NON-POSITIVELY CURVED SQUARED COMPLEXES, APERIODIC TILINGS, AND NON-RESIDUALLY FINITE GROUPS

Daniel T. Wise

Abstract. The first part of this dissertation investigates when every immersed flat plane in a non-positively curved squared complex is the limit of periodic flat planes. This is proven for certain subclasses of the class of non-positively curved squared complexes.

The second part studies the groups which act cocompactly and properly discontinuously on the direct product of two trees. This class of groups turns out to be much richer than one might expect.

An interplay is developed between the immersed flats in the complex and the subgroup separability properties of its fundamental group. This link between algebra and geometry leads to the solution of several open problems concerning the residual properties of automatic groups and small-cancellation groups. In particular an example is given of a compact non-positively curved squared complex whose fundamental group has no finite quotients. This example is also interesting because its universal cover is isomorphic to the direct product of two trees. Other examples include a C(4)-T(4) small-cancellation group which is not virtually torsion-free.

Typeset by $\mathcal{AMS}$-TEX

Acknowledgments

I am grateful to J. Sondow and S. Cappell for helping and encouraging me to become a mathematician. To the Mathematics Department of Princeton University for its continuous and generous support. To my advisor M. Bridson for his helpful guidance and encouragement and many other things. And to L. Mosher for reading this thesis, and for several inspiring conversations. To my fellow students and friends including L. Green, T. Hsu, and T. Tokieda, for their mathematical and non-mathematical friendship. To my Father and Mother for teaching me fractions and crayons, which eventually led to this work. To Lynn, and to David. And to Gidon for writing a computer program which was helpful for exploring CSCs. I am most grateful to my wife Yael for her love and support - this work could not have been done without her. And to my wonderful new baby, Talia Alexandra Wise, this work would have been completed more easily and quickly without her. To these and others I express my warmest thanks.

Contents

Introduction

The theory of small-cancellation groups was developed in the sixties, though its origins may be traced back much further. With the increased interest in geometric group theory in the last 10 years, the various classes of small-cancellation groups have taken their places as examples of automatic and, in some cases, word-hyperbolic groups. In addition, many small cancellation complexes may be naturally provided with metrics of non-positive curvature so their fundamental groups act properly-discontinuously and cocompactly on a CAT(0) space. This class of *non-positively curved groups* form another popular class of groups in geometric group theory which contains many small cancellation groups. Perhaps the simplest examples of 2-dimensional non-positively curved spaces are the *non-positively curved squared complexes.* A non-positively curved squared 2-complex is a 2-complex formed by gluing squares together, which has the property that the link of each vertex contains no cycles of length < 4.

This thesis presents a study of certain aspects of non-positively curved squared complexes. It is divided into two parts. Part I investigates a class of squared complexes called $\mathcal{V}H$-complexes, and in particular, studies the family of immersed flat planes in a non-positively curved $\mathcal{V}H$-complex. Part II is focused on a subclass of $\mathcal{V}H$-complexes called complete squared complexes, with the goal of providing examples of non-positively curved squared complexes with non-residually finite fundamental group. The relationship between these two parts is substantial though not immediately apparent from the proposed subject matter. Since the most important results of the thesis appear in part II, we will introduce part II first and mention its main results before discussing part I. The reader interested in reading part II first, can do so after reading section 1 of Part I.

It has been asked for some time whether all finitely presented small-cancellation groups are residually finite [S], [Wa], [P]. More recently, the same question has been asked about automatic groups [Ge], and groups which act properly discontinously and cocompactly on CAT(0) spaces. In part II of this thesis, I answer these questions by constructing examples of non-residually finite groups belonging to each of these classes. These examples are the fundamental groups of certain natural spaces defined below. The results contained in part II were circulated earlier in [W3].

The principal objects of study in part II are called *complete squared complexes* (or CSCs for short). A CSC is a squared complex with the property that the link of each vertex is isomorphic to a complete bipartite graph. Section 1 introduces CSCs and provides many of the basic theorems and definitions about CSCs that are used throughout part II. For instance, theorem 1.10 asserts that a squared complex is a CSC if and only if its universal cover is isomorphic to the direct product of two trees. It is easy to see that CSCs satisfy the combinatorial non-positive curvature condition for squared complexes, and therefore satisfy the more general C(4)-T(4) small cancellation condition. It thus follows (by [GS]) that the fundamental groups of compact CSCs are biautomatic. A simple proof of this result for the case of CSCs is sketched in 1.12.

The easiest examples of CSCs are the direct products $\Gamma_1 \times \Gamma_2$ of two graphs. Other examples include complexes with a finite cover which is isomorphic to $\Gamma_1 \times \Gamma_2$. In section 2, I describe a more exotic example called X which is the main example of the thesis. It is a CSC formed by gluing together 6 squares. Most of the examples given throughout the rest of the paper are constructed using X in some way, and many of their properties are inherited from X.

In section 3, a type of aperiodically tiled plane which I call an *anti-torus*, is described. It is shown that this aperiodically tiled plane appears in the universal cover of X. This anti-torus lies behind most of the main results of part II.

In section 4, the anti-torus is used to show that the covering spaces of X are of a very limited nature. In algebraic terms, $\pi_1(X)$ is not subgroup separable with respect to a certain subgroup of $\pi_1 X$ corresponding to the fundamental group of a convex subspace V of X contained in $X^{(1)}$. This subspace V, actually corresponds to one of the trees in the isomorphism between $\tilde{X}$ and the direct product of two trees. Theorem 4.9 asserts that $\pi_1 X$ is not virtually $F_n \times F_m$. This is interesting because $\pi_1 X$ acts properly discontinously and cocompactly on $\tilde{X}$, which is the direct product of two trees. Which is surprising, for one might have naturally guessed, that any group acting properly discontinously and cocompactly on the direct product of two trees, is virtually the direct product $F_n \times F_m$ of two free groups.

A group G is said to be residually finite if for any non-trivial element $g \in G$ there exists a finite quotient $G \mapsto Q$ so that g is not mapped to 1_Q. The question of whether or not small-cancellation groups are residually finite, has been noted, for instance, in [Sc] and [Pr]. The question of whether automatic groups are residually finite has been raised in [Ge].

The most important acheivement of part II of this thesis, is to provide examples of non-residually finite groups which are fundamental groups of compact non-positively curved squared complexes. Of course, these are CAT(0) groups, but they are also examples of compact C(4)-T(4) complexes. It is a result of Gersten and Short (see [GS] and [GS2]) that the fundamental groups of such complexes are biautomatic as well.

In section 5, some algebra is used to show that when X is doubled along V, the resulting fundamental group is not residually finite. But since V is a totally geodesic subspace which is also contained in the 1-skeleton of X, this double of X is itself a CSC. Then, an example E of a CSC is constructed with the property that $\pi_1 E$ has no finite quotients (nor any torsion quotients). This is quite surprising, since it means that E and $F_2 \times F_2$ are algebraically very different but geometrically very similar.

In section 6 various examples are given of biautomatic groups and CAT(0) groups which are not virtually torsion-free. Also, it is shown using similar ideas that the question of whether or not word-hyperbolic groups are virtually torsion free is equivalent to the question of whether or not word-hyperbolic groups are residually finite. At the end of section 6 there is a sketch of a method to attempt to produce examples of word-hyperbolic groups which are not residually finite.

In section 7, a compact non-positively curved squared complex is constructed which has a flat plane in its universal cover which is not the limit of periodic flat planes. This answers a question raised by Gromov in [G2]. Later in section 7 it is noted that the geometric angle and (algebraic) Gersten-Stallings angle between a certain pair of subgroups of $\pi_1 X$ disagree. As a consequence, an interesting example of a non-positively curved triangle of groups is constructed.

In section 8 it is proved that a compact CSC has the property that every immersed flat plane is the pointwise limit of a sequence of immersed periodic planes.

Tilings of the plane have been studied substantially (we refer the reader to [GrS] for an account of the extensive results). From a group theoretic point of view, tilings of the plane have recently become more important because of their connection to (non) word-hyperbolic groups.

In particular, an important result about CAT(0) spaces due to Bridson and Gromov, is that a compact non-positively curved space X has word hyperbolic fundamental group if and only if the universal cover $\tilde{X}$ does not contain an isometrically embedded flat plane. (See [Br], BH). In case X is a non-positively curved 2-complex, a flat plane embedded in $\tilde{X}$ has an induced tiling by the 2-cells of X.

A problem suggested by [Gr2] is to understand when the flat subspaces of the CAT(0) space $\tilde{X}$, are the limits of periodic flat subspaces. For instance this was proven to be the case for non-positively curved cubulated 3-manifolds in M. Section 4 of Part I of this thesis, is devoted to showing that every flat plane is the limit of periodic flat planes for a modest subclass of the class of non-positively curved squared complexes which is described below.

$\mathcal{V}H$-complexes, introduced in section 1 are the principal objects of study in part I. A $\mathcal{V}H$-complex is a squared complex for which there is a sensible notion of vertical and horizontal. Briefly, a $\mathcal{V}H$-complex is a squared complex X such that the edges of X are divided into horizontal edges and vertical edges, and such that the squares of X are attached to $X^{(1)}$ in a manner which respects this division. Namely, the sequence of edges in the attaching map of each square are alternately horizontal and vertical. Section 1 contains the basic definitions and theorems about $\mathcal{V}H$-complexes. The vertical-horizontal geometry of a $\mathcal{V}H$-complex yields a decomposition theorem which is expressed in 1.16 and 1.17, and this enables us to easily describe examples in a manner analogous to the description of a group as a graph of groups. In particular, theorem 1.18 asserts that the fundamental group of a non-positively curved $\mathcal{V}H$-complex splits as a graph of free groups with free edge groups.

The class of $\mathcal{V}H$-complexes is a fairly natural class of squared complexes, and in section 3 we observe that a well known class of squared complexes are $\mathcal{V}H$-complexes. Namely, the Dehn complex of any tame knot projection is a $\mathcal{V}H$-complex. We combine this with the result of Weinbaum's [W],[LS], to show that the Dehn complex of a prime alternating knot projection is non-positively curved. This results in theorem 2.9, which provides a decomposition of π_1 of any prime alternating knot complement as a graph of free groups.

Section 3 offers a construction which, given a squared complex X, provides a branched cover $\hat{X} \to X$, such that $\hat{X}$ is a $\mathcal{V}H$-complex. As is explained in section 3, this provides a way of reducing the study of immersed flats in a squared complex X to the study of immersed flats in its $\mathcal{V}H$-branched cover $\hat{X}$.

Section 4 is the main section of part I, and its main goal is to present certain classes of non-positively curved $\mathcal{V}H$-complexes which have the property that every immersed flat plane is the limit of immersed periodic flat planes.

Section 4 is divided into the following subsections:

Section 4.0 explains when an immersed flat plane is said to be the *limit of* *immersed periodic planes*. In addition, we define the *strong limit of cylinders* which is similar, and which arises more naturally in this work.

Section 4.1 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 4.2 presents some lemmas most of which depend on lemma 4.2.1 which states that an immersed line in a finite graph is the limit of immersed periodic lines. This 1-dimensional analogue is easy to prove.

Section 4.3 introduces *malnormal* $\mathcal{VH}$-complexes. The main theorem of 4.3, is theorem 4.3.12 which states that a malnormal $\mathcal{VH}$-complex has the property that every immersed flat plane is the limit of immersed periodic flat planes. Section 4.3 should be thought of as a warm up for section 4.5 below.

Section 4.4 presents the tensor product graph $\Upsilon_1 \otimes \Upsilon_2$ of two graphs $\Upsilon_1 \to \Upsilon$, $\Upsilon_2 \to \Upsilon$ which are mapped to a third graph. This plays an important (but not deep) role in expressing the proofs of section 4.5.

Section 4.5 contains the main results of section 4. *Cyclonormal* $\mathcal{VH}$*-complexes* are a generalization of Malnormal $\mathcal{VH}$-complexes. Main theorem 4.5.3 states that every flat plane in a non-positively curved cyclonormal complex is the limit of immersed periodic flat planes.

Section 4.6 contains a conjectured generalization of the theorems of section 4.5, which will hopefully bring the techniques of this thesis from the realm of $\mathcal{VH}$-complexes, to bear on (cylonormal or malnormal) graphs of word-hyperbolic groups.

Section 5 seeks to foreground some of the connections between parts I and II, by explaining the role that the immersion of a bad plane in X (e.g. one that is not approximable) plays as an obstruction to the subgroup separability of $\pi_1 X$ with respect to certain subgroups.

Part I: $\mathcal{V}H$-complexes and limits of tori

Some Terminology

Let $X \subset G$ be a subset of a group. Then we let $\langle X \rangle$ denote the subgroup generated by X, and we let $\langle\langle X \rangle\rangle$ denote the normal subgroup generated by X, and sometimes we write $\langle\langle X \rangle\rangle_G$ for emphasis. We let $[x, y] = x^{-1}y^{-1}xy$ and $x^y = y^{-1}xy$.

When X is a 2-complex, and $x \in X^{(0)}$ then we let Link_x denote the link of x in X. We let I denote the closed interval $[0, 1]$ with the obvious structure as a graph. More generally we let I_n denote $[0, n]$, which has the obvious structure as a graph with n edges. We let $R = (-\infty, \infty)$ and $R^+ = [0, \infty)$, again with the obvious structures as graphs.

$\mathcal{V}H$ Definitions

1.1 Definition: Squared Complex. A *squared 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 $X^{(1)}$.

1.2 Definition: $\mathcal{V}H$-complex. Let X be a squared complex. X is said to be a $\mathcal{V}H$-complex if the 1-cells of X are partitioned into two classes V and H called *vertical* and *horizontal* edges respectively, and the attaching map of each 2-cell of X alternates between edges in V and H. (See square (i) of figure 2.)

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*.

Given a 0-cell $x \in X^{(0)}$, we let V_x denote the component of V_X containing x. We define H_x similarly.

1.3 Definition: Bipartite Graph. A *bipartite* graph is a graph Γ such that

(1) $\Gamma^{(0)}$ is partitioned into two disjoint classes.
(2) Each 1-cell of Γ connects 0-cells from distinct classes.

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

This motivates the following definition:

1.5 Definition: Locally $\mathcal{V}H$. We will call a complex Y *locally* $\mathcal{V}H$, provided that for each $y \in Y^{(0)}$, there is a chosen bipartite structure on the graph Link_y. Note that as in 1.4, a $\mathcal{V}H$-complex has an induced local $\mathcal{V}H$-structure

Note that Y may admit several $\mathcal{V}H$-structures. This is because a graph which admits a bipartite structure, actually admits 2^c such structures (if we name the classes), where c is the number of connected components of the graph.

1.6 Example: Not Global $\mathcal{V}H$. The simplest example of a squared complex with a local $\mathcal{V}H$-structure which is not consistent with any (global) $\mathcal{V}H$-structure is a loop with one 0-cell and one edge, where the bipartite structure on the link of the 0-cell has one vertex in each class. Note that the underlying graph of this example does have two $\mathcal{V}H$-structures, but both are inconsistent with the local $\mathcal{V}H$-structure that we chose.

An example of a squared complex which has a local $\mathcal{V}H$-strucuture but no $\mathcal{V}H$-structure at all, is given by a square were we identify two of its edges as in figure 1 below.

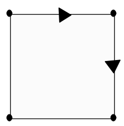

Figure 1: Local but not Global. The complex obtained by identifying two edges of a square according to the gluing pattern in the figure above, has a local $\mathcal{V}H$-structure, but no global $\mathcal{V}H$-structure.

While 1.6 shows that there are many examples of local $\mathcal{V}H$-complexes with no consistent (global) $\mathcal{V}H$-structure, we do have the following theorem which is analogous to the existence of orientable double covers of manifolds.

1.7 Theorem: $\mathcal{V}H$ Double Cover. *Let X be a squared complex which is locally $\mathcal{V}H$, then there is a double cover $\hat{X} \to X$ such that the induced local $\mathcal{V}H$-structure on $\hat{X}$, is consistent with a global $\mathcal{V}H$-structure.*

Proof. In this proof, we attempt to 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 while referring to 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 Link_y. So B_y is a 2 element set, and each of its elements is a (possibly empty) class of vertices of Link_y.

Let $x \in X^{(0)}$ be the base point of X, we will form a homomorphism

$$\phi : \pi_1(X, x) \to \text{Aut}(B_x)$$

Z_2.) We will show that the covering space $\hat{X}$ of X corresponding to ϕ has a $\mathcal{V}H$-structure. More precisely it has an induced local $\mathcal{V}H$-structure, which is consistent with a global $\mathcal{V}H$-structure.

A 1-cell e with initial vertex y and terminal vertex z induces a one-to-one correspondence $\phi_e : B_y \to B_z$. This is done in the obvious way: the end of e at y corresponds to a vertex in Link_y, and the end of e at z corresponds to a vertex in Link_z. We just choose $\phi_e : B_y \to B_z$ 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 $\phi_\sigma : B_{x_0} \to B_{x_n}$. 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 $B_{x_0} \to B_{x_1} \to B_{x_2} \ldots B_{x_{n-1}} \to B_{x_n}$. 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 $\psi : \pi_1(X, x) \to \text{Aut}(B_x)$.

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

Now each 1-cell of $\hat{X}$ 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. And the bipartite structure of each link implies that the condition on the attaching maps of squares of $\hat{X}$ is satisfied. Thus $\hat{X}$ has a global $\mathcal{V}H$-structure consistent with the local $\mathcal{V}H$-structure induced from X.

One can form a similar argument by looking at the universal cover $\tilde{X}$, and observing that it has a $\mathcal{V}H$-structure, One then considers the subgroup of $\mathrm{Aut}(\hat{X})$ which preserves the local $\mathcal{V}H$-structure, and takes the subgroup of index 1 or 2 which preserves the $\mathcal{V}H$-structure, and takes the intersection of $\pi_1(X)$ with that subgroup.

A similar argument works for higher dimensional versions of the class of $\mathcal{V}H$-complexes, but the corresponding homomorphism is to S_n and so may require a cover of higher degree than 2. $\square$

1.8 Definition: Locally CAT(0). The squared complex X is said to be *non-positively curved* or *locally CAT(0)* provided all reduced cycles in the link of each 0-cell of X have length at least 4. A squared complex satisfying this combinatorial link condition is actually locally CAT(0) in the sense of Gromov (see [Gr], [BH]).

1.9 Remark: Locally-$\mathcal{V}H$ and locally CAT(0). If X is a locally-$\mathcal{V}H$ squared complex, then X satisfies the combinatorial non-positive curvature condition if and only if there are no cycles of length 2 in the links of vertices 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.

1.10 Definition: Directed $\mathcal{V}H$-complex.

We view the attaching map of each square of a squared complex X, as a map from the boundary of the standard square $I \times I$ to $X^{(1)}$. We orient both horizontal edges of the standard square from left to right. This is illustrated in the second square of figure 2.

Let X be a $\mathcal{V}H$-complex. Suppose that H_X is a directed graph. We say that the $\mathcal{V}H$-structure on X is *(Horizontally) directed* provided that the attaching map of each square of X is orientation preserving on its horizontal edges. We define *vertically directed* similarly. We will use the term *directed* to mean *horizontally directed.*

1.11 Remark: Subdividing. There is little loss of generality in considering only directed $\mathcal{V}H$-complexes. This is because, given a $\mathcal{V}H$-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 $\mathcal{V}H$-complex. See the third square of figure 2. Similarly, the first barycentric subdivision of X is both vertically and horizontally oriented. See the fourth square of figure 2.

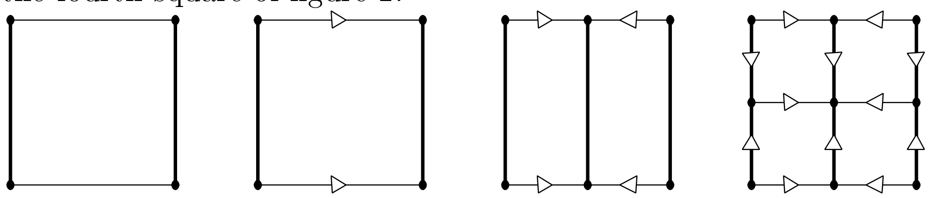

FIGURE 2: SQUARES. The squares in the figure above are meant to suggest (from left to right) a $\mathcal{V}H$-square, a horizontally directed $\mathcal{V}H$-square, a subdivided $\mathcal{V}H$-square (to obtain two horizontally directed subsquares), and the first barycentric subdivision of a $\mathcal{V}H$-square.

1: Decomposition Theorem

1.13 **Vertical Foliation and** V_x

We define a singular vertical foliation on the $\mathcal{V}H$-complex X. The 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. Given a point $x \in X$ (we do not assume that $x \in X^{(0)}$) we define V_x to be the smallest subset of X having the property that $x \in V_x$ and 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 discuss below from a slightly different point of view.

1.14 Remark: Singular Leaves and Directed $\mathcal{V}H$. When x is a point in the interior but not in the center of a horizontal edge, then for y very close to x, V_x is isomorphic to V_y by an isomorphism which is 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 $\mathcal{V}H$-structure. Note that if H_X can be oriented so that V_X is directed then it is easy to see that for any points x and y in the same horizontal edge V_x and V_y are isomorphic by a translation isomorphism.

1.15 Example: Möbius Strip. A moebius strip obtained by identifying the top and bottom horizontal edge of a square with a twisted identification map, has an obvious $\mathcal{V}H$-structure. The circle down the middle of the Möbius Strip is singular.

1.16 Definition: The Decomposition Graph $X \to \Gamma_X$. Given a directed $\mathcal{V}H$-complex X, we define a graph Γ_X and a map $\varrho : X \to \Gamma_X$. The 0-cells of Γ_X correspond to the connected components of V_X which we will refer to as *gates*. Note that each gate is V_x for some 0-cell $x \in X^{(0)}$. The 1-cells of Γ_X correspond to the connected components of $X - V_X$, which we call *corridors*.

It is easy to see that if x and y are in the same corridor, then V_x and V_y are isomorphic graphs and there is a natural isomophism between them. In fact, if x is a point in some corridor 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 squared complex which we will denote by $\bar{C}$. We will also call $\bar{C}$ a corridor.

For each corridor 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 corridor C. Note that these are combinatorial maps and in case X is non-positively curved then they are immersions (i.e. local injections) and so induce π_1-injections.

The components of V_X that the ends of $\bar{C}$ are mapped to correspond to the vertices of Γ_X that the edge of Γ_X corresponding to C is attached to.

1.17 Constructing X from Data. In order to further understand the map $\varrho : X \to \Gamma_X$ we show how X can be built up from the information encoded in Γ_X and the associated data.

Consider a graph Γ_X, and suppose that for each vertex $v \in \Gamma_X^{(0)}$ we have an associated graph Λ_v, and for each edge $e \in \text{Edges}(\Gamma_X)$ we have an associated corridor $C_e = \Lambda_e \times I$ where Λ_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 $\phi_{e0} : \Lambda_e \times 0 \to \Lambda_{e_0}$ and $\phi_{e1} : \Lambda_e \times 1 \to \Lambda_{e_1}$.

