The Hrushovski Property, Finite Height, and Maximal Ascents

Brahim Abdenbi

Department of Mathematics and Statistics
McGill University, Montreal

Fall 2022

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy

Abstract

In Chapter 1, we show that any compact nonpositively curved cube complex Y embeds in a compact nonpositively curved cube complex R where each partial local isometry of Y extends to an automorphism of R. We prove a similar result for compact special cube complexes provided that the partial local isometries satisfy certain conditions.

In Chapter 2, we define the directed height of a mapping of graphs and relate it to the "algebraic" height of subgroups. We show that a map has finite directed height if and only if the corresponding mapping torus has negative immersions. We survey related properties and discuss how they relate to one another.

In Chapter 3, we show that given a bi-order $\succ$ on the free group $\mathcal{F}$, every non-periodic cyclically reduced word $W \in \mathcal{F}$ admits a maximal ascent that is uniquely positioned. This provides a cyclic permutation of W' that decomposes as $W' = AD$ where A is an ascent and D is a descent. We show that if D is not uniquely positioned in W, then it must be an internal subword in A. Moreover, we show that when $\succ$ is the Magnus ordering, $D = 1_{\mathcal{F}}$ if and only if W is monotonic.

Résumé

Dans le chapitre 1, nous démontrons que chaque complexe cubique Y qui est compact et à courbure non-positive, s'intégre dans un complexe cubique R qui est compact et à courbure non-positive, tel que chaque isométrie locale et partielle de Y s'éttend à un automorphisme de R. Nous démontrons un résultat similaire pour les complexes cubiques spéciaux sous conditions que les isométries locales satisfassent certains critères.

Dans le chapitre 2, nous définissons la notion de hauteur dirigée d'une application de graphes et nous étudions la relation entre celle-ci et la notion de hauteur d'un sous-groupe. Nous démontrons qu'une application de graphes a une hauteur dirigée finie si et seulement si l'extension HNN correspondante a des immersions négatives. Nous finissons le chapitre par une analyse de propriétés similaires.

Dans le chapitre 3, nous utilisons l'ordre du groupe libre pour démontrer que chaque mot W qui est non-périodique et réduit se décompose comme un produit de deux sous-mots $W = AD$, dont un est uniquement positioné dans W. En particulier, nous montrons que si l'ordre utilisé est celui de Magnus, le mot D est vide si et seulement si W est un mot monotone.

Acknowledgments

I would like to thank my supervisor Daniel T. Wise for his support, guidance, and encouragement. I also extend my thanks to Piotr Przytycki and Frédéric Haglund for their insightful feedback and corrections. This research was made possible through the generous funding of NSERC and the McGill Department of Mathematics and Statistics.

To my friends and family.

Contents

Introduction

This thesis engages with three topics that are presented in three self contained chapters. In Chapter 1, we establish the Hrushovski property for compact special cube complexes. A well-known theorem of Hrushovski [**Hru92**] states that for any finite graph X, there exists a finite graph Z containing X as an induced subgraph with the property that every isomorphism between induced subgraphs of X extends to an automorphism of Z. Since then, various classes of spaces were shown to have this property which came to be known as the *Extension Property for Partial Automorphisms*, or the *Hrushovski Property*. For example, the Hrushovski property was established for finite metric spaces [**Sol05**], structures of a given finite relational language [**Her95**], and various classes of graphs [**Her98**]. So asking the question of whether compact special cube complexes have this property is a natural one to pursue, especially given the important role these spaces played in recent developments in geometric group theory. The challenge in this case is to preserve specialness and to ensure that the original space embeds as a locally convex subcomplex of its extension. Our approach to this question has similarities to the work of Herwig and Lascar who showed in [**HL00**] that the Hrushovski property for certain spaces is related to the profinite topology of free groups. We make extensive use of this relationship, albeit using a different construction, namely the *horizontal quotients* of graphs of spaces. Several statements in this work could be generalized in various directions, but we focus on compact nonpositively curved cube complexes and partial local isometries. This is arguably a natural generalization of the original statement about graphs. To prove the Hrushovski property for compact special cube complexes, we require that the collection of the partial local isometries be "controlled". For example, crossing hyperplanes cannot be mapped to osculating hyperplanes. This mild condition is necessary in order to avoid creating artificial pathologies that would make specialness fail.

The main results in Chapter 1 are:

THEOREM 0.1. *Let Y be a compact nonpositively curved cube complex and let $\mathcal{O}$ be the set of injective partial local isometries of Y. Then Y embeds in a compact nonpositively curved cube complex R where each $\varphi \in \mathcal{O}$ extends to an automorphism $\Phi \in \mathrm{Aut}(R)$.*

Theorem 0.1 could be generalized to the category of compact nonpositively curved spaces instead of cube complexes. See Remark 5.6. But we focus on a generalization towards special cube complexes as stated in Theorem 0.2. This generalization requires the notion of "controlled" embeddings, which generalizes subgraphs of graphs. See Definition 3.15 and Definition 5.7.

THEOREM 0.2. *Let Y be a compact special cube complex and let $\mathcal{O}$ be a controlled collection of injective partial local isometries of Y. Then there exists a compact special cube complex R containing Y as a locally convex subcomplex such that each $\varphi \in \mathcal{O}$ extends to an automorphism $\Phi \in \mathrm{Aut}(R)$.*

Sections 2 and 3 of Chapter 1 provide definitions and background. Section 4 uses subgroup separability of free groups to find finite covers whose horizontal quotients have certain desired properties. In Section 5 we prove Theorem 5.4 and Theorem 5.10.

In Chapter 2, we study properties of a class of HNN extensions of free groups. Specifically, HNN extensions which conjugate a free factor of a free group to another subgroup. These are typically free-by-cyclic groups, but usually the free group is infinitely generated; in general it can be locally-free–by–cyclic. Such groups have been of continual interest in combinatorial and geometric group theory especially because of works of Gilbert Baumslag [**Bau72**] [**Bau93**] and Feighn-Handel [**FH99**]. Our results introduce a new notion of "directed height" of a mapping which is a directed version of the algebraic notion of height given in [**GMRS98**], which we discuss below. The goal of this chapter is to relate the height to "negative immersion" which is an Euler characteristic condition related to locally quasiconvex and coherent groups.

A 2-complex X has *nonpositive immersions* if for any combinatorial immersion $Y \to X$, with Y compact, connected, and collapsed (meaning Y has no free faces), either $\pi_1 Y$ is trivial or $\chi(Y) \leq 0$. A 2-complex X has *negative immersions* if there exists $c > 0$ such that for any combinatorial immersion $Y \to X$, where Y is compact, connected, collapsed, and with no isolated edges, either $\pi_1 Y$ is trivial or $\chi(Y) \leq -c|Y|_2$ where $|Y|_2$ is the number of 2-cells in Y. Let $\mathcal{F} = \langle S \rangle$ be a free group over a set S and let $\mathcal{H} \xrightarrow{\cong} \mathcal{K}$ be an isomorphism of subgroups of $\mathcal{F}$. The *HNN extension* of $\mathcal{F}$ with *associated subgroups* $\mathcal{H}$ and $\mathcal{K}$ is the group $\langle S, t \ : \ t^{-1}\mathcal{H}t = \mathcal{K} \rangle$, where t is a new generator not in S. The HNN extension is *ascending* if $\mathcal{H} = \mathcal{F}$, and *partially ascending* if $\mathcal{H}$ is a proper free factor of $\mathcal{F}$. A group is *coherent* if each of its finitely generated subgroups is finitely presented. Ascending HNN extensions were proved to be coherent in [**FH99**], and are shown to have nonpositive immersions in [**Wis22**].

Given a subgroup $\mathcal{H}$ of $\mathcal{G}$, the *height* of $\mathcal{H}$ in $\mathcal{G}$, denoted by $\mathrm{Height}(\mathcal{H})$, is the supremal number of distinct cosets $\{\mathcal{H}g_i\}_{i \in I}$ such that $\bigcap\limits_{i \in I} \mathcal{H}^{g_i}$ is infinite. This was introduced in [**GMRS98**]. In particular, the height of $\mathcal{H}$ in $\mathcal{G}$ is 1 precisely when $\mathcal{H}$ is a malnormal subgroup.

In this chapter, we define a closely related notion of "directed height" of a mapping. Let H be a subgraph of a finite graph F and let $\psi : H \to F$ be a cellular immersion. Let X be the corresponding mapping torus. Note that $\pi_1 X$ is a partially ascending HNN extension. The *directed height* of ψ is $\overrightarrow{\mathrm{Height}}(\psi) = \inf\{i \ : \ \psi^{-i}(H) \text{ is a forest}\}$, where $\psi^{-i} = (\psi^i)^{-1}$ whenever the *partial composition* ψ^i is defined. See Definition 4.1 and Definition 4.2. Note that $\overrightarrow{\mathrm{Height}}(\psi) < \infty$ if and only if $\mathrm{Height}(\pi_1 H) < \infty$ in π_1X. See Lemma 4.5 for a proof of this statement.

The main results in Chapter 2 are:

THEOREM 0.3. *Let F be a finite graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Then X has negative immersions if and only if ψ has finite directed height.*

In particular, the simplest nontrivial case of Theorem 0.3 is the following special case which is proven in this text as Theorem 3.6:

COROLLARY 0.4. *Let F be a finite graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Suppose $\psi^{-1}(H)$ is a forest. Then X has negative immersions.*

Motivated by our desire to verify Property 11 of Section 5, we note the following consequence of the preceding statements. This does not prove Property 11 since it does not assert that the edge groups in the splitting of $\mathcal{K}$ equal the intersections of $\mathcal{K} \cap \mathcal{H}^g$, for $g \in \pi_1 X$. The following statement is proved in this text as Theorem 4.15:

THEOREM 0.5. *Let F be a finite graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$ with $\overrightarrow{\text{Height}}(\psi) < \infty$. Let $\mathcal{K} \subset \pi_1 X$ be a finitely generated subgroup. Then $\mathcal{K}$ splits over edge groups with uniformly bounded Euler characteristic.*

Our directed height characterization of negative immersions can be re-interpreted in an attractive way using a natural generalization of fully irreducible partial endomorphisms. A partial endomorphism $\psi : \mathcal{H} \to \mathcal{F}$ is *fully irreducible* if there does not exist $n > 0$, a proper free factor $\mathcal{H}' \subset \mathcal{H}$, and $g \in \mathcal{F}$ such that $\psi^n(\mathcal{H}') \subset g^{-1}\mathcal{H}'g$. See Definition 4.1 for the notion of *generalized composition* explaining ψ^n. The standard notion of fully irreducible endomorphism focuses on the case where $\mathcal{H} = \mathcal{F}$ [**BH92**]. Using this language, we show the following statement proved in the text as Theorem 4.12:

THEOREM 0.6. *Let $\mathcal{H}$ be a proper free factor of a finitely generated free group $\mathcal{F}$, and let $\psi : \mathcal{H} \to \mathcal{F}$ be a monomorphism. Let X be the standard 2-complex of the HNN extension of $\mathcal{F}$ with respect to ψ. Then X has negative immersions if and only if ψ is fully irreducible.*

In Section 2 of Chapter 2, we give some background. In Section 3 we prove a special case of the main theorem, namely, that a partially ascending HNN extensions of free groups with malnormal associated subgroups have negative immersions. In Section 4, we prove the main theorems, and in Section 5 we discuss related properties and state three conjectures.

In Chapter 3 we investigate the relationship between orderability and structural property of *nonperiodic* words. A word is *nonperiodic* if it is not a proper power. Weinbaum conjectured in [**Wei90**] that any nonperiodic word W of length > 1 has a cyclic permutation that is a concatenation UV where each of U and V appear exactly once as a prefix of a cyclic permutation of W and W^{-1}. This conjecture was proved by Duncan-Howie in [**DH92**] using the right-orderability of one-relator groups [**BH72**]. It was also proved in [**HN06**] using the

Critical Factorization Theorem [**CV78**] (also see [**CP91**] and [**HN02**]) and basic properties of Lyndon words [**Lyn54**]. This provided the motivation to investigate this question without the complex machinery used above, but rather using only bi-orderability of the free group. For this purpose, we introduce the notions of *ascents/descents* in a cyclically reduced words. Our result provides a partial proof of Weinbaum's conjecture and gives some additional insights into the structure of cyclically reduced nonperiodic words.

A group $\mathcal{G}$ is *bi-orderable* if there is a total order $\prec$ on $\mathcal{G}$ such that for all $f, g, h \in \mathcal{G}$, if $f \prec g$ then $hf \prec hg$ and $fh \prec gh$. A well known result of Shimbireva [**Shi47**] states that the free group on two generators (and so every non-abelian free group) is bi-orderable. This result is sometimes attributed to Vinogradov and Magnus as well [**DNR14**].

Let $X = \{a, b\}$ and $X^* = X \cup X^{-1}$ where $X^{-1} = \{a^{-1}, b^{-1}\}$. A word W in X^* of *length* n, is a sequence of n letters $W = x_1 \cdots x_n$, with $n \geq 0$ and $x_i \in X^*$. The length of W is denoted by $|W| = n$. When $n = 0$, we call W the *empty* word. The word W is *reduced* if it contains no parts xx^{-1}. It is *cyclically reduced* if it is reduced and $x_1 \neq x_n^{-1}$. We henceforth consider only cyclically reduced words in X^*. A word V is a *subword* of W if $W = SVU$ for some reduced words S and U in X^* with $|W| = |S| + |V| + |U|$. The subword V is a *prefix* of W if S is the empty word and a *suffix* if U is the empty word. If neither S nor U is an empty word, then V is *internal* in W. The word W is *periodic* if there exist a cyclically reduced word U and $n > 1$ such that $W = U^n$. It is *nonperiodic* otherwise. A word W' is a *cyclic permutation* of W if $W' = VU$ where $W = UV$, for some prefix U and suffix V of W. Let R_W be the set of all cyclic permutations of W and W^{-1}. A word V is *uniquely positioned* in W if V is the prefix of exactly one element of R_W.

Following Lyndon-Schupp [**LS77**], let $\mathcal{F} = \mathcal{F}(X)$ be the free group with basis X. Then distinct cyclically reduced words in X^* represent distinct elements of $\mathcal{F}$, and the empty word represents the identity element $1_{\mathcal{F}}$. So a bi-ordering $\prec$ of $\mathcal{F}$ induces a bi-ordering $\prec^*$ on the set of cyclically reduced words in X^* in the following sense: given $U, V \in X^*$ representing $g_u, g_v \in \mathcal{F}(X)$, we have $U \prec^* V \iff g_u \prec g_v$. Moreover, the operation of word concatenation from both the right and the left (followed by cyclic reduction) preserves the order. For simplicity, we denote both bi-orders by $\prec$.

Given a bi-order $\prec$ on $\mathcal{F}$, a cyclically reduced word U in X^* is an *ascent* if every non trivial element of $\mathcal{F}$ represented by a prefix of U is $\succ 1_{\mathcal{F}}$ and every element of $\mathcal{F}$ represented by a suffix of U is $\succ 1_{\mathcal{F}}$. The word U is a *descent* if every element of $\mathcal{F}$ represented by a prefix of U is $\prec 1_{\mathcal{F}}$ and every element of $\mathcal{F}$ represented by a suffix of U is $\prec 1_{\mathcal{F}}$. For a given cyclically reduced word W, we define the *maximal ascent* of W to be the greatest ascent over all subwords in R_W with respect to $\succ$.

The main result of Chapter 3 is:

THEOREM 0.7. *Let $\mathcal{F}(X)$ be a bi-ordered free group on the alphabet $X = \{x_1, x_2\}$. Let $W \in \mathcal{F}(X)$ be a cyclically reduced nonperiodic word of length > 1. Then W has a cyclic permutation $W' = AD$ where:*

(1) A is the uniquely positioned maximal ascent in W and D is a descent whenever $D \neq 1_{\mathcal{F}}$.
(2) If D is not uniquely positioned, then it appears as an internal subword of A.
(3) Using the Magnus ordering on $\mathcal{F}$, we have $D = 1_{\mathcal{F}}$ if and only if W is monotonic.

Theorem 0.7 shows that nonperiodic cyclically reduced words have cyclic permutations that factor as concatenations of maximal ascents and descents. The maximal ascents are always uniquely positioned, whereas the descents are not necessarily so. In this sense, this result provides a partial answer to Weibaum's conjecture since only the maximal ascent is guaranteed to be uniquely positioned. We show that when the descents are not uniquely positioned, they appear as internal subwords of the maximal ascents. Consequently, when the descents are not uniquely positioned, they have shorter lengths than maximal ascents. We also show that when the Magnus bi-ordering is used, we can assert that the ascents are equal to W if and only if W is monotonic, in the sense that it only contains letters in X or X^{-1}, but not both. This is all achieved without using any of the machinery of right-orderability/local indicability of one-relator groups, but rather using only bi-orderability of the free group and some basic combinatorial arguments.

1. Contributions to original knowledge

The main results in this text are:

- Chapter 1: Theorem 5.4 and Theorem 5.10;
- Chapter 2: Theorem 3.6, Theorem 4.10, and Theorem 4.15;
- Chapter 3: Theorem 3.2.

CHAPTER 1

The Hrushovski Property For Compact Special Cube Complexes

1. Introduction

A well-known theorem of Hrushovski [**Hru92**] states that for any finite graph X, there exists a finite graph Z containing X as an induced subgraph such that every isomorphism between induced subgraphs of X extends to an automorphism of Z. Since then, various classes of spaces were shown to have this property which came to be known as the *Extension Property for Partial Automorphisms*, or the *Hrushovski Property.* Of particular interest is the work of Herwig and Lascar who showed in [**HL00**] that the Hrushovski property for certain spaces is related to the profinite topology of free groups. In this paper, we make extensive use of this relationship, albeit using a different construction, namely the *horizontal quotients* of graphs of spaces. Although several statements in this work could be generalized in various directions, we focus on compact nonpositively curved cube complexes and partial local isometries. This is arguably a natural generalization of the original statement about graphs.

Our main results are:

THEOREM 1.1. *Let Y be a compact nonpositively curved cube complex and let $\mathcal{O}$ be the set of injective partial local isometries of Y. Then Y embeds in a compact nonpositively curved cube complex R where each $\varphi \in \mathcal{O}$ extends to an automorphism $\Phi \in \mathrm{Aut}(R)$.*

The next statement, proved in the text as Theorem 5.10, requires the notion of "controlled" embeddings, which generalizes subgraphs of graphs. See Definition 3.15 and Definition 5.7.

THEOREM 1.2. *Let Y be a compact special cube complex and let $\mathcal{O}$ be a controlled collection of injective partial local isometries of Y. Then there exists a compact special cube* *complex R containing Y as a locally convex subcomplex such that each $\varphi \in \mathcal{O}$ extends to an automorphism $\Phi \in \operatorname{Aut}(R)$.*

Sections 2 and 3 provide definitions and background. Section 4 uses subgroup separability of free groups to find finite covers whose horizontal quotients have certain desired properties. In Section 5 we prove Theorem 5.4 and Theorem 5.10.

2. Special Cube Complexes

2.1. Cube Complexes. An *n-cube* is a copy of I^n where $I = [-1, 1] \subset \mathbb{R}$ and $n \geq 0$. Its faces are restrictions of some coordinates to -1 or 1. A *cube complex* is a cell complex built from cubes glued together along their faces. The dimension of a cube complex is the supremum of the dimensions of the cubes contained in it.

Let $v = (\epsilon_i)_{i=1}^n$ be a vertex of I^n; so each $\epsilon_i = \pm 1$. The *v-corner* of I^n is the simplex spanned by $\{w_j\}_{j=1}^n$ where each w_j is obtained from v by replacing ϵ_j by $\frac{\epsilon_j}{2}$.

Let X be a cube complex and $C \subset X$ be the image of a map $I^n \to X$. An *x-corner* of C for $x \in X^0$ is the union of images of v-corners of I^n where $v \mapsto x$.

In general, if $J = \prod_{i=1}^{n} \epsilon_i$ is an m-dimensional subcube of I^n where

$$\epsilon_i \in \left\{ \{-1\}, \{1\}, [-1, 1] \right\}$$

then the *J-corner* of I^n is the simplex spanned by the points $\{w_j\}_{j=1}^{n-m}$ obtained from J as follows:

Given the *center of mass* of J, denoted by $v = (t_k)_{k=1}^n$ where

$$t_k = \begin{cases} 0 & \text{if} \quad \epsilon_k = [-1, 1] \\ 1 & \text{if} \quad \epsilon_k = \{1\} \\ -1 & \text{if} \quad \epsilon_k = \{-1\} \end{cases},$$

the point w_j is obtained from v by replacing the j^{th} nonzero coordinate t with $\frac{t}{2}$. Note that each point $w_j \in \{w_j\}_{j=1}^{n-m}$ corresponds to a cube containing J as a codimension-1 subcube.

Let D be a subcube of an n-cube C of X. A *D-corner* of C is the image of a J-corner of I^n under a map $I^n \to X$, where $(I^n, J) \to (C, D)$. The *link* of D in X, denoted by $\operatorname{link}_X(D)$ is the union of all D-corners of cubes containing D. Note that $\operatorname{link}_X(D)$ is a simplex complex and it is a subspace of X but not a subcomplex. We write $\operatorname{link}(D)$ instead of $\operatorname{link}_X(D)$ when X is clear from context.

A cube complex X is *simple* if the link of each cube in X is simplicial.

