $|x_i| + |y_i| \geq |x_i - y_i|$, we have

$$\sum_{i=1}^{n} |x_i| + |y_i| > \sum_{i=1}^{n} |x_i - y_i|,$$

which would contradict $(*)$. Therefore, the number of edges entering and exiting each vertex in Ω are equal, so there exists a *directed Eulerian trail* γ that is a cycle traversing each edge in Ω precisely once. Therefore,

$$|R([\gamma])| = \prod_{i=1}^{n} \frac{|x_i - y_i|^{2|x_i - y_i|}}{(2|x_i|)^{2|x_i|}(2|y_i|)^{2|y_i|}},$$

and the claim will be proven by showing that $|R([\gamma])| < 1$. Considering each factor in the product separately, we want to show that

$$\frac{|x - y|^{2|x-y|}}{(2|x|)^{2|x|}(2|y|)^{2|y|}} \leq 1$$

with equality only when $x = -y$. Since $v_1 \neq a^n b^{-n}$, the first term in the product is strictly less than one, so $|R([\gamma])| < 1$.

This is immediate in the trivial case where $|x - y| = 0$, $|x| = 0$, or $|y| = 0$. After taking the square root and applying the logarithm, the nontrivial case is equivalent to showing that

$$|x - y| \log |x - y| \leq |x| \log(2|x|) + |y| \log(2|y|) \qquad (**)$$

with equality only when $x = -y$. Since $z \log(z)$ is strictly convex for $z > 0$, the following inequality holds with equality when $p = q$;

$$\frac{|p + q|}{2} \log \left(\frac{|p + q|}{2} \right) \leq \frac{|p|}{2} \log |p| + \frac{|q|}{2} \log |q|.$$

Thus, $(**)$ holds by letting $p = 2x$ and $q = -2y$.

EXAMPLE 5.2. The right angled Artin group $A = \langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1 \rangle$ is not subgroup separable (see [5]). From the alternative presentation $A = \langle a_i, b_i \mid [a_i, b_i], a_1 = a_2, b_2 = b_3; 1 \leq i \leq 3 \rangle$ it is clear that A is a fundamental group of a tubular group given by three tori and two cylinders. Consider an immersed wall with the following equitable set: $\{a_1 b_1^2, b_1\}$, $\{a_2 b_2^{-1}, a_2 b_2^2\}$, and $\{a_3^2\}$; and the arcs given in Figure 5. The diagram shows the underlying graph Ω with the associated edge weightings. There is a single simple closed path with $R([\gamma]) = (2) \cdot (1/2) \cdot (1/2)^{-1} \cdot (1) = 2$; therefore, $\widetilde{\Lambda}$ is dilated, and $\mathrm{Stab}(\widetilde{\Lambda})$ is not separable (see discussion in Section 1). Furthermore, by Theorem 4.10 there must exist an infinite cube in the associated cubulation.

EXAMPLE 5.3. The group $G = \langle a, b, t \mid [a, b], t^{-1}at = b \rangle$ from Example 2.1 demonstrates that the choice of arcs connecting circles can change whether or not
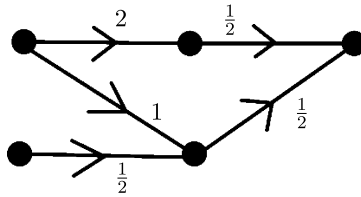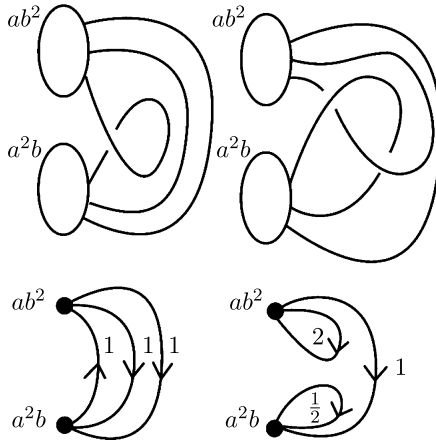
Figure 5   The graph Ω from Example 5.2

Figure 6   Example 5.3. The immersed walls Λ_1, Λ_2 are shown on the top, and the associated quotient Ω_i is shown below with its orientation and weighting. The immersed wall on the left is nondilated, and the immersed wall on the right is dilated

a cubulation is finite dimensional. Consider the equitable set $\{a^2b, ab^2\}$, which extends to both a dilating wall and a nondilating wall shown in Figure 6.

