Canad. J. Math. Vol. **57** (2), 2005 pp. 416–448

Approximating Flats by Periodic Flats in CAT(0) Square Complexes

Daniel T. Wise

Abstract. We investigate the problem of whether every immersed flat plane in a nonpositively curved square complex is the limit of periodic flat planes. Using a branched cover, we reduce the problem to the case of $\mathcal{VH}$-complexes. We solve the problem for malnormal and cyclonormal $\mathcal{VH}$-complexes. We also solve the problem for complete square complexes using a different approach. We give an application towards deciding whether the elements of fundamental groups of the spaces we study have commuting powers. We note a connection between the flat approximation problem and subgroup separability.

1 Introduction

A word-hyperbolic group cannot contain "poisonous" subgroups of the form $\langle a, b \mid ab = ba^n \rangle$ for $n \neq 0$, and in particular cannot contain a $\mathbb{Z}^2$ subgroup. There do exist groups without poisonous subgroups which are not word-hyperbolic because they fail to satisfy finiteness conditions; indeed, an infinitely generated free group is such an example, and higher dimensional failures were exhibited in [16, 3]. Nevertheless, it is surprisingly still unknown whether a group with a finite Eilenberg-MacLane space and no poisonous subgroups is word-hyperbolic.

While it is widely believed that there are counterexamples in general, there are classes of groups for which a dichotomy between word-hyperbolicity and a $\mathbb{Z}^2$ subgroup is strongly suspected. For instance, the *weak hyperbolization conjecture* which is a coarse version of Thurston's geometrization conjecture, proposes that every finitely generated 3-manifold group is either word-hyperbolic or contains a $\mathbb{Z}^2$ subgroup.

For a compact nonpositively curved manifold M, a result of Eberlein shows that $\pi_1 M$ is word-hyperbolic unless $\widetilde{M}$ contains an isometrically embedded Euclidean plane [7]. Moreover the results of [8, 14] show that any $\mathbb{Z}^n$ subgroup of $\pi_1 M$ acts freely and cocompactly on an isometrically embedded copy of $\mathbb{E}^n$ in $\widetilde{M}$. Furthermore, both these statements have direct generalizations to compact metric spaces with nonpositive curvature [4]. Thus the possible dichotomy between word-hyperbolicity and $\mathbb{Z}^2$ subgroups is very suggestive.

Current research has focused on the possibility that a flat plane in $\widetilde{M}$ implies that $\mathbb{Z}^2 \subset \pi_1 M$, and there have been various positive results. For instance, if $\widetilde{M}$ contains a flat plane then $\mathbb{Z}^2 \subset \pi_1 M$ when M is a closed 3-manifold with nonpositive sectional curvature [6, 17], when M is a nonpositively curved cubulated 3-manifold [15] or even when the 3-manifold M merely has a Riemannian metric such that $\widetilde{M}$ contains a flat plane [13].

Received by the editors March 25, 2003.
AMS subject classification: 20F67,20F06.
Keywords: CAT(0), periodic flat planes.

While several results have been obtained in the manifold case, much less is known for nonpositively curved metric spaces that are not manifolds. As observed by Gromov in [9], the existence of finite sets of tiles which can only tile the plane aperiodically, casts substantial doubt on the possibility that this dichotomy holds in general. The first such aperiodic sets were square tiles with matching conditions constructed by Burger [1], and the study of various finite sets of such tiles eventually culminated in Penrose's kite and dart example. We refer the reader to [11] for an account of the extensive results about aperiodic tilings.

A *square complex* X is a 2-complex formed by attaching squares to a graph by identifying their sides with edges in the graph. The square complex X is *nonpositively curved* if the link of each 0-cell has no cycles of length < 4. The universal cover of a nonpositively curved square complex X can be metrized so that each edge is isometric to a unit interval and each square is isometric to a square, and moreover, the metric satisfies the CAT(0) condition [4]. Thus the combinatorial link condition implies that X is itself locally CAT(0) or "nonpositively curved".

We have focused on the class of nonpositively curved square complexes because they are the simplest examples for which the flat/$\mathbb{Z}^2$ problem is intractable, and thus afford a good test case for the general situation. Moreover, as mentioned above, without the nonpositive curvature assumption, a set of squares can already lead to aperiodic tilings of the plane. On the other hand, there are numerous natural examples of nonpositively curved square complexes.

In [10], Gromov raised the question of which semihyperbolic spaces have the property that their flats are limits of periodic flats. A flat in $\widetilde{X}$ is *periodic* if it is stabilized by some $\mathbb{Z}^2$ subgroup of $\pi_1 X$. And a flat $E \hookrightarrow \widetilde{X}$ is the *limit of periodic flats* if there is a sequence $E_i \hookrightarrow \widetilde{X}$ of periodic flats which converge pointwise to X.

Our results, which are most of part I of [24], provide conditions on a compact nonpositively curved square complex X which guarantee that $\pi_1 X$ contains a $\mathbb{Z}^2$ subgroup if and only if $\widetilde{X}$ contains an isometrically embedded flat plane. Our hypotheses are admittedly restrictive, yet for square complexes satisfying our hypotheses we are able to reach the significantly stronger conclusion that every flat plane is the limit of periodic flat planes. The possible range of positive results is limited by the existence of an example of a compact nonpositively curved square complex X such that $\widetilde{X}$ contains a flat plane which is not the limit of periodic flats [27]. It has thus become an intriguing problem to find the appropriate hypotheses guaranteeing that flats are limits of periodic flats.

The principal objects of our study are $\mathcal{VH}$*-complexes* which are introduced in Section 2.1. A $\mathcal{VH}$-complex is a square complex whose 1-cells are partitioned into *vertical* and *horizontal* classes so that each square 2-cell is attached along two vertical and two horizontal 1-cells in a manner respecting the partition. Nonpositively curved $\mathcal{VH}$-complexes arise quite naturally. For instance, following an idea of Weinbaum's [20], we observed in [23] that the Dehn complex of any prime alternating link is a nonpositively curved $\mathcal{VH}$-complex.

In Section 3 we show that every nonpositively curved square complex X has a branched double cover $\widehat{X}$ which is a $\mathcal{VH}$-complex, and moreover every flat in $\widetilde{X}$ is the limit of periodic flats if and only if every flat in $\widetilde{\widehat{X}}$ is the limit of periodic flats. We thus

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

reduce the problem to the study of compact nonpositively curved $\mathcal{VH}$-complexes.

The vertical-horizontal geometry of a $\mathcal{VH}$-complex X yields a decomposition of X as a graph of spaces which is described in Definition 2.14 and Construction 2.15. This enables us easily to describe examples in a manner analogous to the description of a group as a graph of groups. In fact, Theorem 2.16 asserts that the fundamental group of a nonpositively curved $\mathcal{VH}$-complex splits as a graph of free groups with free edge groups. By placing various hypotheses on the attaching maps of the edge spaces of the graph of spaces, we obtain various classes of tractable $\mathcal{VH}$-complexes. For instance, the motivating cases are *clean* $\mathcal{VH}$-complexes where all the attaching maps are embeddings.

Section 4 contains the main technical conditions which are used to prove that every immersed flat plane is the limit of immersed periodic flat planes. Section 4.1 defines when an immersed flat plane is the *limit of immersed tori* which is equivalent to the lifted flat in the universal cover being the limit of periodic flats. In addition, we define the *strong limit of cylinders* which is similar but stronger, and which arises more naturally in this work. Section 4.2 presents some of the key terminology and lemmas involving trajectories that will be used later on. Roughly speaking, the *trajectory* of an immersed flat plane is a partial description of the way in which the flat plane travels through the complex. Section 4.3 presents some lemmas centered around Lemma 4.15 which states that an immersed line in a finite graph is the limit of immersed periodic lines. This is a 1-dimensional analogue of our problem which is easy to prove. Our criterion for showing that flats are strong limits of cylinders is Lemma 4.16 which uses trajectories to provide a bridge from our 2-dimensional problem to the easier 1-dimensional analogue.

Section 5 presents a sequence of conditions on the attaching maps of the decomposition of X which imply that the criterion of Lemma 4.16 is met. Section 5.1 presents the fiber product $\Upsilon_1 \otimes \Upsilon_2$ of maps of graphs which is useful for presenting the proofs of Section 5.3. Section 5.2 introduces *malnormal* $\mathcal{VH}$-complexes. The main theorem of 5.2 is Theorem 5.12, which states that a malnormal $\mathcal{VH}$-complex has the property that every immersed flat plane is the strong limit of immersed cylinders. Section 5.2 should be thought of as a warm up for Section 5.3, which contains the strongest results of Section 5. *Cyclonormal $\mathcal{VH}$-complexes* are a generalization of Malnormal $\mathcal{VH}$-complexes. Theorem 5.17 states that every immersed flat plane in a nonpositively curved cyclonormal complex is the strong limit of immersed cylinders. Section 5.4 contains a conjectured generalization of the theorems of Section 5.3, which will hopefully bring the techniques used here to bear on certain graphs of word-hyperbolic groups and thus augment the combination theorem of Bestvina and Feighn [2].

Section 6 presents a surprising connection between the geometric property of approximable flats in X and the algebraic property of separability with respect to certain subgroups of $\pi_1 X$. An immersion of a nonapproximable plane in X is an obstruction to the separability of edge groups of $\pi_1 X$. This idea was exploited in [24] to construct compact nonpositively curved square complexes X such that $\pi_1 X$ is not residually finite, which resolved a longstanding open problem about $C(4)$-$T(4)$ small-cancellation groups. We raise the question of whether conditions involving flats are the only obstructions to virtual cleanliness, and we describe some progress we have made on this problem.

A *Complete Square Complex*, or *CSC*, is a square complex X such that $\mathrm{Link}(x)$ is a complete bipartite graph for each $x \in X^0$. In Section 7 we prove that a compact CSC has the property that every immersed flat plane is the limit of periodic flats. This result is of a different nature from the results of Section 5, because the method of Section 4, enables us to show that flat planes are strong limits of cylinders, whereas it seems unlikely that such a strong result holds for compact CSC. While the results of Section 4 and 5 stem largely from the scarcity of flats with a given trajectory, the CSC result derives from a density of flats.

Elements g and h of a group G *nearly commute* if $g^n h^n = h^n g^n$ for some $n \geq 1$. In Section 8 we describe an algorithm to decide if elements nearly commute for fundamental groups of complexes satisfying the conditions of Sections 4 and 5. We ask whether such an algorithm exists for the fundamental group of a CSC.

2 $\mathcal{VH}$-complexes

2.1 $\mathcal{VH}$ Definitions

Notation 2.1 Throughout this paper, we let $\mathbb{R} = (-\infty, \infty)$ and $\mathbb{R}^+ = [0, \infty)$ with the usual structures as graphs, with 0-cells at each $n \in \mathbb{Z}$ and open 1-cells at each $(n, n+1)$. We let $I_n \subset \mathbb{R}$ denote the subgraph $[0, n]$, and we let $I = I_1 = [0, 1]$.

Definition 2.2 (Square Complex $\mathcal{VH}$-Complex) A *square complex* X is a combinatorial 2-complex whose 2-cells are attached by combinatorial paths of length 4. Thus, we think of each 2-cell as a square attached to X^1.

A square complex X is a $\mathcal{VH}$*-complex* if the 1-cells of X are partitioned into two classes V and H, called *vertical* and *horizontal* edges respectively, and as in square (i) of Figure 2, the attaching map of each 2-cell of X alternates between edges in V and H.

We let $V_X = V \cup X^0$ denote the *vertical* 1*-skeleton* and $H_X = H \cup X^0$ denote the *horizontal* 1*-skeleton*. For a 0-cell $x \in X^0$, we let V_x denote the component of V_X containing x. We define H_x similarly.

Remark 2.3 (Bipartite Links) Let $x \in X^0$ where X is a 2-complex. We let $\mathrm{Link}(x)$ denote the *link* of x in X which is a graph whose vertices and edges correspond to the ends of 1-cells and corners of 2-cells incident with x. Note that $\mathrm{Link}(x)$ is topologized so that it looks like the ϵ-sphere about x in X.

Recall that a graph Γ is *bipartite* if Γ^0 is partitioned into two disjoint classes such that each edge of Γ connects vertices from distinct classes.

Let X be a $\mathcal{VH}$-complex and let $x \in X^0$. The partition of the 1-cells of X into two classes V and H, induces a partition of the vertices of $\mathrm{Link}(x)$. Furthermore, since attaching maps of the squares of X alternate between 1-cells in V and 1-cells in H, we see that the edges of $\mathrm{Link}(x)$ connect vertices from different classes. Therefore, the $\mathcal{VH}$-structure on X induces a bipartite structure on $\mathrm{Link}(x)$ for each $x \in X^0$.

This motivates the following definition:

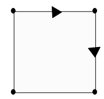

Figure 1: Local but not global: The complex obtained by identifying two sides of a square as indicated above has a local $\mathcal{VH}$-structure but no global $\mathcal{VH}$-structure.