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

$$V_X = \{ \bigcup_{v \in \Gamma_X^{(0)}} \Lambda_v \}.$$

And form X by attaching C_e along its ends to V_X for each edge e of Γ_X, so that

$$X = \{ V_X \bigcup_{e \in \text{Edges}(\Gamma_X)} C_e \}.$$

1.18 Theorem: Graph of Free Groups. *Suppose that X is a non-positively curved $\mathcal{V}H$-complex. Then the map $\varrho : X \to \Gamma_X$ determines a splitting of $\pi_1 X$ as a graph of free groups. Where for each vertex $v \in \Gamma_X$, $\pi_1 v = \pi_1 \Lambda_v$ and for each edge e of Γ_X, $\pi_1 e = \pi_1 \Lambda_e$. For each edge e, attached to the vertices e_0 and e_1, the inclusion $\pi_1 \Lambda_e \to \pi_1 \Lambda_{e_i}$ is induced by the maps ϕ_{ei} described above.*

Proof. This is easy once one notes that the maps ϕ_{ei} are immersions and hence local-isometries and therefore they induce π_1-injections. □

1.19 Definition: Map of $\mathcal{V}H$-complexes. Let X and Y be $\mathcal{V}H$-complexes, and let $\phi : X \to Y$ be a cellular map. Then ϕ is said to me a $\mathcal{V}H$-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 $\mathcal{V}H$-map provided it maps vertical edges to vertical edges and horizontal edges to horizontal edges.

Note that if ϕ is a $\mathcal{V}H$-map, then it is easy to see that there is an induced map $\phi_\star : \Gamma_X \to \Gamma_Y$.

2: Dehn Complexes of Alternating Links

Consider a tame embedding of a link into three-dimensional space, $\rho : M^1 \to \mathbb{R}^3$. We will identify M^1 with $\rho(M^1)$. In this section we are interested in the fundamental group of the link complement, $\pi_1(S^3 - M^1)$. We may assume that M^1 is in general position relative to a projection plane $\mathcal{E}$, and let Π denote its projection onto that plane. (So there are only finitely many multiple points of the projection all of which are isolated double points.) It is convenient to imagine $\mathcal{E}$ as dividing M^1 into a top and bottom, with the overpasses lying above the plane and the underpasses lying below it, and M^1 lying very near Π. By the *regions* of Π we mean the connected components of $\mathcal{E} - \Pi$. By the *crossing points* of Π we mean the double points of the projection. By the *overpass* and *underpass* at a crossing point we mean the usual thing. Of course normally a link projection is drawn so that each overpass is drawn completely, wereas each underpass is missing a neighbourhood of the crossing point. We call this subgraph of Π, Π'.

Because of the following well known fact we are able to choose a checkerboard coloring of the tessellation of $\mathcal{E}$ induced by Π. In other words, we color the regions of $\mathcal{E} - \Pi$ so that they are alternately black and white, so that two regions which meet along an edge of Π are of different colors.

2.1 Lemma: The Checkerboard Coloring. *Consider an embedding of an even-valence graph Λ in the sphere, S^2. Then the regions $S^2 - \Lambda$ can be colored black and white in such a way that any two regions meeting at a common edge of λ are colored differently.*

proof. Consider the dual, Λ', of the induced 2-complex structure on S^2. Note that each 2-cell has an even number of 1-cells in its boundary. There is therefore a groupoid homomorphism $\pi_1(\Lambda') \to Z_2$ induced by sending each edge to a generator. We obtain the coloring by choosing some base vertex to be white, and letting the others be black or white depending on whether they are an odd or even distance away from the base vertex. This is well defined because of the homomorphism.

□

2.2 Definition: The Dehn Complex. The Dehn Complex, $\mathcal{D}_\Pi$ of Π is a squared complex consisting of the following cells:

(1) There are two 0-cells called top and bot .
(2) The 1-cells are in correspondence with the regions of Π. Each 1-cell is connected to both 0-cells. It is convenient to orient 1-cells corresponding to white regions so that they are directed from top to bot . We orient black 1-cells from bot to top .
(3) The 2-cells are in correspondence with the crossings of Π. Let the regions surrounding a crossing X be in counter-clockwise order, a, b, c, d and suppose the overpass of X divides a, d from b, c. Then the attaching map of the 2-cell corresponding to this crossing is the closed path: *abcd* using the orientations of the 1-cells that we have already specified. See figure 3.

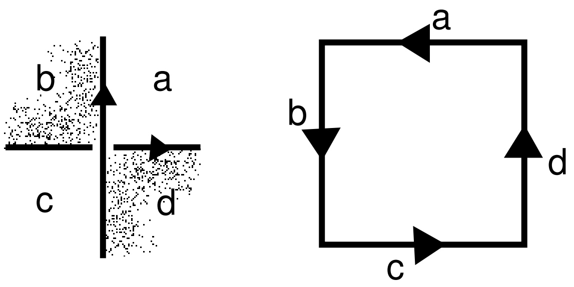

FIGURE 3: A 2-CELL IN THE DEHN COMPLEX. The Figure on the right indicates the attaching map for the 2-cell corresponding to the crossing on the left. Note that the 1-cells corresponding to the white regions on the left are oriented into the page, and the 1-cells corresponding to the black regions are oriented out of the page.

2.3 Lemma: Checkerboard $\rightarrow$ $\mathcal{VH}$-complex. *The Checkerboard Coloring of* Π *induces a corresponding coloring of the edges of* $\mathcal{D}_\Pi$. *Furthermore the checkerboard coloring of regions around a crossing implies that the boundary of each square is alternatingly vertical and horizontal and so the condition on attaching maps of 2-cells for* $\mathcal{VH}$*-complexes is satisfied. Therefore* $\mathcal{D}_\Pi$ *is a* $\mathcal{VH}$*-complex.*

2.4 Construction: Reading The Vertex Links.

There is an easy way to visualize the link of each vertex of $\mathcal{D}_\Pi$directly from the projection Π. See the illustration of figure 4. Assume that we have drawn Π in the usual way - so that a neighbourhood of each underpass at the crossing point is deleted. Then Link_{top} can be thought of as being embedded in $\mathcal{E} - \Pi'$ in the following way:

The 0-cells of the top link correspond to the regions, and imagine one such 0-cell in each region. The 1-cells of the the link correspond to the top corners of the 2-cells in $\mathcal{D}_\Pi$. They can therefore be thought of as curves between adjoining regions that pass through the gaps at the crossing. The bot link is clearly formed from Π in the same way only with all overcrossings switched to undercrossings and vice-versa.

Another way of saying this is that the 1-cells of the Link of top are in correspondence with the connected components of $\Pi - \Pi'$ and they are attached to the 0-cells according to the way these components border regions. Using this correspondence there is a natural way to identify the link with a graph embedded in $\mathcal{E} - \Pi'$.

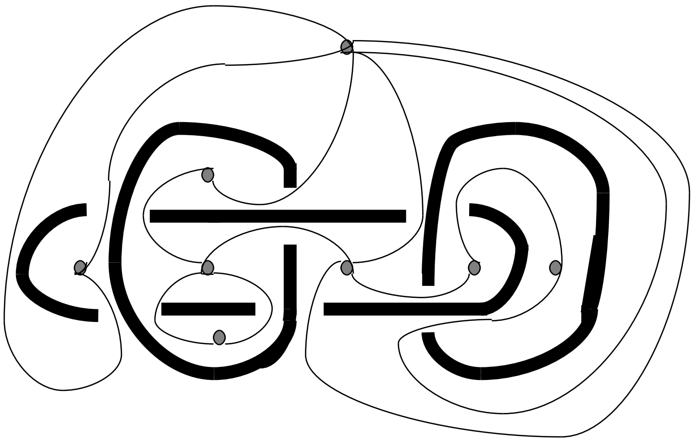

FIGURE 4: READING THE LINK. The graph with large nodes and thin curvy edges which is embedded in the complement of the link projection above, is isomorphic to $\mathrm{Link}_\top$, the top vertex of the Dehn complex of the projection.

2.5 Definition: Alternating Projection. We say that a planar projection of a Link is *alternating* if someone travelling along the projection of each circle in M^1 meets the crossing points alternatingly as overpasses and underpasses.

2.6 Definition: Prime Projection. We say that a projection of a link Π is *prime*, if for any simple closed curve γ in $\mathcal{E}$, transverse to Π, which intersects Π in two points (neither of which are crossing points), one of the two components of $\mathcal{E} - \gamma$ intersects Π in a simple arc.

It is easy to show that a prime alternating link has a prime alternating projection.

The following theorem is due to Weinbaum [We], see [LS] for a more detailed discussion.

2.7 Theorem: Prime Alternating Projections. *$\mathcal{D}_\Pi$ of a prime alternating projection satisfies the combinatorial non-positive curvature condition.*

proof. We showed above that $\mathcal{D}_\Pi$is a $\mathcal{V}H$-complex so it is sufficient to exclude cycles of length 2. Suppose there was a cycle of length 2 in one of the links of vertices of $\mathcal{D}_\Pi$, say in Link_{top}. Consider this cycle in the embedding of Link_{top} in $\mathcal{E} - \Pi'$. It is a simple closed curve γ, that intersects Π in two points. Since Π is prime, one side of the complement of γ intersects Π in a simple arc. But then this arc is an underpass on both ends, which contradicts the hypothesis that Π is alternating. $\square$

2.8 Corollary: Alternating knot is π_1 of a locally-CAT(0) complex. *The fundamental group of an alternating knot complement is the fundamental group of a non-positively curved 2-complex.*

proof. The previous lemma tells us that we can express an alternating link complement as a connected (knot) sum of finitely many link with alternating projections. But each such connected sum, is equivalent on the level of groups, to a free product with amalgamation over a cyclic subgroup. It is easy to see that the subgroup can be chosen to be represented by a local-isometry of a copy of S^1 into each complex.

Now the free product with amalgamation can be realized on a geometric level by gluing the ends of $S^1 \times I^1$ to each complex. □

One can obtain 2.8 by folding away the positive curvature in the Dehn Complex. Also, since the attaching maps are all π_1-injective, we can obtain a version of 2.9 below as well. The following theorem gives some partial results towards a conjecture of Neuwirth [Ne].

2.9 Theorem: Graph of Free Groups. *The fundamental group of a prime alternating link complement has a decomposition as*

(1) *either an HNN extension of a free group of rank r associating two subgroups of rank r.*
(2) *or a free product of two free groups of rank r amalgamating a subgroup of rank $2r - 1$ which is of index 2 in one of the factors.*

2.10 Remark: Similar. Note that these two types of splittings are geometrically and algebraically very similar.

proof. The splitting as a graph of free groups follows from theorem 1.18 because $\mathcal{D}_\Pi$is a $\mathcal{V}H$-complex. In order to obtain more precise information about the splitting we analyze the $\mathcal{V}H$-structure of $\mathcal{D}_\Pi$more carefully.

Let w, and b denote the number of white and black regions. Let c denote the number of crossings of Π.

Let x denote the point at the center of some black (horizontal) edge of $\mathcal{D}_\Pi$. Consider the graph V_x. It has an edge for each 2-cell of $\mathcal{D}_\Pi$and a vertex for each horizontal 1-cell of $\mathcal{D}_\Pi$. So $\chi(V_x) = b - c$.

Note that the vertical skeleton, $V_X = V_{top}$, is connected. It has as many edges as there are white regions, and has exactly 2 vertices. So $\chi(V_{top}) = 2 - w$.

If we think of Π as a graph whose vertices are the crossing points. we see that there are twice as many edges as crossing points. Examining the tessellation of S^2 we have,

$$\chi(S^2) = (c) - (2c) + (b + w)$$

$$2 = b + w - c$$

and we see that

$$\chi(V_{top}) = 2 - w = b - c = \chi(V_x).$$

Thus letting $r = -\chi(V_x) + 1 = w - 1$ we see that both V_{top} and V_x have π_1 which is free of rank r. In case, V_x is a singular leaf, we obtain an amalgamated free product, where the amalgamated subgroup is of index 2 in $\pi_1(V_x)$ and so has index $2r - 1$.

In case V_x is not singular, we obtain an HNN extension where the associated subgroups are isomorphic to $\pi_1 V_x$ and are hence of rank r. □

2.11 Remark: Even ⇔ HNN. It is easy to show that the HNN extension case above arises exactly when all the white regions have an even number of sides.

2.12 Presentations of torus knot groups. We leave as an exercise to the reader to show that two of the usual presentations of torus knot groups arise from the vertical and horizontal decompositions of the Dehn complex.

$\mathcal{V}H$ Branched Cover of Squared Complex

3: Reducing to the $\mathcal{V}H$-case

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

An *immersion* is a combinatorial map which is locally injective or equivalently, a map which induces an injection on the links of vertices. Below, by a *flat* $\mathcal{V}H$-complex, we mean a $\mathcal{V}H$-compex which is locally Euclidean, or equivalently, the link of each vertex is either a cycle of length 4 or a bouquet of 2 edges.

3.1 Theorem: $\mathcal{V}H$-Branched Covering

Let X be a squared complex. Then there exists a 1 or 2 fold branched cover $\rho : \hat{X} \to X$ such that

(1) $\hat{X}$ is a $\mathcal{V}H$-complex.
(2) Every immersion $\phi : F \to X$ from a flat $\mathcal{V}H$-complex F to X lifts to $\hat{X}$.
(3) If X is non-positively curved then every immersion $\phi : F \to \hat{X}$ of a flat $\mathcal{V}H$-complex F, projects to an immersion in X.
(4) If X is non-positively curved then up to covering space transformations, there is a one-to-one correspondence between immersed flat spaces in $\hat{X}$ and immersed flat $\mathcal{V}H$-complexes in X.

3.2 Remark: Using the one-to-one correspondence of theorem 3.1

Theorem 3.1 enables us to restrict our study of immersed flat planes in non-positively curved squared complexes to the immersions in non-positively curved $\mathcal{V}H$-complexes. The main benefit is that non-positively curved $\mathcal{V}H$-complexes can be easily described in terms of their decompositions (see 1.16). In particular, we can describe classes of non-positively curved $\mathcal{V}H$-complexes in terms of properties of their decompositions. For instance, in the next section we show that certain classes of $\mathcal{V}H$-complexes, have the property that any immersed flat plane is the limit of tori.

The theorem above shows that if $\hat{X}$ 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 $\hat{X}$ has the property (1), then we can conclude that X has property (1) as well. This is because any immersed flat plane in X lifts to an immersion in $\hat{X}$, if there is such an immersed flat plane, then there is an immersed flat torus $\phi : T^2 \to \hat{X}$ in $\hat{X}$.

But again because of the above theorem, the composition

$$T^2 \xrightarrow{\phi} \hat{X} \xrightarrow{\rho} X$$

is an immersion T^2 in X.

Likewise, if $\hat{X}$ has the property that every immersed flat plane E is the limit of a sequence of immersed tori, $\{T_n\}$, then $\phi(E)$ is the limit of the sequence of tori

$\{\phi(T_n)\}$. And so X has this property as well. One argues similarly for the strong limit of cylinders.

3.3 Definition: Branched Coverings

Given a (combinatorial) 2-complex X, we may 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 handles similarly 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 intersections with e. For each 1-cell e of X, let e' denote its barycenter. We now obtain a subspace $X_\star \subset X$ which is defined to be:

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

Note that while $X_\star$ is not a subcomplex of X, it is a subcomplex of an obvious subdivision of X. Also, it is easy to see that $\pi_1(X_\star)$ is free.

3.4 Lemma:

$X - X^{(0)}$ 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

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

Finally one notes that for each $v \in X^{(0)}$, $\overline{\text{Star}(v)} - v$ deformation retracts onto $\text{Link}(v)$. $\square$

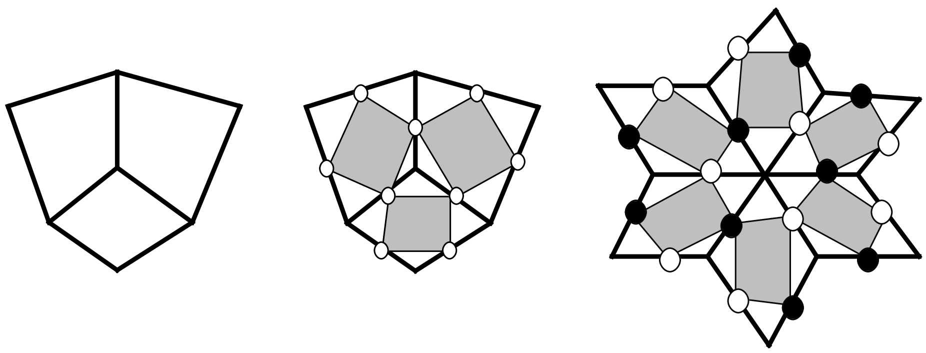

FIGURE 5. The figure on the left is a squared complex X formed by gluing three squares together. The figure in the center is a subdivided

copy of X, containing $X_\star$ as a subspace. The figure on the right depicts the branched cover $\hat{X}$ of X. Note that it contains a double cover of $X_\star$, and that the coloring of the vertices of $X_\star$ induces a $\mathcal{V}H$-structure on $\hat{X}$

Proof of Theorem 3.1.1. We now use the (squared complex) subspace $X_\star \subset X$ defined above, to choose the desired branched cover.

Notice that $X_\star$ may not be connected. In the argument that follows, we will assume that it is connected. The case where it 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 $\mathcal{V}H$-structures for the space $\hat{X}$ obtained below is 2^c where c is the number of connected components $X_\star$.

The homomorphism $\pi_1(X_\star^{(1)}) \to Z_2$ induced by amalgamating the vertices and then sending all edges to generators of 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 $\hat{X}_\star \to X_\star$ and therefore a branched covering space $\hat{X} \to X$

This covering space has (and may be characterized by being the smallest with) the property that all cycles in $\hat{X}_\star^{(1)}$ have even length. It follows that choosing one vertex to be black (in each component of $\hat{X}_\star$), the homomorphism induces a coloring of the vertices of $\hat{X}_\star$ as alternatingly black and white. The 1-cells of $\hat{X}$ inherit this coloring, and furthermore the edges of each square 2-cell of $\hat{X}$ are colored alternatingly black and white. This corresponds to the square polygon inside each 2-cell of $\hat{X}$ whose vertices are colored alternatingly. Thus we have obtained a $\mathcal{V}H$-structure on the branched cover $\hat{X}$ of X. $\square$

proof of 3.1.2. Consider an immersed flat plane $\phi : E \to X$. There is an induced map $\phi_\star : E_\star \to X_\star$. All closed loops in $E_\star$ have even length. It follows that $E_\star$ lifts to $\hat{X}_\star$ and therefore E lifts to $\hat{X}$. $\square$

proof of 3.1.3. Consider an immersed flat torus in $\hat{X}$. The simple cycles in the links of its vertices all have length 4. Now, for any $v \in \hat{X}$, the induced map $\text{Link}(v) \to \text{Link}(\rho(v))$ is a covering map. Consequently, a reduced cycle in $\text{Link}(v)$ is sent to a reduced cycle in $\text{Link}(\rho(v))$. It follows that the projection of the torus is an immersion, since there are no cycles of length < 4 in $\text{Link}(\rho(v))$. $\square$

3.5 Remark: Cycles of length 3. It is interesting to note that the last two conclusions actually only use that every simple cycle in the links is of length ≥ 3.

3.6 Remark: Atoroidal branched covers. It is very tempting to try and form a compact non-positively 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 of in such a way that some of the flat planes of this complexes will lift, but none of the periodic flat planes will lift. This seems tricky 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: Finding Tori

4.0.1: Definition: Limit of Tori. Let X be a non-positively curved space. Let $\phi : E \to X$ be a local isometry of a flat plane to X. We say that E is the *limit of tori*, or more precisely, we say that $\phi : E \to X$ is the limit of tori, if there is a sequence $\phi_n : E \to X$ of local isometries of periodic planes, such that $lim_{n\to\infty}\phi_n = \phi$. That is, for each $x \in E$, $lim_{n\to\infty}\phi_n(x) = \phi(x)$. Note that we say that a map $E \to X$ is *periodic* if it factors through a torus $E \to T^2 \to X$.

4.0.2: Remark: In dimension 2. Suppose that X is of dimension 2, in particular suppose that X is a non-positively curved 2-complex. Then an immersed flat plane $E \to X$ is the limit of tori if and only if for every compact neighbourhood $B \subset E$, the restriction $\phi : B \to X$ of ϕ to B, can be extended to an immersed flat torus $T^2 \to X$. In particular if X is a non-positively curved squared complex, then $\phi : E \to X$ is the limit of tori, if and only if the restriction of ϕ to every $n \times n$ rectangle contained in E can be extended to a flat torus.

4.0.3: Definition: Strong limit of cylinders. Let X be a non-positively curved 2-complex. Given an immersed flat plane $\phi : E \to X$ we say that it is the *strong limit of cylinders* provided that E is the increasing union of flat strips (e.g. $\{[-n, n] \times R\}$) each of which can be extended to a map of a cylinder $S^1 \times R \to X$. This is a stronger property for E to satisfy than being the limit of tori.

To put some of the results that will be obtained in the remainder of section 4 into perspective, we give a simple (but very strict) condition on a space which guarantees that all flats are periodic. Some similar results were obtained independently by P. Brown [Bro].

4.0.4 Proposition: Isolated $\Rightarrow$ Periodic. *Suppose that the group G acts cocompactly and properly discontinously on the CAT(0) space $\tilde{X}$. Suppose that there is a constant K such that $\tilde{X}$ has the property that at most finitely many flat planes pass through any ball of radius K in $\tilde{X}$. Then every flat in X is periodic.*

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

The condition of 4.0.4 arises quite naturally. For example, it is satisfied by any compact C(6) small-cancellation property, and in particular, it is satisfied by any compact non-positively curved 2-complex whose 2-cells are regular hexagons. To see this note that a plane is determined uniquely by a pair of hexagons which meet at an edge.

While 4.0.4 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 II.8 were it is the abundance of planes which provides the result.

The plan is to investigate the properties of immersed flat planes in non-positively curved $\mathcal{V}H$-complexes, by analyzing the entire set of immersed flat strips.

4.1.1 Definition: Flat Strips. By a *bi-infinite flat strip* we mean either $[0,n]\times R$ or $[0,\infty)\times R$ or $(-\infty,\infty)\times R$. Similarly by an *infinite flat strip* we mean either $[0,n]\times R^+$ or $[0,\infty)\times R^+$ or $(-\infty,\infty)\times R^+$. Often we will just refer to any of these possibilities as a flat strip, and either it will be clear from the context which (or possibly all) of these possibilities is meant.

Any integer length (or infinite) interval $\subset(\infty,\infty)$ has an obvious structure as a graph, where each edge is a length 1 subinterval. Thus any flat strip has the structure of a squared complex since it is the product of two graphs. Furthermore any flat strip has the structure of a $\mathcal{V}H$-complex, where we regard the edges parallel to the R or R^+ factor as being horizontal.

4.1.2 Definition: Trajectory. Given a $\mathcal{V}H$-immersion $\phi:F\to X$ of a flat strip to a $\mathcal{V}H$-complex X, we call the induced map $\phi_\star:\Gamma_F\to\Gamma_X$ the *trajectory* of F in X.

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

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

4.1.3 Remark: Boundary Point. The role of a trajectory can be replaced with a point in the boundary. (This will become clear below.) In principle this would provide a more sensitive approach (which also avoids the notion of $\mathcal{V}H$-complex). However, the information about immersed flat strips representing the same boundary point is contained in the connected components of the graph Θ_τ defined below, and so there is little difference for our purposes.

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

4.1.6 Definition: parallel paths. We say that two paths, $\sigma:I_n\to X$ and $\lambda:I_n\to X$, are parallel provided that they bound an immersed rectangle. More precisely, for some m, there exists an immersion $\phi:I_n\times I_m\to X$ such that the restriction of ϕ to $I_n\times\{0\}$ is σ and the restriction of ϕ to $I_n\times\{m\}$ is λ.

4.1.7 Definition: 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 $\alpha:I\times R\to X$ of Θ_τ is attached to the vertices of Θ_τ which are the restrictions of α to $\{0\}\times R$ and $\{1\}\times R$.

Note that 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 refer to a path in Θ_τ which corresponds to a τ-extension of the path σ in Λ_τ.