6. Dilated Walls Are Not Quasi-Isometrically Embedded

Using Theorem 4.10, Theorem 1.3 can be restated in the following form.

THEOREM 6.1. *If $\widetilde{\Lambda}$ is dilated, then $\mathrm{Stab}(\widetilde{\Lambda}) \leq \pi_1 X$ is not quasi-isometrically embedded.*

The following example motivates the proof of Theorem 6.1 and illustrates that a nondilated wall can fail to be quasi-isometrically embedded.

EXAMPLE 6.2. Consider the tubular group $G = \langle a, b, t \mid [a, b], t^{-1}at = b \rangle$ from Example 2.1. Let Λ be the immersed wall given by the equitable set $\{a^{-1}b\}$, which intersects each end of the edge space exactly once (see Figure 7). Observe that Λ is embedded and therefore nondilated. Let $H = \mathrm{Stab}(\widetilde{\Lambda}) \leq G$. Note that H
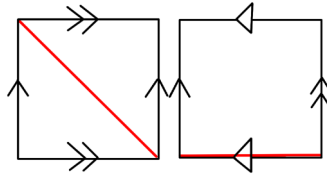
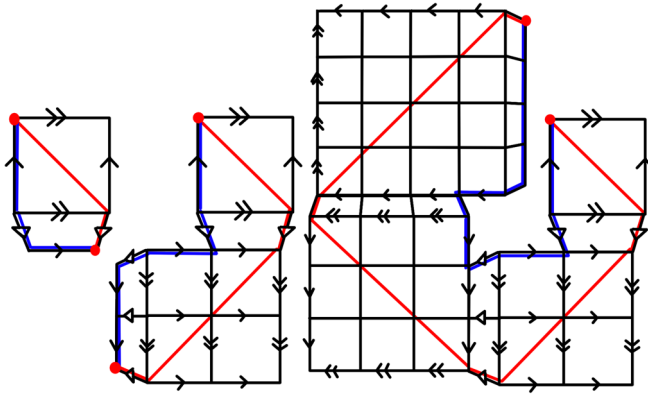**Figure 7**   The tubular space and immersed wall in Example 6.2

Figure 8   The paths corresponding to α_1, β_1, to α_2, β_2, and to α_4, β_4

is a rank 2 free group with basis $\{a^{-1}b, t\}$. We will show that H is quadratically distorted in G.

Define $\alpha_0 = 1$, and for each $n \geq 1$, define $\alpha_n = t(a^{-1}b)^n \alpha_{n-1}$. This sequence of elements has the property that $|\alpha_n|_H = \sum_{i=1}^{n}(i+1) = \frac{1}{2}n(n+1) + n$. The element $\beta_n = a^n(ta^{-1})^n \in G$ satisfies $|\beta|_G \leq 3n$. It can be verified by induction that $\alpha_n = \beta_n$. Therefore, H is a quadratically distorted subgroup of G. Figure 8 illustrates the paths these elements correspond to in the universal cover: α_n spirals outward increasing in length quadratically, whereas β_n cuts through inside, increasing linearly in length before traveling out by another linear factor.

Before embarking on the proof of Theorem 6.1, we outline the strategy employed.

(1) Inspired by the geometric interpretation of Example 6.2, construct a sequence of spiraling paths γ_m in Λ determined by a sequence of parameters.
(2) Choose the parameters of γ_m such that the length of γ_m in Λ is a quadratic function of m.
(3) Double up γ_m, as illustrated in Figure 10, to produce a second sequence of paths ρ_m that spiral out before spiraling back in, with the length growing quadratically in Λ. This stage is necessary because in Example 6.2 the rate at which γ_m spiraled outward was linear.

(4) Specify a path homotopy from ρ_m to ρ'_m, where the length of ρ'_m in X is a linear function of m.

Proof of Theorem 6.1. Let Λ be a dilated immersed wall in X. Let $z \in \pi_1 \Lambda$ satisfy $|R(z)| \neq 1$, and let γ be a closed immersed path in Λ representing z.