Definition 2.4 (Locally $\mathcal{VH}$) The square complex X is *locally* $\mathcal{VH}$ if for each $x \in X^0$ there is a chosen bipartite structure on the graph $\mathrm{Link}(x)$. As in Remark 2.3, a $\mathcal{VH}$-complex has an induced local $\mathcal{VH}$-structure.

Note that X may admit several $\mathcal{VH}$-structures because a graph admiting a bipartite structure actually admits 2^c such structures, where c is the number of connected components of the graph.

Example 2.5 (Not Global $\mathcal{VH}$) The simplest example of a square complex with a local $\mathcal{VH}$-structure which is not consistent with any (global) $\mathcal{VH}$-structure is a loop with one 0-cell and one 1-cell, where the bipartite structure on the link of the 0-cell has one vertex in each class. While the underlying complex in this example does have two $\mathcal{VH}$-structures, both are inconsistent with the local $\mathcal{VH}$-structure that we chose.

A square complex which has a local $\mathcal{VH}$-structure but no (global) $\mathcal{VH}$-structure, is obtained from a square by identifying two of its sides as in Figure 1.

While Example 2.5 shows that there are many examples of local $\mathcal{VH}$-complexes with no consistent (global) $\mathcal{VH}$-structure, we do have the following theorem which is analogous to the existence of orientable double covers of manifolds.

Theorem 2.6 ($\mathcal{VH}$ Double Cover) *Let X be a square complex which is locally $\mathcal{VH}$, then there is a double cover $\widehat{X} \to X$ such that the induced local $\mathcal{VH}$-structure on $\widehat{X}$ is consistent with a global $\mathcal{VH}$-structure.*

Proof In this proof, we use the words 0-cell and 1-cell when referring to 0-cells and 1-cells of X, while reserving the words vertices and edges for the vertices and edges in the links of 0-cells of X.

For each 0-cell $y \in X^0$, let B_y denote the pair of classes of vertices comprising the bipartite structure of $\mathrm{Link}(y)$. So B_y is a two-element set, and each of its elements is a (possibly empty) class of vertices of $\mathrm{Link}(y)$.

Let $x \in X^0$ be the basepoint of X, we will form a homomorphism

$$\phi\colon \pi_1(X, x) \to \mathrm{Aut}(B_x).$$

We will show that the covering space $\widehat{X}$ of X corresponding to ϕ has a $\mathcal{VH}$-structure. More precisely, it has an induced local $\mathcal{VH}$-structure which is consistent with a global $\mathcal{VH}$-structure.

A 1-cell e with initial vertex y and terminal vertex z induces a one-to-one correspondence $\phi_e\colon B_y \to B_z$. This is done in the obvious way: the end of e at y corresponds to a vertex in $\mathrm{Link}(y)$, and the end of e at z corresponds to a vertex in $\mathrm{Link}(z)$.

We just choose $\phi_e\colon B_y \to B_z$ to be the correspondence which associates the classes corresponding to the the ends of e.

Similarly, a combinatorial path $\sigma \to X^1$ which begins at x_0 and ends at x_n induces a bijection $\phi_\sigma\colon B_{x_0} \to B_{x_n}$. This is defined by breaking σ up into the concatenation of edges, $\sigma = e_1 \cdot e_2 \cdots e_n$, and letting $B_{x_0} \to B_{x_n}$ be the corresponding composition: $\phi_{e_n} \circ \cdots \circ \phi_{e_2} \circ \phi_{e_1}$, or in other words, the obvious map $B_{x_0} \to B_{x_1} \to B_{x_2} \cdots B_{x_{n-1}} \to B_{x_n}$. In particular, if σ is a closed path based at y, then it induces an automorphism of B_y.

Now observe that if σ is a (closed) path of length 4 starting at y which travels around the boundary of some square of X, then ϕ_σ induces the identity map on B_y. From this it is easy to conclude that if two paths σ_1 and σ_2 with the same endpoints y and z are homotopic, then $\phi_{\sigma_1} = \phi_{\sigma_2}$. It follows that there is a homomorphism $\psi\colon \pi_1(X, x) \to \mathrm{Aut}(B_x)$.

Let $\widehat{X}$ denote the covering space of X corresponding to the kernel of ψ. We claim that the local $\mathcal{VH}$-structure on $\widehat{X}$ induced by the covering map is consistent with a global $\mathcal{VH}$-structure. To form this global $\mathcal{VH}$-structure, we choose a basepoint $\hat{x}$ of $\widehat{X}$ which is a preimage of x. Then we declare the pair of classes in $B_{\hat{x}}$ of vertices of $\mathrm{Link}(\hat{x})$, to be Vertical and Horizontal (choose either way). Now for each $\hat{y} \in \widehat{X}^0$, a path in $\widehat{X}^1$ from $\hat{x}$ to $\hat{y}$ determines a one-to-one correspondence $B_{\hat{x}} \cong B_{\hat{y}}$, and thus determines a vertical and horizontal name for each class of $B_{\hat{y}}$. Since $\pi_1\widehat{X}$ is the kernel of ψ, the one-to-one correspondence does not depend on the path, and so this way of labeling the classes of $B_{\hat{y}}$ for each $\hat{y}$ is well defined.

Now each 1-cell of $\widehat{X}$ can be declared either horizontal or vertical according to the vertices that its ends correspond to in the links of the 0-cells that the 1-cell is attached to. The bipartite structure of each link implies that the condition on the attaching maps of squares of $\widehat{X}$ is satisfied. Thus $\widehat{X}$ has a global $\mathcal{VH}$-structure consistent with the local $\mathcal{VH}$-structure induced from X. ■

A similar argument works for higher dimensional versions of the class of $\mathcal{VH}$-complexes, but the corresponding homomorphism is to S_n and so may require a larger degree cover.

Definition 2.7 (Nonpositive Curvature) A square complex X is *nonpositively curved* if all immersed cycles in the link of each 0-cell of X have length at least 4. A square complex satisfying this combinatorial link condition admits a locally CAT(0) metric (see [9] or [4]).

Remark 2.8 (Locally-$\mathcal{VH}$ and Curvature) If X is a locally-$\mathcal{VH}$ square complex, then X satisfies the combinatorial nonpositive curvature condition if and only if there are no cycles of length 2 in the links of 0-cells of X. This is because all the cycles in a bipartite graph have even length, and so we may rule out the short cycles of length 1 or 3.

Definition 2.9 (Directed $\mathcal{VH}$-Complex) We view the attaching map of each square of a square complex X, as a map from the boundary of the unit square $I \times I$ to X^1. We orient both horizontal edges of the unit square from left to right as illustrated in the

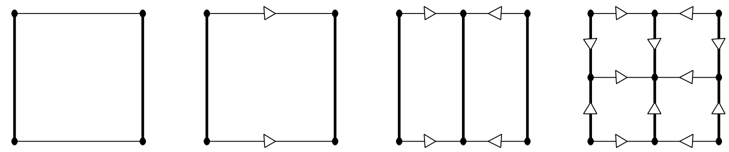

Figure 2: Squares: The squares in the figure above are meant to suggest (from left to right) a $\mathcal{VH}$-square, a horizontally directed $\mathcal{VH}$-square, a subdivided $\mathcal{VH}$-square (to obtain two horizontally directed subsquares), and the first barycentric subdivision of a $\mathcal{VH}$-square.

second square of Figure 2. Let X be a $\mathcal{VH}$-complex and suppose that H_X is a directed graph. The $\mathcal{VH}$-structure on X is *horizontally directed* if the attaching map of each square of X is orientation preserving on its horizontal edges. We define *vertically directed* similarly. We shall use the term *directed* to mean *horizontally directed*.

Remark 2.10 (Subdividing) There is little loss of generality in considering only directed $\mathcal{VH}$-complexes. This is because, given a $\mathcal{VH}$-complex, we may subdivide H_X and subdivide each square of X by adding a vertical edge connecting the centers of its horizontal edges. If we orient all horizontal edges towards the new 0-cells then we obtain a directed $\mathcal{VH}$-complex. See the third square of Figure 2. Similarly, we can subdivide each square both vertically and horizontally so that X is both vertically and horizontally oriented. See the fourth square of Figure 2.

2.2 Decomposition Theorem

Definition 2.11 (Vertical Foliation and V_x) We now define a "singular vertical foliation" on the $\mathcal{VH}$-complex X. The unit square $I \times I$ is foliated by vertical line segments. Similarly, the image of a square in X is foliated by *vertical segments* parallel to the pair of vertical edges on its boundary. For an arbitrary point $x \in X$ we define the *leaf* V_x to be the smallest subset of X having the property that $x \in V_x$ and that V_x contains any vertical segment which intersects it. This definition of V_x is consistent with the definition given earlier for $x \in X^0$.

Thinking of X as being foliated by these vertical subspaces, it is natural to take the quotient of X in which each leaf is identified to a point. This quotient is a graph denoted by Γ_X which we shall discuss in Definition 2.14.

Remark 2.12 (Singular Leaves and Directed $\mathcal{VH}$) When x is a point in the interior but not in the center of a horizontal edge, then for y very close to x, the leaf V_x is isomorphic to the leaf V_y, by an isomorphism induced by sliding it along in X. However, when x is the center of a horizontal edge, the leaf V_x may be different from the surrounding leaves, in which case they correspond to double covers of it. This situation can occur only if X is not directed. It is convenient to add all such singular leaves V_x to the vertical 1-skeleton of X. This corresponds to subdividing certain squares of X. The resulting complex has a directed $\mathcal{VH}$-structure. Note that if H_X can be oriented so that V_X is directed, then for any points x and y in the same horizontal edge, the leaves V_x and V_y are isomorphic by a translation isomorphism.

Example 2.13 (Möbius Strip) A Möbius strip obtained by identifying the top and bottom horizontal edge of a square with a twisted identification map, has an obvious $\mathcal{VH}$-structure. The circle at the center of the Möbius Strip is singular.

Definition 2.14 (The Decomposition Graph $X \to \Gamma_X$) Given a directed $\mathcal{VH}$-complex X, we define a graph Γ_X and a map $\varrho\colon X \to \Gamma_X$. The vertices of Γ_X correspond to the connected components of V_X, which are called *vertex spaces*. Note that each vertex space arises as V_x for some 0-cell $x \in X^0$. The edges of Γ_X correspond to the connected components of $X - V_X$, which are called *edge spaces*.

If x and y are in the same edge space, then V_x and V_y are isomorphic graphs and there is a natural isomorphism between them. Indeed, if x is a point in some edge space C, then $C \cong V_x \times (0, 1)$. It is natural to think of C as a subspace of $V_x \times [0, 1]$ which is a square complex that we denote by $\bar{C}$. We will also refer to $\bar{C}$ as an edge space.

For each edge space C, the inclusion $C \hookrightarrow X$ uniquely extends to a combinatorial map $\bar{C} \to X$. Since $\bar{C} \cong V_x \times I$, we obtain induced maps $V_x \times \{0\} \to V_X$ and $V_x \times \{1\} \to V_X$, which we call the attaching maps of the edge space C. Note that these are combinatorial maps, and if X is nonpositively curved then they are immersions, *i.e.*, local injections, and are therefore π_1-injective [19]. The components of V_X to which the ends of $\bar{C}$ are mapped correspond to the vertices of Γ_X to which the edge of Γ_X corresponding to C is attached.

Finally, the map $X \to \Gamma_X$ is the quotient map induced by identifying each vertical leaf of Definition 2.11 to a point.

Construction 2.15 (Constructing X from Data) In order to further understand the map $\varrho\colon X \to \Gamma_X$ we show how X can be built up from the information encoded in Γ_X and the associated data. This is a special case of the notion of a graph of spaces [18].

Consider a graph Γ_X, and suppose that for each vertex $v \in \Gamma_X^0$ we have an associated *vertex space* X_v which is a graph, and for each edge $e \in \text{Edges}(\Gamma_X)$ we have an associated *edge space* $X_e \times I$ where X_e is a graph. Suppose that for each edge e which is attached to the vertices e_0 and e_1, there are corresponding maps $\phi_{e0}\colon X_e \times \{0\} \to X_{e_0}$ and $\phi_{e1}\colon X_e \times \{1\} \to X_{e_1}$.

Using this data we may construct a $\mathcal{VH}$-complex X as follows: Let V_X be the disjoint union of the set of vertex spaces, that is

$$V_X = \Big\{ \bigcup_{v \in \Gamma_X^0} X_v \Big\}.$$

We form X by attaching $X_e \times I$ along its ends to V_X for each edge e of Γ_X, so that

$$X = \Big\{ V_X \bigcup_{e \in \text{Edges}(\Gamma_X)} X_e \times I \Big\}.$$

Theorem 2.16 (Graph of Free Groups) *Suppose that X is a nonpositively curved $\mathcal{VH}$-complex. Then the map $\varrho\colon X \to \Gamma_X$ determines a splitting of $\pi_1 X$ as a graph of free groups. Specifically, for each vertex $v \in \Gamma_X$, $\pi_1 v = \pi_1 X_v$ and for each edge e of Γ_X, $\pi_1 e = \pi_1 X_e$. For each edge e attached to the vertices e_0 and e_1, the inclusion $\pi_1 X_e \to \pi_1 X_{e_i}$ is induced by the maps ϕ_{ei} described above.*

Proof The nonpositive curvature hypothesis implies that the attaching maps ϕ_{ei} are immersions. They therefore induce π_1-injections [19]. ■

We close this section with the following:

Definition 2.17 (Map of $\mathcal{VH}$-Complexes) Let X and Y be $\mathcal{VH}$-complexes, and let $\phi\colon X \to Y$ be a cellular map. Then ϕ is a $\mathcal{VH}$-*map* provided that $\phi(V_X) \subset V_Y$ and $\phi(H_X) \subset H_Y$. We will be dealing with combinatorial maps, and so ϕ is a $\mathcal{VH}$-map provided it maps vertical edges to vertical edges and horizontal edges to horizontal edges. Note that for a $\mathcal{VH}$-map ϕ, there is an induced map $\phi_\star\colon \Gamma_X \to \Gamma_Y$.