Note that in case X is non-positively curved, there is a natural immersion $\Theta_\tau \to \Lambda_\tau$ 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. We also note that this immersion induces a labelling of Θ_τ by the labels occurring in Λ_τ.

4.1.8 Remark: Θ_τ is a Large Graph. In general Θ_τ is a very large graph with uncountably many components. However if the connected components of Θ_τ are finite, then we can make assertions about τ-flats. For instance, in 4.1.9 we show that if the components are finite then any τ-flat is the strong limit of τ-cylinders. And in 4.1.12 we show that if there is a uniform bound on diameters of components then L_τ is a regular language.

4.1.9 Lemma: Finitely many extensions $\Rightarrow$ Compact Components. *Suppose that Λ_τ is a finite graph (in particular, this is true when X is finite), and let τ be a trajectory. Suppose that 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. Note that Θ_τ is locally finite. This is because, since X is non-positively curved, the map $\Theta_\tau \to \Lambda_\tau$ is an immersion, and since X is finite, the graph Λ_τ is finite (and hence locally finite). It follows that 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 immersed paths in Θ_τ of length $\geq D = J(HL+1)$. Consider a path $\hat{\sigma}$ of length $J(HL+1)$ in Θ_τ. We divide $\hat{\sigma}$ into $(HL+1)$ subpaths each of length J. But there are at most HL different paths of length J in Θ_τ, therefore by the pigeon-hole principle, 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. $\square$

4.1.10 Lemma: Finitely many injective τ-extensions $\Rightarrow$ compact components.
Suppose that there are numbers J and L such that each path σ of length J, has at most L distinct τ-extensions which correspond to injective paths in Θ_τ, then there is a number D such that each component of Θ_τ has diameter $\leq D$.

Note that Θ_τ can have uncountably many non-tree components even if (X, τ) satisfies the hypothesis of lemma 4.1.10. Furthermore, note that this is not the case if we impose the slightly more stringent hypothesis of lemma 4.1.9.

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. $\square$

4.1.11 Definition: The τ-Language. By the language L_τ of τ-extendible paths, we mean the set of all finite immersed τ-extendible paths in Λ_τ. If we label and orient each edge of Λ_τ (as in 4.1.7) then we obtain a language L_τ of words corresponding to the τ-extendible paths.

4.1.12 Lemma: Uniformly Bounded Diameters $\Rightarrow$ L_τ is regular. *If there is a constant D such that each component of Θ_τ has diameter $\leq D$ then L_τ is a regular language.*

Proof. The graph Θ_τ is labelled by the immersion $\Theta_\tau \to \Lambda_\tau$ mentioned in 4.1.7. If we consider two components of Θ_τ to be equivalent if they are labelled the same way, then there are only finitely many components up to equivalence and we obtain a finite labelled graph. Now letting each vertex be a start and accept state we obtain a finite state automata accepting the language L_τ of τ-extendible paths. $\square$

4.1.13 Remark: Language of lines appearing in flats is regular. If there is a universal bound on the diameters of the components of Θ_τ for each trajectory τ, then arguments similar to the finiteness argument above show that:

(1) The language of vertical words appearing in some flat strip is regular.
(2) The language of vertical words appearing in some flat plane is regular.
(3) The language of vertical words extending to a flat cylinder is regular.

4.2.1 Lemma: Limit of circles. *Consider a graph Θ with finitely many vertices. Suppose that there is an immersion $\phi : R \to \Theta$ where $R = (-\infty, \infty)$. Then $\phi : R \to \Theta$ is the limit of immersed circles. More precisely, for each finite connected subinterval $J = [a, b] \subset R$ there is a subdivision of S^1, and an immersion $\psi : S^1 \to \Theta$ and an embedding $J \subset S^1$ so that ψ is an extension of ϕ restricted to J.*

Proof. This is easy. 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, it is easy to see that 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 therefore a limit of the corresponding circles as well. $\square$

4.2.2 Lemma: Compact Components $\Rightarrow$ τ-flats are approximable by cylinders. *Suppose that each component of Θ_τ is compact, then every τ-flat is the strong limit of τ-cylinders.*

proof. A τ-flat is just a bi-infinite path $\hat{\omega}$ in Θ_τ. But the component of Θ_τ containing $\hat{\omega}$ is compact, therefore by lemma 4.2.2, 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 theorem shows that if an immersed plane is the strong limit of cylinders then it is the limit of tori.

4.2.3 Theorem: Cylinder is the limit of tori. *Suppose that $\phi : C \to X$ is an immersion of a cylinder to a $\mathcal{V}H$-complex, where $C = S^1 \times R$, for some subdivision of S^1. Then for any finite interval $J \subset R$, the restriction of ϕ to $S^1 \times J$ extends to an immersion of a torus.*

Proof. We form a graph Θ as follows: Fix a combinatorial structure on S^1 as above. The vertices of Θ correspond to the distinct $\mathcal{V}H$-immersions $S^1 \to X$ where we declare S^1 to be vertical. The edges of Θ correspond to the distinct $\mathcal{V}H$-immersions $S^1 \times I \to X$, where S^1 is vertical and I is horizontal. Each edge $\phi : S^1 \times I \to X$ 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, Θ is finite as well. Furthermore, an immersion of a circle in Θ corresponds to an immersion of a torus in X.

Now, the map $\phi : C \to X$ corresponds to an immersion $R \to \Theta$. Likewise the restriction of ϕ to $S^1 \times J$, corresponds to an immersion $J \to \Theta$ of a subinterval $J \subset (-\infty, \infty) = R$.

Applying lemma 4.2.3 to the map $J \to \Theta$ we obtain an immersion of a circle in Θ which extends $J \to \Theta$, and which corresponds to an immersion of a torus in X which extends the map $S^1 \times J \to X$. $\square$

4.2.4 Lemma: Finite $\Rightarrow$ Finite. *Let $\tau_1 : (-\infty, 0] \to \Gamma_X$ and $\tau_2 : [0, \infty) \to \Gamma_X$ be trajectories, and suppose that they induce a bi-infinite trajectory $\tau : (-\infty, \infty) \to \Gamma_X$. If any length K path σ 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 of σ is n_i, then it is easy to see that the number of τ-extensions is $\leq n_1 n_2$. $\square$

4.3.1 Definition: Malnormal. We say that H is a malnormal subgroup of G, provided that for $x \in G - H$, $xHx^{-1} \bigcap H = \{1_G\}$, the trivial subgroup.

4.3.2 Remark: Geometric meaning. That H is malnormal in G can be interpreted geometrically as a statement about the uniqueness of lifts of closed loops.

If $\hat{G}$ denotes a complex with $\pi_1(\hat{G}) = G$ and $\hat{H}$ denotes the cover corresponding to $H \hookrightarrow G$ then H is malnormal in G exactly if any closed essential loop in $\hat{G}$ has at most one closed lift in $\hat{H}$.

4.3.3 Definition: r-malnormal. We say H is an r-malnormal subgroup if for any set of elements $g_1, \dots, g_{r+1}$ of elements representing distinct right cosets of H, the intersection $\cap_{i=1}^{r+1} g_i^{-1} H g_i$ of the $r+1$ conjugates of H is the trivial group. Note that the only 0-malnormal subgroup of G is $\{1_G\}$. 1-malnormal is the same as malnormal. 2-malnormal is the only other case that we will have use for.

4.3.4 Remark: geometric definition. As above, the geometric meaning of r-malnormal is that any essential closed path has at most r distinct closed lifts.

4.3.5 Definition: r-Malnormal map of Graphs. We say that a map of graphs $\rho : \Lambda \to \Upsilon$ is r-malnormal if every non-trivial reduced cycle in Υ has at most r closed lifts to Λ. We will just say malnormal instead of 1-malnormal. Note that we do not assume that Λ is connected

From now on we assume that all maps between graphs are immersions, in particular all paths are reduced.

When ρ is an immersion then it arises as the restriction of some covering map to a subspace. That is, it comes from group theory anyway and the two notions of r-malnormal are roughly equivalent under this correspondence. The following lemma requires the definitions and terminology of section 4.4

4.3.6 Lemma: unique lifts of long paths. *Let $\rho : \Lambda \to \Upsilon$ be a malnormal map of graphs. Suppose Λ is finite so that $\|(\Lambda \otimes \Lambda)^{(0)}\| < \infty$. Let σ be a reduced path in Υ. 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 two 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 reduced closed path in Λ. But the only component of $\Lambda \otimes \Lambda$ which isn't a tree is the diagonal component. Therefore the two lifts of σ are identical. $\square$

We are now able to give easily verifiable conditions on a $\mathcal{V}H$-complex X which insure that X has the property that there exists k such that for each trajectory τ in Γ_X, paths of length k have finitely many (or even at most one) τ-extension.

4.3.7 Definition: Malnormal $\mathcal{V}H$-Complex. We say that a graph of groups is *Malnormal* if each edge group is a malnormal subgroup of each of its associated vertex groups. A $\mathcal{V}H$-complex is said to be malnormal if the associated graph of groups is malnormal. This just means that the attaching map of each corridor is a malnormal map of graphs.

4.3.8 Definition almost malnormal $\mathcal{V}H$-Complex. We say that a directed non-positively curved $\mathcal{V}H$-complex X is *almost malnormal* if it has the following property:

(1) The attaching map of each corridor is 2-malnormal.
(2) For any pair of distinct attaching maps of corridors $\Lambda_e \to \Lambda_v$ and $\Lambda_f \to \Lambda_v$, the disjoint union $(\Lambda_e \cup \Lambda_f) \to \Lambda_v$ is 2-malnormal.

4.3.9 Remark: 2-malnormal and singular leaves. The most common examples of 2-malnormal subgroups are subgroups of index 2. These turn up quite often in the decompositions of directed $\mathcal{V}H$-complexes, because any singular leaf in a $\mathcal{V}H$-complex corresponds to a graph with a single corridor connected to it by a double cover. This is the source of many almost malnormal $\mathcal{V}H$-complexes.

4.3.10 Theorem: k-uniqueness for Malnormal $\mathcal{V}H$-Complexes. *Let X be a malnormal $\mathcal{V}H$-complex, then there is a constant k so that for each trajectory $\tau : R \to \Gamma_X$, any immersed path of length $\geq k$ has at most one τ-extension by an immersed flat strip.*

Proof. By lemma 4.2.4 it is sufficient to prove this for a trajectory $\tau : R^+ \to \Gamma_X$.

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 $\sigma : I_K \to \Lambda_\tau$ be an immersed path which is 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 corridor, has at most one way to exit via a second corridor. This is because exiting Λ_v to a corridor corresponds to lifting a path in Λ_v to Λ_e (where e is the edge corresponding to the corridor), but since the strip has width $\geq k$, the path we are lifting has width $\geq k$ and so our choice of k and lemma 4.3.6 shows that there is at most one such lift and hence at most one extension. □

4.3.11 Theorem: k-finiteness for almost Malnormal $\mathcal{V}H$-Complexes. *Let X be an almost malnormal $\mathcal{V}H$-complex, then there is a constant k so that for each trajectory $\tau : R \to \Gamma_X$, any immersed path of length $\geq k$ has finitely many τ-extensions by an immersed flat strip.*

Sketch. The proof is similar to the proof of 4.3.10. We first prove the analogous statement for a trajectory $\tau : R^+ \to \Gamma_X$, and then apply lemma 4.2.4. Namely, we show that there are at most 2 τ-extensions. The idea is the same as above, except that at the initial extension there might be 2 choices, but for each succesive elementary extension there is at most one extension by an immersion. □

4.3.12 Theorem: Limits in almost malnormal. *Let X be an almost malnormal $\mathcal{V}H$-complex. Then every immersed flat in X is the strong limit of immersed cylinders and hence the limit of immersed tori.*

Proof. This follows by combining 4.3.11 with 4.2.2 and 4.2.3. □

4.4.1 Definition: The product $\Upsilon_1 \otimes \Upsilon_2$ of Graphs. Given a pair of immersions, $\phi_1 : \Upsilon_1 \to \Upsilon$ and $\phi_2 : \Upsilon_2 \to \Upsilon$. We define a graph $\Upsilon_1 \otimes \Upsilon_2$ and a pair of projection maps $\pi_1 : \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_1$ and $\pi_2 : \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_2$.

First consider the natural map

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

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

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

It is easy to see that $\Upsilon_1 \otimes \Upsilon_2$ consists of a set of 0-cells and diagonals of 2-cells of $\Upsilon_1 \otimes \Upsilon_2$ and therefore it is a graph. For each i we define the map $\pi_i : \Upsilon_1 \otimes \Upsilon_2 \to \Upsilon_i$ in the obvious way. See figure 6 for an illustration of $\Upsilon_1 \otimes \Upsilon_2$ in an easy special case.

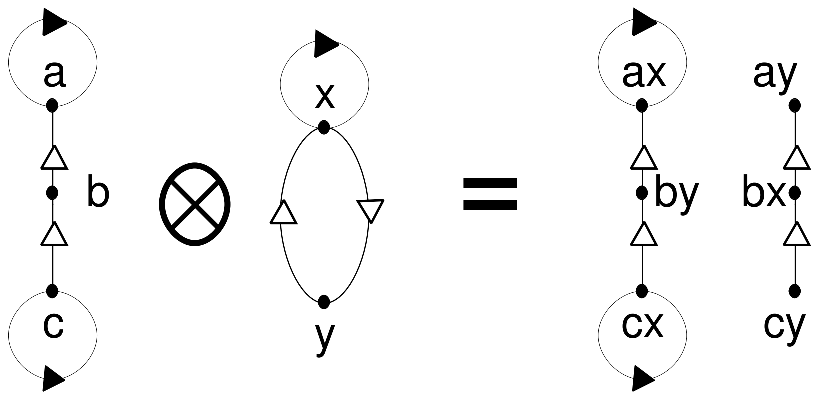

FIGURE 6: AN EASY $\Upsilon_1 \otimes \Upsilon_2$. Let Υ denote a bouquet of two circles, whose edges are labelled by by a white arrow and a black arrow. Let Υ_1 and Υ_2 denote the two labelled graphs on the left of the diagram. Note that their labellings induce maps $\Upsilon_1 \to \Upsilon$ and $\Upsilon_2 \to \Upsilon$. The graph $\Upsilon_1 \otimes \Upsilon_2$ is depicted on the righthand side of the equal sign. Note that $\Upsilon_1 \otimes \Upsilon_2$ has two components. The vertices of Υ_1 are labelled by a, b, c, and the vertices of Υ_2 are labelled by x, y. Each vertex of $\Upsilon_1 \otimes \Upsilon_2$ is labelled by the pair of vertices of Υ_1 and Υ_2 that it corresponds to.

It is easy to see that 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 σ_1 and σ_2 to Υ_1 and Υ_2 of a path σ in Υ. Given two lifts $\sigma_1 : I_n \to \Upsilon_1$ and $\sigma_2 : I_n \to \Upsilon_2$ of the path $\sigma : I_n \to \Upsilon$, there is a map $\sigma_1 \times \sigma_2 : I_n \to \Upsilon_1 \times \Upsilon_2$ whose image is contained in the $\Upsilon_1 \otimes \Upsilon_2$ subspace.

If we let $\bar{\Upsilon}_i \to \Upsilon$ be the completion of Υ_i to a covering space of Υ (one just adds trees) then, $\Upsilon_1 \otimes \Upsilon_2 \hookrightarrow \bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$. The components of $\bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$ are the covers of Υ which correspond to the intersections of conjugates of $\pi_1\Upsilon_1$ and $\pi_1\Upsilon_2$. And the subspace $\Upsilon_1 \otimes \Upsilon_2$ corresponds to a deformation retract of $\bar{\Upsilon}_1 \otimes \bar{\Upsilon}_2$.

Now, in case $\phi_1 = \phi_2$ then we have the space $\Upsilon_1 \otimes \Upsilon_1$. We will refer to the component of $\Upsilon_1 \otimes \Upsilon_1$ which is the diagonal of $\Upsilon_1 \times \Upsilon_1$ as the *diagonal* of $\Upsilon_1 \otimes \Upsilon_1$.

The graph $\Upsilon_1 \otimes \Upsilon_1$ has a labelling induced by the map $\Upsilon_1 \otimes \Upsilon_1 \to \Upsilon_1$.

4.4.2 Definition: Cyclonormal Subgroup. We say that $H \subset G$ is a *cyclonormal* subgroup provided that for any $g \in G - H$, the intersection $H \cap gHg^{-1}$ is cyclic or trivial. We define *r-cyclonormal* similarly (compare 4.3.3). We say that an immersion of graphs $\phi : \Gamma \to \Upsilon$ is cyclonormal provided that the image of $\pi_1\Gamma$ is a cyclonormal subgroup of $\pi_1\Upsilon$.

4.4.3 Remark: $\otimes$ and Malnormal and Cyclonormal. 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 $\chi \geq 0$. In other words, each non-diagonal component is either a tree or a simple cycle with several trees attached to it (each along a single point).

4.4.4 Lemma: Long path in Cyclonormal is ST^rU. *If $\Upsilon_1 \to \Upsilon$ is cyclonormal and finite, then there is a finite set of triples (S, T^d, U) of words, such that the label of any path in a non-diagonal component of $\Upsilon_1 \otimes \Upsilon_1$ can be expressed as $S(T^d)^rU$*

for some triple and some r. Note that S, T, U may be empty words. We will also think of the words S, T, and U as paths in $\Upsilon_1 \otimes \Upsilon_1$ and thus by projection they represent paths in Υ_1 and in Υ. We also assume that T is not a proper power. To see that finitely many triples suffice note that for a finite labelled graph Γ with $\chi(\Gamma) = 0$, we can let T^d correspond to the labels of closed simple paths in Γ. And we can let S and U correspond to the injective paths which end (respectivly begin) on the simple cycle of Γ.

4.4.5 Notation: $\bar{S}$, $\bar{T}$, $\bar{U}$. *We let $\bar{S}$, $\bar{T}$, and $\bar{U}$ denote the maximums of $|S|$, $|T|$, and $|U|$ respectively, over all triples (S, T^d, U) associated to the complex X.*

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

4.5.0 Theorem: Cylindrical $\mathcal{V}H$-complexes. *Let X be a $\mathcal{V}H$-complex, and assume that each corridor of X is a cylinder, that is, for each edge e of Γ_X, $\Lambda_e \cong S^1$. 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 corridor corresponding to e, and let $N = \mathrm{LCM}\big(\{\Lambda_e^{(0)}\}\big)$ be the least common multiple of the perimeters of corridors 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 itself cylindrical, it is obviously the strong limit of immersed cylinders.

4.5.1 Definition: Cyclonormal $\mathcal{V}H$-complex. We say that a graph of groups is *cyclonormal* provided each edge group is a cyclonormal subgroup of its respective vertex groups. Similarly, we say that a directed $\mathcal{V}H$-complex is *cyclonormal* provided that the attaching maps of each corridor are cyclonormal.

We will call a directed $\mathcal{V}H$-complex *almost cyclonormal* provided that for each $v \in \Gamma_X^{(0)}$ the following two conditions are satisfied:

(1) For any corridor 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.

(These two conditions can be combined into one but it isn't worth it.) Note that the class of cyclonormal $\mathcal{V}H$-complexes, includes both malnormal $\mathcal{V}H$-complexes as well as the class of cylindrical $\mathcal{V}H$-complexes mentioned in theorem 4.5.0.

Note that the class of cyclonormal $\mathcal{V}H$-complexes, includes both malnormal $\mathcal{V}H$-complexes as well as the class of cylindrical $\mathcal{V}H$-complexes mentioned in theorem 4.5.0.

4.5.2 Example: Figure 8 knot group. The Dehn complex of the usual projection of the Figure 8 knot is almost cyclonormal. This appears to be true for the Dehn complex of any alternating link. First note that the Dehn Complex is embedded in the link complement which is a 3-manifold. First note, that the direct product of a bouquet of three edges with itself cannot embed in a 3-manifold (because the bipartite 3-3 graph doesn't embed in a sphere). The almost cyclonormal property follows because if some intersection had rank ≥ 2 then there would have to be an embedded direct product as above.

4.5.3 Main Theorem: Cyclonormal $\Rightarrow$ Flat Planes are limits of Tori. *Any immersed flat plane in a non-positively curved almost cyclonormal $\mathcal{V}H$-complex is the strong limit of immersed cylinders and hence the limit of immersed tori.*

Proof. Consider a flat $F \to X$ with trajectory τ. By Theorem 4.5.4 we see that components of Θ_τ are compact, and then by lemma 4.2.2 we see that F is the strong limit of τ-cylinders. And by lemma 4.2.3 we see that F is the limit of tori as well. $\square$

4.5.4 Theorem: almost cyclonormal $\Rightarrow$ Θ_τ has compact components. *Let X be a non-positively curved almost cyclonormal $\mathcal{V}H$-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 main lemma 4.5.5 there is a number J such that for any trajectory τ, any path σ of length J in Λ_τ has finitely many (at most L) τ-extensions $\hat{\sigma}$ which are injective paths in Θ_τ. Now applying lemma 4.1.10, we obtain a bound D to the diameter of components of Θ_τ, and this bound is independent of τ. $\square$

4.5.5 Main Lemma: Finitely many extensions of length J path. *Let X be an (almost) cyclonormal non-positively curved $\mathcal{V}H$-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. The proof is broken up into the lemmas below. By lemma 4.2.4, it is sufficient to prove the statement of main lemma 4.5.5 for a trajectory $\tau : R^+ \to \Gamma_X$. Let J be chosen as in lemma 4.5.9 and consider a path λ in Λ_τ which is of length J. We show that if there are more than one (or two) τ-extensions of λ, then no τ-extension of λ corresponds to an injective path in Θ.

If there are more than one (or two) τ-extensions, then lemma 4.5.6 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 4.5.9 that λ' has a subpath $\sigma' = \gamma^p$, which satisfies the conditions of lemma 4.5.10. Therefore lemma 4.5.10 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. $\square$

4.5.6: Lemma: A bifurcation subpath. *Let X be a non-positively curved (almost) cyclonormal $\mathcal{V}H$-complex and let $\tau : R^+ \to \Gamma_X$ be a trajectory. If λ is an immersed path in Λ_τ which has more than one (or two) τ-extension then there exists a path λ' which is parallel to λ, such that λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$.*

Proof. We prove this for the case of cyclonormal $\mathcal{V}H$-complex, the proof for an almost cyclonormal $\mathcal{V}H$-complex is similar. If there are two τ-extensions of λ then we can consider the vertical path λ' where they bifurcate. See figure 7. 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). $\square$

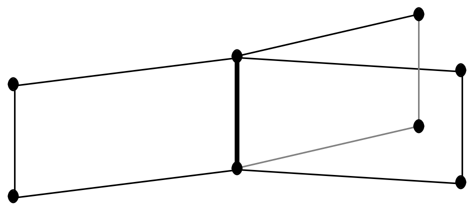

FIGURE 7: BIFURCATION. The figure above is meant to suggest a pair of τ-extensions of the vertical path γ on the left. The two τ-extensions bifurcate at the path γ'. indicated by the bold vertical path.

4.5.7 Lemma: 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 Z_p. In particular, T is a subword of T^2 in only the obvious ways.*

proof. We think of the cyclic word as a labelled 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 an unusual way, so that $T^p = XTY$ where $X \neq T^k$. Then because of the action of 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 labelling 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. $\square$

4.5.8 Lemma: locally T^3 implies T^p. *Let T be a cyclically reduced word which is not a proper power. Suppose that W is a cyclic word, such that each vertex in W is in a strict neighbourhood of T^3. By this we mean that the bi-infinite path W^∞ has the property that each vertex of W^∞ is contained in the middle third of a T^3 subpath. Then for some p, $W = T^p$.*

Proof. Since by the previous lemma, $T \subset T^2$ only in the obvious ways. We see that the T^3 patches must line up. $\square$

4.5.9: Lemma: A long highly periodic subpath. *There exists a number J, such that for every immersed path λ' in X_V which is of length J, if λ' lifts to a non-diagonal component of $\Lambda_e \otimes \Lambda_e$ (for some e), then λ' has a subpath σ' whose label is γ^p with γ not a proper power. And 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 4.4.4 if λ' lifts to a non-diagonal component then λ' can be written 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).

thus assuring that $r \geq 1$.)