Stage 1: Decompose γ as

$$\gamma = x_0 \cdot y_1 \cdot x_1 \cdot y_2 \cdots x_{\ell-1} \cdot y_\ell,$$

where each x_i is a path in a circle of Λ immersed in a vertex space X_{u_i}, and y_i is an arc of Λ embedded in an edge space X_{e_i}. Since γ is a closed path, the terminal point of y_ℓ is the initial point of x_0. The subpath x_i is *diagonal* if $\overrightarrow{X_{e_i}}$ and $\overleftarrow{X_{e_{i+1}}}$ are not parallel in X_{u_i}, and the subpath $y_i \cdot x_i \cdots x_{i+r} \cdot y_{i+r+1}$ is *straight* if $x_i, x_{i+1}, \ldots, x_{i+r}$ are not diagonal. The subscripts are considered modulo ℓ, so that subpaths of γ containing the initial point are considered. Thus, after cyclically parameterizing, γ decomposes as

$$\gamma = a_0 \cdot b_1 \cdot a_1 \cdot b_2 \cdots a_{n-1} \cdot b_n,$$

where each a_i is diagonal, and each b_i is straight. For notational purposes, these subscripts are later considered modulo n.

CLAIM. *There is at least one diagonal subpath in γ.*

Proof. Let $\widetilde{\Lambda}$ be a wall in $\widetilde{X}$ stabilized by $\pi_1 \Lambda$, and $\widetilde{X}_{\widetilde{u}_0}$ be a vertex space intersected by $\widetilde{\Lambda}$. Let $\widetilde{u}_1, \ldots, \widetilde{u}_{\ell-1}$ be the vertices in $\widetilde{\Gamma}$ on the geodesic between $\widetilde{u}_0$ and $\widetilde{u}_\ell = z\widetilde{u}_0$. Let $\widetilde{e}_i$ be the edge between $\widetilde{u}_{i-1}$ and $\widetilde{u}_i$. Let $\langle g_i \rangle = G_{\widetilde{e}_i}$ for $1 \leq i \leq \ell$, and let $g_0 \in G_{\widetilde{u}_i}$ be perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{u}_0}$ and $g_{\ell+1} = zg_0{}^{-1}$. Let $\langle \rho_i \rangle = \text{Stab}(\widetilde{\Lambda}) \cap G_{\widetilde{u}_i}$ for $0 \leq i \leq \ell$. Suppose that the claim is false. This would imply that exist a primitive $h_i \in G_{\widetilde{u}_i}$ and $p_i, q_i \in \mathbb{Z} - \{0\}$ such that $g_i = h_i^{p_i}$ and $g_{i+1} = h_i^{q_i}$ for $1 \leq i \leq \ell - 1$ and $g_\ell = zh_0^{p_\ell}z^{-1}$ and $g_1 = h_0^{q_\ell}$. Since G acts freely on a CAT(0) cube complex, it cannot contain a subgroup isomorphic to $\langle r, t \mid t^{-1}r^n t r^m \rangle$ where $n \neq \pm m$ [11]. Therefore, $p_1 \cdots p_\ell / (q_1 \cdots q_\ell) = \pm 1$, so Lemma 4.2 says

$$R(z) = \prod_{i=0}^{\ell} \frac{\#[\rho_i, g_{i+1}]}{\#[\rho_i, g_i]} = \frac{\#[\rho_0, h_0^{q_\ell}]}{\#[\rho_0, g_0]} \frac{\#[\rho_\ell, zg_0z^{-1}]}{\#[\rho_\ell, zh_0^{p_\ell}z^{-1}]} \prod_{i=1}^{\ell-1} \frac{\#[\rho_i, h_i^{q_i}]}{\#[\rho_i, h_i^{p_i}]}$$

$$= \frac{q_1 \cdots q_\ell}{p_1 \cdots p_\ell} = \pm 1$$

since $\#[\rho_0, h_0] = \#[\rho_\ell, zh_0z^{-1}]$ and $\#[\rho_0, g_0] = \#[\rho_\ell, zg_0z^{-1}]$. This contradicts our choice of γ. □