3 $\mathcal{VH}$ Branched Cover of Square Complex

In this section we show how to use a branched covering space in order to restrict our analysis of flat planes in square complexes to $\mathcal{VH}$-complexes.

A *flat* $\mathcal{VH}$-complex, is a $\mathcal{VH}$-complex that is locally Euclidean in the sense that the link of each 0-cell is either a cycle of length 4 or isomorphic to I_2.

Note that since we are working in the combinatorial category, a map is an *immersion* if and only if the induced maps on the links of 0-cells are injective.

The following is the main result of this section. Its proof is deferred until after Lemma 3.4.

Theorem 3.1 ($\mathcal{VH}$-Branched Covering) *Let X be a square complex. Then there exists a* 1 *or* 2 *fold branched cover $\rho\colon \overline{X} \to X$ such that:*

(1) *$\overline{X}$ is a $\mathcal{VH}$-complex.*
(2) *Every immersion $\phi\colon F \to X$ from a flat $\mathcal{VH}$-complex F to X lifts to $\overline{X}$.*
(3) *If X is nonpositively curved, then every immersion $\phi\colon F \to \overline{X}$ of a flat $\mathcal{VH}$-complex F projects to an immersion in X.*
(4) *If X is nonpositively curved, then up to covering space transformations there is a one-to-one correspondence between immersed flat $\mathcal{VH}$-complexes in $\overline{X}$ and immersed flat $\mathcal{VH}$-complexes in X.*

Remark 3.2 The one-to-one correspondence of Theorem 3.1 enables us to restrict our study of immersed flat planes in nonpositively curved square complexes to the immersions in nonpositively curved $\mathcal{VH}$-complexes. The main benefit is that nonpositively curved $\mathcal{VH}$-complexes can be easily described in terms of their decompositions (see Definition 2.14). In particular, we can describe classes of nonpositively curved $\mathcal{VH}$-complexes in terms of properties of their decompositions.

We apologize to the reader for invoking Definitions 4.2 and 4.4 at this point. Theorem 3.1 shows that if $\overline{X}$ has any of the following properties, then so does X.

(1) If there is an immersed flat plane then there is an immersed flat torus.
(2) Every immersed flat plane is the limit of immersed tori.
(3) Every immersed flat plane is the strong limit of immersed cylinders.

For instance, if $\overline{X}$ has property (1), then we can conclude that X has property (1) as well. Indeed, by Theorem 3.1.2, any immersed flat plane in X lifts to an immersion in $\overline{X}$, but since $\widehat{X}$ has property (1), there is an immersed flat torus $\phi\colon T \to \overline{X}$. Finally, by Theorem 3.1.3, the composition $T \xrightarrow{\phi} \overline{X} \xrightarrow{\rho} X$ is an immersion of T to X.

Likewise, if $\widehat{X}$ has the property that every immersed flat plane E is the limit of a sequence $\{T_n\}$ of immersed tori, then $\phi(E)$ is the limit of the sequence $\{\phi(T_n)\}$, and so X has this property as well. One argues similarly for the strong limit of cylinders.

Construction 3.3 (Branched Coverings) Given a combinatorial 2-complex X, we form the branched covering spaces of X where the branching takes place over the 0-cells of X by considering genuine covering spaces of $X - X^0$ and then extending these covering spaces to branched covering spaces of X by adding copies of the removed 0-cells.

$X - X^0$ deformation retracts onto the subspace $X_\star$ defined as follows: first, it is convenient to regard the 2-cells of X as unit, regular, Euclidean polygons. (One similarly handles a 2-cell which is a 1-gon or 2-gon.) Given such a 2-cell c, we let c' denote the convex closure of the centers of the edges of c.

We will also denote by c' the image of $c' \to c \to X$. Thus, if c_1 and c_2 denote 2-cells with a common 1-cell e on their boundaries, then c_1' and c_2' have a common 0-cell at their intersection with e. For each 1-cell e of X, let e' denote its center. We now obtain a subspace $X_\star \subset X$ which is defined to be:

$$X_\star = \Big\{ \bigcup_{e \in \text{ 1-cells}(X)} e' \Big\} \cup \Big\{ \bigcup_{c \in \text{ 2-cells}(X)} c' \Big\}.$$

While $X_\star$ is not a subcomplex of X, it is a subcomplex of an obvious subdivision of X. Finally, observe that $\pi_1(X_\star)$ is free.

Lemma 3.4 *$X - X^0$ deformation retracts onto $X_\star$.*

Proof This is easy to see by first subdividing X so that $X_\star$ is a subcomplex. We first subdivide each 1-cell e, by adding the 0-cell e' at its barycenter. We then subdivide each 2-cell c in the same way as the Euclidean polygon c is subdivided by the inclusion $c' \subset c$. Having done this, it is easy to see that

$$X_\star = X - \bigcup_{v \in X^0} \operatorname{Star}(v).$$

Finally, Closure(Star(v)) $- v$ deformation retracts onto Link(v) for each $v \in X^0$. ■

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

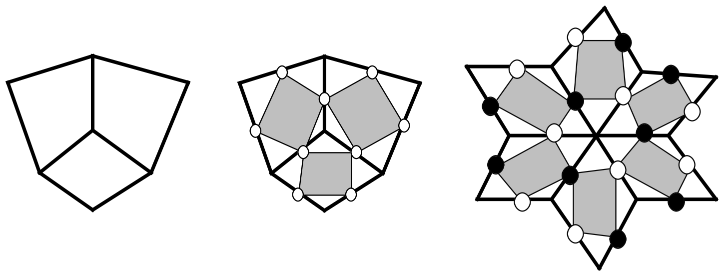

Figure 3: The figure on the left is a square complex X formed by gluing three squares together. The figure in the center is a subdivided copy of X, containing $X_\star$ as a subspace. The figure on the right depicts the branched cover $\overline{X}$ of X. Note that it contains a double cover of $X_\star$, and that the coloring of the vertices of $X_\star$ induces a $\mathcal{VH}$-structure on $\overline{X}$

Proof of Theorem 3.1 (1) We now use the subspace $X_\star \subset X$ defined above, to choose the desired branched cover. In the following argument we will assume that $X_\star$ is connected. The case where $X_\star$ is not connected may be handled by choosing covering spaces as below for each component of $X_\star$ and proceeding analogously. Notice that the number of different $\mathcal{VH}$-structures for the space $\overline{X}$ obtained below is 2^c where c is the number of connected components in $X_\star$.

The homomorphism $\pi_1(X_\star^1) \to \mathbb{Z}_2$ induced by amalgamating the vertices and then sending all edges to generators of $\mathbb{Z}_2$, extends to a homomorphism of $\pi_1(X_\star)$ because all the 2-cells of $X_\star$ have an even number of sides (they are all squares). This homomorphism corresponds to a (possibly trivial) covering space $\widehat{X_\star} \to X_\star$ and therefore a branched covering space $\overline{X} \to X$. This covering space has the property that all cycles in $\widehat{X_\star}^1$ have even length. (It is characterized by being the smallest such covering.) It follows that choosing one vertex to be black in each component of $\widehat{X_\star}$, the homomorphism induces a coloring of the vertices of $\widehat{X_\star}$ as alternatingly black and white. The 1-cells of $\overline{X}$ inherit this coloring, and furthermore the edges of each square 2-cell of $\overline{X}$ are colored alternatingly black and white. This corresponds to the square polygon inside each 2-cell of $\overline{X}$ whose vertices are colored alternatingly. Thus we have obtained a $\mathcal{VH}$-structure on the branched cover $\overline{X}$ of X.

(2) Consider an immersed flat plane $\phi\colon E \to X$. There is an induced map $\phi_\star\colon E_\star \to X_\star$. All closed loops in $E_\star$ have even length. It follows that $E_\star$ lifts to $\widehat{X}_\star$ and therefore E lifts to $\overline{X}$.

(3) Consider an immersed flat torus in $\overline{X}$. The simple cycles in the links of its vertices all have length 4. Now, for any $v \in \overline{X}$, the induced map $\mathrm{Link}(v) \to \mathrm{Link}(\rho(v))$ is a covering map. Consequently, an immersed cycle in $\mathrm{Link}(v)$ is sent to an immersed cycle in $\mathrm{Link}(\rho(v))$. It follows that the projection of the torus is an immersion, since there are no cycles of length < 4 in $\mathrm{Link}(\rho(v))$. ■

Remark 3.5 (Atoroidal Branched Covers) It is tempting to try and form a compact nonpositively curved complex with an immersed flat plane but no immersed torus in the following manner: we begin with a complex X which admits immersed tori. Then we form a branched covering space such that some of the flat planes of this complex lift, but none of the periodic flat planes lift. This difficulty arises because the short cycles which we hope will not lift, are generally in the normal closure of the set of cycles that have length 2π which must be lifted to the branched cover in order that the flat plane lifts. In any case, the branched cover would have to correspond to a non-normal finite index subgroup.

4 Approximating by Tori

4.1 Limits

We preface this section with the following observation which justifies our interest in immersed flats in X rather than embedded flats in $\widetilde{X}$.

Lemma 4.1 *Let F be a flat $\mathcal{VH}$-complex and X be a nonpositively curved square complex. Let $\phi\colon F \to X$ be a map and let $\widetilde{\phi}\colon \widetilde{F} \to \widetilde{X}$ be a lift of this map to the universal covers. Then ϕ is an immersion if and only if $\widetilde{\phi}$ is an embedding.*

Proof If $\widetilde{\phi}$ is an embedding then ϕ is obviously an immersion, so we shall concentrate on the converse. An immersion $\phi\colon Y \to X$ of nonpositively curved complexes is a *local-isometry* if for each $y \in Y^0$ the induced map $\phi\colon \mathrm{Link}(y) \to \mathrm{Link}(\phi(y))$ preserves adjacency, where two vertices are *adjacent* if they are joined by an edge. It follows easily from the definitions that if Y is a flat $\mathcal{VH}$-complex then any immersion $Y \to X$ is a local-isometry. A map which is a local-isometry in this combinatorial sense is actually a local-isometry in the usual metric sense if we give Y and X the metrics of nonpositive curvature consistent with the Euclidean squares. A local isometry between nonpositively curved metric spaces lifts to an isometric embedding of their universal covers (see for instance, [4]). ■

Definition 4.2 (Limit of Tori) Let X be a nonpositively curved space. Let $\phi\colon E \to X$ be a local isometry of a flat plane to X. We say that $\phi\colon E \to X$ is the *limit of periodic flats* or the *limit of tori* if there is a sequence $\phi_n\colon E \to X$ of local isometries of periodic planes such that $\lim_{n\to\infty} \phi_n = \phi$, in the sense that $\lim_{n\to\infty} \phi_n(x) = \phi(x)$ for each $x \in E$. Note that $E \to X$ is a *periodic flat* if it factors through a covering map of an immersed torus $E \to T \to X$.

Remark 4.3 When X is a 2-dimensional nonpositively curved complex, then an immersed flat plane $E \to X$ is the limit of tori if and only if for every compact set $B \subset E$, the restriction $\phi\colon B \to X$ extends to a local-isometry $T \to X$ from a flat torus to X. In particular, if X is a nonpositively curved square complex then $\phi\colon E \to X$ is the limit of periodic flats if and only if for each n the restriction of ϕ to $[-n, n] \times [n, n]$ extends to a local isometry from a flat torus to X.

Definition 4.4 (Strong limit of cylinders) Let X be a nonpositively curved 2-complex. A local isometry $\phi\colon E \to X$ from a flat plane to X is the *strong limit of cylinders* if E is the increasing union of flat strips, *e.g.,* $\{[-n, n] \times \mathbb{R}\}$, each of which can be extended to an immersion of a cylinder $S^1 \times \mathbb{R} \to X$.

Remark 4.5 As we shall see in Lemma 4.16, if X is compact, a plane which is the strong limit of cylinders is also the limit of tori. I believe that the strong limit of cylinders is actually a stronger property. For instance, as proven in Theorem 7.8, every immersed flat plane in a compact CSC is the limit of tori. However, it appears likely that there is an anti-torus in a CSC as in Example 7.3 which is not the strong limit of cylinders.

To put some of the results that will be obtained in the remainder of Section 4 into perspective, we first give a simple but very strict condition on a space which guarantees that all flats are periodic.