LEMMA 2.1. *A cube complex X is simple if the link of each cube of X has no loops and no bigons.*

PROOF. Let $D \subset X$ be an n-cube. Let σ_1 and σ_2 be distinct m-simplices in $\operatorname{link}(D)$ with $\sigma_1 \cap \sigma_2 \neq \emptyset$ and $m \geq 1$. If σ_1 is not embedded then $\operatorname{link}(D)$ has a loop. If $\partial\sigma_1 = \partial\sigma_2$, then there exists an $(n+m-1)$-cube $Y \supset D$ such that $\operatorname{link}(Y)$ contains a bigon. Indeed, the case $m = 1$ corresponds to $Y = D$ with a bigon in $\operatorname{link}(D)$. For $m \geq 2$, the m-simplices σ_1 and σ_2 are D-corners of distinct $(n+m+1)$-cubes A_1 and A_2 intersecting along their faces. An $(m-2)$-simplex $\Delta \subset \sigma_1 \cap \sigma_2$ is a D-corner of an $(n+m-1)$-cube $Y \supset D$. Moreover, two distinct $(m-1)$-simplices containing Δ are D-corners of distinct $(n+m)$-cubes $B \supset Y$ and $B' \supset Y$ that are shared faces of A_1 and A_2. We can see that the Y-corners of B and B' are 0-simplices that are boundaries of the 1-simplices corresponding to the Y-corners of A_1 and A_2. □

2.2. Nonpositive curvature. A simple cube complex X is nonpositively curved if it satisfies Gromov's no-$\triangle$ property **[Gro87]**, which requires that 3-cycles in $\operatorname{link}(D)$ bound 2-simplices for each cube $D \subset X$. An equivalent criterion for nonpositive curvature states that a cube complex is nonpositively curved if the links of its 0-cubes are flag. A simplicial complex is *flag* if any collection of $(n+1)$ pairwise adjacent 0-simplices spans an n-simplex.

2.3. Local Isometries. A subcomplex K of a simplicial complex L is *full* if any simplex of L whose 0-simplices lie in K is itself in K. A subcubecomplex $A \subset B$ is *locally convex* if $\operatorname{link}_A(x) \subset \operatorname{link}_B(x)$ is a full subcomplex for every 0-cube $x \in A$.

A map $X \to Y$ of cube complexes is *combinatorial* if open cells are mapped homeomorphically to open cells, where each homeomorphism is an isometry. It is *cubical* if for each $k \leq \dim(X)$, the k-skeleton of X is mapped to the k-skeleton of Y. A combinatorial map $\Phi : X \to Y$ is an *immersion* if the restriction $\operatorname{link}(x) \to \operatorname{link}(\Phi(x))$ is an embedding for each 0-cube $x \in X$. If X and Y are nonpositively curved and $\operatorname{link}(x)$ embeds as a full subcomplex of link $(\Phi(x))$ then Φ is a *local isometry*. Equivalently, a combinatorial locally injective map $\Phi : X \to Y$ of nonpositively curved cube complexes is a local isometry if Φ has *no missing squares* in the sense that if two 1-cubes a_1, a_2 at a 0-cube x map to $\Phi(a_1), \Phi(a_2)$ that bound the corner of a 2-cube at $\Phi(x)$, then a_1, a_2 already bound the corner of a 2-cube at x.

2.4. Immersed Hyperplanes. A *midcube* of an n-cube is the subspace obtained by restricting one coordinate to 0. Note that a midcube of an n-cube is isometric to an $(n-1)$-cube. An *immersed hyperplane* H in a nonpositively curved cube complex X is a component of the cube complex ${}^{M}/_{\sim}$ where M denotes the disjoint union of midcubes of X and $\sim$ is the equivalence relation induced by identifying faces of midcubes under the inclusion map into X. A 1-cube of X is *dual* to H if its midcube is a 0-cube of H. We note that the edges dual to H form an equivalence class generated by *elementary parallelisms* of 1-cubes, where two 1-cubes are *elementary parallel* if they appear on opposites sides of a 2-cube. The *carrier* of H, denoted by $N(H)$, is the cubical neighborhood of H formed by the union of the closed cubes whose intersection with H is nonempty.

2.5. Special Cube Complexes. An immersed hyperplane H in X *self-crosses* if it contains two distinct midcubes from the same cube. It is *two-sided* if the combinatorial immersion $H \to X$ extends to $H \times I \to X$. In this case, the 1-cubes dual to H can be oriented in such a way that any two dual 1-cubes lying in the same 2-cube are oriented in the same direction. An immersed hyperplane that is not two-sided is *one-sided*. H *self-osculates* if it is dual to two oriented 1-cubes that share the same initial or terminal 0-cube and do not form a corner of a 2-cube. Two distinct immersed hyperplanes, H, H', *cross* if they contain distinct midcubes of the same cube. They *osculate* if they are dual to two 1-cubes that share a 0-cube and the 1-cubes do not form a corner of a 2-cube. Two distinct immersed hyperplanes *inter-osculate* if they both cross and osculate. See Figure 1. A nonpositively curved cube complex is *special* if it satisfies the following:

(1) No immersed hyperplane self-crosses;
(2) No immersed hyperplane is one-sided;
(3) No immersed hyperplane self-osculates;
(4) No two immersed hyperplanes inter-osculate.

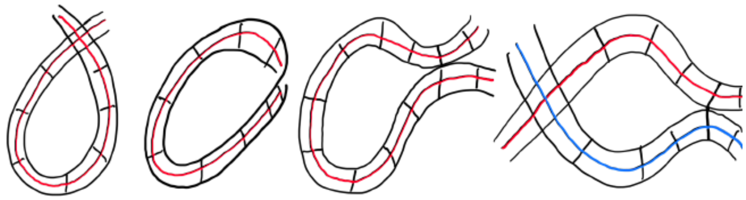

FIGURE 1. From left to right: Self-crossing, one-sidedness, self-osculation, and inter-osculation.

3. Horizontal Quotient of a Graph of Spaces

3.1. Graph of Spaces. An *undirected graph* $\Gamma(V,E)$ is a 1-dimensional CW-complex whose *vertices* and *edges*, denoted by V and E, are the 0-cells and open 1-cells, respectively. There exist two *incidence* maps $\tau_1,\tau_2 : E \to V$ mapping each edge $e \in E$ to its *boundary vertices*, $\tau_1(e)$, $\tau_2(e)$ called initial and terminal vertex, respectively. A *graph of spaces* X with underlying graph $\Gamma(V,E)$, *vertex-spaces* $\{X_v\}_{v\in V}$, and *thick edge-spaces* $\{X_e \times I\}_{e\in E}$ is a topological space X obtained as a quotient of $\{X_v\}_{v\in V}$ and $\{X_e \times I\}_{e\in E}$ in the following manner: for each edge $e \in E$ with boundary vertices $v_1 = \tau_1(e), v_2 = \tau_2(e)$, the corresponding thick edge-space $X_e \times I$ is attached to the vertex-spaces X_{v_1}, X_{v_2} via *attaching maps* which are also denoted by $\tau_1 : X_e \times \{-1\} \to X_{v_1}$ and $\tau_2 : X_e \times \{1\} \to X_{v_2}$. In this text, we always assume the attaching maps of edge-spaces are injective and combinatorial. For simplicity, the subspaces $X_e \times \{-1\} \subset X_{v_1}$ and $X_e \times \{1\} \subset X_{v_2}$ are referred to as *edge subspaces* $X_e \subset X_{v_1}$ and $X_e \subset X_{v_2}$, respectively. The graph $\Gamma(V,E)$ is the quotient of X obtained by mapping X_v to v and $X_e \times (-1,1)$ to e for each $v \in V$ and $e \in E$. We will henceforth denote a graph of spaces X with underlying graph Γ_X by the corresponding canonical quotient map $X \to \Gamma_X$.

3.2. Horizontal Quotient. Let $X \to \Gamma_X$ be a graph of spaces and let E be the edge set of Γ_X. Given an edge $e \in E$, let $\sim_e$ be the equivalence relation on $X_e \times I$ where for all $s,t \in [-1,1]$, we have $(x,t) \sim_e (y,s)$ if and only if $x = y$. Let $X^e = X/\sim_e$ be the corresponding quotient. The *horizontal quotient* of X along the edge e, denoted by $q_e : X \to X^e$, is the quotient map $X \to X^e = X/\sim_e$. In general, if $E' = \{e_1,\dots,e_n\} \subset E$, then the horizontal quotient of X along E' is the quotient $X \to X^{E'} = X/\sim_{E'}$ where $\sim_{E'}$ is the equivalence relation spanned by $\sim_e$ for $e \in E'$. When $E' = E$, we call X^E *the seamless graph of spaces* associated to X and the corresponding map is the *horizontal quotient* which we denote by $q : X \to X^E$. (This terminology was introduced in [**HW22**]). Note that the letter E in X^E is generic in the sense that it refers to the set of all edges of a given graph. For example, given two graphs of spaces $X \to \Gamma_X$ and $Y \to \Gamma_Y$, their horizontal quotients will be denoted by X^E and Y^E, respectively, even when $\Gamma_X \neq \Gamma_Y$. The horizontal quotient q is *strict* if the restriction of q to each vertex-space is an embedding. The *E-parallelism class* of a subset $A \subset X$ is $q^{-1}(q(A))$, that is, the set of all points of X mapping to $q(A)$. When A is a point in an edge-space, $q^{-1}(q(A))$ is the *horizontal graph* associated to A. Note that the restriction of the map $X \to \Gamma_X$ to a horizontal graph in X is an immersion since the attaching maps of edge-spaces are embeddings. In particular, if $X \to \Gamma_X$ is a tree of spaces, then q is strict and the horizontal graphs are trees that intersect each vertex-space of X in at most one point. When X is a graph of cube complexes, an n-cube $C \subset X$ is *vertical* if $q(C)$ is also an n-cube.

REMARK 3.1. In the case of a graph of cube complexes X, we make the following observations:

(1) The quotient X^E is not necessarily a cube complex as cubes of X may be quotiented to simplices in X^E.
(2) When q is strict, it corresponds to an orthogonal projection of cubes (of thick edge-spaces) onto their faces. Then q is cubical and X^E is a cube complex.
(3) When X is a nonpositively curved cube complex and q is strict, the horizontal quotient X^E is not necessarily nonpositively curved.
(4) If $X_1 \to X_2$ is a cover of graphs of cube complexes, then there is an induced map $X_1^E \to X_2^E$ that is not a cover in general.

LEMMA 3.2. *Let $X \to \Gamma_X$ be a graph of cube complexes with a strict horizontal quotient. Then for each immersed hyperplane $U \xrightarrow{f} X^E$, there exists an immersed hyperplane $V \xrightarrow{g} X$, with $f(U) = (q \circ g)(V)$. Furthermore,*

(1) if V is two-sided then so is U;
(2) if $U \xrightarrow{f} X^E$ self-crosses, then $V \xrightarrow{g} X$ self-crosses.

Consequently, if the hyperplanes of X are two-sided/embedded then so are the hyperplanes in X^E.

PROOF. Since q is strict, it is cubical and so X^E is a cube complex. Let $U \xrightarrow{f} X^E$ be an immersed hyperplane. Then the parallelism class of 1-cubes dual to U lifts to a parallelism class of 1-cubes in X. The latter corresponds to an immersed hyperplane $V \xrightarrow{g} X$ that maps onto U, and so $f(U) = (q \circ g)(V)$.

Now suppose $V \xrightarrow{g} X$ is a two-sided immersed hyperplane. If $g(V) \subset X_v$ for some vertex-space X_v, then $q(g(V))$ is two-sided since q is a strict horizontal quotient and thus restricts to an embedding on each vertex-space. If on the other hand, $g(V)$ has nonempty intersection with some edge-space $X_e \times I$ attached to vertex-spaces X_{v_1}, X_{v_2}, then there exist vertical 1-cubes $A_1 \in X_{v_1}$ and $A_2 \in X_{v_2}$ dual to $g(V)$ that lie on opposite sides of a 2-cube $B \subset X_e \times I$. Since V is two-sided, there is a consistent way of orienting A_1 and A_2 so that their initial points lie on the same 1-cube of B. Taking the horizontal quotient along the edge-space $X_e \times I$, induces an orientation on $q(A_1) = q(A_2)$ consistent with the orientation of the vertical 1-cubes of $q(g(V))$. By taking consecutive quotients along all the edge-spaces intersecting $g(V)$, the two-sidedness is preserved at each stage and the claim follows.

Finally, suppose $U \xrightarrow{f} X^E$ is not injective. Then there exists a 2-cube $S \subset X^E$ where $f(U)$ self-intersects. The preimage of S contains a 2-cube where the immersed hyperplane $g(V)$ self-intersects. $\square$

REMARK 3.3. Let $X \to \Gamma_X$ be a graph of cube complexes and let $q : X \to X^E$ be the horizontal quotient. Let $V \xrightarrow{g} X$ be an immersed hyperplane. Then $(q \circ g)(V)$ is not necessarily the image of an immersed hyperplane in X^E. Indeed, not all midcubes of X map to midcubes of X^E. In particular, each immersed hyperplane $g(V) = X_e \times \{0\} \subset X_e \times [-1, 1]$ projects to a subcomplex $q(g(V)) \subset X^E$ that is not a hyperplane.

DEFINITION 3.4. Let $X \to \Gamma_X$ be a graph of cube complexes and $q : X \to X^E$ be the horizontal quotient. Let $x \in X^E$ be a 0-cube and let $q^{-1}(x)$ be the corresponding horizontal graph. Let $\Gamma_0 \subset \Gamma_X$ be the image of $q^{-1}(x)$ under the quotient $X \to \Gamma_X$. Let V_0 and E_0 be the vertices and edges of Γ_0 and let $\{X_v : v \in V_0\}$ and $\{X_e : e \in E_0\}$ be the corresponding vertex-spaces and edge-spaces in X, respectively. Let $\{x_1, \ldots\}$ be the 0-cubes of $q^{-1}(x)$. *The induced graph of links* of x is the graph of spaces $Y \subset X$ with underlying graph $q^{-1}(x)$, whose vertex-spaces are $\operatorname{link}_{X_{v_i}}(x_i)$ and whose edge-spaces are $\operatorname{link}_{X_{e_{ij}}}(x_i)$, where $X_{v_i} \in \{X_v : v \in V_0\}$ is the vertex-space containing x_i and $X_{e_{ij}} \in \{X_e : e \in E_0\}$ are the edge-spaces containing x_i. Note that taking the quotient $X \to X^E$ induces a quotient $Y \to Y^E$ where $\operatorname{link}_{X^E}(x) = Y^E$.

REMARK 3.5. When the edge-spaces of X are embedded locally convex subcomplexes, the edge-spaces of an induced graph of links are embedded full subcomplexes. However, the vertex-spaces of an induced graph of links are not necessarily connected.

LEMMA 3.6. *Let $Y = A \cup_C B$ where A, B are simplicial complexes and C embeds as a full subcomplex in A and B. Then Y is simplicial and A embeds as a full subcomplex of Y.*

PROOF. We show the nonempty intersection of two simplices is a simplex. Let $\sigma_1, \sigma_2 \subset Y$ be simplices with $\sigma_1 \cap \sigma_2 \neq \emptyset$. Each simplex of Y is either in A or in B. Suppose $\sigma_1 \subset A$, $\sigma_2 \subset B$ with $\sigma_1 \not\subset B$ and $\sigma_2 \not\subset A$. Let Z be the set of 0-simplices of $\sigma_1 \cap \sigma_2$ and note that $Z \subset C$. Then Z spans simplices $\delta_1 \subset A$ and $\delta_2 \subset B$. Since C is full in A and B, we see that δ_1 and δ_2 are the same simplex of C. That is, $\sigma_1 \cap \sigma_2$ is a simplex.
To show $A \hookrightarrow Y$ is full, we show that whenever a set of 0-simplices $S \subset A$ spans a simplex Δ, we have $\Delta \subset A$. Indeed, suppose $\Delta \subset B$, then $S \subset C$. But C is full in B and so $\Delta \subset C \subset A$. $\square$

LEMMA 3.7. *Let $Y = A \cup_C B$ where A, B are flag complexes and C embeds as a full subcomplex in A and B. Then Y is flag and A embeds as a full subcomplex of Y.*

PROOF. Y is simplicial by Lemma 3.6. To show flagness, let $K \subset Y$ be an n-clique. We claim that $K \subset A$ or $K \subset B$. We proceed by induction on n. The base case $n = 0$ is trivial. Suppose the claim holds for all cliques of size $\leq n$ and let K be an $(n+1)$-clique. By induction, every proper subclique of K lies in either A or B. Without loss of generality, let $\sigma_1 \in K^0$ be a 0-simplex with $\sigma_1 \notin A$. Then $\sigma_1 \in B$ and for any 0-simplex $\sigma_2 \in K^0$, the 1-simplex $\sigma_1\sigma_2$ lies in B. Indeed, if $\sigma_1\sigma_2$ lies in A, then σ_1 lies in A which is a contradiction. Therefore, $\sigma_2 \in B$ and so $K^0 \subset B$. Moreover, given 0-simplices σ_2 and σ_3 in K^0, the 1-simplex $\sigma_2\sigma_3$ lies in B. To see this, suppose $\sigma_2\sigma_3 \in A$. Then and σ_2 and σ_3 lie in $A \cap B = C$.

But C is full in A and so $\sigma_2\sigma_3 \in C \subset B$. Since B is flag, K bounds a simplex.
Let $K \subset Y$ be a clique such that $K^0 \subset A$. Then by the previous part, $K^1 \subset A$ and it spans a simplex $\Delta \subset A$. Hence A embeds as a full subcomplex of Y. □

LEMMA 3.8. *Let Y be a tree of spaces where each vertex-space is a flag complex and each edge-space embeds as a full subcomplex in its vertex-space. Then Y^E is flag.*

PROOF. Any failure of flagness arises in a quotient of a finite subtree. Therefore, it suffices to prove the claim for finite trees. This follows by induction from Lemma 3.7. Note that a full subcomplex of a full subcomplex is full. □

COROLLARY 3.9. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a tree of nonpositively curved cube complexes where the attaching maps of edge-spaces are injective local isometries. Then $\widehat{X}^E$ is nonpositively curved.*

PROOF. Let x be a 0-cube in $\widehat{X}^E$ and let $Y \to \Gamma_Y$ be the corresponding induced graph of links with underlying graph $q^{-1}(x)$. Since $q^{-1}(x)$ immerses in $\Gamma_{\widehat{X}}$, it is a tree. Then Γ_Y is a tree of flag complexes with embedded full edge-spaces. By Lemma 3.8, the horizontal quotient Y^E is flag, and so is $\operatorname{link}_{\widehat{X}^E}(x)$. □

DEFINITION 3.10. Let X be a graph of cube complexes with horizontal quotient $q : X \to X^E$. Let G be a connected subgraph of a horizontal graph in X. Then:

(1) A hyperplane U *osculates with* G if U is dual to a vertical 1-cube whose initial or terminal 0-cube lies in G.
(2) A two-sided hyperplane U *self-osculates at* G if U is dual to oriented vertical 1-cubes a and b whose initial (or terminal) 0-cubes t_a and t_b lie in G, where $q(a)$ and $q(b)$ are not consecutive 1-cubes of a 2-cube in X^E, and $q(a) \neq q(b)$. When $t_a \neq t_b$, the hyperplane U *remotely self-osculates at* G, in which case we say X has *remote self-osculation*.
(3) A pair of distinct crossing hyperplanes U and V *inter-osculate at* G if there are vertical 1-cubes a and b, with a dual to U and b dual to V, with boundary 0-cubes t_a and t_b lying in G, but $q(a)$ and $q(b)$ are not consecutive 1-cubes of a 2-cube in

X^E. When $t_a \neq t_b$, the hyperplanes U and V *remotely inter-osculate at* G in which case we say X has *remote inter-osculation.*

Note that Definition 3.10 agrees with the definitions in Section 2.5 when $t_a = t_b$.

REMARK 3.11. Remote self-osculations and inter-osculations in X are not actual self-osculations and inter-osculations, but they project to self-osculations/inter-osculations under the horizontal quotient $q : X \to X^E$ whenever q is cubical.

LEMMA 3.12. *Let X be a graph of cube complexes and suppose X is special. If the horizontal quotient $q : X \to X^E$ is cubical and X^E has self-osculation/inter-osculation then X has remote self-osculation/inter-osculation.*

PROOF. Let $U \xrightarrow{f} X^E$ be a self-osculating hyperplane. By Lemma 3.2, there is a hyperplane $V \xrightarrow{g} X$ with $q \circ g\,(V) = f\,(U)$. Since X is special, g and (hence) f are embeddings, and so we can identify U and V with their images. Since the hyperplanes of X are 2-sided, and q is orientation-preserving, the 1-cubes of X^E can be oriented consistently with the orientations of 1-cubes of X. Let a_u and b_u be distinct oriented 1-cubes dual to U that share the 0-cube t where the self-osculation occurs. We can assume without loss of generality that t is the terminal 0-cube of a_u and b_u. Let a_v and b_v be oriented 1-cubes dual to V and mapping to a_u and b_u, respectively. Let $G = q^{-1}\,(t)$ be the horizontal graph mapping to t. Let t_a and t_b be terminal points of a_v and b_v. See Figure 2. Then t_a and t_b lie in G and since X is special, $t_a \neq t_b$. Since $q\,(a_v) = a_u \neq b_u = q\,(b_v)$, the hyperplane V remotely self-osculates at G.

Let U_1 and U_2 be inter-osculating hyperplanes in X^E, and let V_1 and V_2 be the crossing hyperplanes in X mapping to U_1 and U_2, respectively. Suppose the inter-osculation occurs at 1-cubes a_{u_1} and b_{u_2} dual to U_1 and U_2 and meeting at a 0-cube t. Let a_{v_1} and b_{v_2} be 1-cubes dual to V_1 and V_2 and mapping to a_{u_1} and b_{u_2}, respectively. Since X is special, $G = q^{-1}\,(t)$ is nontrivial and contains the distinct 0-cubes t_a and t_b of a_{v_1} and b_{v_2}. Moreover, since a_{u_1} and b_{u_2} do not form a consecutive pair of edges of a 2-cube, V_1 and V_2 remotely inter-osculate at G. $\square$

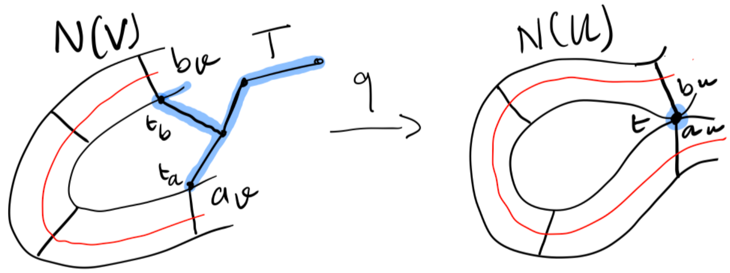

FIGURE 2. The hyperplane V osculates with $G = q^{-1}(t)$ at two points t_a and t_b.