Parameterize the circle in Λ containing a_i as an immersed path c_i with $c_i(0) = c_i(1) = a_i(1)$ and such that if $a_i(0) \neq a_i(1)$, then there is an immersed path a'_i satisfying $c_i^{k_i} = (a'_i)^{-1} \cdot a_i$ for some $k_i \geq 1$. If $a_i(0) = a_i(1)$, then let $c_i^{k_i} = a_i$, and

let $a_i' = a_i^{-1}$. Note that $a_i(0) = a_i'(0)$ and $a_i(1) = a_i'(1)$. Define the sequence of immersed paths $\{\gamma_m : I \to \Lambda\}$ inductively by setting $\gamma_{-1} = a_0(0)$ and

$$\gamma_m = \gamma_{m-1} \cdot \prod_{i=0}^{n-1} \hat{a}_{mn+i} \cdot c_i^{s(mn+i)} \cdot b_{i+1},$$

where $\hat{a}_{mn+i} \in \{a_i, a_i'\}$ and $s(mn+i) \in \mathbb{Z} - \{0\}$. Both $\hat{a}_{mn+i}$ and $s(mn+i)$ will be defined inductively to replicate the spiralling effect in Example 6.2. Note that for γ_m to be an immersed path, if $\hat{a}_{mn+i} = a_i$, then $s(mn+i) > 0$, and if $\hat{a}_{mn+i} = a_i'$, then $s(mn+i) < 0$.

Stage 2: Let X_{v_i} be the vertex space containing the image of a_i. Observe that for each b_i, there is a cylinder $Y_i \hookrightarrow X$ that b_i factors through, where Y_i is a union of cylinders that cover edge spaces and cylinders that immerse in the vertex spaces. Let $\overleftarrow{Y_i}$ and $\overrightarrow{Y_i}$ denote the boundary components of Y_i with $\overleftarrow{f_i} : \overleftarrow{Y_i} \to X_{v_{i-1}}$ and $\overrightarrow{f_i} : \overrightarrow{Y_i} \to X_{v_i}$ denoting the corresponding restrictions of the immersion. When γ traverses every vertex space diagonally, then each Y_i is an edge space.

We choose $\hat{a}_{mn+i} \in \{a_i, a_i'\}$ inductively by first setting $\hat{a}_1 = a_1$. Assuming that $\hat{a}_{mn+i-1}$ has been chosen, we specify that $s(mn+i-1) \in \mathbb{Z} - \{0\}$ is positive or negative so that $\hat{a}_{mn+i-1} \cdot c_i^{s(mn+i-1)}$ is a locally geodesic path in X_{v_i}. Choose $\hat{a}_{mn+i} \in \{a_i, a_i'\}$ such that the terminal point of $\hat{a}_{mn+i-1} \cdot c_i^{s(mn+i-1)}$ meets $\overleftarrow{Y_i}$ with an acute angle on the same side as the initial point of $\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}$ leaves $\overrightarrow{Y_i}$. This ensures that γ_m spirals in a consistent direction.

Inside the torus X_{v_i}, the paths $\overrightarrow{f_i}$, $\overleftarrow{f_{i+1}}$, and c_i determine a triangle with angles θ_i between the $\overrightarrow{f_i}$ and c_i sides, ξ_i between the $\overleftarrow{f_{i+1}}$ and c_i sides, and ζ_i between the $\overrightarrow{f_i}$ and $\overleftarrow{f_{i+1}}$ sides. The rule of sines states

$$\frac{Z_i}{\sin \zeta_i} = \frac{\Xi_i}{\sin \xi_i} = \frac{\Theta_i}{\sin \theta_i},$$

where the value of the numerator is the length of the side opposite the angle in the denominator (see Figure 9). Let $Z_{mn+i} = |\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|$. The subscripts of ζ_i, ξ_i, and θ_i are considered modulo n, whereas the subscripts of Z_i, Ξ_i, and Θ_i are not. Note that we chose $\hat{a}_{mn+i}$ and the sign of $s(mn+i)$ so that $0 < \xi_i, \theta_i < \pi/2$.

Let $r_i = (1/|\overleftarrow{f_i}|) \sin \xi_i / \sin \zeta_i$ and $t_i = (1/|\overrightarrow{f_i}|) \sin \theta_i / \sin \zeta_i$ with $r_{i+n} = r_i$ and $t_{i+n} = t_i$. Define $s(mn+i) \in \mathbb{Z} - \{0\}$ inductively by setting $s(0) = 1$, and assuming that $s(mn+i-1)$ is defined, choose $|s(mn+i)|$ to be large enough that

$$1 \leq r_i Z_{mn+i} - t_{i-1} Z_{mn+i-1}$$
$$= r_i(|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|) - t_{i-1}(|\hat{a}_{mn+i-1} \cdot c_{i-1}^{s(mn+i-1)}|) \tag{†}$$