Proposition 4.6 (Isolated $\Rightarrow$ Periodic) *Suppose that the group G acts cocompactly and properly discontinuously on the* CAT(0) *space $\widetilde{X}$. Suppose that for a fixed constant K, finitely many flat planes pass through any ball of radius K in $\widetilde{X}$. Then every flat in $\widetilde{X}$ is periodic.*

Sketch of Proof It is easy to deduce from the hypothesis that for each R, there are finitely many flats passing through any ball of radius R, and that the number of such flats depends only on R. Now, let $E \subset \widetilde{X}$ be a flat. Choose a point $x \in E$, and consider the orbit of x in $\widetilde{X}$. For some R, every point of E is within R of some translate of x. Now the pigeon hole principle shows that E is periodic. ∎

The condition of Proposition 4.6 arises quite naturally. For example, it is satisfied by any compact nonpositively curved 2-complex whose 2-cells are regular Euclidean hexagons. Indeed, any plane is determined uniquely by a pair of hexagons meeting along an edge. We refer to [22] for a generalization of this.

While Proposition 4.6 uses the scarcity of planes, the more general approach in this section will use the scarcity of planes in a certain direction. This is in contrast to the result of Section 7 were it is the abundance of planes which provides the result.

4.2 Trajectories

Our plan is to investigate the properties of immersed flat planes in a nonpositively curved $\mathcal{VH}$-complex by analyzing the entire set of immersed flat strips which share a common "trajectory".

Definition 4.7 (Flat Strips) By a *bi-infinite flat strip* we mean either $I_n \times \mathbb{R}$ or $\mathbb{R}^+ \times \mathbb{R}$ or $\mathbb{R} \times \mathbb{R}$. Similarly by an *infinite flat strip* we mean either $I_n \times \mathbb{R}^+$ or $\mathbb{R}^+ \times \mathbb{R}^+$ or $\mathbb{R} \times \mathbb{R}^+$. We will refer to any of these possibilities as a flat strip, and it will be clear from the context which (if not all) of these possibilities is meant.

Note that any flat strip is a $\mathcal{VH}$-complex since it is the product of two graphs. We shall regard the 1-cells parallel to the second factor as horizontal.

Definition 4.8 (Trajectory) Given a $\mathcal{VH}$-immersion $\phi\colon F \to X$ of a flat strip to a $\mathcal{VH}$-complex X, we call the induced map $\phi_\star\colon \Gamma_F \to \Gamma_X$ the *trajectory* of F in X.

It will be convenient to refer to a combinatorial map $\tau\colon \mathbb{R} \to \Gamma_X$ as *a trajectory*. Furthermore, we will say that $\phi\colon F \to X$ has trajectory τ, provided that the maps $\phi_\star\colon \Gamma_F \to \Gamma_X$ and $\tau\colon \mathbb{R} \to \Gamma_X$ are identical if we identify Γ_F and $\mathbb{R}$. We will use similar language for trajectories $\tau\colon \mathbb{R}^+ \to \Gamma_X$

Note that trajectories are not always immersions. In particular, there are examples where $X \to Y$ is an immersion but the induced map $\Gamma_X \to \Gamma_Y$ is not.

We let 0 be the basepoint of $\mathbb{R}$, and we let Λ_τ denote *the base graph* of τ, namely $\Lambda_\tau = \Lambda_{\tau(0)} \subset X$. We adopt similar conventions with $\mathbb{R}^+$ in place of $\mathbb{R}$.

Definition 4.9 (τ-Extension of a Path) Consider a trajectory $\tau\colon \mathbb{R} \to \Gamma_X$ and a vertical combinatorial path $\sigma\colon I_n \to \Lambda_\tau \subset V_X \subset X$. We say σ is *τ-extendible* to an infinite strip $\phi\colon F \to X$ provided $F = I_n \times \mathbb{R}$ is an infinite strip with trajectory τ, and $\phi\colon F \to X$ is an extension of $\sigma\colon I_n \to X$ provided we identify I_n with $I_n \times \{0\} \subset F$. We make a similar definition with $\mathbb{R}^+$ in place of $\mathbb{R}$.

Definition 4.10 (Θ_τ) Given a trajectory τ in Γ_X, we define a graph Θ_τ as follows:

(1) The vertices of Θ_τ correspond to τ-extensions of vertices of Λ_τ.
(2) The edges of Θ_τ correspond to τ-extensions of edges of Λ_τ.
(3) Each edge $\alpha\colon I \times \mathbb{R} \to X$ of Θ_τ is attached to the vertices of Θ_τ which are the restrictions of α to $\{0\} \times \mathbb{R}$ and $\{1\} \times \mathbb{R}$.

There is an obvious correspondence between paths in Θ_τ and τ-extensions of paths in Λ_τ. It will be convenient to use $\hat{\sigma}$ to denote a path in Θ_τ corresponding to the τ-extension of a path σ in Λ_τ.

Lemma 4.11 *Let X be a nonpositively curved $\mathcal{VH}$-complex and let τ be a trajectory in Γ_X. The map $\Theta_\tau \to \Lambda_\tau$ induced by sending $\hat{v} \to v$ and $\hat{e} \to e$ for each τ-extension of each vertex v and edge e of Λ_τ is an immersion.*

Proof Since the argument for $\mathbb{R}$ is similar, we shall assume that τ is an $\mathbb{R}^+$ trajectory. Suppose that $\Theta_\tau \to \Lambda_\tau$ is not an immersion. Then there is a pair of 1-cells, $\hat{e}_1$ and $\hat{e}_2$, meeting along a 0-cell $\hat{v}$ in Θ_τ such that $\hat{e}_1$ and $\hat{e}_2$ fold to the same 1-cell e of Λ_τ. We shall regard $\hat{e}_1$ and $\hat{e}_2$ as immersed flat strips.

Let s_1 and s_2 be the first squares in $I \times \mathbb{R}^+$ such that $\hat{e}_1$ and $\hat{e}_2$ disagree (so s_i correspond to $I \times [n, n+1]$ for some n). Let x be the 0-cell in X corresponding to $\{0\} \times \{n\}$. Then Link(x) would have a length 2 cycle corresponding to the pair of edges $I \times \{n\}$ and $\{0\} \times [n, n+1]$. This contradicts that X is nonpositively curved. ■

Remark 4.12 (Θ_τ Is Large) In general, Θ_τ may have uncountably many components. However if the connected components of Θ_τ are finite then we can make assertions about τ-flats. For instance, in Lemma 4.13 we show that if the components are finite, then any τ-flat is the strong limit of τ-cylinders.

Lemma 4.13 (Finitely Many τ-Extensions $\Rightarrow$ Compact Components) *Let X be a nonpositively curved directed $\mathcal{VH}$-complex. Let τ be a trajectory whose base graph Λ_τ is finite. Suppose there are numbers J and L such that each immersed path of length J in Λ_τ has at most L distinct τ-extensions. Then there is a number D such that every component of Θ_τ has diameter $\leq D$. Furthermore all but finitely many components of Θ_τ are trees.*

Proof By Lemma 4.11, the map $\Theta_\tau \to \Lambda_\tau$ is an immersion, and so Θ_τ is locally finite because Λ_τ is locally finite. Therefore, if each component of Θ_τ has bounded diameter then each component is compact.

Let H be the number of paths of length J in Λ_τ. We show that there are no injective paths in Θ_τ of length $\geq D = J(HL+1)$. Suppose there is a path $\hat{\sigma}$ of length $J(HL+1)$ in Θ_τ. We divide $\hat{\sigma}$ into $(HL+1)$ subpaths each of length J. Since there are at most HL different paths of length J in Θ_τ, the pigeon-hole principle shows that two of these must be the same and $\hat{\sigma}$ is not injective.

To see that all but finitely many of the components are trees, note that there are at most HL components of Θ_τ containing an immersed path of length J. ∎

Lemma 4.14 (Finitely Many Injective τ-Extensions $\Rightarrow$ Compact Components) *Let X be a nonpositively curved directed $\mathcal{VH}$-complex, and let τ be a trajectory with Λ_τ finite. Suppose there are numbers J and L such that each length J path σ has at most L distinct τ-extensions corresponding to injective paths in Θ_τ. Then there is a number D such that each component of Θ_τ has diameter $\leq D$.*

Proof Let H be the number of paths of length J in Λ_τ. We show that there are no injective paths in Θ_τ of length $\geq D = J(HL+1)$. Consider a path $\hat{\sigma}$ of length D in Θ_τ. We may divide $\hat{\sigma}$ into $(HL+1)$ subpaths each of length J. But each of these subpaths of $\hat{\sigma}$ is an extension of a path of length J in Λ_τ. Since there are only H paths of length J in Λ_τ, we see that by the pigeon-hole principle, there must be $L+1$ subpaths of $\hat{\sigma}$ which are extensions of the same path of length J in Λ_τ. But then since there are at most L distinct injective extensions of a path σ, either one of these paths is not injective, or by the pigeon hole principle, at least two of these are the same, and so either way, σ is not injective. ∎

While Θ_τ may have uncountably many non-tree components under the hypothesis of Lemma 4.14, this is impossible under the slightly more stringent hypothesis of Lemma 4.13.

4.3 Limits of Circles

Let Θ be a graph. An immersed line $\phi_i\colon \mathbb{R} \to \Theta$ is *periodic* if it factors as $\mathbb{R} \to S^1 \to \Theta$ where $\mathbb{R} \to S^1$ is a covering map. An immersed line $\phi\colon \mathbb{R} \to \Theta$ is the *limit of periodic lines* if for each finite interval $[-n, n] \subset \mathbb{R}$, the restriction $\phi\colon [-n, n] \to \Theta$ extends to a periodic line.

Lemma 4.15 (Limit of Circles) *Let Θ be a graph with finitely many vertices. Any immersed line $\phi\colon \mathbb{R} \to \Theta$ is the limit of periodic lines.*

We note that this holds when Θ contains no edge e such that both components of $\Theta - e$ have infinite diameter.

Proof We will show that for each interval $J[-n, n] \subset \mathbb{R}$ there is a subdivision of S^1 and an immersion $\psi\colon S^1 \to \Theta$ and an embedding $[-n, n] \subset S^1$ so that ψ is an extension of ϕ restricted to $[-n, n]$.

We first choose a maximal tree $\Upsilon \subset \Theta$ and form the quotient graph Θ/Υ. Now observe that there is a one-to-one correspondence between the immersed lines in Θ and immersed lines in Θ/Υ. Similarly, there is a one-to-one correspondence between immersed circles in Θ and in Θ/Υ.

Since Θ/Υ is a bouquet of circles, any path in Θ/Υ can be extended to a circle (by possibly adding one edge). Therefore, every immersed line is the limit of immersed circles. The corresponding line in Θ is thus a limit of the corresponding circles as well. ■

A τ-*flat* in X is an immersed flat $\mathcal{VH}$-complex $F \to X$ whose trajectory is τ. We define a τ-cylinder similarly.

Lemma 4.16 *Let X be a nonpositively curved directed $\mathcal{VH}$-complex and let τ be a trajectory. If each component of Θ_τ is compact then every τ-flat is the strong limit of τ-cylinders.*

Proof A τ-flat is a bi-infinite path $\hat{\omega}$ in Θ_τ. But the component of Θ_τ containing $\hat{\omega}$ is compact. Therefore by Lemma 4.15, any finite subpath $\hat{\sigma}$ of $\hat{\omega}$ is contained in a closed immersed path of Θ_τ. But this corresponds to a τ-cylinder which extends the flat strip corresponding to $\hat{\sigma}$. ■

The following shows that if an immersed plane is the strong limit of cylinders then it is the limit of tori.

Lemma 4.17 (Cylinder Is the Limit of Tori) *Let $C = S^1 \times \mathbb{R}$ for some subdivision of S^1 and let $\phi\colon C \to X$ be an immersion where X is a compact nonpositively curved $\mathcal{VH}$-complex. Then for any finite interval $J \subset \mathbb{R}$, the restriction of ϕ to $S^1 \times J$ extends to an immersion of a torus.*

Proof We form a graph Ω as follows: fix a combinatorial structure on S^1 as above. The vertices of Ω correspond to the distinct $\mathcal{VH}$-immersions $S^1 \to X$ where we declare S^1 to be vertical. The edges of Ω correspond to the distinct $\mathcal{VH}$-immersions $S^1 \times I \to X$, where S^1 is vertical and I is horizontal. Each edge $\phi\colon S^1 \times I \to X$ is connected to the vertices corresponding to the restrictions of ϕ to the two boundary components of $S^1 \times I$.

Observe that since X is finite, Ω is finite as well. Furthermore, an immersed circle in Ω corresponds to an immersed torus in X.

The map $\phi\colon C \to X$ corresponds to an immersion $\mathbb{R} \to \Omega$. Likewise the restriction of ϕ to $S^1 \times J$ corresponds to an immersion $J \to \Omega$ of a subinterval $J \subset \mathbb{R}$.