LEMMA 3.13. *Let X be a graph of cube complexes and let G be a horizontal graph in X. Suppose X is special.*

(1) If a hyperplane U of X remotely self-osculates at G, then $G \cap N(U)$ is disconnected.
(2) If crossing hyperplanes U and V of X remotely inter-osculate at G, then $G \cap \big(N(U) \cup N(V)\big)$ is disconnected.

PROOF OF (1). Let U be a remotely self-osculating hyperplane in X. Let a and b be the oriented 1-cubes dual to U with terminal 0-cubes t_a and t_b in G, as in Definition 3.10. Then $t_a, t_b \in N(U)$ and so $G \cap N(U) \neq \emptyset$. We claim that t_a and t_b lie in distinct components of $G \cap N(U)$. Suppose otherwise. Since $t_a \neq t_b$, there is a nontrivial horizontal path $\gamma \to N(U)$ from t_a to t_b. Express γ as a concatenation of horizontal 1-cubes, $\gamma = e_1 \cdots e_n$, where e_1 contains t_a and e_n contains t_b. Since the attaching maps of edge-spaces are injective, any horizontal 1-cube in $N(U)$ lies in a 2-cube whose opposite 1-cube is also horizontal. Since X is special, the hyperplane U does not self-osculate, and so there is a 2-cube $S_1 \subset N(U)$ that contains a and e_1. Let a_1 and e_1' be the 1-cubes in S_1 opposite to a and e_1, respectively. Then $e_1' \subset N(U)$ is horizontal and intersects a in in its initial 0-cube i_a. Furthermore, the 1-cubes $a_1 \subset S_1$ and $e_2 \subset \gamma$ share a 0-cube. By the same argument, there is a 2-cube $S_2 \subset N(U)$ containing a_1 and e_2, where the opposite 1-cube of e_2 is a horizontal 1-cube e_2' that shares a common 0-cube with e_1' and a_1. By induction, there is a sequence of horizontal 1-cubes $e_1', \ldots, e_n'$ in $N(U)$ where e_1' intersects a in its initial 0-cube i_a and where e_n' intersects b in its initial 0-cube i_b. We distinguish two cases. See Figure 3.

<u>Case 1</u>: There is a sequence $e_1', \ldots, e_n'$ that forms a connected horizontal path from i_a to i_b.

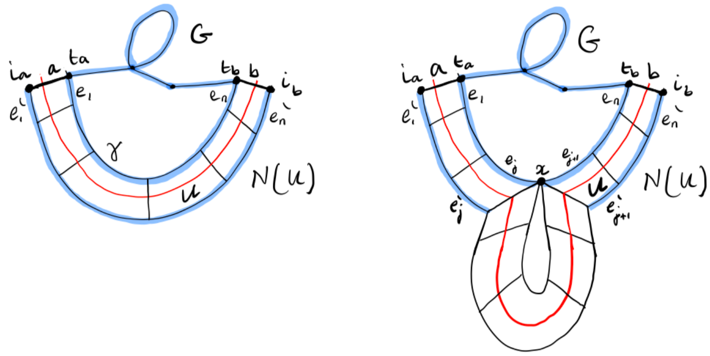

FIGURE 3. Case 1 on the left. Case 2 on the right.

In this case there is a ladder from a to b showing that $q(a) = q(b)$ which is a contradiction.

Case 2: No sequence $e'_1, \ldots, e'_n$ forms a horizontal path from i_a to i_b. Then there is a sequence $e'_1, \ldots, e'_n$ and consecutive 1-cubes e_j and e_{j+1} of γ meeting at a 0-cube x, where the corresponding horizontal 1-cubes e'_j and e'_{j+1} do not intersect. Then x is a point of self-osculation for U which is a contradiction. $\square$

PROOF OF (2). Let U and V be remotely inter-osculating hyperplanes in X. Let a and b be the vertical 1-cubes dual to U and V, respectively, with boundary 0-cubes $t_a \neq t_b$ in G, as in Definition 3.10. We claim that t_a and t_b lie in distinct components of $G \cap (N(U) \cup N(V))$. Suppose otherwise. Then there is a nontrivial horizontal path $\gamma \to (N(U) \cup N(V))$ from t_a to t_b. Let $\gamma = \gamma_u \cdot \gamma_v$, where $\gamma_u \to N(U)$ and $\gamma_v \to N(V)$, and suppose without loss of generality that γ_u is nontrivial. Let $x \in \gamma_u \cap \gamma_v$ and let a_x and b_x be the vertical 1-cubes dual to U and V with boundary 0-cube x. Let $\gamma_u = e_1 \cdots e_n$ be the horizontal path from t_a to x. As in part (1), there is a sequence $e'_1, \ldots, e'_n$ that forms a path in $N(U)$ since otherwise, U self-osculates which is a contradiction. So, a and a_x lie in the same parallelism class.
Similarly, if γ_v is nontrivial, the 1-cubes b and b_x are in the same E-parallelism class. If γ_v is trivial, then $x = t_b$ and $b = b_x$. So we have shown that both a and b are in the same E-parallelism classes as the consecutive 1-cubes a_x and b_x. By assumption, U and V remotely inter-osculate, and so a_x and b_x do not bound a corner of a square. But this means that U and V inter-osculate at x which is a contradiction. $\square$

DEFINITION 3.14. Let X be a cube complex. A subcomplex $X' \subset X$ *self-osculates* if there is a hyperplane U' of X' that extends to a hyperplane U of X dual to a 1-cube whose intersection with X' consists of 0-cubes.

DEFINITION 3.15. A graph of cube complexes is *controlled* if for each thick edge-space $X_e \times I$ attached to vertex-spaces X_{v_1} and X_{v_2}, the following hold for each $i \in \{1, 2\}$:

(1) distinct hyperplanes of X_e extend to distinct hyperplanes of X_{v_i} (wall-injectivity);
(2) non-crossing hyperplanes of X_e extend to non-crossing hyperplanes of X_{v_i} (cross-injectivity);
(3) the edge-space X_e is non self-osculating in X_{v_i}.

LEMMA 3.16. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a controlled tree of cube complexes and suppose each vertex-space of $\widehat{X}$ has embedded hyperplanes. Then each hyperplane U of $\widehat{X}$ dual to a vertical 1-cube splits as a tree of spaces $U \to \Gamma_U$ so that the following diagram commutes:*

$$\begin{array}{ccc} U & \longrightarrow & \widehat{X} \\ \downarrow & & \downarrow \\ \Gamma_U & \longrightarrow & \Gamma_{\widehat{X}} \end{array}$$

Moreover, $\Gamma_U \to \Gamma_{\widehat{X}}$ is an embedding, each hyperplane splits as a tree of connected spaces, each of which embeds in $\widehat{X}$, and consequently, U embeds in $\widehat{X}$ and $U \cap X_v$ is connected for each vertex-space $X_v \subset \widehat{X}$.

PROOF. Let $U \to \Gamma_U$ be a graph of spaces decomposition induced by $\widehat{X} \to \Gamma_{\widehat{X}}$. Since U is dual to a vertical 1-cube, U has nonempty intersection with at least one vertex-space. The vertex-spaces of U are the components of intersections with the vertex-spaces of $\widehat{X}$, and likewise for edge-spaces. Wall-injectivity implies that $U \cap X_v$ is a single hyperplane for each vertex-space X_v intersecting with U. So $\Gamma_U \to \Gamma_{\widehat{X}}$ is an immersion and thus an injection. Therefore, Γ_U is a tree and $U \to \widehat{X}$ is an embedding. $\square$

LEMMA 3.17. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a controlled tree of cube complexes and let X_e be an edge-space in a vertex-space X_v. Let $U \subset \widehat{X}$ be an embedded hyperplane dual to a vertical 1-cube $a \in X_v$. If $a \cap X_e$ consists of 0-cubes then $U \cap X_e = \emptyset$. See Figure 4.*

PROOF. By Lemma 3.16, the intersection $U \cap X_v$ is connected. Since X_e is not self-osculating in X_v, we have $U \cap X_e = \emptyset$. $\square$

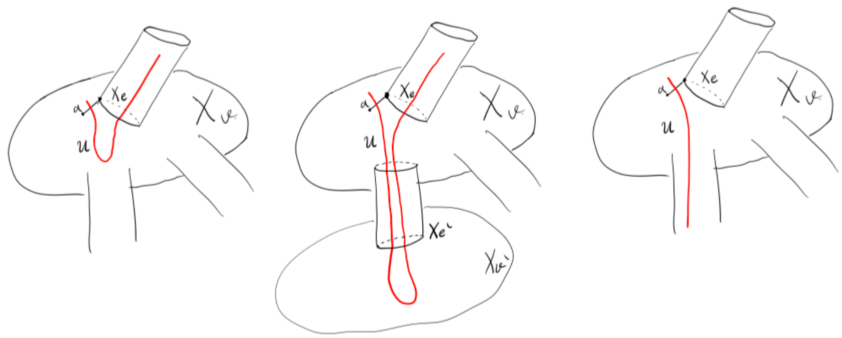

FIGURE 4. The edge-space X_e osculates with the hyperplane U. If $X_e \cap U \neq \emptyset$, then either X_e self-osculates (left) or wall-injectivity fails in some edge-space $X_{e'}$ (middle). $X_e \cap U = \emptyset$ (right).

LEMMA 3.18. *Let $\widehat{X} \xrightarrow{p} \Gamma_{\widehat{X}}$ be a controlled tree of cube complexes with embedded hyperplanes. Then $\widehat{X}$ has no remote self-osculation/inter-osculation.*

PROOF. The horizontal graphs of $\widehat{X}$ are trees that intersect each vertex-space of $\widehat{X}$ in at most one 0-cube. Suppose U is a hyperplane that remotely self-osculates at a horizontal tree T. By Lemma 3.13, $T \cap N(U)$ is not connected. Let K_1 and K_2 be components of $T \cap N(U)$. Let $t_1, t_2 \in T$ be the closest 0-cubes in T with $t_1 \in K_1$ and $t_2 \in K_2$. Let a_1 be the 1-cube dual to U and containing t_1. Let $\gamma = e_1 \cdots e_n$ be the shortest horizontal path in T from t_1 to t_2, where each 1-cube e_i is in the edge-space X_{e_i}. Note that γ is nontrivial since $t_1 \neq t_2$. The 1-cube e_1 with initial 0-cube t_1 does not lie in $N(U)$ for otherwise, the terminal 0-cube of e_1 is in K_1 and is closer to t_2. Then a_1 is not in X_{e_1}. By Lemma 3.17, $U \cap X_{e_1} = \emptyset$. On the other hand, since U splits as a graph of spaces $U \to \Gamma_U$ where Γ_U is a subtree of $\Gamma_{\widehat{X}}$, the image $\left(\gamma \to \Gamma_{\widehat{X}}\right) \hookrightarrow \Gamma_U$ and so $U \cap X_e \neq \emptyset$ which is a contradiction.
Suppose U and V are hyperplanes that remotely inter-osculate at a horizontal tree T. By Lemma 3.13, $T \cap (N(U) \cup N(V))$ is not connected. Let $t_1 \in N(U)$ and $t_2 \in N(V)$ be the closest 0-cubes lying in distinct components of $T \cap (N(U) \cup N(V))$. Let $\gamma_1 = e_1 \cdots e_n$ be the nontrivial horizontal path from t_1 to t_2, where each 1-cube e_i lies in X_{e_i}. Let a_1 and a_2 be the 1-cubes dual to U and V and containing t_1 and t_2, respectively. As in part (1), we have $a_1 \notin X_{e_1}$ and $a_2 \notin X_{e_n}$, and so by Lemma 3.17, we have $U \cap X_{e_1} = \emptyset$ and $V \cap X_{e_n} = \emptyset$. Since $\widehat{X}$ is a tree of spaces, each pair of vertex-spaces is joined by at most one edge-space. Thus,

$U \cap X_{e_1} = \emptyset$ implies $U \cap X_{e_i} = \emptyset$ for all $1 \leq i \leq n$. Similarly, $V \cap X_{e_1} = \emptyset$ for all $1 \leq i \leq n$. Since U crosses V, there is a 0-cube $x \in N(U) \cap N(V)$, and a path $\gamma_2 = f_1 \cdots f_m$ from t_2 to t_1 passing through x, where $f_j \in (N(U) \cup N(V))$. The concatenation $\gamma_1 \cdot \gamma_2$ projects to a closed path in the tree $\Gamma_{\widehat{X}}$. Since γ_1 is horizontal, $\gamma_1 \xrightarrow{p} \Gamma_{\widehat{X}}$ is an embedding. Hence there is a 1-cube $f_j \in \gamma_2$ so that $p(f_j) = p(e_1)$. If $f_j \in N(U)$, then $U \cap X_{e_1} \neq \emptyset$ and if $f_j \in N(V)$, then $V \cap X_{e_1} \neq \emptyset$, both leading to contradictions. $\square$

PROPOSITION 3.19. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a controlled tree of nonpositively curved cube complexes with embedded locally convex edge-spaces. Let $q : \widehat{X} \to \widehat{X}^E$ be the horizontal quotient. If $\widehat{X}$ is special then so is $\widehat{X}^E$.*

PROOF. By Corollary 3.9, $\widehat{X}^E$ is nonpositively curved. Since $\widehat{X}$ is a tree of spaces, the horizontal quotient $q : \widehat{X} \to \widehat{X}^E$ is strict. By Lemma 3.2, each hyperplane of $\widehat{X}^E$ is embedded and two-sided. By Lemma 3.12, self-osculation/inter-osculation in $\widehat{X}^E$ arise from remote self-osculation/inter-osculation in $\widehat{X}$. By Lemma 3.18, $\widehat{X}$ has no remote self-osculation/inter-osculation. $\square$

4. Subgroup Separability

The collection of finite index cosets of a group F forms a basis for the *profinite topology* on F. The multiplication and inversion are continuous with respect to this topology. A subset $S \subset F$ is *separable* if it is closed in the profinite topology. A subgroup $H \subset F$ is separable if and only if H is the intersection of finite index subgroups.

THEOREM 4.1 (Ribes-Zalesskii [**RZ93**]). *Let $H_1, \ldots, H_m$ be finitely generated subgroups of a free group F. Then $H_1 H_2 \cdots H_m$ is closed in the profinite topology.*

It follows that $g_1 H_1 g_2 H_2 \cdots g_m H_m$ is also closed in the profinite topology, for finitely generated subgroups $H_i \subset F$ and $g_i \in F$ with $1 \leq i \leq m$.
Starting with a tree of nonpositively curved cube complexes $\widehat{X} \to \Gamma_{\widehat{X}}$ and using separability properties of the free group action on $\widehat{X}$, we find compact quotients $\widehat{X} \to \overline{X}$ where the the horizontal quotient $\overline{X} \to \overline{X}^E$ is cubical, $\overline{X}^E$ is nonpositively with well-behaved hyperplanes whenever $\widehat{X}$ is controlled and special.

Lemma 4.2. *Let $X \to \Gamma_X$ be a compact graph of cube complexes with one vertex-space Y. Then X has a finite regular cover $\overline{X}$ such that:*

(1) $\overline{X}$ is a graph of spaces whose vertex-spaces are isomorphic to Y;

(2) The restriction of the horizontal quotient $\overline{X} \to \overline{X}^E$ to each vertex-space is injective.

Proof. We find a covering space that splits as a graph of cube complexes with vertex-spaces isomorphic to Y and whose horizontal quotient is strict.

The underlying graph $X \to \Gamma_X$ is a bouquet of circles. Let $\widetilde{\Gamma_X} \to \Gamma_X$ be the universal covering map and let $\widehat{X} \to X$ be the corresponding covering map so that the following diagram commutes:

$$\begin{array}{ccc} \widehat{X} & \longrightarrow & \Gamma_{\widehat{X}} = \widetilde{\Gamma_X} \\ \downarrow & & \downarrow \\ X & \longrightarrow & \Gamma_X \end{array}$$

Then $\pi_1\Gamma_X$ acts freely and cocompactly on $\widehat{X}$. Let $N \subset \pi_1\Gamma_X$ be a finite index normal subgroup, and let $N\backslash\widehat{X} = \overline{X} \to X$ be the covering map induced by $N\backslash\Gamma_{\widehat{X}} = \Gamma_{\overline{X}} \to \Gamma_X$ so that the following diagram commutes:

$$\begin{array}{ccc} \overline{X} & \longrightarrow & \Gamma_{\overline{X}} \\ \downarrow & & \downarrow \\ X & \longrightarrow & \Gamma_X \end{array}$$

Then $\overline{X}$ is a graph of cube complexes where each vertex-space is isomorphic to Y.

We need to choose $\overline{X}$, and thus N, so that no vertex-space has two points in the same parallelism class. In our cubical setting, it is sufficient to ensure that no two 0-cubes of a vertex-space of $\overline{X}$ lie in the same parallelism class. Recall that the attaching maps of edge-spaces are assumed to be injective.

By compactness, there are finitely many 0-cubes $\{C_i\}_{i=1}^n \subset X^0$. Fix a 0-cube C_i and let K_i be the subgroup generated by the horizontal closed paths based at C_i. Then K_i is finitely generated since X is compact. Moreover, since horizontal paths immerse in the underlying graph, the map $X \to \Gamma_X$ induces an injective homomorphism $K_i \to \pi_1\Gamma_X$. Identify K_i with its image. Let $\{\gamma_{ij}\}_{j=1}^m$ be the set of all embedded non-closed horizontal paths between 0-cubes C_i and C_j. Each γ_{ij} maps to an essential closed path in Γ_X and thus represents a nontrivial element $w_{ij} \in \pi_1\Gamma_X$. Furthermore, $w_{ij} \notin K_i$. Indeed, since the attaching maps of edge-spaces are injective, the horizontal graphs immerse in Γ_X, and so the elements represented by γ_{ij} are distinct from elements of K_i. In particular, the products of finitely many cosets $K_i w_{ij} K_j w_{jt} K_t \cdots$ does not contain the identity element. Note that there are finitely many such products of cosets. By Theorem 4.1, there exists a finite index normal subgroup $N \trianglelefteq \pi_1\Gamma_X$ that is disjoint from all such multiple cosets. Let $p : \overline{X} \to X$ be the covering map corresponding to N. Let $Z \subset \overline{X}$ be a vertex-space and $\overline{C_i}, \overline{C_j} \in Z$ be 0-cubes mapping to 0-cubes $C_i, C_j \in Y$. Then $\overline{C_i}$ and $\overline{C_j}$ are not in the same parallelism class of $\overline{X}$. Indeed, if $\overline{\gamma}$ is a horizontal path in $\overline{X}$ from $\overline{C_i}$ to $\overline{C_j}$, then $p\left(\overline{\gamma}\right)$ is a horizontal path γ in X which represents an element in $K_i w_{is} K_s w_{st} K_t \cdots w_{rj} K_j$ where $w_{is}, w_{st}, \ldots, w_{rj}$ are the elements of $\pi_1\Gamma_X$ representing non closed embedded paths between 0-cubes $C_i, C_s, C_t \cdots, C_r, C_j$, respectively. But N contains no such elements, and thus $q : \overline{X} \to \overline{X}^E$ is strict. $\square$

REMARK 4.3. In the proof of Lemma 4.2, we found a covering map $\overline{X} \to X$ that corresponds to a finite index normal subgroup $N \subset \pi_1\Gamma_X$. Note that any normal finite index subgroup $N' \subset N$ induces a finite cover $\overline{X}' \to \overline{X} \to X$ with the same properties as $\overline{X}$. That is, $\overline{X}'$ splits as a graph of spaces with vertex-spaces isomorphic to Y and horizontal quotient $\overline{X}' \to \overline{X}'^E$ is strict.

LEMMA 4.4. *Let $X \to \Gamma_X$ be a graph of nonpositively curved cube complexes and $q : X \to X^E$ be a strict horizontal quotient, where X^E is nonpositively curved. Let Y be a vertex-space of X. If X has no inter-osculating hyperplanes, then $q\left(Y\right) \subset X^E$ is a locally convex subcomplex.*

PROOF. It suffices to show that $q\left(Y\right)$ has no missing squares in X^E. To do so, we show that for each 0-cube $y \in q\left(Y\right)$, the inclusion $\text{link}_{q(Y)}\left(y\right) \subset \text{link}_{X^E}\left(y\right)$ is full.

Let $y \in q\left(Y\right)$ be a 0-cube, and let $e \in \text{link}_{X^E}\left(y\right)$ be a 1-simplex whose boundary 0-simplices x_1 and x_2 lie in $\text{link}_{q(Y)}\left(y\right)$ with $e \notin \text{link}_{q(Y)}\left(y\right)$. Since q is strict, there are consecutive 1-cubes $a_1, a_2 \in q\left(Y\right)$ containing y that are identified with consecutive 1-cubes of a 2-cube $S_e \not\subset q\left(Y\right)$. Since X^E is nonpositively curved, e is the only 1-simplex containing x_1 and x_2 and so a_1 and a_2 are not consecutive 1-cubes of a 2-cube in $q\left(Y\right)$. Then the preimage $q^{-1}\left(S_e\right)$

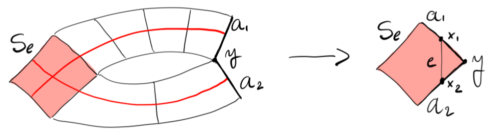

FIGURE 5. Inter-osculation arising from consecutive 1-cubes not bounding a 2-cube in Y, attached to ladders meeting at S_e.

contains inter-osculating hyperplanes since $q^{-1}(a_1)$ and $q^{-1}(a_2)$ contain ladders attached to a 2-cube, contradicting the assumption. See Figure 5. $\square$

The strategy for obtaining $\overline{X}^E$ that is special, is to use multiple coset separability properties of F acting on $\widehat{X}$ to obtain a compact special cube complex $\overline{X}$ whose horizontal quotient $\overline{X}^E$ is special. The property that hyperplanes are embedded and 2-sided is preserved under the map $\overline{X} \to \overline{X}^E$. However, non-inter-osculation and non-self-osculation are not necessarily preserved by $\overline{X} \to \overline{X}^E$. We are therefore forced to revisit and prove a more powerful form of Theorem 5.8 , that provides an intermediate cover $\overline{X}$ for which $\overline{X}^E$ retains all desired properties.