and small enough that

$$r_i(|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|) - t_{i-1}(|\hat{a}_{mn+i-1} \cdot c_{i-1}^{s(mn+i-1)}|) \leq 1 + r_i|c_i^{1+k_i}|. \tag{‡}$$
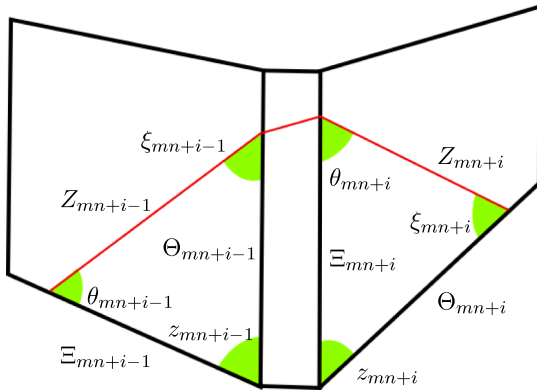
Figure 9 A portion of γ_m traversing $X_{v_{i-1}}$, Y_i, and X_{v_i} is illustrated schematically with the corresponding side lengths labeled

Combining (†) and (‡), replacing $|\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}|$ with Z_{mn+i}, and applying the rule of sines produce the following inequality:

$$1 \le \frac{\Xi_{mn+i}}{|\overleftarrow{f_i}|} - \frac{\Theta_{mn+i-1}}{|\overrightarrow{f_i}|} \le 1 + r_i|c_i^{1+k_i}|. \qquad (\star)$$

Geometrically, this means that $s(mn + i)$ is defined so that consecutive triangles in the sequence have adjacent sides increasing at a bounded rate (see Figure 9). The $|\overleftarrow{f_i}|$ and $|\overrightarrow{f_i}|$ accounts for the different lengths of the attaching maps of Y_i. To estimate a lower bound on the length of γ_m, first rewrite (†) as

$$Z_{mn+i} \ge r_i^{-1}[1 + t_{i-1}Z_{mn+i-1}].$$

Without loss of generality, $\prod_{i=0}^{n-1} r_i^{-1}t_i = \prod_{i=0}^{n-1}(\sin\theta_i/\sin\xi_i)|\overleftarrow{f_i}|/|\overrightarrow{f_i}| \ge 1$, otherwise replace γ with γ^{-1}, which switches θ_i for ξ_i and $\overleftarrow{f_i}$ for $\overrightarrow{f_i}$. Applying this formula recursively produces

$$Z_{mn+i} \ge r_i^{-1}[1 + t_{i-1}r_{i-1}^{-1}[1 + t_{i-2}r_{i-2}^{-1}[\cdots[1 + t_{i-n}Z_{mn+i-n}]\cdots]]];$$

hence,

$$Z_{mn+i} \ge \left[\prod_{i=0}^{n-1} r_i^{-1}t_i\right]Z_{m(n-1)+i} + \cdots \ge Z_{m(n-1)+i} + D,$$

where $D > 0$ is a lower bound on the lower order terms, which are independent of any value of Z_{mn+i} and the value of i. Therefore, there is a lower bound on the length of γ_m:

$$|\gamma_m| \ge \sum_{j=0}^{m(n+1)-1} Z_j \ge \sum_{j=0}^{m} jD \ge \frac{D}{2}m(m+1).$$

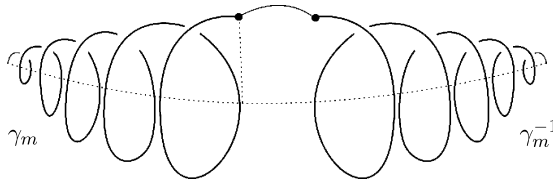Hence, the length of γ_m grows at least quadratically.