Applying Lemma 4.17 to the map $J \to \Omega$, we obtain an immersed circle in Ω which extends $J \to \Omega$ and which corresponds to an immersed torus in X which extends the map $S^1 \times J \to X$. ■

Lemma 4.18 *Let $\tau_1\colon (-\infty, 0] \to \Gamma_X$ and $\tau_2\colon [0,\infty) \to \Gamma_X$ be trajectories, and suppose their union is a bi-infinite trajectory $\tau\colon (-\infty,\infty) \to \Gamma_X$. If a path in Λ_τ has finitely many τ_1-extensions and finitely many τ_2-extensions, then it has finitely many τ-extensions.*

Proof If the number of τ_i-extensions is n_i then the number of τ-extensions is at most $n_1 n_2$. ■

The following summarizes what we have done in Section 4

Theorem 4.19 *Let X be a nonpositively curved $\mathcal{VH}$-complex. Suppose Θ_τ has compact components for each trajectory τ. Then each immersed flat plane in X is the strong limit of cylinders. Moreover, if X is compact, then it is the limit of tori.*

5 Malnormal and Cyclonormal Criteria

In this section we give verifiable conditions on a nonpositively curved directed $\mathcal{VH}$-complex X ensuring that the trajectory criterion of Theorem 4.19 applies and so flats are strong limits of cylinders. It will be useful to first recall the notion of a fiber product of maps of graphs [19].

5.1 Fiber Products of Graphs

Definition 5.1 (The Fiber Product $\Upsilon_1 \otimes \Upsilon_2$) For a pair of immersions, $\phi_1\colon \Upsilon_1 \to \Upsilon$ and $\phi_2\colon \Upsilon_2 \to \Upsilon$, we define a graph $\Upsilon_1 \otimes \Upsilon_2$ and projection maps $\pi_1\colon \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_1$ and $\pi_2\colon \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_2$. Consider the induced map

$$\phi_1 \times \phi_2\colon \Upsilon_1 \times \Upsilon_2 \to \Upsilon \times \Upsilon.$$

Consider the diagonal subspace $D \subset \Upsilon \times \Upsilon$ where $D = \{(x,x) \mid x \in \Upsilon\}$. We define

$$\Upsilon_1 \otimes \Upsilon_2 = (\phi_1 \times \phi_2)^{-1}(D).$$

Observe that $\Upsilon_1 \otimes \Upsilon_2$ consists of a set of 0-cells and diagonals of 2-cells of $\Upsilon_1 \times \Upsilon_2$ and is therefore a graph. For each i we define $\pi_i\colon \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_i$ as the restriction of $\pi_i\colon \Upsilon_1 \times \Upsilon_2 \to \Upsilon_i$. See Figure 4 for an illustration of $\Upsilon_1 \otimes \Upsilon_2$ in an easy case.

The triple $(\Upsilon_1 \otimes \Upsilon_2, \pi_1, \pi_2)$ is also characterized by the property that paths in $\Upsilon_1 \otimes \Upsilon_2$ are in one-to-one correspondence with pairs of lifts $\sigma_1 \to \Upsilon_1$ and $\sigma_2 \to \Upsilon_2$ of a path $\sigma \to \Upsilon$. Indeed, given two lifts $\sigma_1\colon I_n \to \Upsilon_1$ and $\sigma_2\colon I_n \to \Upsilon_2$ of the path

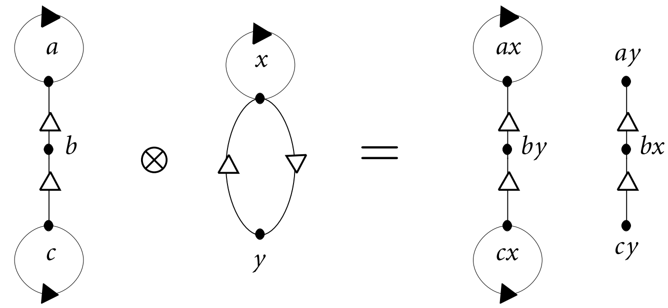

Figure 4: $\Upsilon_1 \otimes \Upsilon_2$. Let Υ denote a bouquet of two circles, whose edges are labeled by a white arrow and a black arrow. Let Υ_1 and Υ_2 denote the two labeled graphs on the left of the diagram. Note that their labelings induce maps $\Upsilon_1 \to \Upsilon$ and $\Upsilon_2 \to \Upsilon$. The graph $\Upsilon_1 \otimes \Upsilon_2$ is depicted on the right-hand side of the equal sign. Note that $\Upsilon_1 \otimes \Upsilon_2$ has two components. The vertices of Υ_1 are labeled by a, b, c, and the vertices of Υ_2 are labeled by x, y. Each vertex of $\Upsilon_1 \otimes \Upsilon_2$ is labeled by the pair of vertices of Υ_1 and Υ_2 to which it corresponds.

$\sigma\colon I_n \to \Upsilon$, there is a map $\sigma_1 \times \sigma_2\colon I_n \to \Upsilon_1 \times \Upsilon_2$ whose image is contained in the $\Upsilon_1 \otimes \Upsilon_2$ subspace.

If we let $\bar{\Upsilon}_i \to \Upsilon$ be the extension (by adding trees) of Υ_i to a covering space of Υ then there is an embedding $\Upsilon_1 \otimes \Upsilon_2 \hookrightarrow \bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$. The components of $\bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$ are the covers of Υ corresponding to intersections of conjugates of $\pi_1\Upsilon_1$ and $\pi_1\Upsilon_2$, and the subspace $\Upsilon_1 \otimes \Upsilon_2$ is a deformation retract of $\bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$.

When $\phi_1 = \phi_2$ we have the space $\Upsilon_1 \otimes \Upsilon_1$. The *diagonal* of $\Upsilon_1 \otimes \Upsilon_1$ is the component of $\Upsilon_1 \otimes \Upsilon_1$ which equals the diagonal of $\Upsilon_1 \times \Upsilon_1$.

We record the following connection between the fiber-product and malnormal and cyclonormal subgroups which are defined in Sections 5.2 and 5.3.

Remark 5.2 (Malnormal, Cyclonormal and $\otimes$) Given an immersion $\Upsilon_1 \to \Upsilon$, properties of the subgroup $\pi_1\Upsilon_1 \hookrightarrow \pi_1\Upsilon$ are reflected in the graph $\Upsilon_1 \otimes \Upsilon_1$.

(1) $\pi_1\Upsilon_1$ is malnormal $\Leftrightarrow$ each non-diagonal component of $\Upsilon_1 \otimes \Upsilon_1$ is a tree.
(2) $\pi_1\Upsilon_1$ is cyclonormal $\Leftrightarrow$ each non-diagonal component of $\Upsilon_1 \otimes \Upsilon_1$ has $\chi \geq 0$.

5.2 Malnormal $\mathcal{VH}$-Complexes

Definition 5.3 (Malnormal) A subgroup H of G is *malnormal* if the intersection $xHx^{-1} \bigcap H$ is trivial for all $x \in G - H$. More generally, H is *r-malnormal* if for any set of elements $g_1, \dots, g_{r+1}$ of elements representing distinct right cosets of H, the intersection $\bigcap_{i=1}^{r+1} g_i^{-1}Hg_i$ of the $r+1$ conjugates of H is the trivial group.

The only 0-malnormal subgroup of G is $\{1_G\}$. Being 1-malnormal is the same as malnormal, and 2-malnormal is the only other case for which we will have use.

Remark 5.4 (Geometric Meaning) Malnormality can be interpreted geometrically in terms of unique lifts of closed essential paths. Let Y be a based space and let $\widehat{Y}$ be a

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

based connected covering space. Then $\pi_1\widehat{Y}$ is r-malnormal if and only if every closed essential path in Y has at most r closed lifts to $\widehat{Y}$. This explains part (1) of Remark 5.2.

Definition 5.5 (r-Malnormal Map of Graphs) A map of graphs $\rho\colon \Lambda \to \Upsilon$ is *r-malnormal* if every nontrivial immersed cycle in Υ has at most r closed lifts to Λ. Note that we do not assume that Λ is connected. As usual we will use the term malnormal instead of 1-malnormal.

When $\rho\colon \Lambda \to \Upsilon$ is an immersion then ρ arises as the restriction of some covering map to a subspace. The two notions of r-malnormal are essentially equivalent under this correspondence.

We use the notation $|A^0|$ for the number of 0-cells in a complex A.

Lemma 5.6 (Unique Lifts of Long Paths) *Let $\rho\colon \Lambda \to \Upsilon$ be a malnormal immersion of graphs. Suppose Λ is finite so that $|(\Lambda \otimes \Lambda)^0| < \infty$. Let $\sigma \to \Upsilon$ be an immersed path. If $|\sigma| \geq |(\Lambda \otimes \Lambda)^0|$ then σ has at most one lift to Λ. Similarly, if ρ is r-malnormal then there are at most r lifts of any sufficiently long path.*

Proof Consider a pair of lifts of σ to Λ. This determines a lift of σ to a component of $\Lambda \otimes \Lambda$. But since $|\sigma| \geq |(\Lambda \otimes \Lambda)^0|$, some subpath of σ lifts to a closed immersed path in Λ. Since nondiagonal components of $\Lambda \otimes \Lambda$ are trees, the closed immersed path lies in the diagonal, and so the two lifts of σ are identical. ■

Definition 5.7 (Malnormal $\mathcal{VH}$-Complex) A graph of groups is *malnormal* if each edge group is a malnormal subgroup of each of its associated vertex groups. A directed nonpositively curved $\mathcal{VH}$-complex is *malnormal* if the associated graph of groups is malnormal, or equivalently, the attaching maps of each edge space are malnormal maps of graphs.

Definition 5.8 (Almost-Malnormal $\mathcal{VH}$-Complex) A directed nonpositively curved $\mathcal{VH}$-complex X is *almost-malnormal* if it has the following property:

(1) The attaching map of each edge space is 2-malnormal.
(2) For any pair of distinct attaching maps of edge spaces $\Lambda_e \to \Lambda_v$ and $\Lambda_f \to \Lambda_v$, their union $(\Lambda_e \cup \Lambda_f) \to \Lambda_v$ is 2-malnormal.

Remark 5.9 (2-Malnormal and Singular Leaves) The most common examples of 2-malnormal subgroups are subgroups of index 2. These turn up often because any singular leaf in a $\mathcal{VH}$-complex corresponds (in the associated directed $\mathcal{VH}$-complex obtained by subdividing) to a vertex space with a single edge space attached to it by a double cover.

Theorem 5.10 (k-Uniqueness for Malnormal $\mathcal{VH}$-Complexes) *Let X be a malnormal compact nonpositively curved directed $\mathcal{VH}$-complex. Then there is a constant k so that for each trajectory $\tau\colon \mathbb{R} \to \Gamma_X$, any immersed path of length $\geq k$ has at most one τ-extension.*

Proof By Lemma 4.18 it is sufficient to prove this for a trajectory $\tau\colon \mathbb{R}^+ \to \Gamma_X$.

Choose k to be larger than the number of vertices in any graph $\Lambda_e \otimes \Lambda_e$ obtained from $\Lambda_e \to \Lambda_v$, where e is an edge of Γ. Let $\sigma\colon I_K \to \Lambda_\tau$ be an immersed path of length $K \geq k$.

There is at most one way to extend σ because an immersed strip of width $\geq k$ which enters some vertex graph Λ_v via some edge space, has at most one way to exit via a second edge space. This is because exiting Λ_v to an edge space corresponds to lifting a path in Λ_v to Λ_e (where e is the edge corresponding to the edge space), but since the strip has width $\geq k$, the path we are lifting has length $\geq k$ and so our choice of k and Lemma 5.6 shows that there is at most one such lift and hence at most one extension. ■

Theorem 5.11 (k-Finiteness for Almost-Malnormal $\mathcal{VH}$-Complexes) *Let X be a compact nonpositively curved almost-malnormal directed $\mathcal{VH}$-complex. There is a constant k so that for each trajectory $\tau\colon \mathbb{R} \to \Gamma_X$, any immersed path in Λ_τ of length $\geq k$ has finitely many τ-extensions.*

Sketch of Proof The proof is similar to the proof of Theorem 5.10. We first prove the analogous statement for a trajectory $\tau\colon \mathbb{R}^+ \to \Gamma_X$, and then apply Lemma 4.18. Namely, we show that there are at most two τ-extensions. The idea is the same as above, except that at the initial extension there might be two choices, but for each successive elementary extension there is at most one extension by an immersion. ■

Theorem 5.12 (Limits in Almost-Malnormal) *Let X be a compact almost-malnormal nonpositively curved $\mathcal{VH}$-complex. Then every immersed flat in X is the strong limit of immersed cylinders and hence the limit of tori.*

Proof This follows by combining Theorem 5.11 with Lemma 4.16 and Lemma 4.17. ■

5.3 Cyclonormal Complexes

As motivation for the main theorem of this section we first prove the following special case:

Theorem 5.13 (Cylindrical $\mathcal{VH}$-complexes) *Let X be a $\mathcal{VH}$-complex, and assume that each edge space of X is a cylinder, that is, $\Lambda_e \cong S^1$ for each edge e of Γ_X. Then every immersed flat in X is the strong limit of immersed cylinders in X.*

Proof Think of $|\Lambda_e^0|$ as the perimeter of the edge space corresponding to e, and let $N = \mathrm{LCM}\{\Lambda_e^0\}$ be the least common multiple of the perimeters of edge spaces of X. Then each immersed flat plane $F \to X$ is actually periodic in the vertical direction, and the period divides N. To see this, note that it is true for each infinite vertical flat strip of F of horizontal width 1. Since F is vertically periodic, it is obviously the strong limit of immersed cylinders. ■

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

Definition 5.14 (Cyclonormal Subgroup) A subgroup H of G is *cyclonormal* if the intersection $H \cap gHg^{-1}$ is cyclic or trivial for any $g \in G - H$. As in Definition 5.3, we define *r-cyclonormal* similarly. An immersion of connected graphs $\phi\colon \Gamma \to \Upsilon$ is *cyclonormal* if the image of $\pi_1\Gamma$ is a cyclonormal subgroup of $\pi_1\Upsilon$.