LEMMA 4.5. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a controlled tree of compact nonpositively curved cube complexes with isomorphic vertex-spaces. Let F be a free group acting freely and cocompactly on $\Gamma_{\widehat{X}}$ and $\widehat{X}$, so that $\widehat{X} \to \Gamma_{\widehat{X}}$ is F-equivariant. Suppose $\widehat{X}$ is special. Then there is a finite index normal subgroup $N \subset F$ and a covering map $\widehat{X} \to N\backslash\widehat{X} = \overline{X}$ where $\overline{X}$ splits as a graph of cube complexes whose horizontal quotient $\overline{X}^E$ contains no self-osculating hyperplanes and no inter-osculating hyperplanes.*

PROOF. Since $\widehat{X}$ has no self-crossing hyperplanes, we can identify each immersed hyperplane with its image in $\widehat{X}$. We first find a finite graph of cube complexes $\overline{X}$ whose horizontal quotient has no inter-osculating hyperplanes. We do so by finding an appropriate finite index subgroup $N \subset F$ and taking the quotient $N\backslash\widehat{X} = \overline{X}$. Note that Lemma 4.2 allows us to pass to a finite cover, if necessary, to ensure that the horizontal quotient is a cube complex. By Lemma 3.12, the horizontal quotient $\overline{X}^E$ has inter-osculation if $\overline{X}$ has remote inter-osculation. Remote inter-osculation in $\overline{X}$ occurs if there are crossing hyperplanes A, B of $\widehat{X}$ and an element $g \in F$ such that gB and A osculate with a horizontal graph T in $\widehat{X}$.

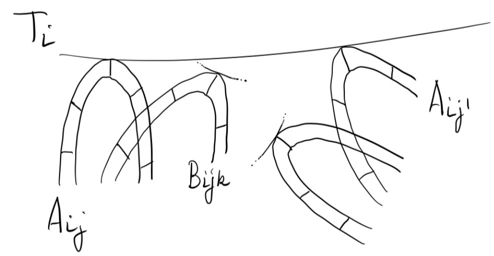

FIGURE 6. The hyperplane B_{ijk} crosses A_{ij} which osculate with the horizontal graph T_i. The element g maps B_{ijk} to A_{ir} which also osculates with T_i

Such an element is called a *remote inter-osculator* at T. Let $\mathcal{R} \subset F$ be the set of remote inter-osculators. We characterize the elements of $\mathcal{R}$ and use subgroup separability to find a finite index subgroup of F that is disjoint from the set $\mathcal{R}$.
By F-cocompactness, there are finitely many F-orbits of horizontal graphs. Let $\{T_i\}_{i=1}^m$ be their representatives. For each tree $T_i \in \{T_i\}_{i=1}^m$ there are finitely many $\text{Stab}\,(T_i)$-orbits of hyperplanes that osculate with T_i. Let $\{A_{ij}\}_{j=1}^{r_i}$ be their representatives. Similarly, for each hyperplane $A_{ij} \in \{A_{ij}\}_{j=1}^{r_i}$, there are finitely many $\text{Stab}\,(A_{ij})$-orbits of hyperplanes crossing A_{ij}. Let $\{B_{ijk}\}_{k=1}^{s_{ij}}$ be their representatives. See Figure 6.

For each B_{ijk} and A_{ir}, if there is an element h_{ijkr} mapping B_{ijkr} to A_{ir}, then the set of all elements g with $gB_{ijk} = A_{ir}$ is:

$$\mathcal{O}_{ijkr} = \text{Stab}\,(A_{ir})\, h_{ijkr}\, \text{Stab}\,(B_{ijk})$$

Furthermore, by precomposing $g \in \mathcal{O}_{ijkr}$ with elements of $\text{Stab}\,(A_{ij})\,\text{Stab}\,(T_i)$, postcomposing g with elements of $\text{Stab}\,(T_i)$, and then taking the union over j, k, r, we obtain the set of remote inter-osculators at T_i:

$$\mathcal{O}_i = \bigcup_{jkr} \text{Stab}\,(T_i)\,\text{Stab}\,(A_{ir})\, h_{ijkr}\,\text{Stab}\,(B_{ijk})\,\text{Stab}\,(A_{ij})\,\text{Stab}\,(T_i)$$

Let $\mathcal{O} = \bigcup_i \mathcal{O}_i$. Each horizontal graph T is a translate of some T_i. Thus each remote inter-osculator at T is conjugate to an element of $\mathcal{O}$. By assumption, $\widehat{X}$ contains no inter-osculating hyperplanes. By Lemma 3.18, $\widehat{X}$ has no remote inter-osculation and thus, $1_F \notin \mathcal{O}$. By Theorem 4.1, the set $\mathcal{O}$ is closed in the profinite topology, and so there exists a finite index normal subgroup N disjoint from $\mathcal{O}$, and hence disjoint from $\mathcal{R}$. Then the horizontal quotient of $N\backslash\widehat{X} \to \left(N\backslash\widehat{X}\right)^E$ has no inter-osculating hyperplanes.

Similarly, to find $\overline{X} \to \overline{X}^E$ with no self-osculating hyperplanes, we use the same method and follow the steps sketched below.

An element $g \in F$ gives rise to self-osculation in $\overline{X}^E$ if $gA = A'$ where A and A' are hyperplanes osculating with the same horizontal graph T. Such elements are called *remote self-osculators* at T. The set of remote self-osculators at T_i is:

$$\mathcal{S}_i = \bigcup_{jr} \operatorname{Stab}(T_i) \operatorname{Stab}(A_{ir})\, h_{ijr} \operatorname{Stab}(A_{ij}) \operatorname{Stab}(T_i)$$

Then any remote self-osculator is conjugate to an element of $\mathcal{S} = \bigcup_i \mathcal{S}_i$. By Lemma 3.18, we have $1_F \notin \mathcal{S}$. Then there exists a finite index normal subgroup $N' \subset F$ such that $N'\backslash\widehat{X} \to \left(N'\backslash\widehat{X}\right)^E$ has no self-osculating hyperplanes and the following diagram commutes:

$$\begin{array}{ccc} \widehat{X} & \longrightarrow & \Gamma_{\widehat{X}} \\ \big\downarrow & & \big\downarrow \\ \overline{X} & \longrightarrow & \Gamma_{\overline{X}} \end{array}$$

The map $\widehat{X} \to (N \cap N')\backslash\widehat{X} = \overline{X}$ provides the desired covering map. $\square$

Remark 4.6. By taking double covers, if necessary, we can ensure that the hyperplanes in $\overline{X}$ are two-sided, which, by Lemma 3.2, means that the hyperplanes of $\overline{X}^E$ are two sided as well.

Up until this point, we have shown how to find a compact quotient where the pathologies precluding specialness do not appear in the horizontal quotients. In the remainder of this section, we show how to ensure that the horizontal quotient is nonpositively curved.

Definition 4.7 (k-corners). For $k \in \{1, 2, 3\}$, a *k-cycle of squares* is a planar complex S_k formed by gluing k squares around a vertex v. A k-cycle of squares has k hyperplanes $\{\alpha_i \mid 1 \leq i \leq k\}$ and k codimension-2 hyperplanes $\{\beta_j \mid 1 \leq j \leq k\}$. Recall that a codimension-2 hyperplane is the intersection of two pairwise intersecting hyperplanes, and the carrier of a codimension-2 hyperplane is the cubical neighborhood containing the intersection. See Figure 7.

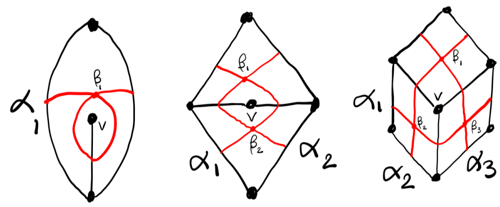

FIGURE 7. 1-cycle, 2-cycle, and 3-cycle of squares with their dual curves.

Let X be a cube complex and $D \subset X$ be an n-cube. An $(n+2)$-dimensional *k-corner of X at D* is a combinatorial immersion $(Z_k, I^n) \to (X, D)$ where $Z_k = S_k \times I^n$ and I^n is identified with $\{v\} \times I^n$ in Z_k.

A k-corner is *empty* if $(Z_k, I^n) \to (X, D)$ does not extend to $(I^{n+3}, I^n) \to (X, D)$. Note that 1-corners and 2-corners are always empty.
We write $Z_k \to X$ when the map $I^n \to D$ is clear from the context.
Note that under the immersion $Z_k \to X$, hyperplanes map to hyperplanes and crossing hyperplanes map to crossing hyperplanes.

REMARK 4.8. Nonpositive curvature can be expressed in terms of k-corners. Specifically, a cube complex is nonpositively curved if it has no empty k-corners. Indeed, if $\mathrm{link}_X(D)$ has a loop [a bigon] then X has a 1-corner [a 2-corner] at D. Furthermore, if the no-$\triangle$ property fails at D, then X has an empty 3-corner at D.
We also note that if X has an empty k-corner at D, then $\mathrm{link}_X(x)$ is not flag for each 0-cube x of D.

DEFINITION 4.9 (k-precorners). Let $X \to \Gamma_X$ be a graph of cube complexes and let $q : X \to X^E$ be the horizontal quotient where q is cubical. Let $Z_k \xrightarrow{\varphi} X^E$ be an $(n+2)$-dimensional k-corner and let $\{A_i = \alpha_i \times I^n \mid 1 \le i \le k\}$ be hyperplanes of $Z_k = S_k \times I^n$ where $\{\alpha_i \mid 1 \le i \le k\}$ are the hyperplanes of S_k. Let $\{B_j = \beta_j \times I^n \mid 1 \le j \le k\}$ be codimension-2 hyperplanes of $Z_k = S_k \times I^n$ where $\{\beta_j \mid 1 \le j \le k\}$ are the codimension-2 hyperplanes of S_k. Let $\left\{H_i \xrightarrow{h_i} X \mid 1 \le i \le k\right\}$ be the immersed hyperplanes of X such that $\varphi(A_i) \subset (q \circ h_i)(H_i)$, and let $N(H_i) \to X$ be their immersed carriers.
The $(n+2)$-dimensional *k-precorner* P_k over the $(n+2)$-dimensional k-corner Z_k is the disjoint union of the corresponding immersed carriers $N(H_i) \to X$ amalgamated along the

carriers of the codimension-2 hyperplanes of H_i that contain the preimages $h_i^{-1}\left(q^{-1}\left(B_j\right)\right)$. See Figure 8. Note that there is a *global* map $h : P_k \to X$ that restricts to h_i on each immersed hyperplane H_i.

A k-precorner $P_k \xrightarrow{h} X$ over a k-corner $Z_k \xrightarrow{\varphi} X^E$ is *empty* if $Z_k \xrightarrow{\varphi} X^E$ is empty. $P_k \xrightarrow{h} X$ is *trivial* if if φ lifts to a combinatorial map $Z_k \to X$ and such that the following diagram commutes:

$$\begin{tikzcd} P_k \arrow[r, "h"] & X \arrow[d, "q"] \\ Z_k \arrow[ur] \arrow[r, "\varphi"] & X^E \end{tikzcd}$$

REMARK 4.10. The map $P_k \xrightarrow{h} X$ induces a splitting of P_k as a graph of spaces as in the following commutative diagram:

$$\begin{tikzcd} P_k \arrow[r] \arrow[d, "h"'] & \Gamma_{P_k} \arrow[d] \\ X \arrow[r] & \Gamma_X \end{tikzcd}$$

Specifically, the vertex-spaces of P_k are the components of the preimages of vertex-spaces of X and the edge-spaces of P_k are the components of the preimages of edge-spaces of X. The graph Γ_{P_k} is the quotient of P_k obtained by identifying vertex-spaces and edge-spaces of P_k with vertices and edges of Γ_{P_k}, respectively. The composition $P_k \to X \to \Gamma_X$ induces a graph morphism $\Gamma_{P_k} \to \Gamma_X$ that maps vertices to vertices and open edges to open edges.

LEMMA 4.11. *Let $\widehat{X} \to \Gamma_{\widehat{X}}$ be a tree of nonpositively curved cube complexes where the attaching maps of edge-spaces are injective local isometries. Let $\widehat{X} \to \widehat{X}^E$ be the horizontal quotient and let $P_k \to \widehat{X}$ be a k-precorner over a k-corner $Z_k \xrightarrow{\varphi} \widehat{X}^E$. Then P_k is trivial and hence nonempty.*

PROOF. Let $T \subset \widehat{X}$ be a minimal connected subtree of spaces containing k cubes $\{C_i \subset P_k\}_{i=1}^k$ that map onto $\varphi\left(Z_k\right)$. Then T is finite since any k cubes mapping onto $\varphi\left(Z_k\right)$ must lie in a finite connected subcomplex of $\widehat{X}$. Note that the minimality is under inclusion and over all possible collections of k cubes mapping onto $\varphi\left(Z_k\right)$. Let $T \to \Gamma_T$ be the underlying tree. We claim that Γ_T is a vertex. Note that if $k = 1$ then there is only one cube that lies in a single vertex-space which by the minimality of T, implies that Γ_T is a vertex. So we

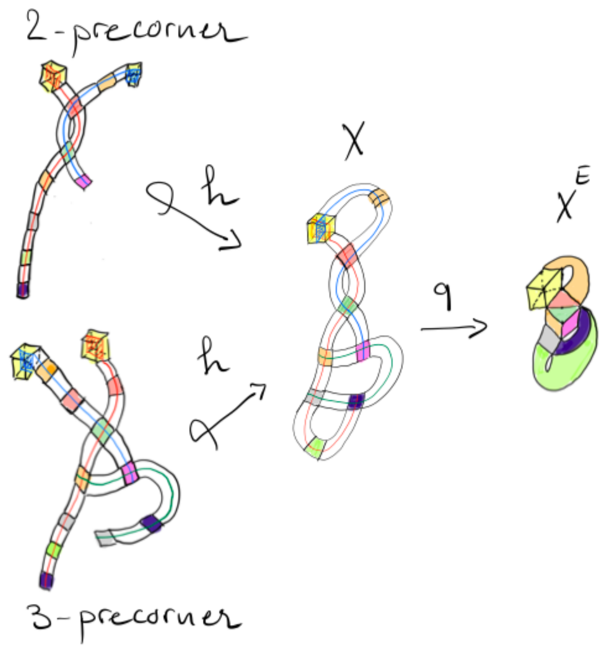

FIGURE 8. A 2-precorner and a 3-precorner.

can assume $2 \leq k \leq 3$. Suppose that Γ_T has a spur e incident on vertices v_1 and v_2, where $\deg(v_1) = 1$. Let T_e be the corresponding edge-space attached to the vertex-spaces T_{v_1} and T_{v_2}. By the minimality of T, we can assume without loss of generality that T_{v_1} contains exactly one cube C_i. There exist distinct immersed hyperplanes $H_1 \to \widehat{X}$ and $H_2 \to \widehat{X}$ that cross in C_i and extend to T_{v_2} through T_e. Since the attaching maps are local isometries, C_i must be in the edge-space. But in that case, the edge-space $T_e \times [-1, 1]$ contains $C_i \times [-1, 1]$ and so the vertex-space T_{v_2} contains $C_i \times \{-1\}$. Therefore, there exists a proper subtree $T' \subset T$ containing k cubes mapping onto $\varphi(Z_k)$, contradicting the minimality of T.
Since T is finite and has no spurs, it is a vertex-space. Moreover, $\widehat{X}$ is a tree of spaces, and so the restriction of the horizontal quotient $q|_T$ in $\widehat{X} \to \widehat{X}^E$ is an isomorphism. This provides the required map $Z_k \to X_v \subset \widehat{X}$, for some vertex-space X_v. So P_k is trivial. By assumption, the vertex-spaces of $\widehat{X}$ are nonpositively curved. By Remark 4.8, Z_k (and hence P_k) is a nonempty k-corner (k-precorner). $\square$

DEFINITION 4.12. Let $X \to \Gamma_X$ be a graph of cube complexes and let F be a group acting on X. Given $k \in \{1, 2, 3\}$, a *k-chain* is an ordered $(k+1)$-tuple of distinct immersed hyperplanes $(H_t)_{t=0}^k$ where H_{t-1} crosses H_t for all $1 \leq t \leq k$. See Figure 9.

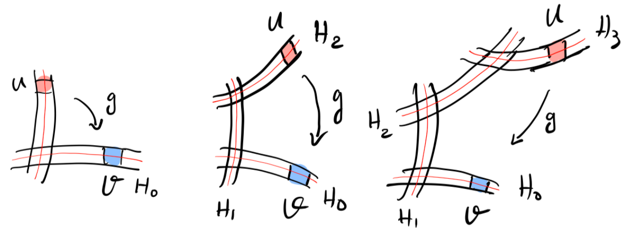

FIGURE 9. From left: 1-chain, 2-chain, and 3-chain.

An element $g \in F$ is a *closing element* if g maps H_k in some k-chain $(H_t)_{t=0}^k$ to H_0 giving rise to an empty k-precorner. We say $(H_t)_{t=0}^k$ is *closed* by g.

REMARK 4.13. Given the setting of Lemma 4.11, let F be a free group acting freely and cocompactly on $\Gamma_{\widehat{X}}$, and thus on $\widehat{X}$, provided that the vertex-spaces are compact and isomorphic. If for some subgroup $G \subset F$, the quotient $G \setminus \widehat{X}$ has an empty k-precorner, then G contains a closing element of some k-chain in X. Note that closing elements map codimension-2 hyperplanes to codimension-2 hyperplanes.

DEFINITION 4.14. Let B be a compact bouquet of circles and let $F = \pi_1 B$ act freely and cocompactly on the universal cover $\widetilde{B}$. Let $\widehat{X} \to \Gamma_{\widehat{X}} = \widetilde{B}$ be a tree of compact isomorphic cube complexes. Then the free cocompact action $F \curvearrowright \widetilde{B}$ induces a free cocompact action $F \curvearrowright \widehat{X}$. We fix a finite collection of immersed hyperplanes $L = L_0 \cup L_1 \cup L_2 \cup L_3$ where:
$L_0 = \{H_1, \ldots, H_{n_0}\}$ are F-representatives of hyperplanes;
$L_1 = \bigcup_i \{H_{i1}, \ldots, H_{in_i}\}$ are $\mathrm{Stab}\,(H_i)$-representatives of hyperplanes crossing H_i, for $1 \leq i \leq n_i$;
$L_2 = \bigcup_{i,j} \{H_{ij1}, \ldots, H_{ijn_{ij}}\}$ are $\mathrm{Stab}\,(H_{ij})$-representatives of hyperplanes crossing H_{ij}, for $1 \leq i \leq n_i$ and $1 \leq j \leq n_{ij}$; and
$L_3 = \bigcup_{i,j,t} \{H_{ijt1}, \ldots, H_{ijtn_{ijt}}\}$ are $\mathrm{Stab}\,(H_{ijt})$-representatives of hyperplanes crossing H_{ijt}, for $1 \leq i \leq n_i$, $1 \leq j \leq n_{ij}$, and $1 \leq t \leq n_{ijt}$.

Let $\mathcal{C}$ be the set of all k-chains of hyperplanes of L. For each hyperplane $A \in L$, there are finitely many $\mathrm{Stab}\,(A)$-representatives of codimension-2 hyperplanes in $N\,(A)$. For each k-chain $C = (A_t)_{t=0}^k$, with $A_t \in L_t$, let J_C be the set of elements of F that map the chosen $\operatorname{Stab}(A_k)$-representatives of codimension-2 hyperplanes of A_k to the chosen $\operatorname{Stab}(A_0)$-representatives of codimension-2 hyperplanes of A_0. Note that J_C is finite.

LEMMA 4.15. *Let B be a compact bouquet of circles and let $X \to B$ be a graph of cube complexes with one compact vertex-space. Let $\widehat{X} \to \Gamma_{\widehat{X}} = \widetilde{B}$ be the tree of cube complexes where $\widetilde{B} \to B$ is the universal covering map such that the following diagram commutes:*

$$\begin{array}{ccc} \widehat{X} & \longrightarrow & \Gamma_{\widehat{X}} = \widetilde{B} \\ \big\downarrow & & \big\downarrow \\ X & \longrightarrow & B \end{array}$$

Let $F = \pi_1 B$ be the free group acting freely and cocompactly on $\widehat{X}$ and let $g \in F$ be a closing element. Then there exist a k-chain $C = (A_t)_{t=0}^k$ with $A_t \in L_t$, and $f \in F$ such that

$$f^{-1}gf \in \operatorname{Stab}(A_0)\, J_C \operatorname{Stab}(A_k) \operatorname{Stab}(A_{k-1}) \cdots \operatorname{Stab}(A_0)$$

PROOF. Let $(B_t)_{t=0}^k$ be the k-chain in $\widehat{X}$ closed by g. Let V and U be codimension-2 hyperplanes in $N(B_0)$ and $N(B_k)$, respectively, such that $gU = V$. The hyperplane B_0 lies in the F-orbit of some representative $A_0 \in L_0$ and so $B_0 = fA_0$ for some $f \in F$. Let $\{V_s \mid 1 \leq s \leq n_0\}$ be $\operatorname{Stab}(A_0)$-representatives of orbits of codimension-2 hyperplanes in $N(A_0)$. Let $a_0 \in \operatorname{Stab}(A_0)$ and $V_s \in \{V_s \mid 1 \leq s \leq n_0\}$ such that $V = fa_0V_s$. Let $A_1 \in L_1$ and $a_0' \in \operatorname{Stab}(A_0)$ such that $B_1 = fa_0'A_1$.