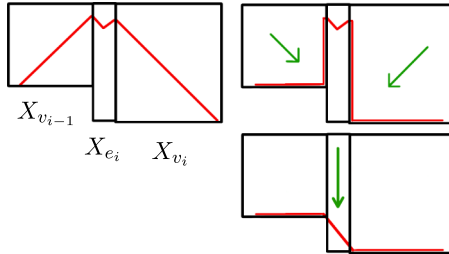**Figure 10** The spiral γ_m is doubled up

Figure 11 The left diagram shows $X_{v_{i-1}}$, Y_i, and X_{v_i} containing $a_{i-1} \cdot c_{i-1}^{s(mn+i-1)} \cdot b_i \cdot a_i \cdot c_i^{s(mn+i)}$. The top right diagram shows the first step of the homotopy, and the bottom right diagram shows the second step

Stage 3: Define $\rho_m := \gamma_m \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot \gamma_m^{-1}$, an immersed path with length growing quadratically with m. See Figure 10.

Stage 4: A path homotopy of ρ_m in X will be described in two stages, and the resulting curve will have length a linear function of m. Compare the following description to Figure 11.

Step 1: Path homotope each diagonal segment $\hat{a}_{mn+i} \cdot c_i^{s(mn+i)}$ in γ_m across the vertex space X_{v_i} to a path $p_{mn+i} \cdot q_{mn+i}$ where p_{mn+i} is a geodesic with image in $X_{v_i} \cap Y_i$ and q_{mn+i} is a geodesic with image in $X_{v_i} \cap Y_{i+1}$ (see the top right diagram in Figure 11). These paths respectively have lengths Ξ_{mn+i} and Θ_{mn+i}. The same path homotopies are made to γ_m^{-1}.

Step 2: Path homotope each $q_{mn+i-1} \cdot b_i \cdot p_{mn+i}$ factoring through Y_i, and the corresponding inverse path, to a local geodesic factoring through Y_i joining the initial and terminal points (see the bottom right diagram in Figure 11). From inequality ($\star$) we can deduce that the resulting path has length at most $|b_{i-1}| + |\overrightarrow{f_i}|(1 + r_i |c_i^{1+k_i}|)$. The segment $q_{mn-1} \cdot b_n \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot b_n^{-1} \cdot q_{mn-1}^{-1}$ can similarly be homotoped since $b_n \cdot a_0 \cdot c_0 \cdot a_0^{-1} \cdot b_n^{-1}$ factors through a cylinder $Y \looparrowright X$.

Therefore, the length of the resulting curve ρ_m' is less than

$$2mn \max_{0 \leq i < n} \{|b_{i-1}| + |\overrightarrow{f_i}|(1 + r_i |c_i^{1+k_i}|)\} + |a_0 \cdot c_0 \cdot a_0^{-1}|,$$

which is a linear function of m. Thus, $\widetilde{\Lambda}$ is distorted. $\square$

References

[1] N. Brady and M. R. Bridson, *There is only one gap in the isoperimetric spectrum,* Geom. Funct. Anal. 10 (2000), no. 5, 1053–1070.

[2] R. G. Burns, A. Karrass, and D. Solitar, *A note on groups with separable finitely generated subgroups,* Bull. Aust. Math. Soc. 36 (1987), no. 1, 153–160.

[3] C. H. Cashen, *Quasi-isometries between tubular groups,* Groups Geom. Dyn. 4 (2010), no. 3, 473–516.

[4] G. C. Hruska and D. T. Wise, *Finiteness properties of cubulated groups,* Compos. Math. 150 (2014), no. 3, 453–506.

[5] G. A. Niblo and D. T. Wise, *The engulfing property for* 3*-manifolds,* The Epstein birthday schrift, pp. 413–418 (electronic), Geom. Topol., Coventry, 1998.

[6] J. H. Rubinstein and S. Wang, π_1*-injective surfaces in graph manifolds,* Comment. Math. Helv. 73 (1998), no. 4, 499–515.

[7] M. Sageev, *Ends of group pairs and non-positively curved cube complexes,* Proc. Lond. Math. Soc. (3) 71 (1995), no. 3, 585–617.

[8] P. Scott, *Subgroups of surface groups are almost geometric,* J. Lond. Math. Soc. (2) 17 (1978), no. 3, 555–565.

[9] D. T. Wise, *Cubular tubular groups,* Trans. Amer. Math. Soc. 366 (2014), no. 10, 5503–5521.

[10] D. J. Woodhouse, *Classifying virtually special tubular groups,* arXiv:1607.06334.

[11] F. Haglund, *Isometries of* CAT(0) *cube complexes are semi-simple* (2007), 1–17.

Department of Mathematics
    and Statistics
McGill University
Montreal, QC
Canada H3A 0B9

daniel.woodhouse@mail.mcgill.ca