Definition 5.15 (Cyclonormal $\mathcal{VH}$-Complex) A graph of groups is *cyclonormal* if each edge group is a cyclonormal subgroup of its associated vertex groups. Similarly, a directed $\mathcal{VH}$-complex is *cyclonormal* if the attaching maps of each edge space are cyclonormal.

A directed $\mathcal{VH}$-complex is *almost-cyclonormal* if for each $v \in \Gamma_X^0$ the following two conditions are satisfied:

(1) For any edge space attaching map $\Lambda_e \to \Lambda_v$, and for any three elements g_1, g_2, g_3 which represent distinct left cosets of $\pi_1\Lambda_e$ in $\pi_1\Lambda_v$, the triple intersection

$$g_1(\pi_1\Lambda_e)g_1^{-1} \cap g_2(\pi_1\Lambda_e)g_2^{-1} \cap g_3(\pi_1\Lambda_e)g_3^{-1}$$

is cyclic or trivial.

(2) For any distinct pair of attaching maps $\Lambda_e \to \Lambda_v$ and $\Lambda_f \to \Lambda_v$, and for any pair of elements g_1, g_2 representing distinct left cosets of $\pi_1\Lambda_e$ in $\pi_1\Lambda_v$, the triple intersection $g_1(\pi_1\Lambda_e)g_1^{-1} \cap g_2(\pi_1\Lambda_e)g_2^{-1} \cap \pi_1\Lambda_f$ is cyclic or trivial.

Note that the class of cyclonormal $\mathcal{VH}$-complexes includes both the malnormal $\mathcal{VH}$-complexes as well as the cylindrical $\mathcal{VH}$-complexes mentioned in Theorem 5.13.

Example 5.16 ($\mathcal{VH}$ and 3-Manifolds) Any nonpositively curved $\mathcal{VH}$-complex which embeds in a 3-manifold is almost-cyclonormal. In particular this holds for the Dehn complexes of prime alternating links (see [23]). The point is that the cartesian product of two bouquets of three edges cannot embed in a 3-manifold, because the complete bipartite graph $K(3,3)$ does not embed in the 2-sphere. The almost-cyclonormal property follows because if some intersection had rank ≥ 2 then there would be a $K(3,3)$ in the link of a 0-cell of the 2-complex.

Main Theorem 5.17 (Cyclonormal $\Rightarrow$ Strong Limits of Cylinders) *Any immersed flat plane in a compact nonpositively curved almost-cyclonormal directed $\mathcal{VH}$-complex is the strong limit of immersed cylinders.*

Proof Consider a flat $F \to X$ with trajectory τ. By Theorem 5.18 we see that components of Θ_τ are compact, and then by Lemma 4.16 we see that F is the strong limit of τ-cylinders. ■

Theorem 5.18 (Almost-Cyclonormal $\Rightarrow$ Θ_τ Has Compact Components) *Let X be a compact nonpositively curved almost-cyclonormal directed $\mathcal{VH}$-complex. Then there is a constant D such that for each trajectory τ, every connected component of the graph Θ_τ has diameter $\leq D$.*

Proof By Lemma 5.20 there is a number J such that for any trajectory τ, any path σ of length J in Λ_τ has finitely many τ-extensions $\hat{\sigma}$ which are injective paths in Θ_τ. Now applying Lemma 4.14, we obtain a bound D to the diameter of components of Θ_τ, and this bound is independent of τ. ■

Definition 5.19 (Parallel Paths) The paths $\lambda\colon I_n \to X$ and $\lambda'\colon I_n \to X$ are parallel if for some m there exists an immersion $\phi\colon I_n \times I_m \to X$ such that the restriction of ϕ to $I_n \times \{0\}$ is λ and the restriction of ϕ to $I_n \times \{m\}$ is λ'.

Main Lemma 5.20 (Finitely Many Extensions of Length J Path) *Let X be an [almost] cyclonormal nonpositively curved $\mathcal{VH}$-complex. There exists a constant J such that for any trajectory τ, any path λ of length J in Λ_τ has finitely many τ-extensions corresponding to injective paths in Θ_τ. In other words, there are finitely many injective paths $\hat{\lambda}$ in Θ_τ corresponding to λ.*

Proof We shall focus on the cylonormal case. The main additional feature in the almost-cyclonormal case is that instead of $\bar{S}, \bar{T}, \bar{U}$ arising from nondiagonal components of $\Lambda_e \otimes \Lambda_e$, they arise from nondiagonal components of $(\Lambda_\otimes \Lambda_e) \otimes \Lambda_f$. This somewhat complicates the statement of Lemma 5.26, but the proofs are similar.

The proof is broken up into the lemmas below. By Lemma 4.18, it is sufficient to prove the statement of Lemma 5.20 for a trajectory $\tau\colon \mathbb{R}^+ \to \Gamma_X$. Let J be chosen as in Lemma 5.26 and consider a path λ in Λ_τ which is of length J. We show that if there is more than one [two] τ-extension of λ, then no τ-extension of λ corresponds to an injective path in Θ.

If there is more than one [two] τ-extension, then Lemma 5.21 shows that there is a path λ' parallel to λ such that an extension of λ bifurcates at λ' and therefore λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ for some e. It follows from Lemma 5.26 that λ' has a subpath $\sigma' = \gamma^p$, which satisfies the conditions of Lemma 5.27. Therefore Lemma 5.27 shows that any extension of σ' is cylindrical. Furthermore, any path parallel to σ' also has the property that any extension of it is cylindrical.

It follows that λ has a subpath σ parallel to σ' with the property that any extension of σ is cylindrical. But that means that if $\hat{\sigma}$ is a path in Θ_τ corresponding to a τ-extension of σ, then $\hat{\sigma}$ is a closed path. It follows that $\hat{\lambda}$ is not injective. ■

Lemma 5.21 (A Bifurcation Subpath) *Let X be a nonpositively curved [almost] cyclonormal $\mathcal{VH}$-complex and let $\tau\colon \mathbb{R}^+ \to \Gamma_X$ be a trajectory. If λ is an immersed path in Λ_τ which has more than one [two] τ-extension, then there exists a path $\lambda' \to X_v$ which is parallel to λ, such that λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ where $\Lambda_e \to \Lambda_v$ is the attaching map of an edge space of X.*

Proof We prove this for a cyclonormal $\mathcal{VH}$-complex; the almost-cyclonormal case is similar. If there are two τ-extensions of λ then as in Figure 5, we can consider the vertical path λ' where they bifurcate. Then λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ where λ' is a path in Λ_v and $\Lambda_e \to \Lambda_v$ is where the bifurcation takes place. ■

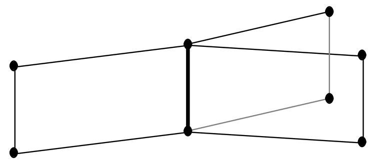

Figure 5: Bifurcation: The figure above is meant to suggest a pair of τ-extensions of the vertical path γ on the left. The two τ-extensions bifurcate at the path γ'. indicated by the bold vertical path.

Lemma 5.22 (Unique T in T^p) *Let T be a cyclically reduced word which is not a proper power. Consider the cyclic word T^p. Then T occurs as a subword of T^p in a unique way up to the cyclic action of $\mathbb{Z}_p$. In particular, T is a subword of T^2 in only the two obvious ways.*

Proof We think of the cyclic word as a labeled graph homeomorphic to S^1, such that the label of the closed cycle is T^p. Suppose that T is a subword of T^p in a nonstandard way, so that $T^p = XTY$ where $X \neq T^k$. Then because of the action of $\mathbb{Z}_p$ on S^1 which is generated by a rotation r_T of length $|T|$, we see that if we read the label of S^1 beginning at the endpoint of the path X above, then we also obtain T^p. It follows that the rotation r_X of length $|X|$ also gives rise to a symmetry of this labeling of S^1, and so, these two elements generate a cyclic subgroup $\langle r_T, r_X \rangle$, which contains $\langle r_T \rangle$ as a proper subgroup. But then the path T is a proper power of the fundamental domain of this action which is a contradiction. ∎

Lemma 5.23 (Locally T^3 Implies T^p) *Let T be a cyclically reduced word which is not a proper power. Let W be a cyclic word such that each vertex in W is in a* strict neighborhood *of T^3 in the sense that each vertex of the bi-infinite labeled edge-path W^∞ is contained in the middle third of a T^3 subpath. Then W is a cyclic permutation of T^p for some p.*

Proof By Lemma 5.22, T occurs in T^2 in only the standard ways, and so the T^3 "patches" must line up. ∎

Lemma 5.24 (Long Path in Cyclonormal is ST^rU) *Let $\Upsilon_1 \to \Upsilon$ be a cyclonormal immersion of graphs where Υ_1 is finite. There is a finite set of triples (S, T^d, U) of (possibly trivial) paths in Υ, such that the projection to Υ of any path in a non-diagonal component of $\Upsilon_1 \otimes \Upsilon_1$ can be expressed as the concatenation $S(T^d)^r U$ for some triple and some r. We also assume that T is not a proper power in the sense that it cannot be expressed as $(T')^m$ for some path T'.*

Proof To see that finitely many triples suffice, note that for a finite labeled graph Γ with $\chi(\Gamma) = 0$, we can let T^d correspond to a simple closed path in Γ, and let S and U correspond to injective paths ending and beginning on the simple cycle of Γ. ∎

Notation 5.25 ($\bar{S}, \bar{T}, \bar{U}$) We let $\bar{S}$, $\bar{T}$, and $\bar{U}$ denote the maximums of $|S|$, $|T|$, and $|U|$, over all triples (S, T^d, U) associated to a cyclonormal $\mathcal{VH}$-complex X.

Lemma 5.26 (A Long Highly Periodic Subpath) *Let X be a compact nonpositively curved cyclonormal directed $\mathcal{VH}$-complex. There is a number J such that for every length J immersed path λ' in X_V, if λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ where $\Lambda_e \to X_V$ is the attaching map of some edge space, then λ' has a nontrivial subpath $\sigma' = \gamma^p$ with γ not a proper power such that p and γ satisfy the following conditions:*

(1) $p \geq 4$.
(2) $p - 1 > 3(\bar{S} + \bar{T} + \bar{U})$ *and so* $|\gamma^{p-1}| > 3(\bar{S} + \bar{T} + \bar{U})$.
(3) $4|T^d|$ *divides* $|\sigma'|$ *for each* T^d *in a triple* (S, T^d, U).

Proof By Lemma 5.24, if λ' lifts to a non-diagonal component then λ' can be expressed as $S(T^d)^r U$ for some triple (S, T^d, U).

First we show that if $|\lambda'|$ is large then r is large. To see this, note that

$$|\lambda'| = |S(T^d)^r U| \leq \bar{S} + r \cdot \bar{T}^d + \bar{U}$$

and so

$$r \geq \frac{|\lambda'| - \bar{S} - \bar{U}}{\bar{T}^d}.$$

Now, choose J so that if $|\lambda'| = J$ then $r \geq q$ where q has the following three properties:

(1) $q \geq 4$.
(2) $q - 1 > 3(\bar{S} + \bar{T} + \bar{U})$.
(3) $4|T^d|$ divides q for each triple (S, T^d, U).

It follows that if λ' is a path which lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ then for some particular triple $(S_1, T_1^{d_1}, U_1)$ we have:

$$\lambda' = S_1 (T_1^{d_1})^r U_1$$

and thus λ' has a subpath σ' of the form $(T_1^{d_1})^q = \gamma^p$ where $\gamma = T_1$ is not a proper power and $p = qd_1$. It is easy to see that the three properties that q satisfies imply the corresponding properties for p. ∎

Lemma 5.27 (Extensions of Long Periodic Paths Are Cylindrical) *Let X be a compact nonpositively curved cyclonormal directed $\mathcal{VH}$-complex. Suppose $\sigma = \gamma^p$ is an immersed path in X_V where γ is closed but not a proper power and suppose the following hold:*

(1) $p \geq 4$.
(2) $|\sigma| \geq 4(\bar{S} + \bar{T} + \bar{U})$ *and so* $|\gamma^{p-1}| > 3\big(|S| + |T| + |U|\big)$.
(3) $4|T^d|$ *divides* $|\sigma|$ *for any triple* (S, T^d, U).