<u>*Case $k=1$*</u>: Let $\{U_r \mid 1 \leq r \leq n_1\}$ be $\operatorname{Stab}(A_1)$-representatives of orbits of codimension-2 hyperplanes in $N(A_1)$. Then $U = fa_0'a_1U_r$ for some $a_1 \in \operatorname{Stab}(A_1)$ and $U_r \in \{U_r \mid 1 \leq r \leq n_1\}$. So $gU = V \Rightarrow gfa_0'a_1U_r = fa_0V_s \Rightarrow \left(a_0^{-1}f^{-1}gfa_0'a_1\right)U_r = V_s$. Therefore $\left(a_0^{-1}f^{-1}gfa_0'a_1\right) \in J_C$, for $C = (A_t)_{t=0}^1$, and so

$$f^{-1}gf \in \operatorname{Stab}(A_0)\, J_C \operatorname{Stab}(A_1) \operatorname{Stab}(A_0)$$

<u>*Case $k=2$*</u>: We have $B_2 = fa_0'a_1A_2$ for some $A_1 \in L_1$ and $a_1 \in \operatorname{Stab}(A_1)$. Then $U = fa_0'a_1a_2U_r$ where $a_2 \in \operatorname{Stab}(A_2)$ and U_r is a $\operatorname{Stab}(A_2)$-representative in $\{U_r \mid 1 \leq r \leq n_1\}$. So, $gU = V \Rightarrow g\,(fa_0'a_1a_2)\,U_r = fa_0V_s$. Therefore,

$$f^{-1}gf \in \operatorname{Stab}(A_0)\, J_C \operatorname{Stab}(A_2) \operatorname{Stab}(A_1) \operatorname{Stab}(A_0)$$

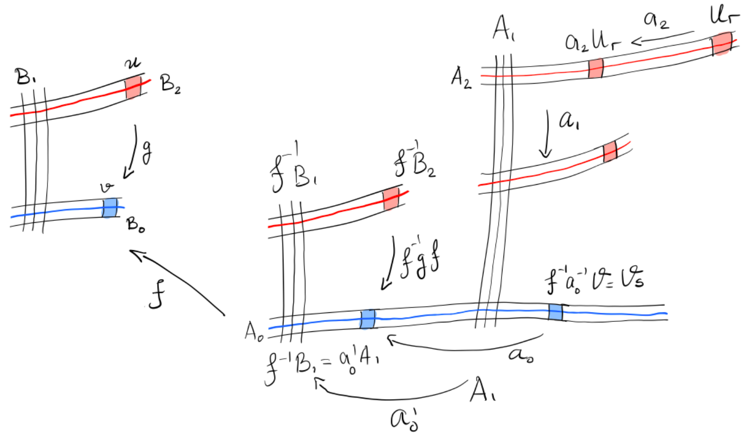

FIGURE 10. Case $k = 2$

Case $k = 3$: Similarly, $U = fa_0'a_1a_2a_3U_r$ where $a_3 \in \text{Stab}(A_3)$ and U_r is a $\text{Stab}(A_3)$-representative in $\{U_r \mid 1 \leq r \leq n_1\}$. Thus, $(gU = V) \Rightarrow g(fa_0'a_1a_2a_3)U_r = fa_0V_s$, and so

$$f^{-1}gf \in \text{Stab}(A_0)\,J_C\,\text{Stab}(A_3)\,\text{Stab}(A_2)\,\text{Stab}(A_1)\,\text{Stab}(A_0)$$

See Figure 10 for case $k = 2$. □

LEMMA 4.16. *Let B be a compact bouquet of circles and let $X \to B$ be a graph of cube complexes with one compact nonpositively curved vertex-space and embedded locally convex edge-spaces. Let $\widehat{X} \to \Gamma_{\widehat{X}} = \widetilde{B}$ be the tree of cube complexes where $\widetilde{B} \to B$ is the universal covering map such that the following diagram commutes:*

$$\begin{tikzcd} \widehat{X} \arrow[r] \arrow[d] & \Gamma_{\widehat{X}} = \widetilde{B} \arrow[d] \\ X \arrow[r] & B \end{tikzcd}$$

Let $F = \pi_1 B$ be the free group acting freely and cocompactly on $\Gamma_{\widehat{X}}$ inducing a free cocompact F-action on $\widehat{X}$. Then there exists a compact graph of cube complexes $\overline{X} \to \Gamma_{\overline{X}}$ and a regular covering map $\widehat{X} \to \overline{X}$ such that the following diagram commutes and the horizontal quotient $\overline{X} \to \overline{X}^E$ is nonpositively curved:

$$\begin{tikzcd} \widehat{X} \arrow[r] \arrow[d] & \Gamma_{\widehat{X}} \arrow[d] \\ \overline{X} \arrow[r] & \Gamma_{\overline{X}} \end{tikzcd}$$

Furthermore, any intermediate covering map $\widehat{X} \to \overline{X}' \to \overline{X}$ *induced by a finite index normal subgroup of* $\pi_1 \Gamma_{\overline{X}}$ *splits as a graph of spaces with nonpositively curved horizontal quotient.*

Proof. Using Lemma 4.2, we can ensure that any finite cover $\overline{X}$ we find below admits a cubical horizontal quotient. Fix collections L and $\mathcal{C}$ as in Definition 4.14. Let

$$\mathcal{O} = \bigcup_{1 \leq k \leq 3} \bigcup_{C \in \mathcal{C}} \left(\operatorname{Stab}\left(A_0\right) J_C \operatorname{Stab}\left(A_k\right) \cdots \operatorname{Stab}\left(A_0\right)\right)$$

where $C = (A_t)_{t=0}^k$ and J_C is as in Definition 4.14. Note that the elements of $\mathcal{O}$ are closing elements by definition. Any empty k-precorner in $\overline{X}$ results from a k-chain in $\widehat{X}$ that is closed by some element $g \in F$. By Lemma 4.15, any closing element in F is conjugate to some element in $\mathcal{O}$. By Lemma 4.11, $\widehat{X}$ admits only trivial k-precorners where each trivial k-precorner is over a k-corner that lies in a single vertex-space of $\widehat{X}$. By assumption, the vertex-spaces of $\widehat{X}$ are nonpositively curved and thus contain only nonempty k-corners. So, $1_F \notin \mathcal{O}$. By Theorem 4.1, there exists a finite index normal subgroup $G \triangleleft F$ that is disjoint from $\mathcal{O}$. Let $\overline{X} = G \backslash \widehat{X} \to G \backslash \Gamma_{\widehat{X}} = \Gamma_{\overline{X}}$ and $\widehat{X} \to \overline{X}$ be the corresponding compact quotient and the regular covering map, respectively. By Remark 4.13, $\overline{X}$ has only trivial nonempty k-precorners, and thus the horizontal quotient $\overline{X}^E$ has no empty k-corners. By Remark 4.8, X^E is nonpositively curved.

Finally, we note that any finite index normal subgroup of G contains no closing elements and so, the corresponding finite covers splits as a graph of spaces with nonpositively curved horizontal quotient. $\square$

5. The Construction

Definition 5.1. Let Y be a compact nonpositively curved cube complex, and let $Y' \subset Y$ be a subcomplex. The map $\varphi : Y' \subset Y \to Y$ is a *partial local isometry* if φ is a local isometry and both Y' and $\varphi\left(Y'\right)$ are locally convex subcomplexes of Y.

Definition 5.2. Let Y be a nonpositively curved cube complex and let $\mathcal{O} = \{\varphi_j : Y_j \subset Y \to Y\}_{j=1}^n$ be a collection of injective partial local isometries of Y where each Y_j is connected. The *realization* of the pair $(Y, \mathcal{O})$ is the cube complex X obtained as the following quotient space (See Figure 11):

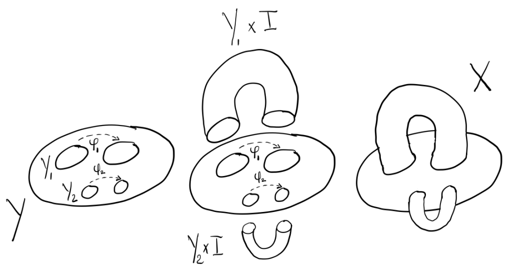

FIGURE 11. The construction of X from Y with two partial local isometries.

$$X = Y \bigsqcup_{j=1}^{n} (Y_j \times I) \Big/ \{(y,0) \sim y,\ (y,1) \sim \varphi_j(y),\ \forall\, y \in Y_j\}_{j=1}^{n}$$

The space X decomposes as a graph of spaces via the map $X \to B$ with $Y \mapsto v$ and $Y_j \times I \mapsto \gamma_j$ where B is the bouquet of n circles $\{\gamma_j\}_{j=1}^{n}$ incident to a vertex v.

LEMMA 5.3. *Let $\overline{X} \to \Gamma_{\overline{X}}$ be a compact graph of cube complexes with a strict horizontal quotient $\overline{X} \to \overline{X}^E$ and isomorphic vertex-spaces. Let $\Phi \in \mathrm{Aut}(\Gamma_{\overline{X}})$ and $\overline{\Phi} \in \mathrm{Aut}(\Gamma_{\overline{X}})$ be automorphisms that map cubes to cubes isometrically. Suppose that the left square of the diagram below commutes. Then there exists an automorphism $\overline{\Phi}^E \in \mathrm{Aut}\left(\overline{X}^E\right)$ such that the right square of the diagram below commutes:*

$$\begin{tikzcd}
\Gamma_{\overline{X}} \arrow[d, "\Phi"'] & \overline{X} \arrow[l] \arrow[r, "q"] \arrow[d, "\overline{\Phi}"'] & \overline{X}^E \arrow[d, dashed, "\overline{\Phi}^E"'] \\
\Gamma_{\overline{X}} & \overline{X} \arrow[l] \arrow[r, "q"] & \overline{X}^E
\end{tikzcd}$$

PROOF. Define $\overline{\Phi}^E : \overline{X}^E \to \overline{X}^E$ by $\overline{\Phi}^E(y) = q\left(\overline{\Phi}\left(q^{-1}(y)\right)\right)$. Then $\overline{\Phi}^E$ is well-defined. Indeed, $q^{-1}(y)$ is either a point or a horizontal graph. By the commutativity of the left square, the automorphism $\overline{\Phi}$ maps points to points and horizontal graphs to horizontal graphs. In both cases, $q\left(q^{-1}(y)\right)$ is a single point. Moreover, for each point $y \in \overline{X}$, we have $\overline{\Phi}^E(q(y)) = q\left(\overline{\Phi}\left(q^{-1}(q(y))\right)\right)$. Since $q(y)$ is a point, $q^{-1}(q(y))$ is either the point y or a horizontal graph containing y. In both cases, $q\left(\overline{\Phi}\left(q^{-1}(q(y))\right)\right) = q\left(\overline{\Phi}((y))\right)$ and thus the right square commutes. By the commutativity of the left square, $\overline{\Phi}$ permutes the vertex-spaces of $\overline{X}$ which makes $\overline{\Phi}^E$ an automorphism of $\overline{X}^E$ that permutes copies of the vertex-spaces. $\square$

THEOREM 5.4. *Let Y be a compact nonpositively curved cube complex and let $\mathcal{O}$ be the set of injective partial local isometries of Y. Then Y embeds in a compact nonpositively curved cube complex R where each $\varphi \in \mathcal{O}$ extends to an automorphism $\Phi \in \mathrm{Aut}(R)$.*

PROOF. We construct a compact graph of spaces $\overline{X}$ whose horizontal quotient $\overline{X}^E = R$ has the desired properties.
Let $\mathcal{O} = \{\varphi_j : Y_j \subset Y \to Y\}_{j=1}^n$ be the collection of injective partial local isometries of Y and let $X \to B$ be the realization of the pair $(Y, \mathcal{O})$. Let γ_j be the loop in B corresponding to φ_j. Let $F = \pi_1 B$ and let $\widehat{X} \to X$ be the covering map induced by the universal covering $\widetilde{B} \to B$ such that the following diagram commutes:

$$\begin{array}{ccc} \widehat{X} & \longrightarrow & \widetilde{B} \\ \downarrow & & \downarrow \\ X & \longrightarrow & B \end{array}$$

Then $\widehat{X} \to \Gamma_{\widehat{X}} = \widetilde{B}$ is a nonpositively curved tree of cube complexes. By Lemma 4.2, Remark 4.3, and Lemma 4.16, there exists a finite regular cover $\overline{X} \to X$ that splits as a graph of spaces according to the following commutative diagram and such that the horizontal quotient $\overline{X} \to \overline{X}^E$ is strict and $\overline{X}^E$ is nonpositively curved. Note that each vertex-space of $\overline{X}$ is a copy of of Y according to some fixed isomorphism.

$$\begin{array}{ccc} \widehat{X} & \longrightarrow & \Gamma_{\widehat{X}} \\ \downarrow & & \downarrow \\ \overline{X} & \longrightarrow & \Gamma_{\overline{X}} \\ \downarrow & & \downarrow \\ X & \longrightarrow & B \end{array}$$

Fix a vertex $v \in \Gamma_{\overline{X}}$ and let $\overline{X}_v$ be the corresponding vertex-space of $\overline{X}$. By subgroup separability of free groups, we can assume that $\Gamma_{\overline{X}}$ has no loops. Thus $\overline{X}_v$ is adjacent to $2n$ vertex-spaces $\{\overline{X}_{v_j}\}_{j=1}^{2n}$ where each $\overline{X}_{v_j}$ is joined to $\overline{X}_v$ by a copy of $Y_j \times I$ attached as follows: $Y_j \times \{0\}$ is identified with a copy of $Y_j \subset \overline{X}_v$ and $Y_j \times \{1\}$ is identified with a copy of $\varphi(Y_j) \subset \overline{X}_{v_j}$. Each $Y_j \times I$ corresponds to a unique map $\varphi_j \in \mathcal{O}$ and thus to a unique circle γ_j in B. The lift of γ_j at v specifies a unique automorphism $\Phi_j \in \mathrm{Aut}(\Gamma_{\overline{X}})$ that maps

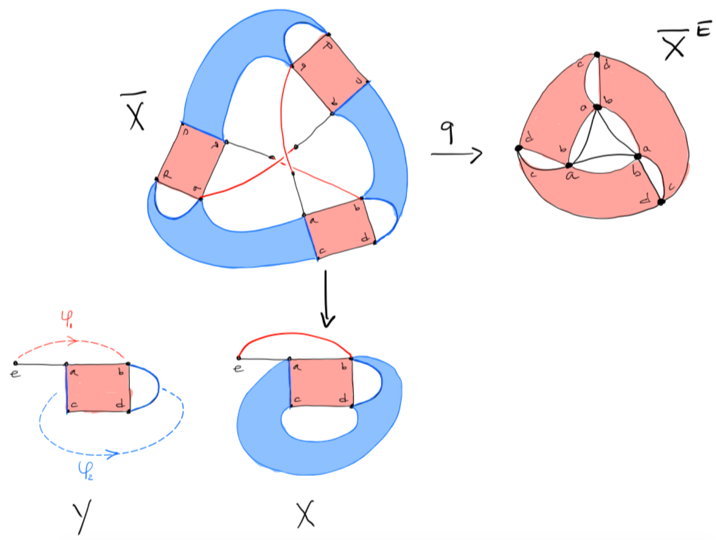

FIGURE 12. The construction of $\overline{X}^E$.

v to v_j. Then there is an automorphism $\overline{\Phi}_j \in \mathrm{Aut}\left(\overline{X}\right)$ that maps $\overline{X}_v$ to $\overline{X}_{v_j}$ such that the following diagram commutes:

$$\begin{tikzcd} \overline{X} \arrow[r, "\overline{\Phi}_j"] \arrow[d] & \overline{X} \arrow[d] \\ \Gamma_{\overline{X}} \arrow[r, "\Phi_j"] & \Gamma_{\overline{X}} \end{tikzcd}$$

In particular, $\overline{\Phi}_j$ maps a copy of $Y_j \subset \overline{X}_v$ to a copy of $\varphi_j\left(Y_j\right) \subset \overline{X}_{v_j}$. By Lemma 5.3, any automorphism $\overline{\Phi} \in \mathrm{Aut}\left(\overline{X}\right)$ induced by an automorphism of the underlying graph $\Phi \in \mathrm{Aut}\left(\Gamma_{\overline{X}}\right)$ descends to an automorphism $\overline{\Phi}^E \in \mathrm{Aut}\left(\overline{X}^E\right)$. So $\overline{\Phi}_j^E\left(q\left(\overline{X}_v\right)\right) = q\left(\overline{\Phi}_j\left(\overline{X}_v\right)\right) = q\left(\overline{X}_{v_j}\right)$. Since $q\left(\overline{X}_v\right) \cong Y$ and $q\left(\overline{X}_{v_j}\right) \cong Y$ are embedded subcomplexes of $\overline{X}^E$ amalgamated along $Y_j = \varphi_j\left(Y_j\right)$, the restriction $\overline{\Phi}_j^E|_{Y_j}$ equals φ_j. See Figure 12. □

REMARK 5.5. Note that $\dim\left(\overline{X}^E\right) = \dim\left(Y\right)$.

REMARK 5.6. Following the *Simple Local Gluing* Lemma in [**BH99**], Theorem 5.4 can be generalized to nonpositively curved metric spaces provided that some finiteness conditions are satisfied and the edge-spaces are locally convex, closed, and complete subspaces.

DEFINITION 5.7. Let Y be a compact nonpositively curved cube complex. A collection of injective partial local isometries $\mathcal{O} = \{\varphi_j : Y_j \subset Y \to Y\}_{j=1}^n$ is *controlled* if the corresponding realization $X \to B$ is a controlled graph of spaces.

THEOREM 5.8 (Haglund-Wise [**HW10**]). *Let X decompose as a finite graph of spaces, where each vertex-space X_v and edge-space X_e is special with finitely many hyperplanes. Then X has a finite special cover provided the attaching maps of edge-spaces satisfy the following:*

(1) the attaching maps $X_e \to X_{\iota(e)}$ and $X_e \to X_{\tau(e)}$ are injective local-isometries;
(2) distinct hyperplanes of X_e map to distinct hyperplanes of $X_{\iota(e)}$ and $X_{\tau(e)}$;
(3) noncrossing hyperplanes map to noncrossing hyperplanes;
(4) no hyperplane of X_e extends in $X_{\iota(e)}$ or $X_{\tau(e)}$ to a hyperplane dual to an edge that intersects X_e in a single vertex.

REMARK 5.9. For a compact special cube complex Y with a controlled collection of partial local isometries $\mathcal{O}$, the realization of the pair $(Y, \mathcal{O})$ satisfies the conditions of Theorem 5.8 and thus has a finite special cover.

THEOREM 5.10. *Let Y be a compact special cube complexe and let $\mathcal{O}$ be a controlled collection of injective partial local isometries of Y. Then there exists a compact special cube complex R containing Y as a locally convex subcomplex such that each $\varphi \in \mathcal{O}$ extends to some automorphism $\Phi \in \mathrm{Aut}\,(R)$.*

PROOF. The claim follows from Remark 5.9 Theorem 5.8, Theorem 5.4 Lemma 4.5, Lemma 4.4, and Lemma 3.2. $\square$

CHAPTER 2

Negative Immersions and Finite Height Mappings

1. Introduction

A 2-complex X has *nonpositive immersions* if for any combinatorial immersion $Y \to X$, with Y compact, connected, and collapsed, either $\pi_1 Y$ is trivial or $\chi(Y) \leq 0$. A 2-complex X has *negative immersions* if there exists $c > 0$ such that for any combinatorial immersion $Y \to X$, where Y is compact, connected, collapsed, and with no isolated edges, either $\pi_1 Y$ is trivial or $\chi(Y) \leq -c|Y|_2$ where $|Y|_2$ is the number of 2-cells in Y. Let $\mathcal{F}$ be a free group and let $\alpha : \mathcal{H} \xrightarrow{\cong} \mathcal{K}$ be an isomorphism of subgroups of $\mathcal{F}$. The *HNN extension* of $\mathcal{F}$ with respect to α is the group presented by $\langle \mathcal{F}, t \ : \ t^{-1}\mathcal{H}t = \mathcal{K} \rangle$. The HNN extension is *ascending* if $\mathcal{H} = \mathcal{F}$, and *partially ascending* if $\mathcal{H}$ is a proper free factor of $\mathcal{F}$. Ascending HNN extensions which were studied in [**FH99**], are shown to have nonpositive immersions in [**Wis22**].

The *height* of a subgroup $\mathcal{H}$ in a group $\mathcal{G}$, denoted by Height $(\mathcal{H})$, is the supremal number of distinct cosets $\{\mathcal{H}g_i\}_{i \in I}$ such that $\bigcap_{i \in I} \mathcal{H}^{g_i}$ is infinite [**GMRS98**]. In this chapter, we define a closely related notion of "directed height" of mappings. Let H be a subgraph of a finite graph F and let $\psi : H \to F$ be a cellular immersion. Let X be the mapping torus representing the partially ascending HNN extension. Then the *directed height* of ψ is $\overrightarrow{\text{Height}}(\psi) = \inf\{i \ : \ \psi^{-i}(H) \text{ is a forest}\}$. We will show in Lemma 4.5 that $\overrightarrow{\text{Height}}(\psi) < \infty$ if and only if $\pi_1 H$ has finite height in π_1X. Our main results are:

Theorem. Let H be a subgraph of a finite graph F. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Suppose $\psi^{-1}(H)$ is a forest. Then X has negative immersions.