It then 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' whose label is $(T_1^{d_1})^q = \gamma^p$ such that $\gamma = T_1$ is not a proper power, and $p = qd_1$. It is easy to see that the 3 properties that q satisfies imply the corresponding properties for p. □

4.5.10: Lemma: Extensions of Long Periodic Words are Cylindrical. *Suppose that $\sigma = \gamma^p$ where γ is not a proper power. and assume that the following conditions are satisfied:*

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

Then any extension of σ is cylindrical.

To see why $|\gamma^{p-1}| > 3(\|S\| + \|T\| + \|U\|)$ in property (2), note that by property (1), $|\gamma^{p-1}| \geq \frac{3}{4}\|\sigma\|$.

Proof. It is sufficient to show that this is true one corridor at a time. Namely, suppose that σ is a path in Λ_α and there is a corridor $\Lambda_\alpha \leftarrow \Lambda_e \rightarrow \Lambda_\beta$ then we show that any lift of σ to Λ_e is closed and we show that the projection σ_1, of this lift to Λ_B, the other side of the corridor, also has the properties listed above that σ satisfied: Namely we will show that $\sigma_1 = (\lambda_1)^{p_1}$ where p_1 satisfies property (1). And 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 doesn't lift to a closed path, then there would be two distinct lifts of γ^{p-1}. And therefore $\sigma = \gamma^p$ lifts to a closed path which is of period p, and so its projection σ_1 on the other side of the corridor, has period $\geq p$.

The case where γ^{p-1} doesn't have a unique lift is a bit trickier and it is here that we use our 3 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 neighbourhood. We will show that every point of the word γ^{p-1} has a strict T^3 neighbourhood.

Since $|\gamma^{p-1}| > 3(\|S\| + \|T\| + \|U\|)$, every point in the middle third of γ^{p-1} has a strict T^3 neighbourhood. By property (1), $p - 1 \geq 3$, and so the middle third of γ^{p-1} contains a γ subpath. Now, if we translate the covering of this γ subpath by strict T^3 neighbourhoods, inside $\sigma = \gamma^p$ (we translate by the action of Z_p on the cyclic word $\sigma = \gamma^p$), then we see that every point of σ has a strict T^3 neighbourhood. It follows from lemma 4.5.8 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 $\frac{\|\sigma\|}{\|T^d\|}$ which is ≥ 4 by hypothesis (3). $\square$

There are results similar to those of section 4.6 for general non-positively curved squared complexes. But 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.

Also, theorem 4.5.4 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 cylclonormal. The idea is very similar to 4.5.4.

The theorems of section 4.5 suggest the following more general results:

4.6.1 Conjecture: Word-hyperbolic generalization. *Let Γ be a graph of word-hyperbolic groups where each of the edge groups is embedded by a quasi-isometry. Suppose that Γ is a cylonormal 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 $Z \times Z$ subgroups.*

Given a group G and a pair of elements $g, h \in G$. we say that g and h *nearly commute* if for some n and m, the commutator $[g^m, h^n]$ is trivial. 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.

It seems likely that there are groups with solvable word problem for which the problem of deciding if two elements nearly commute is recursively unsolvable.

In the case of non-positively curved complexes, the problem of whether two elements nearly commute can be interpereted concretely in terms of tilings. We explain this in the case of a compact non-positively curved squared complex, but the details are similar in the general case.

The idea is that if g and h nearly commute then they have powers g^m and h^n which generate a free abelian subgroup of rank 1 or rank 2.

The case where they generate a cyclic subgroup is not difficult to handle. And in case the elements generate a $Z \times Z$ subgroup they act periodically on a flat plane, see [BH].

Applying the results of the previous section, we can give some positive results about the 'nearly commute?' problem for certain non-positively curved squared 2-complexes.

4.6.2 Proposition: No Anti-torus. *Let X be a non-positively curved cyclonormal $\mathcal{VH}$-complex, then X doesn't 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. □

4.6.3 Theorem: Nearly Commute. *Let X be a non-positively curved cyclonormal $\mathcal{VH}$-complex. There exists an algorithm which determines for any pair of elements $g, h \in \pi_1 X$ if g and h nearly commute. In particular let g and h denote immersed circles in the horizontal and vertical 1-skeletons of X. Then it is decidable whether or not the elements represented by g and h have powers which commute.*

sketch. Just fill in the quarter plane bounded by g^∞ and h^∞ either one obtains a torus at some point, or else the filling procedure eventually fails. Note that the non-positive curvature shows that there is only one way to fill it in. Since there is no anti-torus in X, the filling procedure cannot continue indefinitely without producing a torus. □

5: SUBGROUP SEPARABILITY STUFF

5.1 Definition: Subgroup Separability. Let $H \subset G$ be a subgroup of the group G. Then G is said to be *subgroup separable* with respect to H, or H*-separable*, provided that for each element $g \in G - H$, there exists a finite quotient $G \to Q$ such that $\bar{g} \notin \bar{H}$. A group is said to be *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. Thus, subgroup separability is clearly a much stronger property than residual finiteness.

5.2: Graph of Graphs Consider a 2-complex X such that X is determined by a graph Γ_X and some associated data in the following manner:

(1) For each $v \in \Gamma_X^{(0)}$ there is a graph Λ_v.
(2) For each edge $e \in \Gamma_X$ there is a graph Λ_e, and a *corridor* $C_e = \Lambda_e \times I$,
(3) and a pair of maps $\phi_{e0} : \Lambda_e \times \{0\} \to \Lambda_{e_0}$ and $\phi_{e1} : \Lambda_e \times \{1\} \to \Lambda_{e_1}$, where e_0 and e_1 are the vertices that the edge e is attached to.
(4) X is constructed from the data by taking the disjoint union $\cup_v \Lambda_v$, and then attaching the corridors C_e along their boundaries using the maps ϕ_{e0} and ϕ_{e1}.

We will generally assume that the maps ϕ_{ei} are immersions, though we will not assume that they are cellular). Note that there is an obvious map $\psi : X \to \Gamma_X$.

5.3 Definition: Clean Complex. We call a 2-complex X obtained as in 5.2, a *clean complex* provided that each attaching map $\phi_{ei} : \Lambda_e \times \{i\} \to \Lambda_{e_i}$ is an embedding.

5.4 Remark: Clean $\Rightarrow$ Malnormal. Note that if X is a clean $\mathcal{V}H$-complex, then X is a malnormal $\mathcal{V}H$-complex. This is because an embedding of a subgraph in a graph is obviously malnormal.

Given a 2-complex X as in 5.2 we will say that $\pi_1 X$ is $\pi_1\Lambda_e$-separable provided that for some point $x_e \in \Lambda_e \times \{1/2\} \subset X$, we have $\pi_1(X, x)$ is $\pi_1\big(\Lambda_e \times \{1/2\}, x\big)$-separable. Of course, this is independent of the choice of base point and only depends on the conjugacy class represented by $\Lambda_e \to X$.

5.5 Theorem: Subgroup Separable $\Leftrightarrow$ Virtually Clean. *Let X be as in 5.2, then $\pi_1 X$ is $\pi_1\Lambda_e$-separable for each e, if and only if X has a finite cover $\hat{X}$ which is clean (with respect to the decomposition of $\hat{X}$ induced by $\hat{X} \to X$).*

Proof. For each edge e, and each attaching map $\phi_{ei} : \Lambda_e \times \{i\} \to \Lambda_{e_i}$, we will show that there is a (based) finite cover $\hat{X}_{ei} \to X$ such that Λ_e lifts, and such that the induced attaching map $\hat{\phi}_{ei} : \hat{\Lambda}_e \to \hat{\Lambda}_{e_i}$ is an embedding. Taking a finite regular cover $\tilde{X} \to X$ which factors through $\hat{X}_{ei} \to X$ for each $_{ei}$ gives us a clean complex.

Since ϕ_{ei} is an immersion, ϕ_{ei} is an embedding if and only if ϕ_{ei} is injective on $\Lambda_e^{(0)}$. So for each pair of distinct vertices a and b of Λ_e which are sent to the same vertex of Λ_v we choose a path σ_{ab} in Λ_e connecting a and b. Let $\bar{\sigma}_{ab}$ denote the path which is the image of σ_{ab} in Λ_v. Choose as well, a path λ_{ab} contained in C_e, which travels from x to $\phi_{ei}(a) = \phi_{ei}(b)$. Now let γ_{ab} be the composition $\lambda_{ab} \cdot \bar{\sigma}_{ab} \cdot \lambda_{ab}^{-1}$, which represents an element of $\pi_1(X, x)$. Note that γ_{ab} is not an element of $\pi_1\Lambda_e$.

Now we use the fact that $\pi_1(X,x)$ is $\pi_1(\Lambda_e,x)$-separable to choose a finite index subgroup of $\pi_1(X,x)$ which contains $\pi_1(\Lambda_e,x)$ but does not contain γ_{ab} for any a,b,σ_{ab} as above. In the corresponding based cover $\hat{X}_{ei} \to X$, the induced map $\hat{\phi}_{ei} : \hat{\Lambda}_{ei} \to \hat{\lambda}_v$ is injective, for if it were not injective then some γ_{ab} would lift to a closed loop, which is impossible because we chose the finite index subgroup not to contain any such element. The converse is easy to prove. $\square$

5.6 Lemma: Virtual Limits. *Let X be a non-positively curved space, and let $\hat{X} \to X$ be a finite cover. Then for each of the properties below, $\hat{X}$ 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*

proof. We prove this for property (2), the others are similar. Consider an immersed flat plane $F \to X$. It lifts to an immersion $F \to \hat{X}$. Now suppose that $F \to \hat{X}$ is the limit of tori $T^2_n \to \hat{X}$. Then $F \to X$ is the limit of their projections $T^2_n \to \hat{X} \to X$.

In the other direction, consider an immersed flat plane $F \to \hat{X}$. Its projection $F \to X$ is the limit of tori $T^2_n \to X$. Then each such torus has a finite cover $\hat{T}^2_n$ (which is still a torus) which lifts to $\hat{X}$. The sequence $\hat{T}^2_n \to \hat{X}$, limits to $\hat{E}^2$. $\square$

5.7 Remark: Edge-Separable $\Rightarrow$ Properties. As a corollary of 5.5 and 5.6 we see that if $\pi_1 X$ is subgroup separable with respect to its edge subgroups then X is virtually clean, and so inherits various properties about flats from $\hat{X}$. In particular, since $\hat{X}$ is clean, it has and therefore X has the following properties:

(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 limit of immersed tori.
(4) Each immersed flat is the strong limit of immersed cylinders.

In a later section, we take advantage of this by turning it around. Namely, 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 for some $e \in \Gamma_x$, $\pi_1(X)$ is not Λ_e-separable. This will be the starting point for the solution of several long standing open questions, concerning residual properties of certain groups.

5.8 Question: Malnormal or Cyclonormal $\Rightarrow$ Virtually Clean. *Let X be a $\mathcal{V}H$-complex. If X is Malnormal, then is X virtually clean? If X is almost cylconormal then is X virtually clean? Are there any obstructions besides those involving the properties of flats in clean complexes, which are inherited by complexes that they cover?*

5.9 Remark: Cyclic edge groups. Some partial results in this direction were obtained in [W2], were it was shown that if all the corridors of the non-positively curved $\mathcal{V}H$-complex X, are cylinders, then there is a finite clean cover $\hat{X} \to X$.

5.10 Question: A special case. *Is a pair of free groups amalgamated over a malnormal subgroup: Residually finite? Virtually Clean? Subgroup Separable? (note that these properties are listed in order of increasing strength)*

5.11 Conjecture: Subgroup Separable CAT(0) $\Rightarrow$ Limits. *Let G act cocompactly and properly-discontinously on a CAT(0) space $\hat{X}$. I guess that if G is subgroup-separable then every flat in $\hat{X}$ is the limit of periodic flats. In particular, if there is an immersed flat plane then there is an immersed flat torus. Hopefully more restrictive hypotheses relative to certain subgroups of G would suffice as (as in theorem 5.5).*

5.12 Remark: General Squared Complex. Some evidence to support conjecture 5.11 is that, conjecture 5.11 is true in the case of general compact non-positively curved squared complexes. This is because, if $\pi_1 X$ is subgroup separable with respect to the subgroups corresponding to corridors of the double branched cover $\hat{X}$, then one can deduce that $\hat{X}$ is virtually clean, using arguments similar to theorem 5.5.

In a slightly different direction we ask:

5.13 Question: When does Clean $\Rightarrow$ Subgroup Separable. *Which clean complexes have subgroup separable π_1.*

This seems to be a tricky question. Note that the non-subgroup separable example of [BKS] is π_1 of a virtually clean complex. Note as well that $F_2 \times F_2$ is not subgroup separable but it is the fundamental group of a clean complex.

To obtain some partial results note that if X is a non-positively curved clean $\mathcal{V}H$-complexes, then by using the horizontal decomposition of X, one sees that under certain circumstances $\pi_1 X$ splits as a graph of free groups with cyclic edge subgroups. For instance, this is true when each each edge in V_X appears in at most 2 squares. In particular applying the results of [W2] we see that such groups are subgroup separable.

Continuing in this direction and avoiding the non-subgroup separable examples mentioned above, an optimistic conjecture is that clean $\mathcal{V}H$-complexes are subgroup separable provided the union of the attaching maps at each vertex is 2-cylonormal.

Part II: Complete Squared Complexes

Section 1: Complete Squared Complexes

1.1 Definition: Complete Bipartite Graph. A bipartite graph Γ with disjoint classes of vertices V and H, is said to be *complete* if the attaching map $\text{Edges}(\Gamma) \to V \times H$ is a bijection, that is, each pair of vertices from V and H are joined by a unique edge.

1.2 Definition: Complete Squared Complex. A squared complex X is called a *complete squared complex* (or *CSC* for short) if for each $v \in X^{(0)}$, the graph Link_v is a complete bipartite graph.

1.3 Remark: CSCs are non-positively curved. Note that if X is a CSC then X obviously satisfies the link condition because for each vertex $v \in X^{(0)}$, the cycles in Link_v have even length, and there are no cycles of length 2.

1.4 Remark: CSCs are locally $\mathcal{V}H$. If X is a CSC then X is obviously a locally $\mathcal{V}H$-complex. In fact, if X is connected, then X has a unique local $\mathcal{V}H$-structure.

1.5 Example: Not global-$\mathcal{V}H$. There are examples of CSCs which do not have a global $\mathcal{V}H$-complex. For instance, the standard 2-complex of

$$\langle a, b \mid a^2b^2 \rangle$$

is a CSC but does not have a $\mathcal{V}H$-structure. Note that this is just one of the usual ways of forming the Klein bottle.

1.6 Theorem: CSC has $\mathcal{V}H$double cover. *Let X be a CSC, then there is a double cover $\hat{X} \to X$ such that $\hat{X}$ is a CSC which is also a $\mathcal{V}H$-complex.*

Proof. Since X is obviously a locally $\mathcal{V}H$-complex we may apply theorem I.1.7 to obtain the desired cover.

Note that the isometry group of $\tilde{X}$ has a subgroup of index 2 which preserves the local $\mathcal{V}H$-structures, and so the proof can be made slightly easier in this case. □

1.7 Lemma: Trees. *Let $\tilde{X}$ be a simply connected non-positively curved $\mathcal{V}H$-complex. Then for each point $x \in \tilde{X}$, the spaces V_x and H_x are trees.*

Proof. The map $V_x \to \tilde{X}$ is a local-isometry and therefore induces a π_1-injection. (This is because it sends local-geodesics to local-geodesics, see [BH???], [M]). But since $\pi_1\tilde{X}$ is trivial, it follows that $\pi_1 V_x$ is trivial, and so the graph V_x must be a tree. □

1.8 Lemma: Connecting with vertical and horizontal. *Any two points w and z in a connected simply connected CSC $\tilde{X}$ can be connected by a path $\sigma_v \cdot \sigma_h$, which is the concatenation of a vertical and horizontal path. Similarly, any two points can be connected by a path $\sigma'_h \cdot \sigma'_v$. which is the concatenation of a horizontal and vertical path.*

Proof. By subdividing $\tilde{X}$ we may assume that w and z are in $\tilde{X}^{(0)}$. Now since $\tilde{X}$ is connected we may connect w to z by a path of minimal length σ in $\tilde{X}^{(1)}$. Now we show that σ is homotopic (relative to its endpoints) to a path $\sigma_v \cdot \sigma_h$ which is the concatenation of a vertical path beginning at w, followed by a horizontal path ending at z.

To see this, note that any length 2 subpath say $h_1^{\epsilon_1} v_2^{\epsilon_2}$ of a path λ, determines a unique square (2-cell) s of $\tilde{X}$, whose attaching map can be thought of as $h_1^{\epsilon_1} v_2^{\epsilon_2} = v_3^{\epsilon_3} h_4^{\epsilon_4}$ for some choices of h_i, v_i and ϵ_i. But then we can homotope λ through the square s to a path of the same length by homotoping the subpath $h_1^{\epsilon_1} v_2^{\epsilon_2}$ to $v_3^{\epsilon_3} h_4^{\epsilon_4}$, and leaving everything outside the subpath $h_1^{\epsilon_1} v_2^{\epsilon_2}$ unchanged.

Now, we define the complexity of a path λ in $\tilde{X}^{(1)}$ as follows: For each horizontal edge a in λ, let K_a denote the number of vertical edges in λ coming after a. We let the complexity of λ equal $\sum_{a \in \text{Horizontal Edges}(\lambda)} K_a$.

Now observe that each time we do a square homotopy as above we reduce by 1 the complexity of the path. It follows that after finitely many homotopies through squares we obtain a path of complexity 0 which is therefore of the form $\sigma_v \cdot \sigma_h$. □

1.9 Lemma: $V_w \cap H_z$ is one point. *Let w, z be points in the connected simply connected CSC $\tilde{X}$, then the intersection of V_w and H_z consists of a single point.*

Proof. To see that $V_w \cap H_z$ has at least one point, note that by Lemma 1.8, there is a path $\sigma_v \cdot \sigma_h$ which connects w to z. But σ_v ends at the point where σ_h begins, and so this point is in $V_w \cap H_z$, because $\sigma_v \to V_w$ and $\sigma_h \to H_z$.

To see that $V_w \cap H_z$ has at most one point, note that V_w and H_x are totally geodesic trees, and so if they intersected at x and y, then the geodesics joining x and y in V_w and H_z must be identical, and so since this geodesic is contained in $V_{\tilde{X}}$ and in $H_{\tilde{X}}$ it must be contained in $V_{\tilde{X}} \cap H_{\tilde{X}} = \tilde{X}^{(0)}$ and so it consists of a single point. Therefore $x = y$. □

As motivation for the following theorem, observe that the condition that X is a CSC is equivalent to X being locally isometric to the direct product of two trees.

1.10 Theorem: Tree: X is CSC $\Leftrightarrow \tilde{X} \cong$ Tree $\times$ Tree. *Let X be a squared complex with universal cover $\tilde{X}$. If $\tilde{X}$ is a CSC then X is a CSC. If X is a CSC and $x \in \tilde{X}$, and consider the two trees V_x and H_x. Then $\tilde{X}$ is isomorphic to $V_x \times H_x$.*

Proof. Note that by lemma 1.6 $\tilde{X}$ is a $\mathcal{V}H$-complex, so the statement of the theorem makes sense. We define a map $\phi : \tilde{X} \to V_x \times H_x$ defined by $y \mapsto \big(V_x \cap H_y, H_x \cap V_y\big)$. The map ϕ is well-defined by the previous lemma.

Note that ϕ has an inverse, $V_x \times H_y \to \tilde{X}$ defined by $(a, b) \mapsto H_a \cap V_b$. It is easy to see that these two maps are inverse to each other because $H_a = H_y$ and $V_b = V_y$. It is easy to check that ϕ is cellular as well and so it is an isomorphism and hence induces an isometry as well. □

1.11: Remark: Factoring π_1. *In case x_o is the only 0-cell of X, then one can deduce from lemma 1.8 or theorem 1.10 that $\pi_1(X)$ factors as $\pi_1(X) = \pi_1(V)\pi_1(H) = \pi_1(H)\pi_1(V)$. In general, $\pi_1(V)\pi_1(H) \subset_f \pi_1(X)$, but is not a subgroup. Also, $\pi_1(V) \cap \pi_1(H) = \{1\}$.*

1.12: Theorem: Biautomatic - The tree $\times$ tree language.

Using biautomatic structures for groupoids (see [E]), it is easy to show that the simple language consisting of paths which are, vertical followed by horizontal, is the language of a biautomatic structure for the fundamental group of a compact CSC. Note that it is obvious that this is a language of geodesics which satisfies the fellow traveler property. A similar easy language works for the higher dimensional cases (that is Tree $\times$ Tree $\times$ Tree, etc). Biautomatic structures have been obtained for more general classes of complexes: For finite C(4)-T(4)-complexes this is proven in [GS], and for compact non-positively curved cubulated complexes this is proven in [NR].

1.13 Proposition: CSC $\Leftrightarrow$ attaching maps are covering maps. *A directed $\mathcal{V}H$-complex X is a CSC if and only if each attaching map of each corridor in the decomposition of X is a covering map.*

Proof. If all the attaching maps are covering maps, then it is easy to see that the links are complete bipartite graphs. To see the converse, we note that for cellular maps between graphs, a covering map is just a local homeomorphism. But this local condition is satisfied by the attaching maps of the corridors in the case of CSCs because the links of vertices are complete bipartite graphs. So, in case X is a CSC, all the attaching maps of corridors are covering maps. $\square$

1.14 Remark: Weighted Graph. Because of lemma 1.13, we are able to add some structure to the decomposition graph Γ_X. When X is a directed CSC, we label the ends of edges in Γ_X, with numbers indicating the degrees (as a covering space) of the associated attaching maps.

Section 2: The Main Example: X

2.1: Main Example: X as Six Squares.

We define X to be the complex consisting of the six squares indicated in figure 8. The squares are glued together as indicated by the coloring of edges. It is easy to see that X is a $\mathcal{V}H$-complex and that it is a CSC. Note that X has a unique 0-cell. Let $H = H_X$, denote the subcomplex consisting of the 2 horizontal edges, and let $V = V_X$ denote the subcomplex consisting of 3 vertical edges.

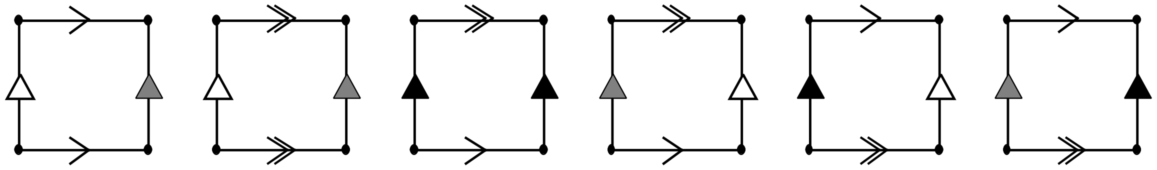

Figure 8: X as six squares. The figure above indicates the gluing pattern for the six squares of X. The three vertical edges colored white, grey, and black are denoted a, b, and c respectively. The two horizontal edges, single and double arrow, are denoted x and y respectively.

2.2: The vertical decomposition of X.