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

Then for any extension of $\sigma \to X$ to an immersion $\sigma \times \mathbb{R}^+ \to X$, the path $\{0\} \times \mathbb{R}^+ \to X$ equals $\{n\} \times \mathbb{R}^+ \to X$ where 0 and n are the endpoints of σ. Moreover, the map $\sigma \times \mathbb{R}^+ \to X$ factors through an immersion $S^1 \times \mathbb{R}^+ \to X$ where $S^1 \to X$ is the immersed subdivided circle corresponding to the closed path σ.

Note that $|\gamma^{p-1}| > 3\big(|S| + |T| + |U|\big)$ in property (2) since $|\gamma^{p-1}| \geq \frac{3}{4}|\sigma|$ by property (1).

Proof It is sufficient to show that this is true one edge space at a time. Namely, suppose that σ is a path in Λ_α and there is an edge space $\Lambda_\alpha \leftarrow \Lambda_e \to \Lambda_\beta$. We shall show that any lift of σ to Λ_e is closed and we show that the projection σ_1 of this lift to Λ_β has the same properties that σ satisfied: specifically, we will show that $\sigma_1 = (\gamma_1)^{p_1}$ where p_1 satisfies property (1), and that properties (2) and (3) hold because $|\sigma_1| = |\sigma|$.

If $\sigma = \gamma^p$ has a lift to Λ_e then so does γ^{p-1}. In case γ^{p-1} has a unique lift to Λ_e, then it is easy to see that the lift of γ corresponding to the initial segment of σ is closed. For if the initial γ subpath of γ^p does not lift to a closed path, then there would be two distinct lifts of γ^{p-1}. And therefore $\sigma = \gamma^p$ would lift to a closed path $\hat{\gamma}^p$, and so its projection σ_1 to Λ_β would be of the form $\gamma_1^{p_1}$ where p_1 is a multiple of p.

The case where γ^{p-1} does not have a unique lift is trickier and it is here that we use our three hypotheses. Since there are two distinct lifts, we see that there is a lift of γ^{p-1} to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$. Thus γ^{p-1} is of the form $S(T^d)^r U$ for some triple (S, T, U).

Every point in the path $S(T^d)^r U$, except for the initial ST and the terminal TU, is contained in the middle third of a T^3 subpath. Thus every point of γ^{p-1} except for the initial length $|ST|$ and terminal length $|TU|$ subpaths has a strict T^3 neighborhood. We will show that every point of the word γ^{p-1} has a strict T^3 neighborhood.

Since $|\gamma^{p-1}| > 3(|S| + |T| + |U|)$, every point in the middle third of γ^{p-1} has a strict T^3 neighborhood. By property (1), $p - 1 \geq 3$, and so each vertex in the closed path γ is identical to a vertex in the middle third of γ^{p-1}, and thus has a strict T^3 neighborhood. Consequently, by translating these strict T^3 neighborhoods around, we find that every vertex of the cyclic word $\sigma = \gamma^p$ has a strict T^3 neighborhood. It follows from Lemma 5.23 that σ is a power of T and therefore the lift of γ^{p-1} extends to a lift of σ.

Now, by property (3), $|\sigma|$ is a multiple of $|T^d|$ and so the lift of σ to $\Lambda_e \otimes \Lambda_e$ is closed and so the lift of σ to Λ_e is closed. Furthermore, the closed lift of σ to Λ_e projects to a closed path in Λ_β which is of period at least $|\sigma|/|T^d|$ which is ≥ 4 by hypothesis (3). ∎

5.4 Generalizations and a Conjecture

The results we have described have several straightforward generalizations. First of all, there are results similar to those of Section 5.4 for general nonpositively curved square complexes. They can be obtained from our point of view by taking the branched cover of Section 3 and using the correspondence between flats described there.

Secondly, Theorem 5.18 can be refined in a straightforward way by allowing a slightly less local approach. Instead of looking at the edge subgroups, one can look at the intersection of edge subgroups corresponding to finite trajectories and ask that these be cyclonormal. The idea is very similar to Theorem 5.18.

We now propose a conjectured generalization of the theorems of Section 5.3:

Conjecture 5.28 (Word-Hyperbolic Generalization) *Let Γ be a finite graph of word-hyperbolic groups where each of the edge groups is embedded by a quasi-isometry. Suppose that Γ is a cyclonormal graph of groups. Then $\pi_1\Gamma$ is word-hyperbolic if and only if it contains no $BS(n,m)$ subgroups. If Γ is a malnormal graph of groups, then $\pi_1\Gamma$ is word-hyperbolic if and only if it contains no $\mathbb{Z}^2$ subgroups.*

6 Connections to Subgroup Separability

Definition 6.1 (Subgroup Separability) Let H be a subgroup of the group G. Then G is *subgroup separable with respect to H*, or *H-separable*, if for each $g \in G - H$ there is a finite quotient $G \to \bar{G}$ such that $\bar{g} \notin \bar{H}$. A group is *subgroup separable* if it is H-separable for each finitely generated subgroup $H \subset G$. Note that G is $\{1_G\}$-separable if and only if G is residually finite.

Definition 6.2 (Clean Complex) A directed $\mathcal{VH}$-complex X is *clean* if all the attaching maps of edge spaces are embeddings.

Remark 6.3 (Clean $\Rightarrow$ Malnormal) Note that if X is a clean $\mathcal{VH}$-complex, then X is a malnormal $\mathcal{VH}$-complex. This is because an embedding of a subgraph in a graph is malnormal. Furthermore, if X is clean then X is nonpositively curved.

Let X be a connected directed $\mathcal{VH}$-complex. Then we say $\pi_1 X$ is $\pi_1\Lambda_e$-separable if $\pi_1(X,x)$ is $\pi_1\left(\Lambda_e \times \{1/2\}, x_e\right)$-separable for some point $x_e \in \Lambda_e \times \{1/2\} \subset X$. This is independent of the choice of basepoint since it depends only on the conjugacy class of subgroup represented by $\Lambda_e \to X$.

The following is proven in [23].

Theorem 6.4 (Separable $\Leftrightarrow$ Virtually Clean) *Let X be a compact nonpositively curved directed $\mathcal{VH}$-complex. Then $\pi_1 X$ is $\pi_1\Lambda_e$-separable for each e, if and only if X has a finite cover $\widehat{X}$ whose induced $\mathcal{VH}$-structure is clean.*

Lemma 6.5 (Virtual Limits) *Let X be a nonpositively curved space, and let $\widehat{X} \to X$ be a finite cover. Then for each of the properties below, $\widehat{X}$ has the property if and only if X has the property.*

(1) *If there is an immersed flat plane, then there is an immersed flat torus.*
(2) *Every immersed flat plane is the limit of tori.*
(3) *Every immersed flat plane is the strong limit of cylinders.*

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

Proof We prove this for property (2); the other proofs are similar. Consider an immersed flat plane $F \to X$. It lifts to an immersion $F \to \widehat{X}$. Now suppose that $F \to \widehat{X}$ is the limit of tori $T_n \to \widehat{X}$. Then $F \to X$ is the limit of their projections $T_n \to \widehat{X} \to X$.

In the other direction, consider an immersed flat plane $F \to \widehat{X}$. Its projection $F \to X$ is the limit of tori $T_n \to X$. Then each such torus has a finite cover $\widehat{T}_n$ which lifts to $\widehat{X}$. The sequence $\widehat{T}_n \to \widehat{X}$ limits to $\widehat{F}$. ■

Remark 6.6 (Edge-Separable ⇒ Properties) A consequence of Theorem 6.4 and Lemma 6.5 is that if $\pi_1 X$ is subgroup separable with respect to its edge subgroups then X inherits the following properties from its finite clean cover $\widehat{X}$.

(1) For each trajectory τ, the graph Θ_τ has compact components; in particular X admits no immersed anti-torus.
(2) If it admits an immersed flat plane, then it admits an immersed torus.
(3) Each immersed flat is the strong limit of immersed cylinders.
(4) Each immersed flat is the limit of immersed tori.

We turn Remark 6.6 to our advantage in [21]. The idea is that if a complex X fails to have one of the above properties, then it cannot have a clean finite cover. Therefore we can conclude (the more algebraic property) that $\pi_1(X)$ is not Λ_e-separable for some $e \in \Gamma_x$. This was the starting point for the solution of several long standing open questions, concerning residual properties of certain groups.

Question 6.7 (Malnormal or Cyclonormal ⇒ Virtually Clean) Let X be a compact nonpositively curved directed $\mathcal{VH}$-complex. Is X virtually clean provided that X is malnormal? Provided that X is almost-cyclonormal? Are there any obstructions to virtual cleanliness besides properties enjoyed by flats in clean complexes, which are inherited by complexes that they cover?

Remark 6.8 During the course of the last seven years since I raised this question, I have made some progress on it in [25] and [26]. In [25] I proved virtual cleanliness for the cylindrical complexes of Theorem 5.13, and in [26] I proved virtual cleanliness for square complexes that satisfy a slightly stronger hypothesis than malnormality.

Conjecture 6.9 (Subgroup Separable CAT(0) ⇒ Limits) *Let G act cocompactly and properly-discontinuously on a* CAT(0) *space $\widetilde{X}$. If G is subgroup-separable, then every flat in $\widetilde{X}$ is the limit of tori. In particular, if there is an immersed flat plane, then there is an immersed flat torus.*

I suspect that as in Theorem 6.4, a separability hypothesis relative to certain subgroups of G would suffice.

Remark 6.10 (General Square Complex) Some evidence to support Conjecture 6.9 is that Conjecture 6.9 is true in the case of general compact nonpositively curved square complexes. This is because, if $\pi_1 X$ is subgroup separable with respect to the subgroups corresponding to edge spaces of the double branched cover $\overline{X}$, then one can deduce that $\overline{X}$ is virtually clean, using arguments similar to Theorem 6.4.

In the search for even stronger results we ask:

Question 6.11 Which clean complexes have subgroup separable π_1?

Some partial results towards this are given in [25] where subgroup separability is proven for cylindrical $\mathcal{VH}$-complexes. Further results were proven in [23] under even more general conditions.

However, not every clean complex has subgroup separable π_1. Indeed, $F_2 \times F_2$ is not subgroup separable but it is the fundamental group of a clean complex. Furthermore, the non subgroup separable Burns-Karrass-Solitar example is π_1 of a virtually clean complex [5]. However, in joint work with Tim Hsu, we proved that all quasiconvex subgroups are separable [12], so perhaps one should only try to show that all quasiconvex subgroups are separable.

7 Limits of Tori in CSCs

In this final section, we show that any immersed flat plane in a compact CSC is the limit of periodic planes. The proof given is easily seen to generalize to higher dimensional CSCs. While the results from Section 4 stem from a scarcity of flat planes, Theorem 7.8, the main result of this section, derives from the density of flat planes.

Definition 7.1 A *complete bipartite graph* is a graph whose vertices are divided into two classes such that there is exactly one edge joining each pair of vertices from distinct classes. A *complete square complex* or *CSC* is a square complex X such that for each $x \in X^0$, the graph $\mathrm{Link}(x)$ is complete bipartite.

Proposition 7.2 *We record some properties of CSCs. See* [24, 21] *for the details.*

(1) *Every CSC X has a (single or) double cover $\widehat{X}$ which is a $\mathcal{VH}$-complex.*
(2) *X is a CSC if and only if its universal cover $\widetilde{X}$ is the product of two trees. In particular, if X is $\mathcal{VH}$ then $\widetilde{X} \cong \widetilde{V}_x \times \widetilde{H}_x$ for $x \in X^0$.*
(3) *Let X be a directed $\mathcal{VH}$-complex, then X is a CSC if and only if all attaching maps in the decomposition of X are covering maps.*

Example 7.3 There is a CSC X consisting of six squares described in [24] and [21] such that $\widetilde{X}$ contains an anti-torus. An *anti-torus* is a flat plane F in $\widetilde{X}$ such that each horizontal and vertical line in F is periodic, but F is not periodic.

I strongly believe that there are CSCs containing flat planes that are not strong limits of tori. By doubling, one sees that this is not true for every anti-torus, but perhaps it is true for typical examples.

Definition 7.4 (Combinatorial Convex Closure) Let $\widetilde{X}$ be a CAT(0) 2-complex and let $S \subset \widetilde{X}$ be a subspace. The *combinatorial convex hull* $\bar{S}$ of S is the intersection of all convex subcomplexes of $\widetilde{X}$ containing S.

Definition 7.5 (Diagonal line) Let $\widetilde{X}$ be the universal cover of a CSC X and let $D \to \widetilde{X}$ be an isometrically embedded line. We say D is a *diagonal line* if $D \nsubseteq X^1$. We will be interested in diagonal lines that pass through a 0-cell of $\widetilde{X}$, and that have a rational slope. We define a *diagonal line segment* similarly.