Theorem. Let F be a finite connected graph and let $H \subset F$ be a subgraph. Let $\psi : H \to F$ be a cellular immersion. Then the mapping torus of ψ has negative immersions if and only if ψ has finite directed height.

THEOREM. Let F be a finite graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$ with $\overrightarrow{\text{Height}}(\psi) < \infty$. Let $\mathcal{K} \subset \pi_1 X$ be a finitely generated subgroup. Then $\mathcal{K}$ splits over edge groups with uniformly bounded Euler characteristic.

In Section 2, we give some background. In Section 3 we prove a special case of the main theorem, namely, that a partially ascending HNN extensions of free groups with malnormal associated subgroups have negative immersions. In Section 4, we prove the main theorem, and in Section 5 we discuss related properties and state two conjectures.

2. Background

We work in the category of CW-complexes. Let Y be a CW-complex. We denote by Y^k the k-skeleton of Y and by $|Y|_k$ the number of k-cells in Y. Given complexes X and Y, a map $Y \to X$ is *cellular* if it maps Y^k into X^k for all k. It is *combinatorial* if it maps open cells of Y homeomorphically onto open cells of X. It is an *immersion* if it is locally injective. A complex is collapsed if it has no free faces. A 1-cell (edge) is *isolated* if it is not a face of a 2-cell. A 2-complex X has *negative immersions* if there is $c > 0$ such that for any combinatorial immersion $Y \to X$ with Y compact, connected, collapsed (containing no free faces), and containing no isolated edges, either $\pi_1 Y$ is trivial or $\chi(Y) \leq -c|Y|_2$ where $|Y|_2$ is the number of 2-cells in Y and $\chi(Y)$ is the Euler characteristic of Y.

A group $\mathcal{G}$ is *coherent* if every finitely generated subgroup of $\mathcal{G}$ is finitely presented. The proof of Theorem 2.1 can be found in [**Wis20**].

THEOREM 2.1. *Let X be a compact 2-complex with negative immersions. Then $\pi_1 X$ is coherent.*

A *graph* F is a 1-dimensional CW-complex whose *vertices* and *edges* are the 0-cells and 1-cells, respectively. There exist two *incidence* maps $\tau_1, \tau_2 : F^1 \to F^0$ mapping each edge $e \in F^1$ to its *boundary vertices*, $\tau_1(e)$, $\tau_2(e)$ called initial and terminal vertex, respectively. Each edge is *oriented* from its initial vertex to its terminal vertex. The *degree* of a vertex v relative to the graph F, denoted by $\deg_F(v)$, is the number of edges in F^1 containing v as an initial or terminal vertex. An edge whose initial and terminal vertices coincide with

v counts twice in $\deg_F(v)$. A *leaf* is a vertex of degree 1 and a *spur* is an edge containing a leaf. A graph is *trivial* if it is a union of vertices. A *tree* is a non-empty graph with no embedded circles and a *forest* is a disjoint union of trees. The *empty graph* is the graph with no edges and no vertices. We consider the empty graph as a forest.

A *graph of graphs* X with underlying graph Γ_X, *vertex-spaces* $\{X_v\}_{v\in\Gamma_X^0}$, and *edge-spaces* $\{X_e\}_{e\in\Gamma_X^1}$ is a topological space X obtained as a quotient of graphs $\{X_v\}_{v\in\Gamma_X^0}$ and $\{X_e \times I\}_{e\in\Gamma_X^1}$ in the following manner: for each edge $e \in \Gamma_X^1$ with boundary vertices $v_1 = \tau_1(e), v_2 = \tau_2(e)$, the edge-space $X_e \times I$ is attached to the vertex-spaces X_{v_1}, X_{v_2} via an *outgoing* attaching map $X_e \times \{0\} \to X_{v_1}$ and an *incoming* attaching map $X_e \times \{1\} \to X_{v_2}$. The Euler characteristic of the resulting space is given by

$$\chi(X) = \sum_{v\in\Gamma_X^0} \chi(X_v) - \sum_{e\in\Gamma_X^1} \chi(X_e)$$

A subgroup $\mathcal{H} \subset \mathcal{G}$ is *malnormal* if $g\mathcal{H}g^{-1} \cap \mathcal{H} = 1_{\mathcal{G}}$ whenever $g \notin \mathcal{H}$. The pair $\mathcal{H}, \mathcal{K} \subset \mathcal{G}$ is malnormal if $g\mathcal{H}g^{-1} \cap \mathcal{K} = 1_{\mathcal{G}}$ for all $g \in \mathcal{G}$. An HNN extension is *malnormal* if the associated subgroups form a malnormal pair.

3. Malnormal Partially Ascending HNN Extension

Definition 3.1. Let H be a subgraph of a graph F. The *boundary* of H in F is

$$\partial H = \left\{ v \in H^0 \ : \ \deg_F(v) > \deg_H(v) \right\}.$$

Lemma 3.2. *Let $H \subset F$ be a subgraph of a finite leafless graph F with no trivial components. Then:*

$$\chi(F) - \chi(H) \ \leq \ \frac{-1}{2}|\partial H|_0$$

Proof. A graph J satisfies $\chi(J) = \sum_{v\in J^0}\left(1 - \frac{\deg(v)}{2}\right)$. We temporarily use χ to denote the number of vertices minus the number of open edges. Let $J = \left(F - H\right) \bigcup_{v\in\partial H^0} S_v^1$ be obtained by removing H and adding a circle at each vertex of ∂H. Then

$$\chi(F) - \chi(H) \ = \ \chi(F - H) \ = \ \chi(J) \ \leq \ -\frac{1}{2}|\partial H|_0. \qquad \square$$

LEMMA 3.3. *Let H be a subgraph of a finite graph F. Let $\psi : H \to F$ be a cellular immersion with $H \subset \psi(H)$. Suppose H has no tree component and $\psi^{-1}(H)$ is homeomorphic to a forest. Then $\psi(T) \cap \partial H \neq \emptyset$ for each component $T \subset \psi^{-1}(H)$. Consequently, there exists $M = M(F, H, \psi) > 0$ such that $|H|_1 \leq M\ |\partial H|_0$.*

PROOF. Note that $\psi^{-1}(H)$ is not necessarily a subgraph of F. Each tree $T \subset \psi^{-1}(H)$ can be subdivided into a tree $\bar{T}$ so that $\psi|_{\bar{T}}$ is combinatorial. Let

$$d = \max\left\{\mathrm{Diam}\left(\bar{T}\right)\ :\ T \subset \psi^{-1}(H),\ \text{where } \bar{T} \text{ is the subdivision of } T\right\}$$

Since H has no tree components, each component $T \subset \psi^{-1}(H)$ has a leaf that maps to ∂H. So $H \subset \bigcup_{v \in \partial H} \mathcal{N}_d(v)$ where $\mathcal{N}_d(v)$ is a ball of radius d centered at v. Let $M = \max\left\{|\mathcal{N}_d(v)|_1\ :\ v \in F^0\right\}$. Then $|H|_1 \leq M|\partial H|_0$. □

DEFINITION 3.4. Let F be a graph and let $H \subset F$ be a subgraph. The *mapping torus* of a map $\psi : H \to F$ is the 2-complex X obtained as follows:

$$X = \left(F \sqcup (H \times [0,1])\right) / \left\{(x,0) \sim x,\ (x,1) \sim \psi(x)\ :\ x \in H\right\}$$

The 2-complex X decomposes as a graph of spaces $X \to \Gamma_X$, where Γ_X is a circle with one vertex v and one edge e. Let $X_v = F$ and $X_e = H \times [0,1]$ be the vertex-space and edge-space, respectively, where X_e is attached to X_v via the maps $H \times \{0\} \to X_v$ and $H \times \{1\} \to X_v$. We refer to the images of $H \times \{0\}$ and $H \times \{1\}$ in X_v as the *outgoing* and *incoming* edge-spaces, respectively. An edge e of X is *vertical* if $e \subset F$, and *horizontal* otherwise. Note that each vertex of H gives rise to a horizontal edge of X, and each edge of H gives rise to a 2-cell of X. Moreover, each horizontal edge and each 2-cell of X arises in this manner.

REMARK 3.5. Let X be the mapping torus of a cellular immersion $\psi : H \to F$, where H is a subgraph of a finite graph F. Let $Y \to X$ be a combinatorial immersion where Y is a nontrivial compact, connected, and collapsed 2-complex with no isolated edges. The decomposition $X \to \Gamma_X$ induces a decomposition $Y \to \Gamma_Y$ whose vertex-spaces are the components of the preimage of F and whose open edge-spaces are the components of the preimage of $H \times (0,1)$. Let Y_v be the disjoint union of the vertex-spaces, and let $Y_e \subset Y_v$ be the disjoint union of the outgoing edge-spaces. Then there is a cellular immersion $\Psi : Y_e \to Y_v$ whose mapping torus is Y and the following diagram commutes:

$$\begin{array}{ccc} Y_e & \xrightarrow{\Psi} & Y_v \\ \downarrow & & \downarrow \\ H & \xrightarrow{\psi} & F \end{array}$$

Define the *boundary* of Y, denoted by ∂Y, as the union of the boundary vertices of Y_e in Y_v. We make the following remarks:

(1) The 2-cells of Y are in correspondence with the edges of Y_e.

(2) Distinct outgoing edge-spaces in a vertex-space are disjoint. This holds since $Y \to X$ is an immersion and the outgoing edge-space in X is an embedding. In particular, each edge of Y_v is in at most one outgoing edge-space.

(3) Since Y is collapsed and has no isolated edges, each edge in Y_v lies in image(Ψ). Indeed, if there is a non-isolated edge $e \not\subset$ image(Ψ), then by Remark (2), e lies in a unique outgoing edge-space. However, outgoing edge-spaces are embedded and so e is a free face, contradicting that Y is collapsed.

(4) No edge-space of Y has a leaf, since a leaf would give rise to a free face.

(5) No edge-space (vertex-space) is a single vertex since otherwise Y would have an isolated edge, a free face, or be trivial.

(6) Outgoing edge-spaces are embeddings and Ψ is an immersion since these mappings pull back from the combinatorial immersion $Y \to X$.

(7) No vertex-space in Y_v has a leaf. Indeed, by Remark (3), each edge of Y_v lies in an incoming edge-space. By Remark (4), no edge-space has a leaf. By Remark (6), the attaching maps of edge-spaces are immersions. Since the image of an immersed leafless graph contains no leafs, the claims holds. Furthermore, by Remark (5), no vertex-space of Y is a tree, and so, $\chi(Y_{v_i}) \leq 0$ for all vertex-spaces Y_{v_i} of Y.

THEOREM 3.6. *Let H be a subgraph of a finite graph F. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Suppose $\psi^{-1}(H)$ is homeomorphic to a forest. Then X has negative immersions.*

PROOF. Let $Y \to X$ be a combinatorial immersion where Y is a nontrivial compact, connected, and collapsed 2-complex with no isolated edges. As in Remark 3.5, let $Y \to \Gamma_Y$

be the induced graph-of-spaces decomposition, and let $\Psi : Y_e \to Y_v$ be the map whose mapping torus is Y. By Remark 3.5.(3), we have $Y_e \subset \text{image}(\Psi)$. By Remark 3.5.(6), the map Ψ projects to ψ and so $\Psi^{-1}(Y_e)$ is homeomorphic to a forest. Each component $T' \subset \Psi^{-1}(Y_e)$ can be subdivided to form a tree $\bar{T}'$ so that $\Psi|_{\bar{T}'}$ is combinatorial. Since $Y \to X$ is a combinatorial immersion, the subdivided trees of $\Psi^{-1}(Y_e)$ embed into the subdivided trees of $\psi^{-1}(H)$ (as in Lemma 3.3), and so for each component $T' \subset \Psi^{-1}(Y_e)$, we have $\text{Diam}\left(\bar{T}'\right) \leq d$, where $d = \max\left\{\text{Diam}\left(\bar{T}\right) : T \text{ is a component in } \psi^{-1}(H)\right\}$. Moreover, since X is compact, there is an upper bound $M = M(d)$ on the number of edges in any d-ball in Y_v. By Remarks 3.5.(4)-(5), Y_e has no tree component. By Lemma 3.3, we have $|Y_e|_1 \leq M|\partial Y_e|_0$. By Lemma 3.2, and Remark 3.5.(1), we have:

$$\chi(Y) = \chi(Y_v) - \chi(Y_e) \leq \frac{-1}{2}|\partial Y_e|_0 \leq \frac{-1}{2M}|Y_e|_1 = \frac{-1}{2M}|Y|_2. \qquad \square$$

4. Finite Height Mappings

DEFINITION 4.1. The *generalized composition* of the functions $\alpha : A \to B$ and $\beta : C \to D$, where $C \subseteq B$, denoted by $\beta \bullet \alpha$, is $\beta \bullet \alpha = \beta \circ \alpha|_{\alpha^{-1}(C)}$.

DEFINITION 4.2. Let F be a connected graph and let $H \subset F$ be a subgraph. Let $\psi : H \to F$ be a cellular immersion. For each $i \geq 0$, let ψ^i denote the generalized composition of ψ with itself i times, where $\psi^0 = id_F : F \to F$. Let $\psi^{-i}(H) = \left(\psi^i\right)^{-1}(H)$.

Let Z_i denote the domain of ψ^i. Then $Z_{i+1} = \{x \in Z_i : \psi^i(x) \in H\} = \psi^{-i}(H)$, for each $i \geq 0$. The *combinatorial domain* D_i of ψ^i is the largest subgraph in Z_i. Note that Z_i is not necessarily a subgraph of F, $Z_{i+1} \subseteq Z_i$, and $D_{i+1} \subseteq D_i$ for all $i \geq 0$. Moreover, Z_i has a part that deformation retracts to D_i and a part that is a disjoint union of closed intervals and singletons. Thus, when Z_i is not homeomorphic to a forest, at least one component of D_i is not a tree. Let $D_\infty \subset H$ be the subgraph whose edges and vertices map into H under all powers of ψ. Note that $\emptyset \subseteq D_\infty \subseteq D_{i+1} \subseteq D_i$.

The *directed height* of ψ is:

$$\overrightarrow{\text{Height}}(\psi) = \inf\left\{i : \psi^{-i}(H) \text{ is a forest}\right\}$$

Note that $\overrightarrow{\text{Height}}\,(\psi) = 0$ if and only if H is a forest. We use the following notation:

$$\|\psi\| = \max\left\{ |\psi\,(e)\,|_1 \ : \ e \subset H^1 \right\}$$

REMARK 4.3. $\overrightarrow{\text{Height}}\,(\psi) = \ell < \infty$ if and only if the length of embedded directed paths in the Bass-Serre tree with infinite stabilizers is bounded by ℓ. Note that the Bass-Serre tree is directed because of the map to the underlying graph of the HNN extension which is a directed loop.

DEFINITION 4.4. The *height* of a subgroup $\mathcal{H}$ in $\mathcal{G}$, denoted by $\text{Height}\,(\mathcal{H})$, is the supremal number of distinct cosets $\{\mathcal{H}g_i\}_{i\in I}$ such that $\bigcap_{i\in I} \mathcal{H}^{g_i}$ is infinite.

LEMMA 4.5. *Let H be a subgraph of a finite graph F and let $\psi : H \to F$ be a cellular immersion. Let X be the mapping torus of ψ. Then $\pi_1 H$ has finite height in $\pi_1 X$ if and only if ψ has finite directed height.*

PROOF. Let $\mathcal{H} = \pi_1 H$ and $\mathcal{F} = \pi_1 F$. Suppose $\mathcal{H}$ has finite height in $\pi_1 X$. Then $\text{Height}\,(\mathcal{H})$ bounds the number of distinct cosets $\{\mathcal{H}g_i\}$ such that $|\mathcal{H}^{g_1} \cap \cdots \cap \mathcal{H}^{g_n}|$ is infinite. So the number of edges in the Bass-Serre tree T with a common infinite stabilizer is likewise bounded. Hence $\text{Height}\,(\mathcal{H})$ bounds the length of embedded paths in T with infinite stabilizer. Thus $\overrightarrow{\text{Height}}\,(\psi) < \infty$.

Suppose $\overrightarrow{\text{Height}}\,(\psi) < \infty$. Then $\overrightarrow{\text{Height}}\,(\psi)$ bounds the length of embedded paths in T with infinite stabilizers. There is a uniform upper bound on the degree of vertices of any subtree $T' \subset T$ with point-wise stabilizer of T' infinite. Indeed, the number of incoming edges of each vertex in T' is bounded by $r = \text{Height}\,(\psi_*\,(H))$ in $\mathcal{F}$, since every finitely generated subgroup of a free group has finite height [**GMRS98**]. Thus T' is a rooted tree of length $\leq \overrightarrow{\text{Height}}\,(\psi)$ and incoming degree $\leq r$. So the number of edges in T' is $\leq r^{\overrightarrow{\text{Height}}(\psi)}$. Any set of cosets of the edge group corresponds to a set of edges in T. The intersection of the corresponding conjugates point-wise stabilizes those edges, and thus point-wise stabilizes the smallest tree T' containing them. Hence the number of cosets is bounded by $\leq r^{\overrightarrow{\text{Height}}(\psi)}$. $\square$

LEMMA 4.6. *Let H be a subgraph of F. Let $\psi : H \to F$ be a cellular immersion with $\overrightarrow{\text{Height}}\,(\psi) = \ell < \infty$. Then D_∞ is a (possibly empty) forest.*

PROOF. We have $D_\infty \subseteq D_{\ell+1}$. So, it suffices to show that $D_{\ell+1}$ is a forest. Suppose $C \subset D_{\ell+1}$ is an embedded circle. Then $\psi^{\ell}(C) \subset H$ and so $\psi^{-\ell}(H)$ is not a forest, contradicting the assumption. $\square$

LEMMA 4.7. *Let F be a connected graph and let $H \subset F$ be a finite subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. If ψ has infinite directed height, then H contains a connected subgraph $D \subset H$ with $\chi(D) \leq 0$ such that $\psi(D) = D$. Consequently, X contains a subcomplex $Y \hookrightarrow X$, where Y is a connected, compact, and collapsed 2-complex with no isolated edges and $\chi(Y) = 0$.*

PROOF. Since ψ has infinite directed height, for each $i \geq 0$, we have $\psi^{-i}(H)$ is not a forest. So each D_i contains an embedded circle. Since H is finite and $D_{i+1} \subseteq D_i$, there is an integer p such that for all $j > p$ we have $D_{j+1} = D_j$. Then D_j contains a component D with $\chi(D) \leq 0$ and $\psi(D) = D$. In particular, since ψ is an immersion, $\psi(D_{\text{core}}) = D_{\text{core}}$, where D_{core} is the core of D. The mapping torus of ψ restricted to D_{core} provides Y. $\square$

DEFINITION 4.8. Let $Q \to \Gamma_Q$ be a graph of spaces where Γ_Q is equal to a subdivided interval $[0, k]$ directed from 0 to 1 $\leq k \leq \infty$. Suppose each vertex-space Q_{v_i} is a tree where Q_{v_0} has exactly one edge f_0. For each edge-space $Q_{e_i} \times I$ there is an *outgoing* attaching map $Q_{e_i} \times \{0\} \to Q_{v_{i-1}}$ and an *incoming* attaching map $Q_{e_i} \times \{1\} \to Q_{v_i}$. When each outgoing attaching map is an embedding onto a single edge f of the vertex-space, then Q is a *ladder* and f is a *connecting edge*. When each attaching map is bijective, then Q is a *fan*. The *rim* of a fan Q, denoted by $\mathrm{Rim}(Q)$, is Q_{v_k}. The *length* of Q is $\mathrm{Length}(Q) = k$. We allow the case $k = \infty$ and say that Q is an *infinite ladder/fan*. We say Q *arises* from f_0, and f_0 *gives rise* to Q.

The space Q is a cell complex as follows: we have already declared each Q_{v_i} is a tree and so it remains to describe the additional 1-cells and 2-cells of Q. Each open edge-space $Q_{e_i} \times (0,1)$ has a product structure induced by the graph Q_{e_i}. See Figure 1. The edges in the vertex-spaces are *vertical* and the remaining ones are *horizontal*. Each vertex in the image of $Q_{e_i} \to Q_{v_{i-1}}$ *gives rise* to a horizontal edge in Q. Each edge f in the image of $Q_{e_i} \to Q_{v_{i-1}}$ *gives rise* to a 2-cell $S \subset Q$. We say S *arises* from f.

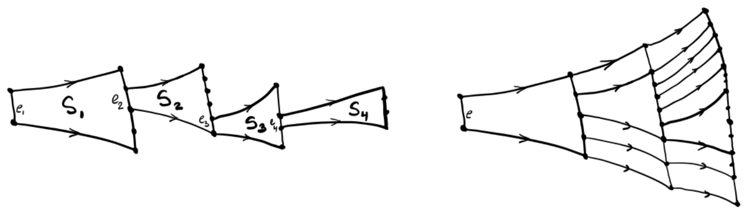

FIGURE 1. Left: A ladder of length 4 emerging from e_1. Right: A fan of length 3 emerging from e.

Let X be a 2-complex with a graph-of-spaces structure whose 1-skeleton is partitioned into horizontal and vertical edges, where the vertical edges are the edges of vertex-spaces, and the horizontal edges are the remaining ones. An *immersed ladder* of X is a combinatorial immersion $\lambda : L \to X$ that maps vertical/horizontal edges of a ladder L to vertical/horizontal edges of X. An *immersed fan* $\varphi : Q \to X$ is defined analogously. An edge $e \subset X$ has a k-ladder (resp. k-fan), if there is an immersed ladder $\lambda : L \to X$ (resp. immersed fan $\varphi : Q \to X$) of length k emerging from e' such that $\lambda(e') = e$ (resp. $\varphi(e') = e$). When X is the mapping torus of $\psi : H \to F$, we require that immersions preserve the orientation of horizontal edges.