Observe that X is both vertically and horizontally directed, and thus X may be described in terms of both its vertical and horizontal decompositions (see I.1.17). Here we describe the vertical decomposition of X.

Let $\rho_L : L \to V$ and $\rho_R : R \to V$ be the covering spaces of V indicated in figure 9. Denote by $\tau : L \to R$ the obvious (translation) graph isomorphism. Note that τ is *not* a covering space isomorphism.

Using the information contained in ρ_L, ρ_R and τ we may form X as follows: Γ_X will be a graph with one vertex v, and one edge e. Let Λ_e denote a graph isomorphic to R and thus isomorphic to L via τ. The vertex $v \in \Gamma_X$ corresponds to $\Lambda_v = V \subset X$. The corridor $\Lambda_e \times I$ corresponds to the edge e in Γ_X. Identify $\Lambda_e \times 0$ and $\Lambda_e \times 1$ with L and R respectively, in such a way that the isomorphism $\Lambda_e \times 0 \cong \Lambda_e \times 1$ is the same as τ. Now, glue each end of $\Lambda_e \times I$ to V using the maps ρ_L and ρ_R, respectively. So we have

$$X = \big(\Lambda_e \times I \,\cup\, V\big)/\big\{(x, 0) \sim \rho_L(x),\ (x, 1) \sim \rho_R(x) \ \mid\ x \in \Lambda_e\big\}$$

Notice that this describes $\pi_1(X)$ as an HNN extension of a free group of rank 3, associating subgroups of index 2. The decomposition in the horizontal direction provides a decomposition of $\pi_1(X)$ as an HNN extension of a free group of rank 2, associating subgroups of index 3.

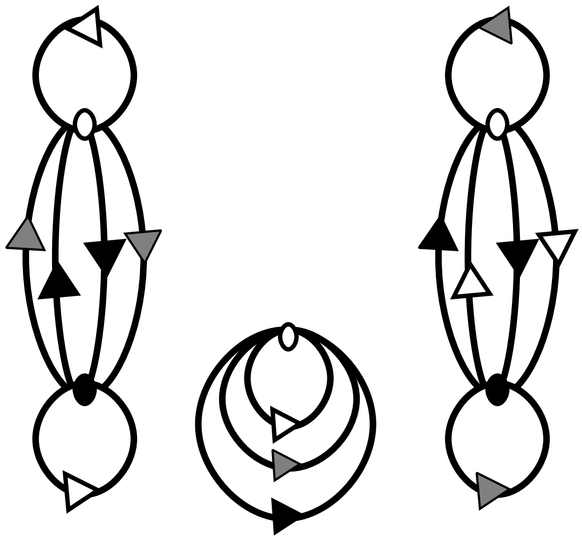

FIGURE 9: VERTICAL DECOMPOSITION OF X. The bouquet of 3 circles in the center is V. The graph on the left and right are denoted by L and R respectively. There are covering maps $\rho_L : L \to V$ and $\rho_R : R \to V$ induced by the colorings of the graphs.

SECTION 3: THE ANTI-TORUS

3.0: Definition: Anti-torus Let Y be a CSC with a $\mathcal{V}H$-structure, and let $\tilde{Y}$ denote the universal cover of Y. Let $y \in Y^{(0)}$ and let $\tilde{y} \in \tilde{Y}$ denote a preimage of y.

Let s_h and s_v denote immersed circles in H_y and V_y which are based at $\tilde{y}$. Let s_h^∞ and s_v^∞ denote the lines in $\tilde{Y}$ based at $\tilde{y}$ which are labelled periodically by the labels of s_h and s_v respectively. That is, the map $s_h^\infty \to \tilde{Y}$ factors through the map $s_h \to Y$. Similarly for s_v.

Let $s_h^\infty \times s_v^\infty$ denote the convex closure of s_h^∞ and s_v^∞ in $\tilde{Y}$. Note that since $\tilde{Y} \cong H_{\tilde{Y}} \times V_{\tilde{Y}}$, the convex closure of $s_h^\infty \cup s_v^\infty$ is the product subspace $s_h^\infty \times s_v^\infty \hookrightarrow H_{\tilde{Y}} \times V_{\tilde{Y}}$.

We say that $s_h^\infty \times s_v^\infty \subset \tilde{Y}$ is an *anti-torus* provided that it is not tiled periodically by the preimages of squares of Y. That is, the map $s_h^\infty \times s_v^\infty \to Y$ doesn't factor through a torus $T^2 \to Y$.

The existence of immersed anti-tori in certain compact CSCs plays a fundamental role in this work.

Remark: Anti-tori in general. Given a (general) non-positively curved space X, and given a local-isometry $E \hookrightarrow \tilde{X}$ from a flat plane to the universal cover. We refer to E as an *anti-torus* provided E contains a pair of intersecting periodic lines.

3.1: The Anti-Torus Π immersed in X. *Let c^∞ denote the infinite periodic vertical line in $\tilde{X}$ which is the component containing the basepoint $\tilde{x}_o \in \tilde{X}^{(0)}$ of the preimage in $\tilde{X}$ of the loop labelled by c in X. Define similarly y^∞. Let Π denote the convex closure in $\tilde{X}$ of the infinite geodesics labelled by c^∞ and y^∞. The plane Π, is tiled by the six orbits of squares in $\tilde{X}$. See figure 10 below.*

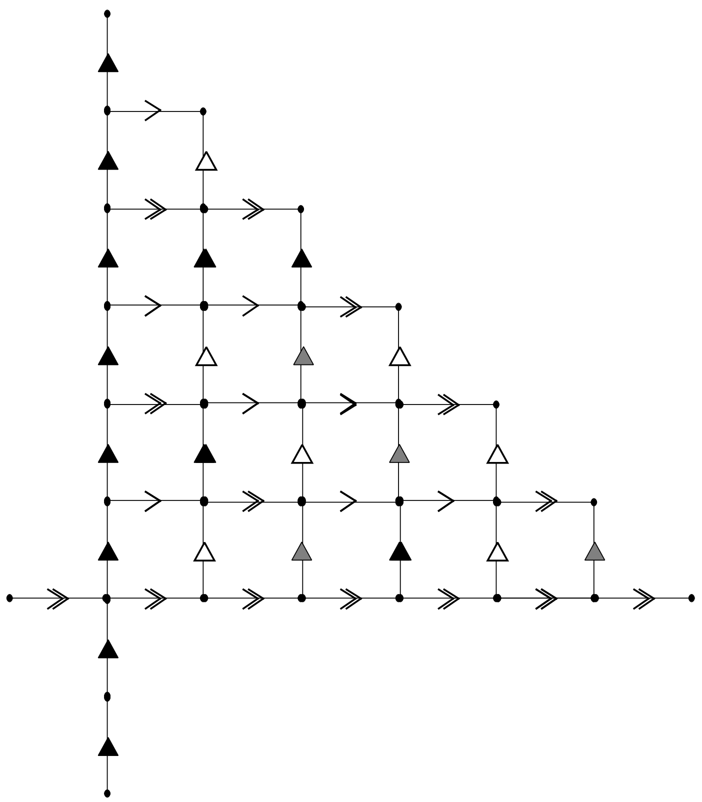

FIGURE 10: THE ANTI-TORUS Π. The plane Π above is the convex closure of two periodically labelled lines in $\tilde{X}$. A small region of the northeast quadrant has been tiled by the squares of X.

3.2: Proof of aperiodicity of Π.

Let the basepoints l_o, x_o, and r_o of L, V, and R be represented by the white vertices. Let the second vertex of L be denoted by l_1.

Given a path $\sigma : I \to V$ with $\sigma(0) = x_o$, let $\tilde{\sigma}^o$ and $\tilde{\sigma}^1$ denote the lifts of σ to L based at l_o, and l_1 respectively.

Define $\phi_i(\sigma) : I \to V$ to be $\phi_i(\sigma) = \rho_r \circ \tau \circ \tilde{\sigma}^i$.

Let $\iota : V \to V$ be the orientation-preserving automorphism that permutes the edges a and c and fixes b.

We are interested in the case were σ is a combinatorial path, and so is naturally represented as a word in $\{a^\pm, b^\pm, c^\pm\}$. As usual, the product of two paths $\sigma_1 \circ \sigma_2$ is represented by the concatenation of words.

Define $\#_a : \pi_1(V) \to \mathbb{Z}$ to be the homomorphism induced by $a \mapsto 1$, $b \mapsto 0$, and $c \mapsto 0$. Note that this merely counts the number of signed occurences of a in the reduced word representing the element. Define similarly $\#_b$ and $\#_c$. Let $\#_{b+c} \overset{\text{def}}{=} \#_b + \#_c$.

3.3: Lemmas.

(1) $\sigma \in \pi_1(L) \Leftrightarrow \#_{b+c}(\sigma) \equiv_2 0$
(2) $\phi_1 = \iota \circ \phi_0$
(3) $\sigma \notin \pi_1(L) \Rightarrow \phi_0(\sigma^2) = \phi_0(\sigma) \cdot \phi_1(\sigma)$

3.4: Lemma: Doubling lemma. $\sigma \notin \pi_1(L) \Rightarrow \phi(\sigma^2) \notin \pi_1(L)$.

Proof.

$$\begin{aligned}&\#_{b+c}(\phi_0(\sigma^2)) \equiv_2 \#_a(\phi_0(\sigma^2)) = \\ &\#_a(\phi_0(\sigma)\cdot\phi_1(\sigma)) = \#_a(\phi_0(\sigma)) + \#_a(\iota\circ\phi_0(\sigma)) = \\ &\#_a(\phi_0(\sigma)) + \#_c(\phi_0(\sigma)) = \#_b(\sigma) + \#_c(\sigma) = \#_{b+c}(\sigma) \equiv_2 1\end{aligned}$$

The first equality follows because $\|\sigma^2\| \equiv_2 0$. The second follows from statement (3) of Lemma 3.3. The third follows from statement (2) of Lemma 3.3 and from the fact that ϕ_0 is a homomorphism. The fourth follows from the definition of ι. The fifth follows because the letters that may be transformed by ϕ_0 into either a or c are the letters b and c (not respectively). The sixth follows by definition. The seventh follows from statement (1) of Lemma 3.3 and the assumption that $\sigma \notin \pi_1(L)$. $\square$

3.4b: Corollary. $\phi^n(c^{2^n}) \notin \pi_1(L)$.

Proof. We use lemma 3.4 to prove this by induction. When $n = 0$, we obviously have $\phi^0(c^{2^0}) \notin \pi_1(L)$, since this just means that $c \notin \pi_1(L)$. Suppose that the statement is true for $n = k$.

By lemma 3.4, $\phi^k(c^{2^k}) \notin \pi_1(L)$ implies that $\phi\left(\left(\phi^k(c^{2^k})^2\right)\right) \notin \pi_1(L)$. But it is easy to check that $\left(\phi^k(c^{2^k}\right)^2 = \phi^k(c^{2^{k+1}})$ and so substituting proves the lemma. $\square$

3.5: Remark:. Notice that $aca^{-1}c^{-1} \xrightarrow{\phi} cac^{-1}a^{-1} \xrightarrow{\phi} aca^{-1}c^{-1}$. Therefore the lemma cannot be strengthened to say that for any σ, there exists n such that, $\phi^n(\sigma) \notin \pi_1(L)$.

3.6: Geometric interpretation Consider the horizontal geodesic in $\tilde{X}$ based at $\tilde{x}_o$ corresponding to y^n. Given a vertical geodesic based at $\tilde{x}_o$ represented by some word σ in $\{a^{\pm1}, b^{\pm1}, c^{\pm1}\}$, the convex closure of the union of these two paths at $\tilde{x}_o$, is a rectangle. It is easy to see that the word $\phi^n(\sigma)$ is simply the label (beginning at the endpoint of y^n) of the side opposite σ in this rectangle. In a sense, all of the 'calculations' in the lemmas above are really occurring inside of Π. Using this point of view, it is easy to show that Π is not periodic, merely as a consequence of 3.4. However, instead of going through that relatively simple argument, we will look more carefully at the tiling of Π, and obtain a somewhat more informative proof of aperiodicity.

3.7: Proposition: Π Period Doubling. *Define $W_n(m)$ to be the horizontal positive path in Π of length n beginning at the endpoint of the vertical path c^m. Thus, $W_n(m)$ determines a word which is the label of the side opposing y^n in the rectangle which is the convex closure of y^n and c^m. (We will also use $W_n(m)$ to denote the associated word.)*

For each n, the words $\{W_n(m) \mid 0 \le m \le 2^n - 1\}$ are all distinct, and thus, every positive word of length n is $W_n(m)$ for some m. (Of course, for each n, there are exactly 2^n distinct positive words in x and y of length n.)

Proof. Given a vertical path, it represents an element of $\pi_1(L)$ if and only if its convex closure at the basepoint with the edge labelled by y or x yields a rectangle whose top edge is labelled by y or x respectively.

Now, arguing by induction, we assume that the theorem is true for $n = k$, so that $W_k(m)$ are distinct for $0 \leq k \leq 2^k - 1$. In particular, it follows that

$$W_k(m) = W_k(m + 2^k)$$

Consequently the first k letters of the words $W_{k+1}(m)$ and $W_{k+1}(m + 2^k)$ are equal, therefore it suffices to show that the last letters of these words are distinct. However, the vertical word connecting the end points of the paths associated with $W_k(m)$ and $W_k(m + 2^k)$ is just a subword of length 2^k of the word $\big(\phi^k(a^{2^k})\big)^2$, so it is a conjugate of $\phi^k(a^{2^k})$ which, by lemma 3.4b, is not an element of $\pi_1(L)$. It follows from the opening remark that the last letters of $W_{k+1}(m)$ and $W_{k+1}(m+2^k)$ are distinct, and we are done. $\square$

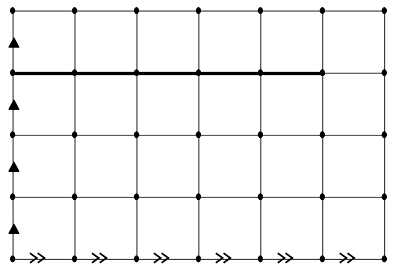

FIGURE 11: COMPUTATIONS IN Π . The figure above represents the upper right quadrant of Π. The bold horizontal line, represents the path corresponding to $W_5(3)$.

3.8: Proposition: Π is an Anti-torus.

Proof. To see that Π is aperiodic, note that we have just shown that the width n, infinite vertical strip in Π bounded on the left by the c^∞ axis has period 2^n. Alternatively, the proof of 3.7 shows that every finite positive word in x and y appears in Π, and so Π cannot possibly be perioidic. $\square$

3.9: Questions: finding anti-tori.

It would be nice if there were an algorithm to determine whether or not two perpendicular elements in a CSC group do not have powers that commute. Or in other words, decide whether the convex closure of their axes is an anti-torus or a periodic plane. See Question I.4.6.2, I.4.6.3. It seems very likely that CSCs typically have immersed anti-tori. I wonder if there is an easy proof of this. See 4.10.

Section 4: The covering spaces of X

4.1 Weight Conservation of $\rho_\star$.

Let X denote a directed $\mathcal{V}H$-complex and let $X \to \Gamma_X$ denote the associated vertical decomposition. As in I.1.19, a map $Y \to X$ of $\mathcal{V}H$-complexes gives rise to a map $\Gamma_Y \to \Gamma_X$. In particular, a covering space $\hat{X}$ of X has an induced directed $\mathcal{V}H$-structure, and the covering map gives rise to a commutative diagram

$$\begin{array}{ccc} \hat{X} & \to & \Gamma_{\hat{X}} \\ \rho\downarrow & & \rho_\star\downarrow \\ X & \to & \Gamma_X \end{array}$$

Note that, as mentioned above, $\rho_\star$ is not usually a covering map, but is instead a submersion of graphs. Let us restrict ourselves to the case where X is a CSC. Then, as mentioned above, it is natural to label each end of each edge of Γ_X with the degree of the associated attaching map in X. So, if the initial vertex of $\varepsilon \in \mathrm{Edges}(\Gamma_X)$ is v, then we label the end of ε at v with the degree of the associated attaching map $\Lambda_\varepsilon \times 0 \to \Lambda_v$ of the corridor corresponding to ε.

Keeping track of this additional information, the map $\rho_\star : \Gamma_{\hat{X}} \to \Gamma_X$ is locally weight conserving. That is, the link of each vertex of $\Gamma_{\hat{X}}$ and Γ_X has a weighting by natural numbers. And, for each $v \in \Gamma_{\hat{X}}$, the map $\mathrm{Link}(v) \to \mathrm{Link}(\rho_\star(v))$ has the property that the sum of the weights of the preimage of a point is equal to the weight of that point. This is a rather simple restriction, but it will enable us to limit the types of covering spaces that may arise.

In order to see why the local weight conservation holds, observe that the same is true for the covers of the mapping cylinder of a covering space.

4.2: Restrictions on Covering Spaces of X.

Since the attaching maps of the sole corridor of X are both of degree 2, (that is, the maps $\rho_L : L \to V$ and $\rho_R : R \to V$ are each of degree 2, as above), there are only four local models for the map $\Gamma_{\hat{X}} \to \Gamma_X$ induced by a covering space $\hat{X} \to X$. These are illustrated in figure 12 below. Using the Anti-Torus II, we will be able to show that under certain circumstances only case (iv) may occur.

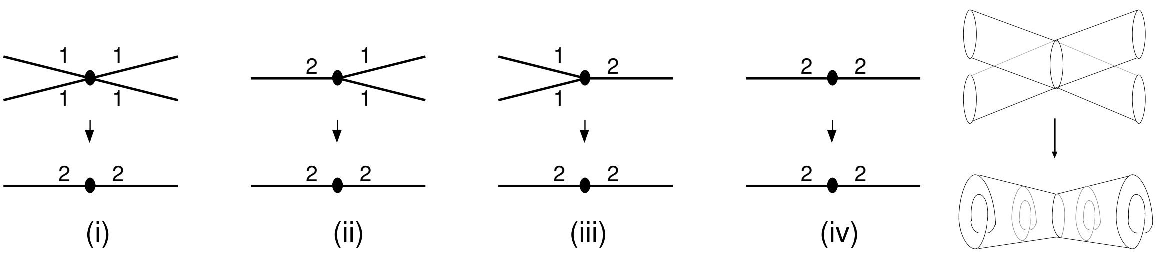

Figure 12: The Four Cases. Diagrams (i), (ii), (iii), and (iv) represent the four different local possibilities for the map $\Gamma_{\hat{X}} \to \Gamma_X$. The figure on the right is meant to suggest the type of local behavior of a map $\hat{X} \to X$ that would produce case (i).

4.3 Lemma: Anti-Torus Obstruction. *Suppose that $\hat{X}$ is a regular covering, and suppose that there exists n such that the closed path c^n lifts to a closed path in*

$\hat{X}$. Then only case (iv) of figure 12 may occur. That is, $\pi_1(V) \cap \pi_1(\hat{X}) \not\subset \pi_1(L)$, and $\pi_1(V) \cap \pi_1(\hat{X}) \not\subset \pi_1(R)$.

It is worth mentioning that in the case where $\hat{X}$ is a finite cover of X, cases (ii) and (iii) are automatically excluded because components of $\rho^{-1}(R)$ and $\rho^{-1}(L)$ must be isomorphic as graphs and so it is impossible for one of them to be isomorphic to a proper cover of $\hat{V}$ while the other one is isomorphic to the trivial cover. To see that lemma 4.3 is true in the finite case, observe that if $\hat{X} \to X$ is of type (i) then $\hat{X}$ is clean. But this is impossible since if $\hat{X}$ is clean, then letting τ denote the trajectory of $\hat{\Pi} \to \hat{X}$, from theorem 3.10 Θ_τ is obviously finite (actually this is trivial special case of the more general result of theorem 3.10 or 4.5.4). However the immersed anti-torus $\hat{\Pi}$ determines an infinite length injective path in Θ_τ, which is a contradiction.

A slightly different proof that there isn't a finite clean cover can be obtained by combining 4.5 and 4.9 below.

Proof of 4.3. Suppose that case (iv) does not occur, then in the decomposition of $\hat{X}$, the attaching maps of corridors corresponding to at least one of $L \to V$ or $R \to V$ have degree 1.

[The idea of the proof is similar to the proof in the finite case mentioned above, but it requires a slightly more sensitive analysis. Essentially, we assume that a cover $\hat{X} \to X$ exists, and produce a contradiction, namely that either the right or left half of Π is periodic (i.e. cylindrical). We do this by showing that $\hat{\Pi}$ factors through an immersion of a cylinder $C \to \hat{X}$, where $C = S^1 \times [0, \infty)$ and where the map $S^1 \to X$ represents c^n and the map $[0, \infty) \to X$ represents $y^{+\infty}$. This seems pretty easy to see, since $C \to X$ is determined by its restriction to these subspaces, but I have included a more detailed explanation below.]

Let us consider the case where the attaching maps of the corridors of $\hat{X}$ corresponding to the map $L \to V$ are of degree 1.

From the point of view of a preimage $\hat{V}$ of V, this means that outgoing horizontal arrows point towards a graph isomorphism between $\hat{V}$ and a cross section of the respective corridor, that is, towards a component of the preimage L. This means that given a closed path in $\hat{V}$ one is able to push it in the direction of the horizontal arrows, towards another copy of $\hat{V}$. More precisely, given a map of a graph $\Gamma \to \hat{V}$ one can extend it through this corridor to a combinatorial map of $\Gamma \times I \to \hat{X}$. Since each y-labelled outgoing edge points towards such a graph isomorphism, we are able to extend the map of Γ so that an edge $\gamma_o \times I \subset \Gamma \times I$ is sent to an outgoing y-labelled edge.

Denote by C the semi-infinite cylindrical squared-complex $S^1 \times [0, \infty)$ where S^1 consists of n edges and $[0, \infty)$ is subdivided into edges in the usual way. Also, orient all the edges of $[0, \infty)$ so that they are pointing away from 0.

We identify S^1 with $S^1 \times \{0\} \hookrightarrow C$ and, letting s be a basepoint 0-cell of S^1, we identify $[0, \infty)$ with $s \times [0, \infty) \hookrightarrow C$. In addition, let $(s, 0)$ denote the basepoint of C.

We will show that there is a $\mathcal{V}H$-immersion $\phi : C \to \hat{X}$, so that $S^1 \hookrightarrow C$ gets sent to a closed path in $\hat{X}$ with label c^n, and so that $[0, \infty) \hookrightarrow C$ gets sent to an infinite path with label y^∞ (that is $yyy\cdots$). This will result in a contradiction, as explained below, it follows that our assumption that there is a covering space $\hat{X}$ falling into case (iii) above is incorrect.

We begin by using our assumption that c^n lifts to a closed path in order to define $\phi : S^1 \to \hat{X}$ to represent such a closed path. Then, as above, we may extend the map ϕ restricted to $S^1 = S^1 \times \{0\}$ to $S^1 \times [0, 1)$ so that the edge $s \times [0, 1]$ is sent to a y labelled outgoing edge. Similarly we extend ϕ to $S^1 \times [0, 2)$, and so on to define ϕ on all of C.

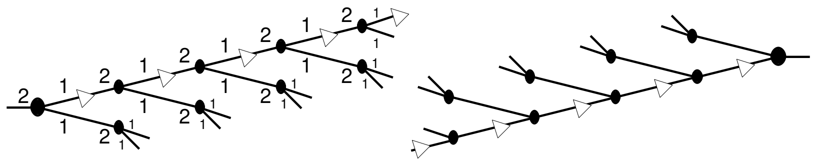

FIGURE 13: TRAJECTORY OF C IN $\hat{X}$. The figures on the left and right represent $\Gamma_{\hat{X}}$ corresponding to cases (ii) and (iii) of Figure 12. The sequence of arrowed edges in the figure on the left, is meant to suggest the image of the map $\Gamma_C \to \Gamma_{\hat{X}}$ that would be induced by the map $C \to \Gamma_X$ constructed in the proof. The sequence of arrowed edges in the figure on the right is meant to suggest the same for a similar proof in case (iii).