Lemma 7.6 (Convex Closure of Diagonal) *The combinatorial convex hull of a diagonal line segment in the universal cover $\widetilde{X}$ of a CSC X is a rectangle. The combinatorial convex hull of a diagonal line in $\widetilde{X}$ is a plane.*

Proof Consider the union of the set of squares intersecting D, and use the completeness of the links of $\widetilde{X}$ to extend outwards.

Alternatively, note that $\widetilde{X}$ is the direct product $V_{\widetilde{x}} \times H_{\widetilde{x}}$, and observe that the combinatorial convex hull of D is the direct product of its projection onto the two factors. ■

A diagonal line $D \hookrightarrow \widetilde{X}$ is *periodic* if the map $D \to X$ factors through a circle.

Lemma 7.7 (Convex Closure Is Periodic Plane) *Let X be a compact nonpositively curved square complex. Let $D \to \widetilde{X}$ be a periodic diagonal line. Then the combinatorial convex hull $\bar{D}$ of D is a periodic plane.*

Proof The cyclic group acting on D must stabilize $\bar{D}$. Consequently each line in $\bar{D}$ parallel to D is periodic and so $\bar{D}$ is cylindrical. Therefore, up to covering space transformation, there are finitely many lines in $\bar{D}$ parallel to D which intersect $\widetilde{X}^0$. Since $\bar{D}$ is the combinatorial convex hull of any such line, we see that $\bar{D}$ admits an additional symmetry by translating along a line orthogonal to D, and so $\bar{D}$ is (doubly) periodic. ■

Theorem 7.8 (Limit of Tori in CSC) *Let X be a compact CSC, and let $E \to X$ be an immersed flat plane. Then E is the limit of tori.*

Proof Without loss of generality, we assume that X is a $\mathcal{VH}$-complex and that X is vertically and horizontally directed. Choose an $n \times n$ region $S_n \subset E$. We will find a periodic diagonal line D whose combinatorial convex hull $\bar{D}$ contains S_n. By Lemma 7.7, $\bar{D}$ is a periodic plane extending S_n, and this proves the theorem.

Imagine that S_n is embedded in the Euclidean plane in the obvious way. Let *sw* denote the southwest square of S_n and let *ne* denote the northeast square of S_n.

Below we will produce a $\mathcal{VH}$-immersion $L \to X$ such that L is the union of S_n and two width 1 flat strips (see Figure 6). Let *sw* and *ne* denote the southwest and northeast squares of S_n. We will choose a horizontal strip from *ne* to a square q and a vertical strip from q to a square sw'. The square sw' will have the property that the maps $sw \to X$ and $sw' \to X$ are identical if we identify *sw* and sw' by translating one to the other through S_n and the flat strips.

Now, the diagonal geodesic segment γ in the rectangle $\bar{L}$ connecting the southwest corner of *sw* to the southwest corner of sw' determines a closed geodesic in X. We

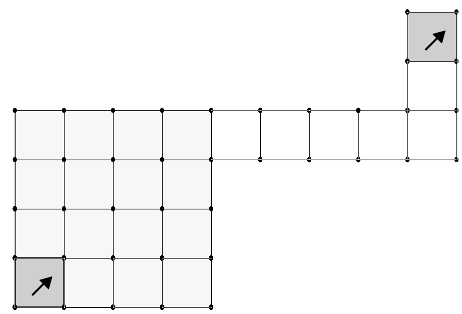

Figure 6: Illustrated above is an example of L. The shaded 4×4 square at the bottom left of L represents S_4. The square *sw* at the bottom left corner of S_n is labeled with an arrow (so that we can keep track of its orientation in L). The square *ne* is at the upper right-hand of S_n. There is a horizontal strip which travels to the right from *ne* and ends at q. There is then a vertical strip which begins at q and travels north to sw' which is oriented in the same way as *sw*.

consider L to be a subcomplex of $\widetilde{X}$, and we let $D = \gamma^\infty$ be the periodic extension of γ in $\widetilde{X}$. Observe that $\bar{L}$ and hence S_n is contained in $\bar{D}$ and we are done.

To construct L, consider the horizontal edge space C^h_{ne} containing *ne* and the vertical edge space C^v_{sw} containing *sw*. Beginning at *ne* we travel in C^h_{ne} until we arrive at a square q of the vertical edge space C^v_{sw} which is oriented horizontally in the same direction as *sw*. We then travel in C^v_{sw} until we pass through the square *sw* in such a way that it is oriented vertically in the same direction as our original *sw*. These two trips determine the horizontal and vertical strips.

To see that the orientations of *sw* and sw' can always be chosen consistently, we note that orientations can be reversed within edge spaces unless they are cylinders. However, if a vertical (respectively horizontal) edge space is a cylinder, then all vertical (respectively horizontal edge spaces) are cylinders, and so there would be no orientation problem to begin with. ■

8 Commuting Powers

A pair of elements g and h of a group G *nearly commute* if g^n and h^n commute for some $n > 0$. The problem of whether two elements nearly commute is a difficult one. Since there are groups with unsolvable word problem, it is easy to see that there are groups for which it is undecidable in general whether two elements nearly commute. Furthermore, it seems likely that there are groups with solvable word problem for which the problem of deciding if two elements nearly commute is recursively undecidable. In fact I suspect the following:

Conjecture 8.1 *There is a compact nonpositively curved square complex X such that the problem of deciding whether elements $g, h \in \pi_1 X$ nearly commute is recursively undecidable. In particular, there is a compact CSC with this property.*

https://doi.org/10.4153/CJM-2005-018-x Published online by Cambridge University Press

The connection between Conjecture 8.1 and the problem addressed in this paper is that for a nonpositively curved 2-complex X, the problem of whether two elements of $\pi_1 X$ nearly commute can be interpreted concretely mostly in terms of periodic tilings. Indeed, if g and h virtually commute then $\langle g^n, h^n \rangle$ is a free abelian subgroup of rank 1 or rank 2. The cyclic case is easy to determine by comparing axes for g and h, and the second case degenerates into the question of whether or not the flat plane containing their axes is periodic.

Let X be a nonpositively curved square complex. The angle between a pair of ends of 1-cells is 0 if they are identical, is $\pi/2$ if they meet along a corner of a square, and is π-otherwise. We define a *turning angle* between two consecutive 1-cells in a path in X^1 analogously. The following is implicit in the proof of Theorem 7.8:

Lemma 8.2 *Let X be a nonpositively curved square complex. Let σ be a closed geodesic path in X^1 with a $\pi/2$ turning angle. Then up to covering translation, the universal cover $\tilde{\sigma}$ of σ is contained in at most one periodic plane $\tilde{\sigma} \subset \widetilde{T} \subset \widetilde{X}$.*

Theorem 8.3 (Nearly Commute) *Let X be a compact nonpositively curved square complex which does not contain an immersed anti-torus. There exists an algorithm which determines if g and h nearly commute for any pair of elements $g, h \in \pi_1 X$.*

Sketch of Proof If g has a $\pi/2$ turning angle, then the centralizer of g is equal to the centralizer of g^n for each n, and this centralizer is a quasiconvex $\mathbb{Z}$ or $\mathbb{Z}^2$ subgroup and by applying Lemma 8.2, it is a simple matter to compute the centralizer and test if h lies in the centralizer. The same argument holds if h has a $\pi/2$ turning angle.

Now assume that neither g nor h has a $\pi/2$ turning angle. Suppose each angle between g and h is π, in the sense that the immersed paths homotopic to gh, gh^{-1}, hg, and hg^{-1} have no $\pi/2$ turning angles. Form the graph $g \vee h \to X^1$ by identifying basepoints, and fold until we obtain an immersion $Y \to X^1$ which is a local isometry and hence π_1-injective. Now g and h are immersed cycles in the graph Y, and hence they commute if and only if they are powers of the same element and this is readily verified by examining the corresponding words.

Finally, suppose that in the sense above, some angle between g and h is $\pi/2$. It is in this last case that we employ the hypothesis that X contains no immersed anti-torus. By possibly inverting one or both of g and h and by possibly simultaneously conjugating, we may assume that a $\pi/2$ angle occurs immediately in the positive direction.

We now algorithmically produce the convex hull of the quarter plane bounded by g^∞ and h^∞ by adding squares whenever two edges at one of their corners appear. Note that the nonpositive curvature implies that there is at most one way to add missing squares at each stage. Eventually, one either obtains the fundamental domain of a torus at some point or else the filling procedure terminates with a missing square at some corner. Indeed, since there is no immersed anti-torus, the filling procedure cannot continue indefinitely without producing a torus. ■

Applying the results of the previous section, we can give some positive results determining whether elements nearly commute for fundamental groups of certain nonpositively curved square 2-complexes. These results are based upon the following fact:

Lemma 8.4 (No Anti-Torus) *Let X be a nonpositively curved cyclonormal $\mathcal{VH}$-complex, then X does not admit an immersed anti-torus.*

Proof An immersed anti-torus determines an injective path of infinite length in Θ_τ, where τ is the trajectory of the anti-torus. ■

Corollary 8.5 *Let X be a compact nonpositively curved square complex. If its $\mathcal{VH}$-branched cover $\overline{X}$ is cyclonormal (after subdividing to make it directed), then it is decidable whether elements of $\pi_1 X$ nearly commute.*

Proof Combine Lemma 8.4, Theorem 8.3 and a variant of Theorem 3.1 showing that $\bar{X}$ has an anti-torus if and only if X has an anti-torus. ■

References

[1] R. Berger, *The undecidability of the domino problem.* Mem. Amer. Math. Soc. **66**(1966).
[2] M. Bestvina and M. Feighn, *Stable actions of groups on real trees.* Invent. Math. **121**(1995), 287–321.
[3] N. Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups.* J. London Math. Soc. **60**(1999), 461–480.
[4] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature.* Grundlehren Math. Wiss. 319, Springer-Verlag, Berlin, 1999.
[5] R. G. Burns, A. Karrass, and D. Solitar, *A note on groups with separable finitely generated subgroups.* Bull. Austral. Math. Soc. **36**(1987), 153–160.
[6] S. V. Buyalo, *Euclidean planes in three-dimensional manifolds of nonpositive curvature.* Mat. Zametki **43**(1988), 103–114, 155.
[7] P. Eberlein, *Geodesic flow in certain manifolds without conjugate points.* Trans. Amer. Math. Soc. **167**(1972), 151–170.
[8] D. Gromoll and J. A. Wolf, *Some relations between the metric structure and the algebraic structure of the fundamental group in manifolds of nonpositive curvature.* Bull. Amer. Math. Soc. **77**(1971), 545–552.
[9] M. Gromov, *Hyperbolic groups.* In: Essays in group theory, Math. Sci. Res. Inst. Publ. 8, Springer, New York, 1987, pp. 75–263.
[10] ———, *Asymptotic invariants of infinite groups.* In: Geometric group theory, Mathematical Society Lecture Note Series 182, Cambridge Univ. Press, Cambridge, 1993.
[11] B. Grünbaum and G. C. Shephard, *Tilings and patterns.* W. H. Freeman, New York, 1987.
[12] T. Hsu and D. T. Wise, *Separating quasiconvex subgroups of right-angled Artin groups.* Math. Z. **240**(2002), 521–548.
[13] M. Kapovich, *Flats in* 3*-manifolds.* Preprint, 1998.
[14] H. B. Lawson, Jr. and S. T. Yau, *Compact manifolds of nonpositive curvature.* J. Differential Geom. **7**(1972), 211–228.
[15] L. Mosher, *Geometry of cubulated* 3*-manifolds.* Topology **34**(1995), 789–814.
[16] E. Rips, *Subgroups of small cancellation groups.* Bull. London Math. Soc. **14**(1982), 45–47.
[17] V. Schroeder, *Codimension one tori in manifolds of nonpositive curvature.* Geom. Dedicata **33**(1990), 251–263.
[18] P. Scott and T. Wall, *Topological methods in group theory.* In: Homological group theory, London Math. Soc. Lecture Note Ser. 36, Cambridge Univ. Press, Cambridge, 1979, pp. 137–203.
[19] J. R. Stallings, *Topology of finite graphs.* Invent. Math. **71**(1983), 551–565.
[20] C. M. Weinbaum, *The word and conjugacy problems for the knot group of any tame, prime, alternating knot.* Proc. Amer. Math. Soc. **30**(1971), 22–26.K, 1971.
[21] D. T. Wise, *Complete squared complexes.* Preprint 2003.
[22] ———, *Cubulating small-cancellation groups.* Geom. Funct. Anal. **14**(2004), 150–214.
[23] ———, *Subgroup separability of the figure* 8 *knot group.* Preprint 1998.

[24] ______, *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*. PhD thesis, Princeton University, 1996.
[25] ______, *Subgroup separability of graphs of free groups with cyclic edge groups.* Q. J. Math. **51**(2000), 107–129.
[26] ______, *The residual finiteness of negatively curved polygons of finite groups.* Invent. Math. **149**(2002), 579–617.
[27] ______, *A flat plane that is not the limit of periodic flat planes.* Algebr. Geom. Topol. **3**(2003), 147–154 (electronic).

Deptartment of Mathematics and Statistics
McGill University
805 Sherbrooke Street West
Montreal, QC
H3A 2K6
e-mail: wise@math.mcgill.ca