Let X be the mapping torus of $\psi : H \to F$. Let $H_i = \overline{D_i - D_{i+1}}$ be the subgraph whose edges give rise to i-fans but not $(i+1)$-fans. When $H_i = \emptyset$, we have $D_i = D_{i+1} = D_\infty$ is the subgraph whose edges give rise to infinite fans. Then D_∞ is ψ-invariant. Let $m = m(\psi)$ denote the supremum of lengths of maximal finite fans in X. Note that when H is finite we have $m < \infty$ since any maximal finite fan is determined by the edge it arises from.

LEMMA 4.9. *Let H be a subgraph of a finite connected graph F. Let X be the mapping torus of a cellular immersion $\psi : H \to F$ with $\overrightarrow{\mathrm{Height}}(\psi) < \infty$. Let $m = m(\psi)$ be the maximal length of immersed finite fans in X. Let $Y \to X$ be a combinatorial immersion, where Y is a nontrivial, compact, connected, and collapsed 2-complex with no isolated edges. Let $Y \to \Gamma_Y$ be the induced graph-of-spaces decomposition and let ∂Y be the associated boundary. Then there exists $M = M(H, F, \psi) > 0$ such that each 2-cell S of Y lies in the image of an immersed ladder $\lambda : L \to Y$ with $\mathrm{Length}(L) \leq m + 1$ emerging from e where $\mathrm{Dist}(\lambda(e), \partial Y) \leq M$.*

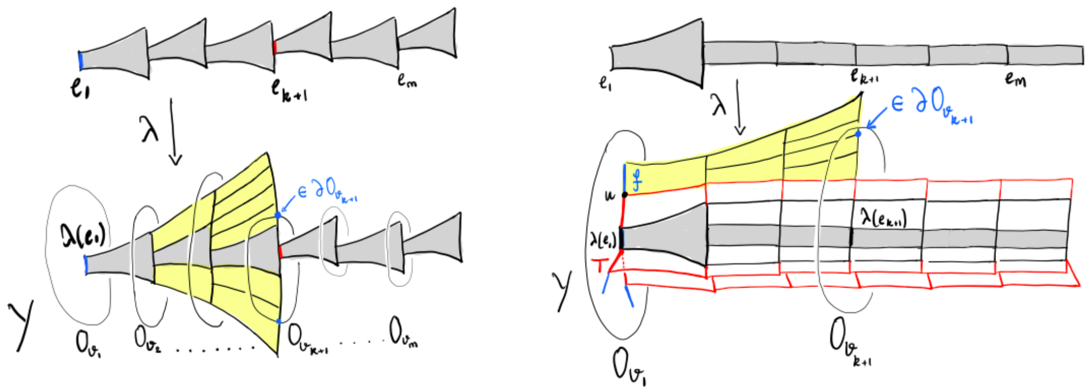

FIGURE 2. On the left: Case 1, and on the right case 2.

PROOF. Let S be a 2-cell of Y. Since Y is collapsed, there is an immersed ladder $\lambda : L \to Y$ with $\text{Length}\,(L) = m+1$ and whose $(m+1)$-th 2-cell maps to S. Let $\{e_1, \dots, e_{m+1}\}$ be the connecting edges of L. For $1 \leq i \leq m+1$, let O_{v_i} be the outgoing edge-space containing $\lambda\,(e_i)$ and let Y_{v_i} be the vertex-space containing O_{v_i}. Let $\varphi : Q \to Y$ be the maximal immersed fan emerging from $\lambda\,(e_1)$. See Figure 2.

Case 1: If $\text{Length}\,(Q) = k \leq m$, then $\lambda\,(e_{k+1}) \subset \varphi\,(\text{Rim}\,(Q)) \cap O_{v_{k+1}}$. Since Q is maximal, image $\left(\text{Rim}\,(Q) \to Y_{v_{k+1}}\right) \not\subset O_{v_{k+1}}$, and so $\varphi\,(\text{Rim}\,(Q)) \cap \partial O_{v_{k+1}} \neq \emptyset$. Since fans in Y project to fans in X, we have $|\,\text{Rim}\,(Q)\,|_1 \leq \|\psi\|^m$. Thus, $\text{Dist}\,(\lambda\,(e_{k+1})\,, \partial Y) \leq \|\psi\|^m$.

Case 2: If $\text{Length}\,(Q) > m$, then the image of $Q \to Y \to X$ is an infinite fan of X. Let $T \subset O_{v_1}$ be the maximal connected subgraph containing $\lambda\,(e_1)$ and whose edges give rise to $(m+1)$-fans in Y. Hence T immerses in D_∞. Since $\overrightarrow{\text{Height}}\,(\psi) < \infty$, it follows from Lemma 4.6 that D_∞ is a forest. So T is a tree with $\text{Diam}\,(T) \leq \text{Diam}\,(D_\infty)$. Let $u \in T$ be a leaf. Since Y is collapsed, outgoing edge-spaces have no leaves. So there is an edge $f \subset O_{v_1}$ containing u with $f \not\subset T$. By maximality of T, the maximal fan $\varphi'\,(Q')$ emerging from f has length $k \leq m$. So $\varphi'\,(\text{Rim}\,(Q')) \cap \partial O_{v_{k+1}} \neq \emptyset$. Hence, $\text{Dist}\,(\lambda\,(e_{k+1})\,, \partial Y) \leq \text{Diam}\,(D_\infty) + \|\psi\|^m$.

The claim follows with $M = \text{Diam}\,(D_\infty) + \|\psi\|^m$. $\square$

THEOREM 4.10. *Let F be a finite connected graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Then X has negative immersions if and only if ψ has finite directed height.*

PROOF. The "only if" direction holds by Lemma 4.7.

Suppose $\overrightarrow{\text{Height}}(\psi) < \infty$. Let $Y \to X$ be a combinatorial immersion where Y is a nontrivial compact, connected, and collapsed 2-complex with no isolated edges. Let $Y \to \Gamma_Y$ be the induced graph-of-spaces decomposition. For each $v \in \Gamma_Y^0$, let Y_v be the corresponding vertex-space and let O_v be the disjoint union of outgoing edge-spaces in Y_v. Let $m = m(\psi)$ be the supremal length of maximal finite fans in X. By Lemma 4.9, there exists $M > 0$ such that each 2-cell of Y lies in a ladder of length $\leq m+1$ emerging from a vertical edge e with $\text{Dist}(e, \partial Y) \leq M$. Let $\partial' Y$ be the set of boundary points of Y that are at a distance $\leq M$ from such edges e. So $\partial' Y \subseteq \partial Y = \bigsqcup_{v \in \Gamma_Y^0} \partial O_v$. Since $Y \to X$ is a combinatorial immersion, there is an upper bound N on the number of edges in an M-ball in the vertex-spaces of Y. Note that $N = N(F, M)$ is a function of F and M. Consider the M-balls centered at vertices of $\partial' Y$. In each such ball, there are at most N edges and each edge gives rise to at most $\|\psi\|^m$ ladders of length $\leq (m+1)$. The number of 2-cells in each ladder is $\leq (m+1)$. Then:

$$|Y|_2 \;\leq\; \sum_{v \in \partial' Y} (m+1)\|\psi\|^m N \;=\; (m+1)\|\psi\|^m N |\partial' Y|_0 \;\leq\; (m+1)\|\psi\|^m N |\partial Y|_0$$

and so

$$\frac{|Y|_2}{(m+1)\|\psi\|^m N} \;\leq\; |\partial Y|_0$$

By Remark 3.5.(7), the vertex-spaces of Y have no leaves. Then the conclusion holds by the following double inequality. Its first equality is straightforward. Its last inequality follows from above, and its middle inequality holds by Lemma 3.2.

$$\chi(Y) \;=\; \sum_{v \in \Gamma_Y^0} \left(\chi(Y_v) - \chi(O_v)\right) \;\leq\; \frac{-1}{2}|\partial Y|_0 \;\leq\; \frac{-1}{2(m+1)\|\psi\|^m N}|Y|_2 \qquad \square$$

DEFINITION 4.11. Let $\mathcal{F}$ be a free group. There is a natural generalization of fully irreducible endomorphisms of free groups to fully irreducible partial endomorphisms. A partial endomorphism $\psi : \mathcal{H} \to \mathcal{F}$ is *fully irreducible* if there does not exist $n > 0$, a proper free factor $\mathcal{H}' \subset \mathcal{H}$, and $g \in \mathcal{F}$ such that $\psi^n(\mathcal{H}') \subset g^{-1}\mathcal{H}' g$. See Definition 4.1 for the notion of *generalized composition* explaining ψ^n. The standard notion of fully irreducible endomorphism focuses on the case where $\mathcal{H} = \mathcal{F}$ [**BH92**].

The *standard 2-complex* associated to a presentation of a group G is a 2-dimensional cell complex formed by a single vertex, one circle at the vertex for each generator of G, and a 2-cell for each relation in the presentation. The attaching maps of the 2-cells are determined by the presentation.

In the language of Definition 4.11, our result shows the following:

THEOREM 4.12. *Let $\mathcal{H}$ be a proper free factor of a finitely generated free group $\mathcal{F}$, and let $\psi : \mathcal{H} \to \mathcal{F}$ be a monomorphism. Let X be the standard 2-complex of the HNN extension of $\mathcal{F}$ with respect to ψ. Then X has negative immersions if and only if ψ is fully irreducible.*

PROOF. The proof follows from Lemma 4.5, Lemma 4.6, Lemma 4.7, and Lemma 4.9. □

REMARK 4.13. If $c = \dfrac{1}{2(m+1)\|\psi\|^m N}$ is the constant in $\chi(Y) \leq -c|Y|_2$, then $0 < c < 1$.

REMARK 4.14. In the proof of Theorem 4.10, we assume that Y has no isolated edges, as required by the definition of Negative Immersions. However, the claim that $\chi(Y) \leq -c|Y|_2$ holds even if we allow Y to have isolated edges. This follows from a simple induction on the number of isolated edges in Y. Indeed, the base case holds by Theorem 4.10. Now, let e be an isolated edge of Y. Then either e is not separating and $Y = Y_1 \cup e$, or e is separating and $Y = Y_1 \cup e \cup Y_2$. In the former case, we have

$$\chi(Y) < \chi(Y - e) = \chi(Y_1) \leq -c|Y_1|_2 = -c|Y|_2$$

where the last inequality holds by induction. In the latter case, we have

$$\chi(Y) = \chi(Y_1) + \chi(Y_2) - 1 < \chi(Y_1) + \chi(Y_2) \leq -c\left(|Y_1|_2 + |Y_2|_2\right) = -c|Y|_2$$

where the last inequality holds by induction.

Motivated by our desire to verify Property 11 of the next section, we note the following consequence of the preceding statements. This does not prove Property 11 since it does not assert that the edge groups in the splitting of $\mathcal{K}$ equal the intersections of $\mathcal{K} \cap \mathcal{H}^g$, for $g \in \pi_1 X$.

THEOREM 4.15. *Let F be a finite connected graph and let $H \subset F$ be a subgraph. Let X be the mapping torus of a cellular immersion $\psi : H \to F$ with $\overrightarrow{\mathrm{Height}}(\psi) < \infty$. Let $\mathcal{K} \subset \pi_1 X$ be a finitely generated subgroup. Then $\mathcal{K}$ splits over edge groups with uniformly bounded Euler characteristic.*

PROOF. By Theorem 4.10, X has negative immersions. By Theorem 2.1, $\pi_1 X$ is coherent. So there is a combinatorial immersion $Y \to X$, with $\pi_1 Y \xrightarrow{\cong} \mathcal{K}$ where Y is compact and connected. We can assume that Y is collapsed since collapsing is a homotopy equivalence. Let $Y \to \Gamma_Y$ be the graph-of-spaces decomposition induced by the decomposition $X \to \Gamma_X$. Let V_Y and O_Y be the disjoint union of vertex-spaces and outgoing edge-spaces of Y, respectively. We show that $\chi(O_Y)$ is uniformly bounded by a function of $\mathrm{rank}(\pi_1 Y)$. In fact, we show that $\chi(O_Y) \geq \dfrac{\chi(Y)}{c}$. By Theorem 4.10, Remark 4.13 and Remark 4.14, there is a constant $c \in (0,1)$ such that $\chi(V_Y) - \chi(O_Y) \leq -c|Y|_2$. So $\chi(O_Y) \geq \chi(V_Y) + c|Y|_2$. We have $\chi(Y) = \chi(V_Y) - E + |Y|_2$, where E are the number of the horizontal edges in Y. Since $c - 1 < 0$, we have

$$\chi(O_Y) \geq \chi(Y) + E - |Y|_2 + c|Y|_2 \geq \chi(Y) + (c-1)|Y|_2 \geq \frac{\chi(Y)}{c}$$

where the last inequality follows by By Theorem 4.10. $\square$

5. Discussion of Related Properties

Let H be a subgraph of a finite connected graph F and let $\mathcal{H} = \pi_1 H$ and $\mathcal{F} = \pi_1 F$. Let X be the mapping torus of a cellular immersion $\psi : H \to F$. Consider the following properties:

(1) $\pi_1 X$ is locally quasiconvex.
(2) $\mathcal{F}$ and $\mathcal{H}$ are quasiconvex.
(3) $\mathcal{F}$ and $\mathcal{H}$ have finite height.
(4) $\pi_1 X$ has the finitely generated intersection property.
(5) X has negative immersions.
(6) $\pi_1 X$ contains no subgroup isomorphic to an ascending HNN extension of a finitely generated free group.
(7) $\pi_1 X$ is hyperbolic.

(8) $\pi_1 X$ contains no Baumslag-Solitar subgroup $\mathrm{BS}(1,m)$ for $m>0$.

(9) $\pi_1 X$ has a quasiconvex hierarchy.

(10) $\pi_1 X$ is virtually special.

(11) $\mathcal{K}\cap\mathcal{H}$ is finitely generated whenever $\mathcal{K}\subset\pi_1 X$ is finitely generated.

(12) Each finitely generated subgroup of $\pi_1 X$ is tamely generated.

(1)$\Rightarrow$ (2) is immediate. When $\pi_1 X$ is hyperbolic, we have (2) $\iff$ (3), where ($\Rightarrow$) holds by [**GMRS98**] and ($\Leftarrow$) holds by [**Mit04**]. A group has the *finitely generated intersection property* (FGIP) if the intersection of any two finitely generated subgroups is also finitely generated. For instance, free groups have the FGIP [**How54**]. (4)$\Rightarrow$(6) holds by [**BW22**] and (1) $\Rightarrow$ (4) holds by [**Sho91**]. (5)$\Rightarrow$(6) since ascending HNN extensions of free groups have Euler characteristic zero, and (6)$\Rightarrow$(3) by Lemma 4.5 and Lemma 4.7. (5) $\iff$ (3) holds by Theorem 4.10, and (5)$\Rightarrow$(8) is a special case of (5)$\Rightarrow$(6). It is well known that (7)$\Rightarrow$(8), e.g [**ABC$^+$91**]. (7) + (9)$\Rightarrow$(10) by [**Wis12**]. (11)$\Rightarrow$(12) since if $\mathcal{K}\cap\mathcal{H}$ is finitely generated for each finitely generated $\mathcal{K}$, then $\mathcal{K}^g\cap\mathcal{H}$ is finitely generated for each g, and so $\mathcal{K}\cap\mathcal{H}^g$ is finitely generated for each g. See [**BW13**] for the definition of "tamely generated". (12)$\Rightarrow$ (1) holds by [**BW13**]. (5) $\Rightarrow$ (9) holds by the following argument: (5) $\Rightarrow$ (6) and by [**CW22**], we have (6)$\Rightarrow$ $\pi_1 X\subset\pi_1 X'$ where X' is the mapping torus of a fully irreducible nonsurjective map of a graph and X' is hyperbolic relative to X. By [**Rey11**], this implies $\pi_1 X'$ is hyperbolic since it contains no $\mathrm{BS}(1,m)$ and so (5)$\Rightarrow$ (7) holds. Since (5)$\Rightarrow$(2), we have (2)+(7) $\Rightarrow$ (9) since $\pi_1 X$ splits along $\mathcal{H}$ and the vertex-group is free.

We end this chapter by stating the following conjectures:

Conjecture 5.1. (5) $\Rightarrow$ (1) and hence (5) $\Rightarrow$ (11).

Conjecture 5.2. If X is a 2-complex with negative immersions, then $\pi_1 X$ has a finite index subgroup that is isomorphic to the fundamental group of a mapping torus of a finite height immersion of graphs $\psi: H\to G$.

Conjecture 5.3. If $\mathcal{G}$ is a locally quasiconvex hyperbolic group, then $\mathcal{G}$ has a finite index subgroup that is isomorphic to the fundamental group of a mapping torus of a finite height immersion of graphs $\psi: H\to G$.

CHAPTER 3

Maximal Ascents

1. Introduction

Weinbaum conjectured in [**Wei90**] that any nonperiodic word W of length > 1 has a cyclic permutation that is a concatenation uv where each of U and V appear exactly once as a prefix of a cyclic permutation of W and W^{-1}. This conjecture was proved by Duncan-Howie in [**DH92**] using the right-orderability of one-relator groups [**BH72**]. This provided the motivation to investigate whether maximal ascents are uniquely positioned in nonperiodic words. Our main result is:

Theorem Let $X = \{x_1, x_2\}$ be an alphabet and let $W \in \mathcal{F}(X)$ be a cyclically reduced nonperiodic word. Then W has a cyclic permutation $W' = AD$ where:

(1) A is the uniquely positioned maximal ascent in W.
(2) If D is not uniquely positioned, then it appears as an internal subword of A.
(3) Using the Magnus ordering on $\mathcal{F}$, we have $D = 1_{\mathcal{F}}$ if and only if W is monotonic.

2. Ascents and Descents

Let X be an alphabet and let $W = y_1 \cdots y_n$ be a word in $X^* = X \cup X^{-1}$. We say W is *reduced* if $y_i \neq y_{i+1}^{-1}$ for all $1 \leq i < n$. It is *cyclically reduced* if W is reduced and $x_1 \neq x_n^{-1}$. We henceforth consider only cyclically reduced words. Each W represents an element g in the free group $\mathcal{F} = \mathcal{F}(X)$, and each $g \in \mathcal{F}$ is represented by a unique reduced word W. For simplicity, we shall use W to denote both the element of $\mathcal{F}$ and its representation in X^*. The *spelling* of W is $W = y_1 \cdots y_n$ where $y_i \in X^*$. The *empty* word is denoted by $1_{\mathcal{F}}$. The *length* of W, denoted by $|W|$, is n if $W = y_1 \cdots y_n$ for $y_i \in X^*$. We also represent a word W as a finite graph denoted by $\overline{W}$, that is linear, directed, and labeled in X^*. The empty word is then represented by a single vertex. A word V is a *subword* of W if $W = SVU$ for some reduced words S and U with $|W| = |S| + |V| + |U|$. Note that each subword V of W

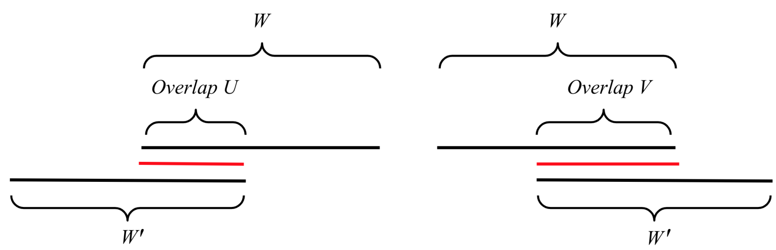

FIGURE 1. We represent words as line segments. Here we see W and W' overlap along U and V.

is represented by a connected subgraph $\overline{V} \subset \overline{W}$ .The subword V is a *prefix* if $S = 1_{\mathcal{F}}$ and a *suffix* if $U = 1_{\mathcal{F}}$. If both $S \neq 1_{\mathcal{F}}$ and $U \neq 1_{\mathcal{F}}$, then V is *internal* in W. Subwords U and V of W are *equivalent* if they have the same spelling and appear in different positions in W.. We write $U \equiv V$. Two subwords of W *overlap* if a nonempty suffix of one is a prefix of the other and neither is a subword of the other. We shall represent words and overlaps diagrammatically as line segments. See Figure 1. W is *periodic* if there is a cyclically reduced word (*period*) $U \in \mathcal{F}$ with $U^k = W$ for some $k > 1$. If U is a prefix of W, with $W = UV$, then the word $W' = VU$ is a *cyclic permutation* of W. Working in the free group, we have $W' = U^{-1}WU$, where W' is reduced if and only if W is cyclically reduced. Let R_W be the set of cyclic permutations of W and W^{-1}. A nontrivial reduced word $U \in \mathcal{F}$ is a *uniquely positioned* in W if U is the prefix of exactly one element of R_W. For example, aa is uniquely positioned in $W = baaba$. The word aba is not uniquely positioned in W since it is a prefix in both $ababa$ and $abaab$.

A *bi-ordered* group is a pair $(\mathcal{G}, \prec)$ where $\mathcal{G}$ is a group and $\prec$ is a total order on $\mathcal{G}$ that is invariant under both right and left group translation. A well known result of Shimbireva [**Shi47**] states that the free group on two generators (and so every non-abelian free group) is bi-orderable. This result is sometimes attributed to Vinogradov and Magnus as well [**DNR14**].

DEFINITION 2.1. Let $\mathcal{F} = \mathcal{F}(X)$ be the free group on an alphabet X. Let $\prec$ be a bi-order on $\mathcal{F}$. A word $U \in \mathcal{F}$ is an *ascent* if each prefix and each suffix of U is $\succ 1_{\mathcal{F}}$, and U

is a *descent* if each prefix and each suffix of U is $\prec 1_{\mathcal{F}}$. The *maximal ascent* A of W is the greatest ascent over all subwords of R_W.

The *peak* (resp. *low*) in W is the largest (resp. smallest) prefix of W with respect to $\prec$.

REMARK 2.2. Let W be a cyclically reduced word representing an element in $\mathcal{F}$. Let $\prec$ be a bi-order on $\mathcal{F}$. Then:

(1) The maximal ascent of W is the largest subword in R_W with respect to $\prec$.
(2) The inverse of an ascent is a descent.
(3) If two ascents in a given word overlap, then the overlap is also an ascent.
(4) Ascents and descents in a given word have no overlaps.
(5) Ascents and descents are always nontrivial. Peaks and lows, on the other hand can be trivial.