Now, since there are only finitely many labels, and because the trajectory of C in $\hat{X}$ is periodic, the induced tiling of C is periodic. But since the anti-torus Π factors through C, or rather, factors through the induced map $C \to X$, it would follow that Π itself is periodic. This would be a contradiction.

Cases (i) and (ii) of figure 12 are handled similarly, so that the only possible covers $\hat{X}$ of X which contain a lift of c^n are of type (iv). $\square$

Note that statements similar to 4.3 can be proven for any compact CSC containing an anti-torus. See also I.5.6 and I.5.7.

4.4 Corollary: $\pi_1(X)$ is not $\pi_1(V)$-separable. *Consider the homomorphism $\psi : \pi_1 X \to \mathbb{Z}$ induced by $a \mapsto 0, b \mapsto 0, c \mapsto 0, x \mapsto 1, y \mapsto 1$. Then $\pi_1 V$ is not separable from any element of Kernel(ψ), in any quotient in which c has finite order. That is, for all n we have the inclusion*

$$Kernel(\psi) \subset \pi_1(V)\langle\langle c^n \rangle\rangle.$$

In particular, $xy^{-1} \in \pi_1(V)\langle\langle c^n \rangle\rangle$ for each $n \geq 1$.

Proof. Let $\hat{X}$ denote the cover of X corresponding to the normal subgroup $\langle\langle c^n \rangle\rangle$. Let $\hat{V}$ denote the component of the preimage of V in $\hat{X}$ containing the basepoint. Consider the decomposition $\hat{X} \to \Gamma_{\hat{X}}$. By 4.3, $\Gamma_{\hat{X}}$ is isomorphic to R. Furthermore $\Gamma_{\hat{X}}$ can be oriented so that $\hat{X} \to \Gamma_{\hat{X}}$ is orientation preserving on the horizontal edges, and of course, the vertical edges are sent to 0-cells. Now, $g \in \pi_1(V)$ exactly when the closed path corresponding to g in X lifts to a path ending at $\hat{V}$ in $\hat{X}$. Looking at the projection $\varrho : \hat{X} \to \Gamma_X$, this is equivalent to requiring that the path corresponding to g being projected to a closed path in Γ_X. But since ϱ is orientation preserving on $H_{\hat{X}}$, the elements of Kernel(ψ) are projected to closed paths and we are done. $\square$

4.5 Proposition: Clean ⇒ Virtually $F_n \times F_m$. *Consider a compact CSC K which is clean, that is all the attaching maps (of the corridors in the decomposition of K) are of degree 1. Then K has a finite cover, $\hat{K} \to K$ so that $\pi_1(K) \cong F_n \times F_m$. In fact, $\hat{K}$ is isomorphic to the direct product of graphs.*

Proof. Choose a basepoint $b \in \Gamma_K^{(0)}$. Any closed loop in Γ_K based at b determines an element of $\text{Aut}(\Lambda_b)$. This is because each edge e of Γ_K attached to the vertices u and v, corresponds to a corridor C_α with attaching maps of degree 1, and therefore induces an isomorphism $\Lambda_u \to \Lambda_v$. Likewise, a closed path based at b in Γ induces an automorphism of Λ_b, which is the composition of isomorphisms.

In this way we obtain a homomorphism

$$\pi_1(\Gamma_K) \to \text{Aut}(\Lambda_b).$$

Composing this with the projection $K \to \Gamma_K$ we obtain (choosing $k \in \Lambda_b$) the homomorphism μ, which is the composition

$$\pi_1(K, k) \to \pi_1(\Gamma_K, b) \to \text{Aut}(\Lambda_b).$$

Since K is compact, Λ_b is a finite graph, so $\text{Aut}(\Lambda_b)$ is finite as well. It follows that the kernel of μ is of finite index in $\pi_1(\Gamma)$, and consequently, the cover $\hat{K}$ that it corresponds to is of finite degree. It is easy to see that

$$\hat{K} \cong \Lambda_b \times \Gamma_{\hat{K}},$$

since we have just chosen it to be a trivial Λ_b bundle over $\Gamma_{\hat{K}}$. □

4.6 Propostion: $F_n \times F_m$ ⇒ Clean. *If K is a CSC, and $\pi_1(K) \cong F_n \times F_m$ then K is clean. That is, the attaching maps of the corridors of K are all degree* 1.

Sketch. Consider first the case where both n and m are ≥ 2. The key point is that a diagonal element, that is, an element which isn't represented by a totally vertical or horizontal closed geodesic path has $Z \times Z$ as its centralizer. This is true because the non-positive curvature condition determines a unique plane containing any given non-vertical non-horizontal line.

From this we conclude that the F_n and F_m factors act vertically and horizontally or vice-versa. And then it is not difficult to show that the complex is clean and in fact the direct product of the two graphs corresponding to F_n and F_m.

In the case where $n \geq 2$ and $m = 1$, using the same idea we can determine that the F_1 subgroup is either vertical or horizontal, and from there it is easy to show that it is clean and in fact a product though the second graph factor might correspond to another F_2 subgroup.

The case of $F_1 \times F_1$ is easily seen to be clean, but here neither factor is necessarily vertical or horizontal.

There are more general results for non-positively curved spaces (see [BH]). □

4.7 Problem: Act on Tree $\times$ Tree closed under quasi-isometry. *Suppose G is quasi-isometric to a Tree $\times$ Tree. Does G have a subgroup of finite index which acts (nicely) on a Tree $\times$ Tree?*

4.8 Remark: c and y don't nearly commute. It follows from general theorems about cocompact actions on CAT(0) spaces, that maximal free-abelian subgroups act on a flat subspace in the same way as lattices on E^n [BH]. Therefore, if c and y had powers that commuted, then there would be a periodic plane with a pair of axis labelled by c^∞ and y^∞. On the other hand, the non-positive curvature condition implies that the (aperiodic) anti-torus Π above, is the unique plane containing the c^∞ and y^∞ axes. Consequently, there are no integers $m, n \geq 1$ such that $[c^m, y^n] = 1_{\pi_1(X)}$. This is strange.

4.9 Corollary: $\pi_1(X)$ is isometric to $F_2 \times F_3$ but not commensurable with it.

Proof. That $\pi_1(X)$ is isometric to $F_2 \times F_3$ follows because $\tilde{X}$ is isometric to the direct product of a regular tree with valence 4 with a regular tree of valence 6 (See 1.10).

To see that $\pi_1(X)$ is not commensurable with the direct product of two free groups, first note that as in 4.6, if some finite cover $\hat{X}$ of X had $\pi_1(\hat{X}) \cong F_n \times F_m$, then F_n and F_m would be vertical and horizontal. This means that they would be subgroups of finite index of $\pi_1(V)$ and $\pi_1(H)$. But then there would be finite integers q and r so that $c^q \in F_n$ and $y^r \in F_m$ and then $[c^q, y^r] = 1$ which cannot be. $\square$

4.10 Conjecture: Virtually $F_n \times F_m \Leftrightarrow$ no anti-torus. *I guess that if a compact CSC X has no immersed anti-torus then there is a finite clean cover $\hat{X} \to X$. This would imply that $\hat{X}$ it has finite cover which is clean. By 4.5, this is equivalent to $\pi_1 X$ being virtually $F_n \times F_m$. Note that as above, if X admits an immersed anti-torus, then X is not virtually clean. See section I.5 for some related material. For instance as in I.5, it is easy to see that X is virtually clean if and only if $\pi_1 X$ is subgroup separable with respect to $\pi_1 \Lambda_x$ for each $x \in X$. Note that there are stronger hypothesis involving Θ_τ of section I.4 that one can place on X in the hope of obtaining that it is virtually $F_n \times F_m$*

Before closing this section, we mention that theorem 4.9 resolves negatively the following question which was posed in [T].

4.11 Question:. *Suppose that K_1 and K_2 are finite simplicial complexes and there is a simplicial complex K which simplicially covers both K_1 and K_2. Must there be a finite simplicial complex which also simplicially covers both K_1 and K_2?*

Theorem 4.9 asserts that both X and the direct product $B_2 \times B_3$ of bouquets of circles, have the same universal cover, but they do not have an isomorphic finite cover. Note that the above quesion has an affirmative answer in case K_1 and K_2 are graphs. See [T].

Section 5: Non-residually finite examples

5.1: Invariant Subgroup.

For use in 5.2, note that if G is finitely generated then the subgroup

$$N = \bigcap_{[G:H]\leq n} H$$

is of finite index and is also fully invariant. To see that it is fully invariant let $\psi : G \to G$ be an endomorphism, then if $[G : H] \leq n$ then $[G : \psi^{-1}(H)] \leq n$ and so

$$N = \bigcap_{[G:H]\leq n} H \subseteq \bigcap_{[G:H]\leq n} \psi^{-1}(H) = \psi^{-1}\Big(\bigcap_{[G:H]\leq n} H\Big) = \psi^{-1}(N)$$

The following theorem (in the case of an automorphism) and its corollary below are due to D. Long and G. Niblo (see [LN]). Their proof is slightly different.

5.2: Theorem: Fixed subgroup of an endomorphism. *Let ϕ be an endomorphism of a finitely generated residually finite group G, and let Fix_ϕ denote the subgroup of elements fixed by ϕ. Then G is Fix_ϕ-separable.*

Proof. Let $a \in G - \mathrm{Fix}_\phi$. Choose $N \not\ni a^{-1}\phi(a)$ to be normal and ϕ-invariant. The map $\phi : G \to G$ induces a map $\bar{\phi} : G/N \to G/N$. By choice of N, the fixed subgroup of $\bar{\phi}$ does not contain $\bar{a}$. It does, however, contain the image of Fix_ϕ, and so we are done. □

5.3: Corollary: Fixed subgroup of Double. *Let $D = G *_{H=\bar{H}} \bar{G}$ be the double of G along H, then if D is residually finite then G is H-separable.*

Proof. Let ι be the involution of D which switches G and $\bar{G}$. It is easy to see that $\mathrm{Fix}_\iota = H$. Now apply the previou theorem. □

5.3b: Remark: Finite order automorphism. Note that if ϕ is an automorphism of finite order, then we do not have to assume that G is finitely generated in theorem 5.2. This is because given a finite index subgroup $H \subset G$, if ϕ has order r, then we can form a ϕ-invariant finite index subgroup N by taking the intersection $\cap_{0\leq i<r}\phi^i(H)$. And the remainder of the proof of theorem 5.2 goes through unchanged. So for instance, corollary 5.3 doesn't require that G is finitely generated.

5.5: Main Theorem: The double of X is not residually finite. *Let D denote the double of X along V. That is, letting $\bar{X}$ denote an isomorphic copy of X, let*

$$D = \big(X \cup \bar{X}\big)/V = \bar{V}.$$

It is easy to see that D is a CSC, see 5.01. The fundamental group $\pi_1 D$ is not residually finite.

Proof. By lemma 4.4 $\pi_1 X$ is not $\pi_1 V$-separable. Therefore by corollary 5.3, $\pi_1 D$ is not residually finite. In particular, the element $(xy^{-1})^{-1}(\overline{xy^{-1}})$ is mapped to the identity in any finite quotient. □

finite index and ι-invariant. Below we obtain a slightly stronger result, but the essential argument is the same. We use $\pi_1(X)'$ to denote an isomorphic copy of $\pi_1(X)$ and denote the action of the isomorphism by $g \mapsto g'$, etc. Also, we use the notation: $d \mapsto \bar{d}$ and $\iota \to \bar{\iota}$ for the effect of the canonical homomorphism $D \to D/M$.

5.4: Theorem: Torsion quotients. *For each $n \geq 1$, $(xy^{-1})^{-1}(xy^{-1})' \in \langle\langle c^n \rangle\rangle$, that is the element $(xy^{-1})^{-1}(xy^{-1})'$ doesn't survive in any quotient of $\pi_1 D$ in which c has finite order.*

Proof. By lemma 4.4 the image of xy^{-1} represents an element of $\pi_1 V$ in the quotient $\pi_1 D \to \pi_1 D / \langle\langle c^n \rangle\rangle$. By symmetry, the element $(xy^{-1})'$ represents the same element of $\pi_1 V$ in the quotient. Therefore $(xy^{-1})^{-1}(xy^{-1})'$ is mapped to the identity in the quotient.

One can make the symmetry mentioned above more precise by using the fixed subgroup proof as in 5.3. Let $\iota : \pi_1 D \to \pi_1 D$ be the obvious involution. And let $\bar{\iota} : \pi_1 D/\langle\langle c^n \rangle\rangle \to \pi_1 D/\langle\langle c^n \rangle\rangle$ be the induced involution. Then $\pi_1 V \langle\langle c^n \rangle\rangle$ is mapped to $\mathrm{Fix}_{\bar{\iota}}$ because both $\bar{\pi_1 V}$ and $\langle\langle \bar{c^n} \rangle\rangle$ are contained in $\mathrm{Fix}_{\bar{\iota}}$.

But if $(xy^{-1})^{-1}(xy^{-1})' \notin \langle\langle c^n \rangle\rangle$ then $(xy^{-1})\langle\langle c^n \rangle\rangle \neq (xy^{-1})'\langle\langle c^n \rangle\rangle$ which means that $\overline{(xy^{-1})} \notin \mathrm{Fix}_{\bar{\iota}}$ which contradicts lemma 4.4. This is because $\pi_V \langle\langle c^n \rangle\rangle$ is mapped to $\mathrm{Fix}_{\bar{\iota}}$ because both $\overline{\pi_1 V}$ and $\langle\langle \bar{c^n} \rangle\rangle$ are contained in $\mathrm{Fix}_{\bar{\iota}}$. □

5.01 Lemma: Gluing CSCs. *Let X and Y be CSCs with a $\mathcal{V}H$-structure, let V_x and V_y be connected components of V_X and V_Y and suppose that there is an isomorphism $\phi : V_x \to V_y$. Then the space $Z = X \cup Y/\phi$ obtained by identifying V_x with V_y using ϕ, is itself a CSC with a $\mathcal{V}H$-structure.*

Proof. To see that Z is a CSC, note that we only have to check this locally. The required fact is that if B and C are complete bipartite graphs, whose vertices are divided into the classes, B_0, B_1 and C_0, C_1 respectively, then if there is a bijection $B_0 = C_0$, then the union of $B \cup C/(B_0 \cong C_0)$ is also a complete bipartite graph.

The fact that the resulting complex is a $\mathcal{V}H$-complex follows because the gluing identifies vertical edges. And so the $\mathcal{V}H$-structure of Z is inherited from X and Y. Slightly more general results about constructing new CSCs from old ones by gluing them together along certain gluing maps can be obtained, the details are similar. □

5.7: Remark. At this stage, it is not difficult to obtain (by starting with X and using several HNN extensions and amalgamated free products) a non-positively curved squared complex which has no finite quotients. However, it seems worthwhile to get a bit more involved and obtain a CSC with no finite covers, this being somewhat remarkable because it is geometrically very similar to $F_n \times F_m$, and yet is algebraically very different.

5.8: Example A We first give an example of a CSC called A which admits a local isometry $X'' \to A$ (where X'' is the first barycentric subdivision of X). Both the vertical and horizontal decompositions of A are amalgamated free products. We describe A through its vertical decomposition.

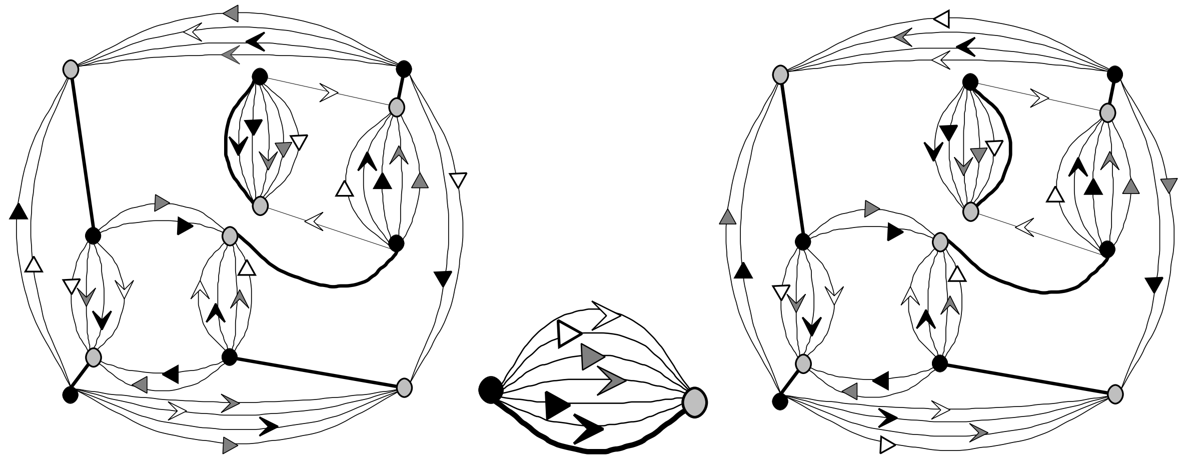

FIGURE 14: INTIMIDATING COMPLEX A. Denote by $\rho_L : L \to B$ and $\rho_R : R \to B$ the covering maps of the central graph B by the large graphs on the left and right. The complex A has underlying graph Γ_A consisting of one edge and two vertices. The vertex graphs Λ_0 and Λ_1 are each isomorpic to B. The single corridor is isomorphic to $L \times I$ and the attaching maps on either side are modelled on ρ_L and ρ_R.

Let $\rho_L : L \to B$ and $\rho_R : R \to B$ denote the covering maps from the left and right colored graphs to copies of the smaller graph in the center of Figure 14. Denote by $\tau : L \to R$ the graph isomorphism between R and L induced by the obvious translation. We form a CSC called A as follows: The underlying graph Γ_A is isomorphic to I. Let the vertex graphs corresponding to the vertices of I, denoted by V_0 and V_1, be isomorphic copies of B. Let V_e denote a graph isomorphic to L and thus to R via τ. The attaching maps of the corridor $V_e \times I$ to each of V_0 and V_1 are just the compositions

$$V_e \times 0 \to L \to B \cong V_0 \qquad \text{and} \qquad V_e \times 1 \to R \to B \cong V_1$$

5.9 The horizontal decomposition of A.

The horizontal decomposition of A is of the form $A \to \Gamma_A^h \cong I$. The horizontal graphs H_0 and H_1 correspond to the two vertices of Γ_A^h. H_0 and H_1 are the two components of H_A. We let H_0 correspond to the horizontal edges in $V_e \times I$ corresponding to grey vertices of V_e, and we let H_1 correspond to the horizontal edges of $V_e \times I$ corresponding to black vertices of V_e. From this it is easy to see that both H_0 and H_1 are bipartite graphs with two vertices and 6 edges.

The unique corridor $H_e \times I$ in the horizontal decomposition of A, has attaching maps represented by $\rho_n : H_e \to H_0$ and $\rho_s : H_e \to V_1$. (These attaching maps are each degree 7 covering spaces).

5.9a Lemma: Normal Closure of H_e. *The normal closure of the image of $\pi_1(H_e)$ in $\pi_1(H_1)$ is all of $\pi_1(H_1)$. Consequently, the normal closure of $\pi_1(H_0)$ in $\pi_1(A)$ is all of $\pi_1(A)$.*

Proof. The covering map $H_e \to H_1$ has degree 7. Since the corresponding subgroup has index 7, it follows that the only subgroups containing it are itself and all of

$\pi_1(H_1)$. Therefore, we will show that $\pi_1(H_e)$ is not a normal subgroup of $\pi_1(H_1)$ and thus since the normal closure $\langle\langle \pi_1(H_e) \rangle\rangle_{\pi_1(H_1)}$ properly contains $\pi_1(H_e)$, we see that the normal closure must be $\pi_1(H_1)$.

To see that the subgroup corresponding to $\pi_1(H_1)$ is not normal, we will show that the associated covering space is not regular. One way of seeing this is to note that there is a closed path of length 2 in H_1 which has both a lift of degree 1 and a lift of degree 6 to the covering space $H_e \to H_1$. The two edges of this closed path correspond to the circled black vertices of figure 15. The two edges of the degree 1 lift of the closed path correspond to the ends of edges marked by a bold **1** in figure 15. The twelve edges in the degree 6 lift of the closed path correspond to the ends of edges marked by a bold **6** in figure 15.

Now to see that $\pi_1(A)$ is the normal closure of $\pi_1(H_0)$, note that

$$\pi_1(A) = \pi_1(H_0) *_{\pi_1(H_e)} \pi_1(H_1).$$

□

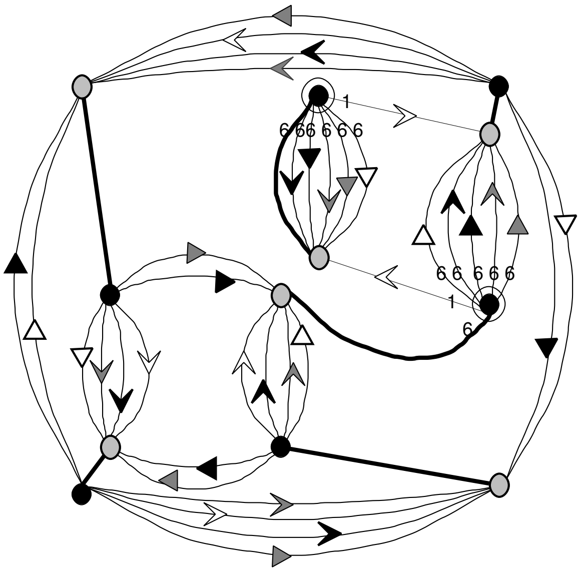

FIGURE 15: $H_e \to H_1$ IS NOT REGULAR. The graph above represents a cross section (near one end) of the vertical corridor $V_e \times I$. The horizontal graph H_1 corresponds to the black vertices in the graph above. The two circled black vertices determine a closed cycle of length 2 in H_1. It is easy to check that the two ends of edges labelled by a bold **1** determine a degree 1 lift of the closed cycle, and that the twelve ends of edges labelled by bold **6** determine a degree 6 lift of the closed cycle.

5.10 Lemma: Local isometry $X \to A$. *Let X'' denote the first barycentric subdivision of X. There is a $\mathcal{VH}$-immersion $X'' \to A$. Furthermore, the path xy^{-1} in X, gets sent to a path in $H_0 \subset A$ which is the concatenation of four distinct edges.*

Let X' denote the vertical subdivision of X, and let X'' denote the horizontal subdivision of X' which is the first barycentric subdivision of X.

In order to describe the $\mathcal{VH}$-immersion $X'' \to A$, we first describe a $\mathcal{VH}$-immersion $X' \to W$, where W is a certain CSC described below. Afterwards, we describe a $\mathcal{V}H$-immersion $W' \to A$ where W' denotes the horizontal subdivision of W. The desired $\mathcal{V}H$-immersion $X'' \to A$ is then obtained by taking the composition $X'' \to W' \to A$.

We first subdivide X vertically by subdividing V_X and adding a horizontal edge in the middle of each square of X. We call the resulting complex X'. The vertical decomposition of the complex X' is illustrated in figure 16. Note that the edges of X' are oriented away from the new vertices of X'.