LEMMA 2.3. *Let W be a cyclically reduced word of $(\mathcal{F}(X), \prec)$. Let A be the maximal ascent in W. Let M and m be the peak and low of W. If A is a subword of W, then $W \succ 1_{\mathcal{F}}$ and $M = mA$. Consequently, W has exactly one subword equivalent to A and no subword of W^{-1} is equivalent to A. Moreover, if $W = AD$ with $D \neq 1_{\mathcal{F}}$, then D is a descent.*

PROOF. Let $W = PAQ$, where P and Q are subwords of W. If $W \prec 1_{\mathcal{F}}$, then $A \prec P^{-1}Q^{-1}$. But $P^{-1}Q^{-1}$ is an initial subword of a conjugate of W^{-1}. Indeed, we have $W^{-1} = Q^{-1}A^{-1}P^{-1}$. Then $AQ\,(Q^{-1}A^{-1}P^{-1})\,Q^{-1}A^{-1} = P^{-1}Q^{-1}A^{-1}$ is a cyclic permutation of W^{-1}. This leads to a contradiction since A is the maximal ascent in W. Thus, if A is a subword of W, then $W \succ 1_{\mathcal{F}}$ and so A is not a subword of conjugates of $W^{-1} \prec 1_{\mathcal{F}}$.

Let $W = y_1 \cdots y_n$ and let $g_i = y_1 \cdots y_i$. Suppose $g_j = g_i A$ for some $0 \leq i, j \leq n$. Then $g_j \preceq M$ and $g_i \succeq m$. So $g_i^{-1} \preceq m^{-1}$. Since $\prec$ is a bi-order, $A = g_i^{-1} g_j \preceq m^{-1}M$.

We now show that $m^{-1}M$ is a subword of W. Note that $m^{-1}M$ is a subword of W whenever $|m| < |M|$. Let $m = y_1 \cdots y_s$ and $M = y_1 \cdots y_t$. Then $s \neq t$. Suppose $s > t$. Then $m^{-1}M = y_s^{-1} \cdots y_{t+1}^{-1}$ is a subword in W^{-1}. By the maximality of A, we have $A \preceq m^{-1}M \Rightarrow A = m^{-1}M$. So A is a subword of W^{-1}, which is a contradiction. Thus, $s < t$ and $A = m^{-1}M$ is a subword of W.

It remains to show that A does not appear twice in W. Following [**DH92**], if $g_j = g_i A$, then:

$$g_j \;=\; g_i \left(g_i^{-1} g_j\right) \;\succeq\; m \left(g_i^{-1} g_j\right) \;=\; m(m^{-1}M) \;=\; M$$

and so $g_j = M$ and $g_i = m$. Thus A appears exactly once in W.

Suppose $W = AD$ where A is the maximal ascent in W and $D \neq 1_{\mathcal{F}}$. Then $D \prec 1_{\mathcal{F}}$ since otherwise $AD > A$ contradicting the maximality of A. If D has a prefix $U \succ 1_{\mathcal{F}}$, then $AU \succ A$ which is a contradiction. If D has a suffix $U \succ 1_{\mathcal{F}}$ then $UA \succ A$ contradicting the maximality of A. Thus D is a descent. $\square$

Remark 2.4. When $W = AD$, the peak M is A and the low m is $1_{\mathcal{F}}$.

Each cyclic permutation of $W = AD$ appears as a subword of W^n with $n \geq 2$. To show that A is uniquely positioned in W, it suffices to show that W^2 has exactly 2 subwords equivalent to A. That is, A appears as a subword of W^2 in only the expected positions which are $W^2 = ADAD$. In general, there are $n \geq 2$ occurrences of the subword A in $W^n = \underbrace{AD \cdots AD}_{n \text{ times}}$. To prove the main theorem, we will show that W^2 contains exactly two occurrences of the maximal ascent A.

Definition 2.5. Let $W = AD$ and $W' = AD'$ be overlapping cyclic permutations in W^n where $W' = U^{-1}WU$ and $n \geq 2$. A *cascade* in W^n induced by the *shift* U is a sequence of concatenated subwords $\{U_i\}_{i \geq 0}$ in W where $U_i \equiv U$ for each i. See Figure 2.

Proposition 2.6. *Let $W \in F(X)$ be a cyclically reduced nonperiodic word. Then the maximal ascent in W is uniquely positioned.*

Proof. Let A be the maximal ascent in W and suppose without loss of generality that $W = AD$. If $D = 1_{\mathcal{F}}$, then $W = A$ and so A is uniquely positioned in W since W is nonperiodic and thus all its cyclic permutations are distinct. If $D \neq 1_{\mathcal{F}}$, then by Lemma 2.3 D is a descent. Following Remark 2.4, suppose $W^2 = ADAD$ contains a third occurrence of A. Since A appears exactly once in each W, the third occurrence of A begins in the first factor W and ends in the second one. Let $W' = AD'$ be a cyclic permutation of W starting with A. Then D' is a descent with $|D'| = |D|$. Since W is nonperiodic, it has distinct

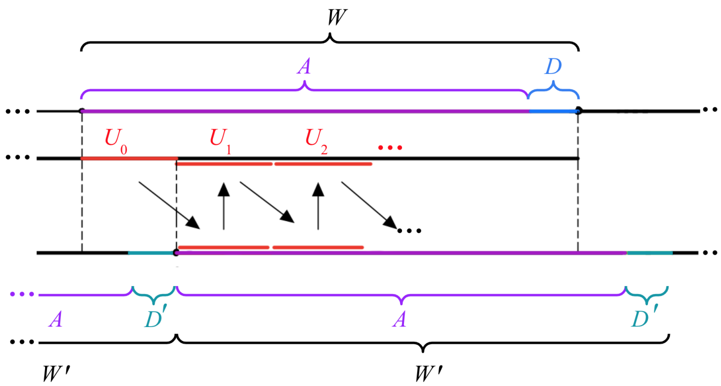

FIGURE 2. Since U is a prefix in AD, it is a prefix in AD'. The rightward shift of AD' by U induces a cascade $U = U_0, U_1, U_2, \ldots$ in AD, where the U_i concatenate and $U_i \equiv U$ for each i.

conjugates, and so $D' \not\equiv D$. Moreover, W has a prefix U_0 such that $W' = U_0^{-1} W U_0$. The shift U_0 induces a cascade $U_0, U_1, U_2, \ldots$ in $W = AD$.

Claim : $|D| < |U_0| < |A|$ and $|U_0|$ is not a divisor of $|W|$.

PROOF OF CLAIM. We have $W' = AD' = U_0^{-1} A D U_0$ with $|U_0| < |W| = |A| + |D|$. If $|U_0| > |A|$, then W', and thus A, begins in the interior of D. Since ascents and descents do not overlap, the ascent A is internal in D, and so A is not unique in W, which contradicts Lemma 2.3. If $|U_0| < |D|$, then the ascent A appearing in W' ends in the interior of D which is a contradiction. If $|U_0| = |D|$, then A and D have a common suffix which is impossible.

Since the subwords U_i are concatenated, if $\dfrac{|W|}{|U_0|} = k \in \mathbb{N}$, then $W = U_0^k$ which is a contradiction. $\square$

Note that $U_0 = A_1 D'$ where A_1 is an ascent. Indeed, U_0 is a prefix of A (in W) and so each prefix of A_1 is $\succ 1_{\mathcal{F}}$; and A_1 is a suffix of A (in W') and so each suffix of A_1 is $\succ 1_{\mathcal{F}}$.

Consider the cascade induced by $U_0 = A_1 D'$. Then $W = AD$ is a proper subword of U_0^n for some $n > 1$. Indeed, the cascade ensures that the ascent A in W appears as a subword of U_0^n for some $n > 1$. However, since $U_0 \equiv A_1 D'$ where A_1 is an ascent and D' is a descent, and A_1 does not overlap with D in $W = AD$, the only possibility is for A_1 to be long enough so that the last occurrence of A_1 in $W = AD$ must begin in A, contain D, and end in the

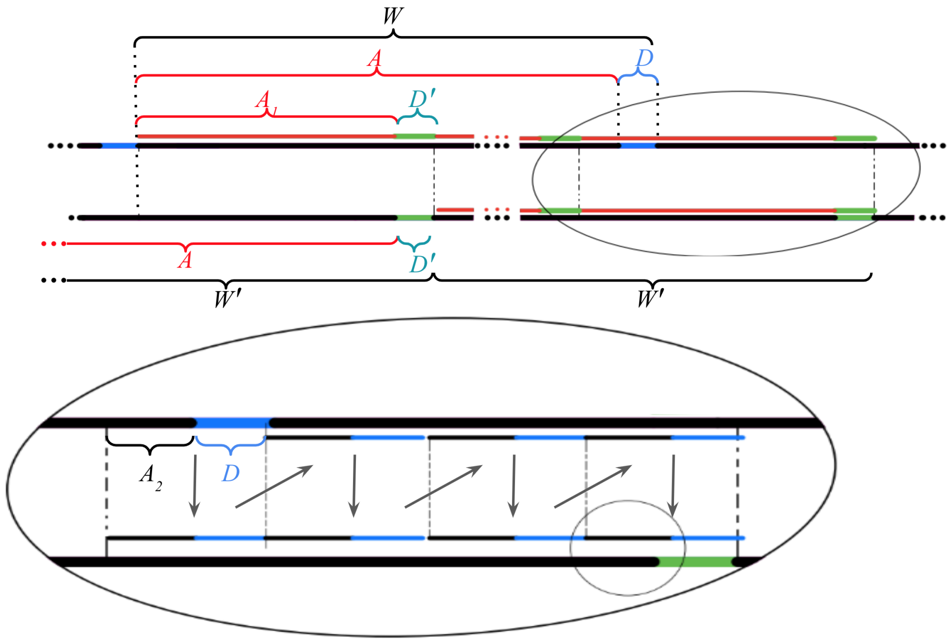

FIGURE 3. The cascade induced by $U_0 = A_1D'$ terminates with D appearing as an internal subword of A_1. See the region inside the ellipse. So A_1 has a prefix $U_2 = A_2D$ which itself is a shift that induces a cascade in A_1D'. In this example, the cascade induced by A_2D immediately leads to a contradiction since it forces an overlap between the ascent A_2 and the descent D'. See the region inside the circle.

interior of the second W factor in W^2. Let U_n be the first term that is not in W. Then $U_n \equiv A_1D'$ and $W = AD$ overlap. So D appears as a subword of U_0. For D to appear as a subword of U_0, it is necessary that D appears as an internal subword of A_1, since D is not equivalent to D' and D does not overlap with the ascent A_1. So we have a new overlap of two subwords equivalent to U_1. See Figure 3. The shift $U_2 = A_2D$ induces a new cascade that follows the same pattern as above with the difference being that $U_2 = A_2D$ is a concatenation of an ascent A_2, with $|A_2| < |A_1| < |A|$, and D instead of D'. Note that by the above Claim, $|D'| = |D| < |U_2| < |A_1|$. Once again, the cascade of copies of A_2D requires that the copies of D must not coincide with D' and cannot overlap with A_2. So D appears as a subword of A_2. Thus A_2 contains a subword $U_3 = A_3D'$ where $|A_3| < |A_2|$ is an ascent. As this process repeats, the shift U_j will be a concatenation of an ascent A_j and the descent D if j is even, and U_j will be a concatenation of an ascent A_j and the descent D' if j is odd. For each cascade, the ascent A_j in the shift $U_j = A_j D$ (or $U_j = A_j D'$) has shorter length. Since $|W| < \infty$ and by induction, this process terminates with either an overlap of an ascent and a descent, or by forcing $D = D'$. Both lead to contradictions. $\square$

3. Magnus Bi-order

The following describes an explicit bi-order on free groups due to Magnus [**MKS66**]. Let $\mathcal{F} = \mathcal{F}(x_1, x_2)$ be the free group on generators x_1, x_2. Let $\Lambda = \mathbb{Z}[[X_1, X_2]]$ be the ring of formal power series in the non-commuting variables X_1 and X_2, one for each generator of $\mathcal{F}$. Define the multiplicative homomorphism $\mu : \mathcal{F} \to \Lambda$ as:

$$\mu : \begin{cases} x_i & \mapsto 1 + X_i \\ x_i^{-1} & \mapsto 1 - X_i + X_i^2 - X_i^3 + \cdots \end{cases}$$

For example:

$$\begin{aligned} \mu\left(x_1 x_2^{-1}\right) &= (1 + X_1)(1 - X_2 + X_2^2 - X_2^3 + \cdots) \\ &= 1 + X_1 - X_2 + O(2) \end{aligned}$$

where $O(n)$ refers to the sum of all terms of order $\geq n$. Then μ is injective and $\mathcal{F}$ embeds in the group of units $1 + O(1) \subset \Lambda$. Order the elements of Λ as follows. First adopt the convention of writing the elements of Λ in standard form starting from lower degree terms in an increasing order. Then, order the terms with the same degree lexicographically where $X_1 \succ X_2$. Compare two elements of Λ according to the coefficients of the first term at which they differ. For example, $1 + X_1 + 3X_2 + O(2) \succ 1 + X_1 + X_2 + O(2)$ since the first term at which they differ is X_2, and the coefficient of X_2 in the first element is greater than the coefficient of X_2 in the second one. Under this order, $1 + O(1) \subset \Lambda$ is a bi-ordered group.

Define an ordering $\succ$ on $\mathcal{F}$ by:

$$v \succ w \iff \mu(v) \succ \mu(w)$$

It is readily verified that $\succ$ is both left and right invariant.

Definition 3.1. A word $W = y_1 \cdots y_n \in \mathcal{F}(X)$ is *monotonic* if either $y_i \in X$ for each $1 \le i \le n$, or $y_i \in X^{-1}$ for each $1 \le i \le n$.

Theorem 3.2. *Let $X = \{x_1, x_2\}$ be an alphabet and let $\mathcal{F} = \mathcal{F}(X)$ be the free group on X equipped with a bi-order $\prec$. Let $W \in \mathcal{F}$ be a cyclically reduced nonperiodic word of length > 1. Then W has a cyclic permutation $W' = AD$ where:*

(1) A is the uniquely positioned maximal ascent in W.

(2) If D is not uniquely positioned, then it appears as an internal subword of A.

(3) Using the Magnus ordering on $\mathcal{F}$, we have $D = 1_{\mathcal{F}}$ if and only if W is monotonic.

Proof. Let A be the maximal ascent in W. By Proposition 2.6, A is uniquely positioned.

Let $W' = AD$ where D is not uniquely positioned in W. If $D \neq 1_{\mathcal{F}}$, then by Lemma 2.3, D is a descent. Suppose $D' \equiv D$ is a subword in AD. By Remark 2.2, D' has no overlap with A, and so it appears as an internal subword of A. Moreover, if $D = 1_{\mathcal{F}}$, then $W = A$ and D is the empty word between any concatenated subwords of W. Thus D appears as an internal, albeit trivial, subword of W. Note that by assumption, $|W| > 1$ and so W has at least two subwords.

Choose the Magnus bi-ordering corresponding to $x_1 \succ x_2 \succ 1_{\mathcal{F}}$. Then any nonempty word in X is $\succ 1_{\mathcal{F}}$. If W is monotonic, then so is each cyclic permutation of W. Suppose without loss of generality that $W' = y_1 \cdots y_n$ with $y_i \in X$, for $1 \le i \le n$. By the maximality of A, if $D \neq 1_{\mathcal{F}}$, then it is a descent, which is impossible since all monotonic words are $\succ 1_{\mathcal{F}}$.

Suppose $D = 1_{\mathcal{F}}$. Then W' is the maximal ascent. Suppose W' is not monotonic. Then $W = Ux_1^{-1}V$ for some words $U, V \in \mathcal{F}$. The case $W = Ux_2^{-1}V$ is similar. Let

$$\mu(U) = 1 + M_1X_1 + M_2X_2 + O(2) \quad \text{and} \quad \mu(V) = 1 + N_1X_1 + N_2X_2 + O(2)$$

where $M_i, N_i \in \mathbb{Z}$. Then

$$\mu(VU) = 1 + (M_1 + N_1)X_1 + (M_2 + N_2)X_2 + O(2)$$

Moreover, we have

$$\begin{aligned} \mu(W') &= (1 + M_1X_1 + M_2X_2 + O(2))(1 - X_1 + O(2))(1 + N_1X_1 + N_2X_2 + O(2)) \\ &= 1 + (M_1 + N_1 - 1)X_1 + (M_2 + N_2)X_2 + O(2) \end{aligned}$$

Hence $W' \prec VU$. But VU is a prefix of the cyclic permutation $x_1U^{-1}W'Ux_1^{-1} = VUx_1^{-1}$ contradicting the maximality of W'. $\square$

Bibliography

[ABC$^+$91] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 3–63. World Sci. Publ., River Edge, NJ, 1991. Edited by H. Short.

[Bau72] Gilbert Baumslag. A non-cyclic, locally free, free-by-cyclic group all of whose finite factor groups are cyclic. *Bull. Austral. Math. Soc.*, 6:313–314, 1972.

[Bau93] Gilbert Baumslag. *Topics in combinatorial group theory*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 1993.

[BH72] R. G. Burns and V. W. D. Hale. A note on group rings of certain torsion-free groups. *Canad. Math. Bull.*, 15:441–445, 1972.

[BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

[BW13] Hadi Bigdely and Daniel T. Wise. Quasiconvexity and relatively hyperbolic groups that split. *Michigan Math. J.*, 62(2):387–406, 2013.

[BW22] Jacob Bamberger and Daniel T. Wise. Failure of the finitely generated intersection property for ascending HNN extensions of free groups. *Internat. J. Algebra Comput.*, 32(5):885–893, 2022.

[CP91] Maxime Crochemore and Dominique Perrin. Two-way string-matching. *J. Assoc. Comput. Mach.*, 38(3):651–675, 1991.

[CV78] Yves Césari and Max Vincent. Une caractérisation des mots périodiques. *C. R. Acad. Sci. Paris Sér. A-B*, 286(24):A1175–A1177, 1978.

[CW22] William Chong and Daniel T. Wise. Embedding of one-sided clean hnn extensions in ascending hnn extensions. 32(5), 2022.

[DH92] Andrew J. Duncan and James Howie. Weinbaum's conjecture on unique subwords of nonperiodic words. *Proc. Amer. Math. Soc.*, 115(4):947–954, 1992.

[DNR14] Bertrand Deroin, Andrés Navas, and Cristóbal Rivas. Groups, orders, and dynamics. 2014.

[FH99] Mark Feighn and Michael Handel. Mapping tori of free group automorphisms are coherent. *Ann. of Math. (2)*, 149(3):1061–1077, 1999.

[GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Her95] Bernhard Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15(3):365–371, 1995.

[Her98] Bernhard Herwig. Extending partial isomorphisms for the small index property of many ω-categorical structures. *Israel J. Math.*, 107:93–123, 1998.

[HL00] Bernhard Herwig and Daniel Lascar. Extending partial automorphisms and the profinite topology on free groups. *Trans. Amer. Math. Soc.*, 352(5):1985–2021, 2000.

[HN02] Tero Harju and Dirk Nowotka. Density of critical factorizations. *Theor. Inform. Appl.*, 36(3):315–327, 2002.

[HN06] Tero Harju and Dirk Nowotka. On unique factorizations of primitive words. *Theoret. Comput. Sci.*, 356(1-2):186–189, 2006.

[How54] A. G. Howson. On the intersection of finitely generated free groups. *J. London Math. Soc.*, 29:428–434, 1954.

[Hru92] Ehud Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12(4):411–416, 1992.

[HW10] Frédéric Haglund and Daniel T. Wise. Coxeter groups are virtually special. *Adv. Math.*, 224(5):1890–1903, 2010.

[HW22] Frédéric Haglund and Daniel T. Wise. Vh-ification. pages 1–33, 2022.

[LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89. Springer-Verlag, Berlin-New York, 1977.

[Lyn54] R. C. Lyndon. On Burnside's problem. *Trans. Amer. Math. Soc.*, 77:202–215, 1954.

[Mit04] Mahan Mitra. Height in splittings of hyperbolic groups. *Proc. Indian Acad. Sci. Math. Sci.*, 114(1):39–54, 2004.

[MKS66] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory: Presentations of groups in terms of generators and relations.* Interscience Publishers [John Wiley & Sons, Inc.], New York-London-Sydney, 1966.

[Rey11] Patrick Reese Reynolds. *Dynamics of irreducible endomorphisms of F(N).* ProQuest LLC, Ann Arbor, MI, 2011. Thesis (Ph.D.)–University of Illinois at Urbana-Champaign.

[RZ93] Luis Ribes and Pavel A. Zalesskii. On the profinite topology on a free group. *Bull. London Math. Soc.*, 25(1):37–43, 1993.

[Shi47] H. Shimbireva. On the theory of partially ordered groups. *Rec. Math. [Mat. Sbornik] N.S.*, 20(62):145–178, 1947.

[Sho91] Hamish Short. Quasiconvexity and a theorem of Howson's. In É. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 168–176. World Sci. Publishing, River Edge, NJ, 1991.

[Sol05] S. Solecki. Extending partial isometries. *Israel J. Math.*, 150:315–331, 2005.

[Wei90] C. M. Weinbaum. Unique subwords in nonperiodic words. *Proc. Amer. Math. Soc.*, 109(3):615–619, 1990.

[Wis12] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012.

[Wis20] Daniel T. Wise. An Invitation to Coherent Groups. In Dylan Thurston, editor, *What's Next?: The Mathematical Legacy of William P. Thurston*, pages 1–89. Princeton University Press, 2020. To appear.

[Wis22] Daniel T. Wise. Coherence, local indicability and nonpositive immersions. *J. Inst. Math. Jussieu*, 21(2):659–674, 2022.