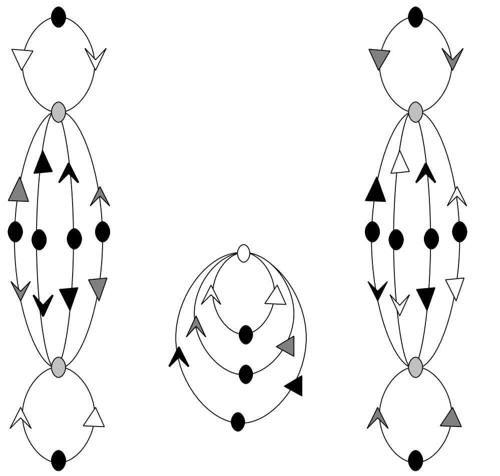

FIGURE 16: THE VERTICAL DECOMPOSITION OF THE VERTICAL SUBDIVISION X' OF X.

The vertical decomposition of W is illustrated in figure 17. There is an obvious $\mathcal{V}H$-immersion $X' \to W$ which preserves the labels of vertical edges of X' and W. One can think of this immersion as the quotient of X' obtained by identifying certain horizontal edges of X'.

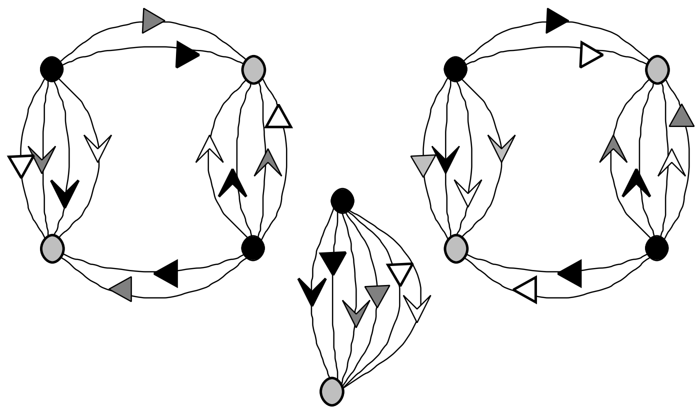

FIGURE 17: THE VERTICAL DECOMPOSITION OF W

We now form the horizontal subdivision W' of W, and describe an immersion $W' \to A$. To understand this immersion, consider the subcomplex A_1 of A obtained by removing the interior of each vertical bold black edge of A, and removing the interior of each square having a bold black vertical edge on its boundary. The vertical decomposition of this subcomplex is depicted in figure 18. Note that there are two corridors each of which is attached at both ends to the subgraphs of V_0 and V_1 not containing any bold black edge. There is an obvious $\mathcal{V}H$-immersion $W' \to A_1$ and thus an immersion $X'' \to W' \to A_1 \to A$.

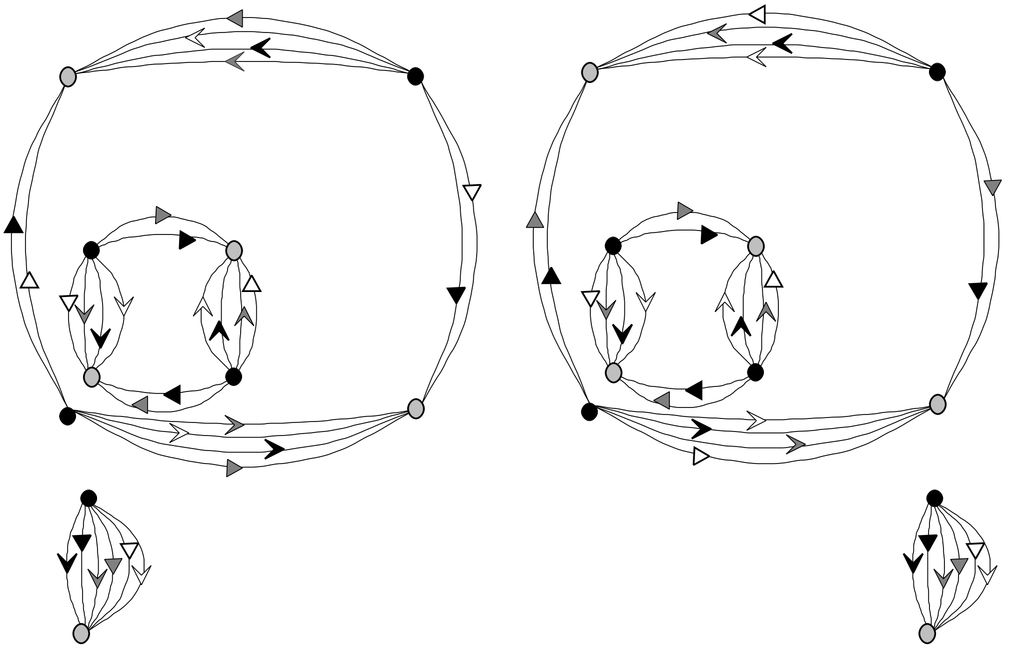

FIGURE 18

Remark: Basepoints. In order to suppress mentioning the basepoints in the arguments concerning π_1 of the various spaces below (all of which are constructed from copies of A), whenever π_1 is mentioned we will assume that an edge of $V_0(A)$ has been contracted and that the resulting vertex is the basepoint.

5.10a: The complex M: Let M denote the CSC which is the double of A along V_0. So

$$M = \big(A \cup \bar{A}\big)/\{V_0 = \bar{V}_0\}.$$

Note that $\Gamma^h_M \cong I$ and that Γ^v_M is a bouquet of 2 edges. See figure 19 for a depiction of the 1-skeleton of M. Let $H_0(M)$ denote the component of H_M which equals $H_0 \cup \bar{H}_0$.

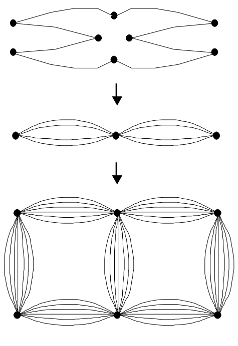

FIGURE 19: THE 1-SKELETON OF M. The lower graph of the figure above represents the 1-skeleton of M. The subgraph $H_0(M)$ of $M^{(1)}$ is represented by the twelve horizontal edges at the top of $M^{(1)}$. The path

ω is represented by an immersion into $H_0(M)$ of the cycle of length 8 illustrated by the top graph of the figure. The immersion $\omega \to H_0(M)$ factors through the subgraph of $H_0(M)$ indicated by the graph in the middle of the illustration, in the obvious way.

Observe that there is a $\mathcal{V}H$-immersion

$$X \bigcup_{V=\bar{V}} \bar{X} \to A \bigcup_{V_0=\bar{V}_0} \bar{A} = M.$$

We let ω denote the path in $H_0(M)$ obtained by taking the projection of the path $(xy^{-1})^{-1}(xy^{-1})'$ in the double of X. Note that by theorem 5.4 the element of $\pi_1 M$ represented by ω is trivial whenever the element a of $\pi_1 M$ has finite order. Furthermore, note for later use that ω consists of the concatenation of two paths in H_0 and $\bar{H}_0$ each consistsing of 4 distinct edges. See figure 19.

5.10b Lemma: Another Normal Closure. $\pi_1 M = \langle\langle \pi_1\big(H_0(M)\big)\rangle\rangle$.

Proof. $\pi_1 M = \langle \pi_1 A \cup \pi_1 \bar{A}\rangle = \big\langle \langle\langle \pi_1 H_0\rangle\rangle_{\pi_1 A} \cup \langle\langle \pi_1 \bar{H}_0\rangle\rangle_{\pi_1 \bar{A}}\big\rangle \subset \langle\langle \pi_1(H_0 \cup \bar{H}_0)\rangle\rangle_{\pi_1 M} = \langle\langle \pi_1(H_0(M))\rangle\rangle_{\pi_1 M}$. $\square$

As in 5.5, M is actually itself a CSC with $\Gamma_M^h \cong I$ (and Γ_M^v is a bouquet of 2 edges). By the $V_0(M)$, we mean the component of V_M containing the basepoint 0 of A. In addition, let $0 \in A \subset M$ be chosen as the basepoint of M.

5.11 Lemma: Generators for $\pi_1(H_0(M))$. *The group $\pi_1(H_0(M))$ is generated by elements represented by closed paths of the form: $ab^{-1}cd^{-1}wx^{-1}yz^{-1}$, where a, b, c, d denote arbitrary distinct edges of H_0 and w, x, y, z denote arbitrary distinct edges of $\bar{H}_0$.*

Proof. Observe that $\pi_1(H_0(M))$ is generated by elements represented by closed paths of length 2 beginning at the basepoint. It is therefore sufficient to show that every closed path of length 2 beginning at the basepoint is the product of two length 8 paths of the form indicated above. By symmetry, it is sufficient to indicate how this is done for one such path of length 2.

$$ab^{-1} = \big(ae^{-1}cd^{-1}wx^{-1}yz^{-1}\big)\big(zy^{-1}xw^{-1}dc^{-1}eb^{-1}\big)$$

and

$$wx^{-1} = \big(wv^{-1}yz^{-1}ab^{-1}cd^{-1}\big)\big(dc^{-1}ba^{-1}zy^{-1}vx^{-1}\big).$$

Note that we are using the fact that H_0 (and hence $\bar{H}_0$) has more than 4 edges. $\square$

5.12: Building the CSC E with no finite quotients.

We will form a CSC by amalgamating several copies of M along their $H_0(M)$ subspaces.

First let K be a graph which is isomorphic to $H_0(M)$ with basepoint corresponding to the basepoint of $H_0(M)$. Let $\{\sigma_1, \dots, \sigma_r\}$ denote the full set of length 8 generators of $\pi_1 K$ which are of the form indicated in lemma 5.11. Choose r copies

$\{M_1, \ldots, M_r\}$ of M. For each i let $H_0(M_i)$ denote the subspace of M_i corresponding to $H_0(M)$. For each i let ω_i be the path in M_i corresponding to the path ω in M.

For each i, choose a graph isomorphism $\psi_i : H_0(M_i) \to K$ which sends ω_i to σ_i.

Now let E be the union of the M_i identified along $H_0(M_i)$ using the maps ψ_i. That is,

$$E = K \cup_{\psi_i} M_i.$$

Note that by lemma 5.11, $\pi_1 K$ is generated by the σ_i and therefore $\pi_1 K$ is generated by the ω_i, because $\sigma_i = \omega_i$ in E. Furthermore, by lemma 5.01 E is clearly a CSC.

5.13: Theorem: E has no finite quotients. *The fundamental group of E has no finite quotients, in fact $\pi_1 E$ has no torsion-quotients.*

Proof. In order to see why this is true, first observe that $\pi_1 E$ is the normal closure of $\pi_1 K$. This is because $\pi_1 E = \langle \cup \pi_1 M_i \rangle$ and by lemma 5.10b for each i, $\pi_1 M_i = \langle\langle \pi_1 H_0(M_i) \rangle\rangle$.

Thus it is sufficient to show that $\pi_1 K$ is trivialized in every finite quotient. But for each i, the element ω_i is trivial in any finite quotient, and so since $\pi_1 K$ is generated by the ω_i, we see that $\pi_1 K$ is trivial in any finite quotient.

The proof that $\pi_1 E$ has no torsion quotients is almost the same, we just use lemma 5.4 instead of lemma 5.5 to show that ω_i is trivialized in any torsion quotient. □

5.14 Question: Does there exists a CSC with (non-trivial) simple π_1.

I guess that one does exist. I can construct examples similar to E above, where (for instance) certain free subgroups embed in every non-trivial quotient.

We close this section with an application of theorem 5.5.

5.15: Theorem: Non-residually finite manifold. *There exists a closed non-positively curved manifold whose fundamental group is not residually finite.*

Proof. In [Hu], it is shown that any compact non-positively curved complex, is a retract of a closed non-positively curved manifold. Applying this to the example D of theorem 5.5 yields the result

□

Something else that should be mentioned is that in [HW], Tim Hsu and myself have applied some of the results of this section to give examples of non-residually finite non-positively curved squares (and triangles) of finite groups.

Section 6: Some Non-Virtually Torsion-Free Examples

6.0 Introduction.

In this section we give some examples of non-virtually torsion-free groups which are biautomatic and/or π_1 of a compact non-positively curved 2-complex. We also show that two well known questions concerning the residual finiteness of word-hyperbolic groups are equivalent. We close this section by outlining a possible approach towards finding a non-residually finite, word hyperbolic group.

6.1 Easy Lemma. *Let $g \in G$ have the property that it belongs to the kernel of any finite quotient, (so G is not residually finite). Suppose that for some n, g is not in the kernel of the quotient $G \to Q = G/\langle\langle g^n \rangle\rangle$. Then Q is not virtually torsion free.*

Proof. We show by contradiction that $\bar{g}$ is contained in every finite index subgroup of Q. Suppose $\bar{g} \notin H$ where $[Q : H] < \infty$, then using the left coset representation, we obtain a finite quotient $Q \to F$ such that $\bar{g}$ is not sent to the identity in F. But now, consideration of the map $G \to Q \to F$ shows that the element $\bar{g} \in Q$ gets sent to the identity. $\square$

The following proposition implies that two well known questions popularized by Gromov, concerning residual properties of word-hyperbolic groups, are equivalent.

6.2: Proposition. *If all word-hyperbolic groups are virtually torsion free, then all word-hyperbolic groups are residually finite.*

Note that the converse is true because a given word-hyperbolic group has only finitely many different torsion elements up to conjugacy. One can deduce this easily from the properties of the action on the Rips Complex of the group (see [Sh]). Thus, if the group is residually finite, then one may separate each of these finitely many conjugacy classes from the identity in a finite quotient, and the kernel of this homomorphism is torsion-free and of finite index.

Before giving the proof, we quote a special case of a result of Delzant, Gromov, Ol'shanskii. See [D],[G1],[O].

6.3: Theorem. *Let G be word-hyperbolic, given an infinite order element $g \in G$, there exists $n > 1$ such that $G/\langle\langle g^n \rangle\rangle$ is itself word-hyperbolic and $\bar{g} \neq 1$.*

Proof of 6.2. What we actually do is show that if one had a non-residually finite word-hyperbolic group, then one could construct a non-virtually torsion-free word-hyperbolic group. So suppose G is word hyperbolic, and suppose that $g \in G$ doesn't survive in any finite quotient of G. Then if g is of finite order, then G itself is not virtually torsion free. If g is of infinite order, then we apply 6.3 to form a quotient $G \to G/\langle\langle g^n \rangle\rangle$ such that $G/\langle\langle g^n \rangle\rangle$ is word hyperbolic, and such that $\bar{g}$ is nontrivial in $G/\langle\langle g^n \rangle\rangle$. But now 6.1 shows that $G/\langle\langle g^n \rangle\rangle$ is not virtually torsion-free.

Or, avoiding 6.3, we may argue as follows: Given a group G with an element $g \in G$ which can't be separated from the identity in any finite quotient, we may construct a new group $\langle G, t \mid [g,t]^n \rangle$ which is easily seen to be word hyperbolic for $n \geq 2$. $\square$

The method used above to construct a non-virtually torsion-free group from a non-residually finite group is a bit trickier to implement in the case of automatic groups or CAT(0) groups which are not word hyperbolic. One difficulty is that a version of 6.3 does not exist in this case.

6.4 Example:. Here we give an example of an automatic group G and an element $b \in G$ such that there are no automatic $G/\langle\langle b^n \rangle\rangle$ Quotients. Consider the group G, presented by:

$$\langle a, b, t \mid [a, b], a^t = a^2 b \rangle.$$

It is isomorphic to

$$\langle a, \beta, t \mid [a, \beta], a^t = \beta \rangle,$$

which is automatic since the standard 2-complex of this presentation is a non-positively curved squared complex (see [GS]).

However, for all $n > 0$ the quotient $Q = G/\langle\langle b^n \rangle\rangle$ is not automatic. One can prove this by first noting that the subgroup generated by $\langle a^n, t \rangle$ is isomorphic to the Baumslag-Solitar group presented by $\langle \alpha, t \mid \alpha^t = \alpha^2 \rangle$. Then by using standard methods, one deduces that Q has an exponential isoperimetric function and is therefore not automatic (see [E]).

Remark:. It is worth mentioning that there is an example of a group G (similar to E of 5.12) which acts properly discontinuously and cocompactly on the direct product of two trees, (so G is biautomatic), such that there is an element $g \in G$ such that for each n the quotient $G/\langle\langle g^n \rangle\rangle$ is trivial. This is worse behavior than example 6.4, and it demonstrates the extent to which word-hyperbolic groups differ from automatic and CAT(0) groups (at least in this sense).

6.5 Modding out by n-th powers for CAT(0) groups. *Suppose that X is a non-positively curved 2-complex, and that $\pi_1 X$ has an element g whose conjugacy class is represented by a closed geodesic path γ which has a point with a large turning angle $\theta > \pi$. Then for large n, the element g survives in the quotient $\pi_1 X \to \pi_1 X/\langle\langle g^n \rangle\rangle$, and $\pi_1 X/\langle\langle g^n \rangle\rangle$ is the fundamental group of a 2-dimensional non-positively orbihedron.*

. We choose n so that for a regular Euclidean polygon P with n sides, and thus interior angle $\frac{(n-2)\pi}{n}$, we would have $\frac{(n-2)\pi}{n} + \theta \geq 2\pi$. In other words, we choose $n \geq \frac{2\pi}{\theta - \pi}$. We first subdivide X so that γ is contained in the 1-skeleton. Now we attach P along γ^n to X, so that each vertex of P corresponds to the point of γ with large turning angle. The group Z_n acts on P in the obvious way, and so since the attaching map of P is Z_n-equivariant we can form the non-positively curved orbihedron $Y = X \cup P/Z_n$. To see that Y is a non-positively curved orbihedron, note that the link condition is still satisfied at the vertex of X where γ has the large turning angle. And that the interior of P, and in particular the cone point, is modeled on P/Z_n. The fundamental group of $X \cup P/Z_n$, is isomorphic to $\pi_1 X/\langle\langle \gamma \rangle\rangle$. (See [H] for information about non-positively curved orbihedra)

In order to use 6.5 to understand the non-virtually torsion free group of example 6.8, we will need the following lemma.

A group G is said to be *Hopfian* if every surjective endomorophism $G \to G$ is an automorphism. It is a theorem of Malćev [LS], that a finitely generated residually finite group is Hopfian. We give a version of this below.

6.7: Lemma: Non-Hopf element. *Let G be a finitely generated group, and let $\phi : G \to G$ be a surjective endomorphism. Suppose that $g \in Kernel(\phi)$, then g is contained in every finite index subgroup of G.*

Proof. Suppose that g where contained outside a finite index subgroup $H \subset G$, then by 5.1 there would be a ϕ-invariant finite index normal subgroup N such that $g \notin N$. But then if we consider the diagram below,

$$\begin{array}{ccc} G & \overset{\phi}{\to} & G \\ \downarrow & & \downarrow \\ G/N & \overset{\bar{\phi}}{\to} & G/N \end{array}$$

since $\bar{\phi}$ is a surjective endomorphism of finite groups, we see that $\bar{\phi}$ is an automorphism, and so $\bar{g}$ is not mapped to the identity by $\bar{\phi}$. But since the diagram is commutative, and we assumed that $\phi(g) = 1_G$, this is a contradiction. □

Here is a relatively simple example which is obtained using 6.5.

6.8 A CAT(0) example.

A relatively simple example of a group which acts properly discontinuously and cocompactly on a CAT(0) space, but is not virtually torsion-free, may be constructed by adapting the following non-Hopfian example of [W1]:

$$G = \langle\ a, b, s, t \ \mid\ [a, b] = 1, a^s = (ab)^2, b^t = (ab)^2\ \rangle.$$

One can see that G is non-Hopfian by considering the endomorphism ψ of G induced by $a \mapsto a^2, b \mapsto b^2, s \mapsto s, t \mapsto t$. A non-trivial element of G which is in the kernel of ψ is $\gamma = \left[(ab)^{s^{-1}}, (ab)^{t^{-1}}\right]$.

Following [W1], it is easy to see that the presentation for G given above is equivalent to the following presentation by using a couple of Tietze transformations:

$$G = \langle\ a, b, c, s, t \ \mid\ c = ab, c = ba, a^s = c^2, b^t = c^2\ \rangle.$$

Now to see that G is π_1 of a non-positively curved 2-complex C, we just use the standard 2-complex of the second presentation given above. We metrize the edges a, b, s, and t as length 1 intervals, and we metrize the edge c as length $\frac{1}{2}$. The cells of the complex C are metrized as Euclidean isosceles triangles and rectangles as in figure 20 below.

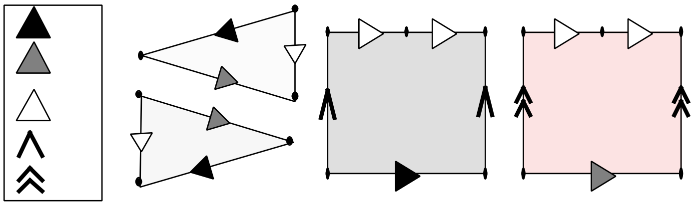

Figure 20: C has a non-positively curved metric

Consider the element $\gamma = \left[(ab)^{s^{-1}}, (ab)^{t^{-1}}\right]$. Subtituting c for ab, we can rewrite γ as $\gamma = \left[c^{s^{-1}}, c^{t^{-1}}\right]$ and thus think of γ as a path in C. Figure 21 illustrates that

γ is homotopic to a closed geodesic, which we will call γ'. To see that γ' is a local geodesic, observe that the turning angle of γ' at each point that is sent to the 0-cell of C is either $\pi + 2\arcsin(1/4) - 2\arctan(1/8) > 3.39 > \pi$ or $\pi + 2\arcsin(1/4) \sim 3.646 > \pi$.

Incidentally, note that the lengths of the edges s and t can be varied independently and the metrics of the corresponding rectangles can be varied accordingly, so that the resulting metric on C is still non-positively curved. As we change the metric on C by increasing the lengths of s and t and the associated rectangles, each of the four turning angles of γ' approaches $\pi + 2\arcsin(1/4) \sim 3.646$.

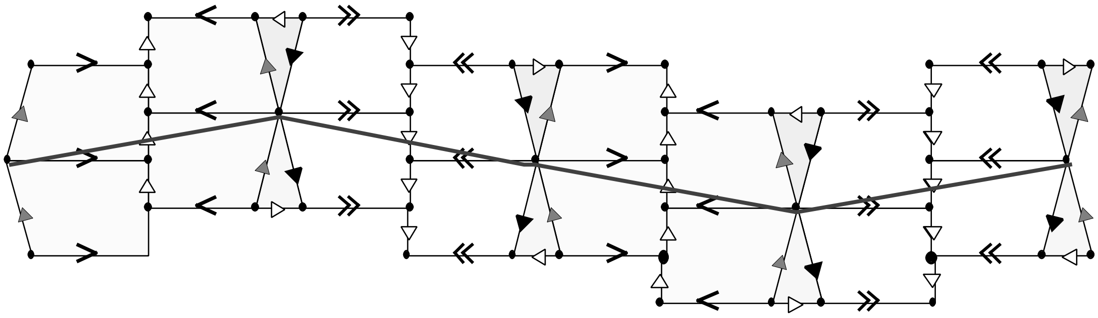

FIGURE 21:

By lemma 6.7, we see that γ is contained in every finite index subgroup of G. By lemma 6.5, we see that the quotient G/γ^{13} is the fundamental group of a non-positively curved 2-complex of groups.

Example 6.9. *Let $n \geq 2$ and let G be as in 6.8, and form the group $J = \langle G, t \mid [g,t]^n \rangle$. Then J is the fundamental group of a non-positively curved orbihedron, and J is not virtually torsion-free.*

Proof. The element $[g,t]$ is contained in every finite quotient because g is. Now we vary the construction of 6.5 slightly to obtain the result. First we add a loop t at the basepoint to C such that t has the same length as γ'. Next we let P be a regular polygon with $4n$ sides, and we attach P to $C \cup t$ along $[\gamma', t]^n$ in the obvious way. It is easy to see that $C \cup t \cup P/Z_{4n}$ is non-positively curved. The resulting orbihedron has $\pi_1 = J$. Since $[g,t]$ survives in the quotient (it has order n) we see by lemma 6.1 that J is not virtually torsion-free. $\square$

The difficulty with the preceding examples is that there are not yet any general results except for certain small cancellation conditions, providing automaticity for groups acting on CAT(0) 2-complexes. Therefore we construct the following example.

6.10: A non virtually torsion-free Biautomatic group. *There is a group G such that*

1. *G is not virtually torsion-free.*
2. *G is a C(4)-T(4) group.*
3. *G acts cocompactly and properly discontinuously on a CAT(0) squared complex.*
4. *G is biautomatic.*

proof. Let D be the CSC of 5.5. And let γ denote the closed horizontal geodesic path $(xy^{-1})^{-1}(x\bar{y}^{-1})$ in the 1-skeleton of D which represents an element of $\pi_1 D$ that is trivial in every finite quotient. We will use γ to refer to the word $(xy^{-1})^{-1}(x\bar{y}^{-1})$.

We form the group $G = \langle \pi_1(D), t / \langle\langle [\gamma, t]^n \rangle\rangle \rangle$. By lemma 5.1, G is not virtually torsion-free.

To see that G is π_1 of a C(4)-T(4) complex, note that G is π_1 of the complex $L = D \cup t, \cup P$ where P is a polygon with $10n$ sides, that is attached to $D \cup t$ along the combinatorial path $[\gamma, t]^n$. It is easy to see that L satisfies the C(4)-T(4) small cancellation condition.

To see that G acts cocompactly on a CAT(0) squared complex we first represent G as the fundamental group of a non-positively curved orbihedron, by taking the quotient $L \to L/Z_n$ where Z_n acts on P as in 6.5. We now subdivide L/Z_n so that it is a squared complex. We subdivide the t edge of L/Z_n so that it is of length 2, and we subdivide the P/Z_n so that it is isomorphic to the direct product $I_2 \times I_4$. The cone point of order n of P/Z_n is therefore at the vertex at its center. The resulting subdivision L' of L, is a non-positively curved squared orbihedron, and therefore its universal cover is a CAT(0) squared complex.

That G is biautomatic follows either by [GS] because G is a C(4)-T(4) group, or by the results of [NR] because G acts properly discontinuously and cocompactly on a CAT(0) squared complex.

6.11: Word-Hyperbolic Suggestion.

One possible way of using the method of this paper to work towards finding a non-residually finite word-hyperbolic group is as follows:

(1) Find a word hyperbolic group G and a malnormal subgroup M such that G is not M-separable.
(2) Prove that M is quasi-convex.
(3) Double G along M to obtain a non-residually finite word hyperbolic group as in 5.3. See [NL], [Pa], and [BF].

Part (3) of the plan is already clear.

Part (2) shouldn't be difficult. My own personal guess is that any finitely generated malnormal subgroup of a word-hyperbolic group is quasi-convex. It is interesting to observe that, were this not the case, then say $M \subset G$ is a f.g. malnormal subgroup of a word-hyperbolic group which is not quasi-convex, then it would have to be distorted, and thus the double $G *_M G$ would be a nice example of a non-word hyperbolic group with no Baumslag-Solitar subgroups.

Part (1) is where the difficulty lies in this approach. To begin with, I know of no (finitely presented) residually finite group G which is not M-separable (for M a f.g., malnormal subgroup). I therefore pose this as as a problem. Examples of word-hyperbolic groups which are not subgroup separable with respect to an arbitrary subgroup are easy to construct, yet in all examples that I am aware of, the subgroup is very far from being malnormal.

6.12: Problem. *Is there a (finitely presented) residually finite group G and a f.g., malnormal subgroup M, such that G is not M-separable?*

Perhaps problem 6.12 is as difficult as the question of residual finiteness for word-hyperbolic groups. Both of there difficulties seem to stem from an avoidance of 'long cylinders'. I hope that some progress on 6.12 will yield some insight. Of course the main goal would be to find a word hyperbolic solution to 6.12

Section 7: A non-approximable Flat

7.0 Intro.

The goal here is to give an example of a compact non-positively curved space Z, which admits an immersed flat plane that is not the limit of a sequence of periodic flat planes. This answers a question of Gromov's (see [G2]), who asked whether or not this might possibly be the case for all CAT(0) spaces.

7.1 Construction of Z.

Let X be the CSC of 2.1 consisting of 6 squares.

We construct a new complex Y as follows: Start with a square s, and then attach four cyclinders, (that is, isomorphic copies of $S^1 \times I$) to s. One such cylinder is attached along each of its edges. The resulting 5 square space is denoted by Y. Note that the complex Y contains immersions of an infinite cross. (That is, an immersion of $[0,1] \times (-\infty, \infty) \cup (-\infty, \infty) \times [0,1] \subset E^2$ with the obvious squared complex structure.)

At each corner of $s \subset Y$, there is a pair of intersecting circles contained in $Y^{(1)}$, which are the boundaries of distinct cylinders. Note that they meet at an angle of $\frac{3\pi}{2}$ in Y. At each of three (NW, SW, & SE) corners of $s \subset Y$ we attach a torus, identifying the pair of circles in the 1-skeleton of a copy of $T^2 = S^1 \times S^1$ with the pair of circles noted above at the respective corner of s. At the fourth (NE) corner of s, we attach a copy of the complex X. Here we identify the pair of circles meeting at the corner of s with the pair of perpendicular circles c and y of X. We name the resulting complex Z. Thus, $Z = T^2 \cup T^2 \cup T^2 \cup Y \cup X$. See figure 22 for a depiction of the 8 squares of $Z - X$ and their gluing patterns.

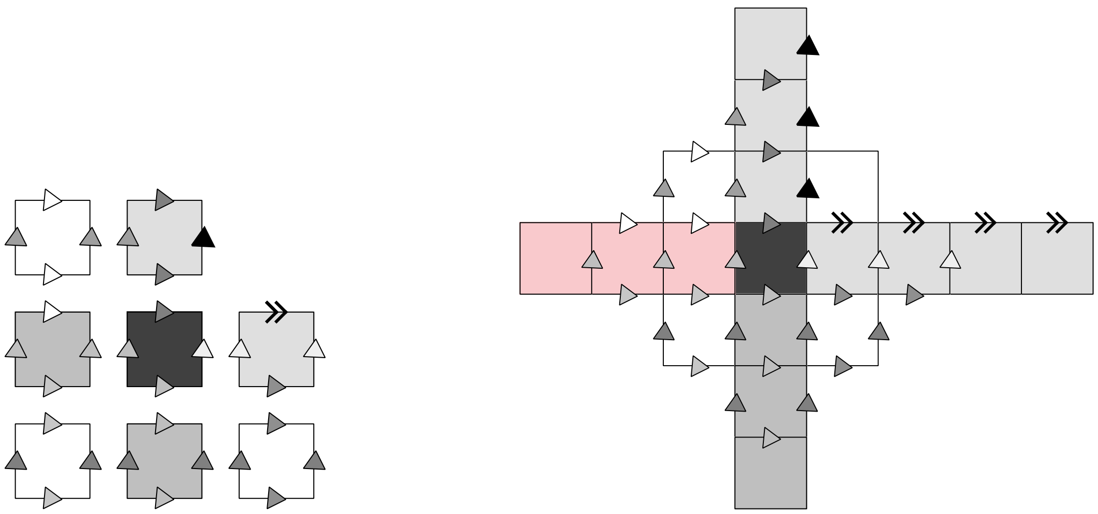

Figure 22: $Z - X$ and Z. The eight squares of the figure on the left are glued together following the gluing pattern to form $Z - X$. To form Z, we add a copy of X at the NE corner, identifying the loops in X labelled by c and y, with the black single and double arrows of the diagram. The figure on the right represents an infinite cross whose convex closure in Z is not approximable by any periodic plane. Note that while the NW, SW, and SE corners of this plane are perioidic (or

rather, quarters of periodic tilings of the plane), the quarter plane on the NE is an aperiodic quarter of Π.

7.2 The planes containing s Note that there are 16 immersed flat planes in Z containing s. These are in correspondence with the 16 different infinite crosses containing s only at their centers, which are immersed in Y. The convex closure in Z of each of these crosses is a flat plane, butz none of these flat planes is periodic. Using an argument similar to the one given below, one can show that these are the *only* flat planes in $\tilde{Z}$ containing s.

7.3 No periodic Approximations. *There is no immersion of a torus $T^2 \to Z$ which contains s. Equivalently, there is no periodic plane in $\tilde{Z}$ containing a copy of s. Thus, no immersed flat planes described in 7.2 is the limit of immersed tori.*

Proof. We argue by contradiction. Suppose that there is an immersed periodic plane Ω containing s. Orienting the tesselation of Ω by the squares of Z appropriately in E^2, one starts at a copy of s and travels north (inside the northern cylinder) until one reaches another copy s_n of s (here we are using that Ω is periodic). Similarly, one travels east from s to reach a square s_e. Traveling north from s_e and east from s_n, one traces out the boundary of a rectangle whose NE corner is a square s_{ne}.

But here we have reached a contradiction, as the inside of this rectangle is tiled by squares in X, yet the boundary of this rectangle reads $\left[a^{\pm n}, y^{\pm m}\right]$. See figure 23. Such a word cannot be trivial in $\pi_1(X)$ because the axes of a and y have convex-closure an anti-torus whereas if they had powers which commuted, their convex-closure would be a periodic plane. (See 4.8 above.) $\square$

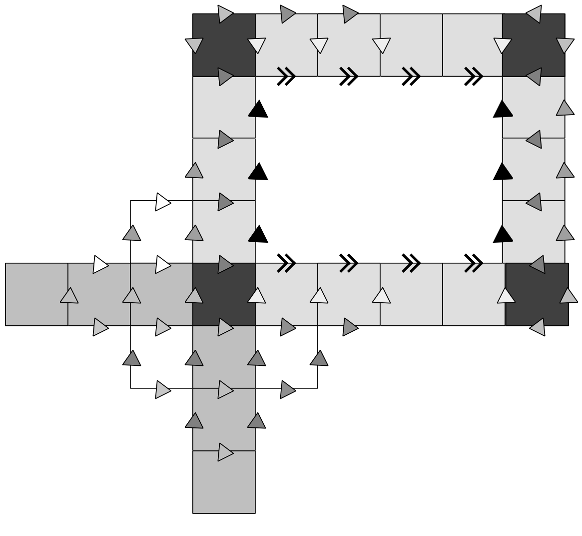

FIGURE 23: CONTRADICTION. The figure above illustrates one of the four possible contradictions which explain why no periodic plane contains the square s.

To show that there is no other plane containing s besides the 16 mentioned above, one looks at the pair of axes intersecting at s in a plane containing s. If this plane is different from each of the 16 mentioned above, then some other s-tile must appear along one of the axes. Looking at the infinite strip in the plane whose corners are the two s tiles, one obtains a similar contradiction to the one obtained above.

7.4 The Algebraic angle vs. Geometric angle.

It is interesting that the elements c and y have axes which intersect perpendicularly in a plane, so that in a sense the geometric angle between these subgroups is $\frac{\pi}{2}$. However, the (algebraic) Gersten-Stallings angle (see [St]) between these subgroups is $\leq \frac{\pi}{3}$. To see this, we need only show that there is no non-trivial relation of the form $c^k y^l c^m y^n = 1_G$.

Now, choosing a basepoint $\tilde{p} \in \tilde{X}^{(0)}$, then, as in 1.5, $\tilde{X}$ is isomorphic with $V_{\tilde{p}} \times H_{\tilde{p}}$. We denote by c^∞ the axis of c in $V_{\tilde{p}}$ (based at $\tilde{p}$), and similary $y^\infty \subset H$. There are unique *orientation preserving* combinatorial retractions $V_{\tilde{p}} \to a^\infty$ and $H_{\tilde{p}} \to y^\infty$. Together they induce an orientation preserving combinatorial retraction

$$\tilde{X} \cong V_{\tilde{p}} \times H_{\tilde{p}} \overset{\rho}{\to} c^\infty \times y^\infty \hookrightarrow \tilde{X}.$$

Since the image of a closed loop in $\tilde{X}$ must be mapped to a closed loop in the plane $c^\infty \times y^\infty$, one sees immediately that the only relations that must be checked are rectangular (i.e., $|k| = |m|$ and $|l| = |n|$), but these are again easily ruled out by the anti-torus Π and the fact that X is both horizontally and vertically directed.

7.5 Remark: Square of Groups & Triangle of groups.

Observe that the complex Z can be thought of in an obvious way as a $K(\pi, 1)$ for a non-positively curved *square of groups* (see [St],[H]) with cyclic edge groups.

Because the algebraic angle between $\langle c \rangle$ and $\langle y \rangle$ is small, it is tempting to form a non-positively curved triangle of groups, where the edge groups are cyclic and the vertex groups are isomorphic to $\pi_1(X)$, and where each edge group is embedded on one (clockwise) side as $\langle c \rangle$ and on the other (counter-clockwise) side as $\langle y \rangle$. This can be done so that the resulting triangle of groups has Z_3 symmetry. This seems to be a curious example, for instance, I wonder if it is the fundamental group of a compact non-positively curved space. It seems as though it might not be, and the the reason why is different from the usual types of problems.

Section 8: 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. P. Brown, [Bro] has obtained some similar results. There are also some generalizations due to J. Meier and myself [MW].

8.1: Definition: Combinatorial Convex Closure. Let $\tilde{X}$ be a CAT(0) 2-complex. And let $S \subset X$ be a subspace. We define the *combinatorial convex closure* $\bar{S}$ of S to be the intersection of all convex subcomplexes of $\tilde{X}$ containing S.

8.2: Definition: Diagonal line. Let $\tilde{X}$ be the universal cover of a CSC X. And let $D \to \tilde{X}$ be an isometrically embedded line. We say that D is a *diagonal line*, provided it is not contained in $X^{(1)}$. We will be interested in diagonal lines which pass through some 0-cell of $\tilde{X}$, and which have a rational slope. We define a *diagonal line segment* similarly.

8.3: Lemma: Convex Closure of Diagonal. *The convex closure of a diagonal line segment in the universal cover $\tilde{X}$ of a CSC X, is a rectangle. The convex closure of a diagonal line in $\tilde{X}$ is a plane.*

proof. This is easy. We take the union of the set of squares intersecting D, and then we use the completeness of the links of $\tilde{X}$ to extend outwards.

Or, we can note that $\tilde{X}$ is the direct product $V_{\tilde{x}} \times H_{\tilde{x}}$, and then note that the convex closure of D is the direct product of its projection onto the two factors. $\square$

8.4: Lemma: Convex closure is periodic plane. *Let X be a compact non-positively curved squared complex. And let $D \to \tilde{X}$ be a periodic diagonal line. Then the convex closure $\bar{D}$ of D is a periodic plane.*

proof. The cyclic group acting on D obviously stabalizes $\bar{D}$, and so $\bar{D}$ is cylindrical. But considering circles in the cylinder parallel to D, or more precisely, lines in $\bar{D}$ parallel to D which intersect $\tilde{X}^{(0)}$, we see that there are only finitely many types. And then since $\bar{D}$ is the convex closure of any such line, we see that $\bar{D}$ admits a symmetry orthogonal to D, and so $\bar{D}$ is periodic. $\square$

8.5: Main Theorem: 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 $\mathcal{V}H$-complex and that X is vertically and horizontally directed. Γ_X^h, inherit a structure as directed graphs from X.

Choose an $n \times n$ region $S_n \subset E$. We will find a periodic diagonal line D whose convex closure $\bar{D}$ contains S_n. By lemma 8.4, $\bar{D}$ is a periodic plane which extends 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 south-west square of S_n, and let *ne* denote the north-east square of S_n. (if $n = 1$ then $sw = ne = S_1$).

In order to find D we first construct an immersion of a certain $\mathcal{V}H$-complex $L \to X$, which is an extension of S_n, that is $S_n \subset L \to X$. The complex L is a subcomplex of the Euclidean plane which consists of the union of S_n and two flat strips of width 1. The first of these strips runs east from ne to some square q, and the second strip runs north from q to a square sw' which is mapped to the same square as sw with the same orientation. (That is, if we identify sw' and sw by the obvious translation, then the two maps to X are identical). See figure 24 for a picture of L.

It is easy to see that the diagonal geodesic segment γ in E connecting the southwest corner of sw to the southwest corner of sw' determines a closed geodesic in X. We consider L to be a subcomplex of $\tilde{X}$, and we let $D = \gamma^{\infty}$, the periodic extension of γ in $\tilde{X}$. It is easy to see that L and hence S_n is contained in $\bar{D}$.

So we only have to construct the immersion $S_n \subset L \to X$ as above. So consider the horizontal corridor C^h_{ne} containing ne and the vertical corridor C^v_{sw} containing sw. Beginning at ne we travel towards a square q of C^h_{ne} until we arrive at a square q of the vertical corridor C^v_{sw} which is oriented horiontally 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 the in the same direction of our original sw.

To see that the orientations of sw and sw' can always be obtained consistently, note that in case a vertical (respectively horizontal) corridor is a cylinder then all vertical (respectively horizontal corridors) are cylinders. $\square$

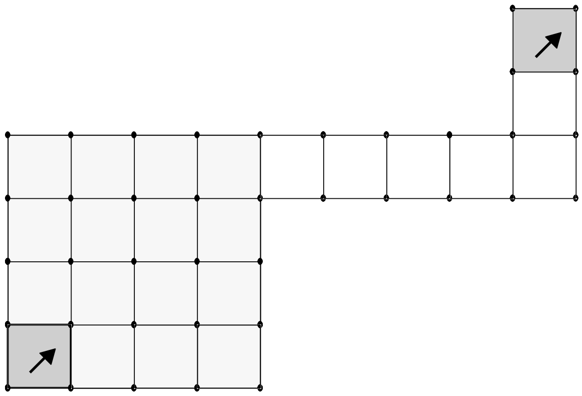

FIGURE 24: THE COMPLEX L. The figure above is meant to suggest the complex L in a particular case. The shaded 4 by 4 square at the bottom left of L, represents S_n. The square sw at the bottom left corner of S_n is labelled with an arrow (so that we can keep track of its orientation in L). The square 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. The complex L, (or rather the infinite periodic extension of its diagonal) determines a periodic plane which contains S_n.

References

[BF] M. Bestvina and M. Feighn, *A combination theorem for Negatively Curved Groups*, J. Differential Geometry **35** (1992), 85-101.

[BH] M.R. Bridson and A. Haefliger, *Metric spaces of nonpositive curvature*, to appear.

[Br] M.R. Bridson, *On the existence of flat planes in spaces of non-positive curvature*, Proceedings of the American Mathematical Society **123** (1994), 223-235.

[Bro] P.R. Brown, *The coarse geometry of non-positively curved polygons of finite groups I: embedded planes*, preprint (1996).

[E] D.B.A. Epstein, J.W. Cannon, D.F. Holt, S.V.F. Levy, M.S. Paterson, and W.P. Thurston, *Word Processing in Groups*, Jones and Bartlett, Boston, MA., 1992.

[Ge] S.M. Gersten, *Problems on Automatic Groups*, Algorithms and Classification in Combinatorial Group Theory (G. Baumslag and C.F. Miller III, eds.), MSRI vol 23, Springer-Verlag, 1992.

[G1] M. Gromov, *Hyperbolic Groups*, Essays in Group Theory (S.M. Gersten, ed.), Springer-Verlag, 1987.

[G2] M. Gromov, *Asymptotic invariants of infinite groups: Geometric Group Theory Volume 2*, London Mathematical Society Lecture Notes 181 (G.A. Niblo and M. Roller, eds.), Cambridge University Press, 1993.

[GrS] Grünbaum and Shepard, *Tilings and patterns*, W.H. Freeman, 1987.

[GS] S.M. Gersten and H. Short, *Small Cancellation Theory and Automatic Groups*, Inventiones Math. **102** (1990), 305–334.

[GS2] S.M. Gersten and H. Short, *Small Cancellation Theory and Automatic Groups: Part II*, Inventiones Math. **105** (1991), 641–662.

[H] A. Haefliger, *Complexes of Groups and Orbihedra*, Group Theory from a Geometrical Viewpoint (E. Ghys, A. Haefliger, and A. Verjovsky, eds.), World Scientific Publishing Co., 1991, pp. 504-542.

[Ha] A. Hatcher, *Algebraic Topology*, Manuscript, 1995.

[Hu] B. Hu, *Retractions of closed manifolds with non-positive curvature*, Proceedings of Ohio topology festival, pp. 136-147.

[HW] T. Hsu and D.T. Wise, *A Non-Residually Finite Non-Positively Curved Square of Finite Groups*, Journal of Pure and Applied Algebra, to appear.

[LN] D.D. Long and G. Niblo, *Subgroup Separability and 3-Manifold Groups*, Mathematische Zeitschrift **207** (1991), 209-215.

[LS] R.C. Lyndon, and P.E. Schupp, *Combinatorial Group Theory*, Springer-Verlag, 1977.

[M] L. Mosher, *Geometry of cubulated 3-manifolds*, Topology (96).

[MW] J. Meier and D.T. Wise, in preparation.

[Ne] L.P. Neuwirth, *Knot Groups, Ann. of Math. Studies*, vol. 56, 1965.

[NR] G. Niblo and L. Reeves, *Non-Positively Curved Cubulated Complexes Automatic*, preprint.

[O] A.Y. Ol'shanskii, *On residualing homomorphisms and G-subgroups of hyperbolic groups*, International Journal of Algebra and Computation (1993), World Scientific Publ. Company.

[Pa] P. Papasoglu, *Geometric Methods in Group Theory*, Columbia U. Thesis (1993).

[Pr] S.J. Pride, *Some Problems in Combinatorial Group Theory*, Lecture Notes in Mathematics 1398 (1989), Springer-Verlag, 147–148.

[R] D. Rolfson, *Knots and Links*, Publish or Perish Inc., Houston, Texas, 1990.

[Sc] P. E. Schupp, *A Survey of Small Cancellation Theory*, Word Problems, North-Holland, 1973.

[SH] H. Short, J. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, and M. Shapiro, *Notes on Word Hyperbolic Groups*, Group Theory from a Geometrical Viewpoint (E. Ghys, A. Haefliger, and A. Verjovsky, eds.), World Scientific Publishing Co., 1991, pp. 3-63.

[St] J. R. Stallings, *Non-Positively Curved Triangles of Groups*, Group Theory from a Geometrical Viewpoint (E. Ghys, A. Haefliger, and A. Verjovsky, eds.), World Scientific Publishing Co., 1991, pp. 491-503.

[T] T.W. Tucker, *Some topological graph theory for topologists: A sampler of covering space constructions*, Topology and Combinatorial Group Theory, proceedings of the fall foliage topology seminar. (P. Latiolais, ed.), Lecture Notes in Mathematics, vol. 1440, Springer-Verlag, pp. 192–207.

[Wa] C.T.C. Wall (ed.), *Homological Group Theory*, Cambridge University Press, Cambridge, 1979, p. 386.

[We] C.M. Weinbaum, *The word and conjugacy problem for the knot group of any tame prime alternating knot*, Proc. Amer. Math. Soc. **22** (1971), 22–26.

[W1] D.T. Wise, *A Non-Hopfian Automatic Group*, Journal of Algebra (96).

[W2] D. T. Wise, *Subgroup separability of cyclic HNN extensions of free groups*, Preprint (1996).

[W3] D.T. Wise, *A non-positively curved squared complex with no finite covers*, Preprint (1995).