Groups Geom. Dyn. 15 (2021), 691–754
DOI 10.4171/GGD/619

Cubulating one-relator products with torsion

Ben Stucky

Abstract. We generalize results of Lauer and Wise to show that a one-relator product of locally indicable groups whose defining relator has exponent at least 4 admits a proper and cocompact action on a CAT(0) cube complex if the factors do.

Mathematics Subject Classification (2020). 20E06, 20E25, 20F05, 20F06, 20F65, 20F67, 57M05, 57M07, 57M10, 57M12, 57M60.

Keywords. Cubulated, one-relator, staggered complex, CAT(0) cube complex, locally indicable, relatively hyperbolic, relatively quasiconvex, van Kampen diagram.

Contents

1 Introduction 691
2 Preliminaries 696
3 Reduced diagrams and extreme 2-cells 702
4 Additional extreme 2-cells 710
5 Geometry of the universal cover 714
6 Relative hyperbolicity 723
7 Walls and trellises 726
8 Walls embed and separate 730
9 Walls are relatively quasiconvex 733
10 Bridges and linear separation 739
11 Existence of the action 747
References 752

1. Introduction

Much effort has been devoted to studying groups which act properly and cocompactly on CAT(0) cube complexes, henceforth referred to as *cubulable groups*, in recent years. Their most famous appearance is in the resolution of the Virtual Haken Conjecture by Agol and Wise, building on work of Bergeron and Wise, Kahn and Markovic, Perelman, Thurston, and others, in which the cubulation of hyperbolic 3-manifold groups is featured prominently [2, 16, 22, 23, 28]. Simply knowing that a group is cubulable is sufficient to conclude a good deal of structural information about it. For instance, these groups satisfy a Tits alternative [25], admit a quadratic-time solution to the word problem [3], and satisfy the Novikov and Baum-Connes conjectures [4, 10]. Cubulable groups which have the stronger property of being *virtually special*, i.e., possess a finite index subgroup which embeds into a right-angled Artin group, enjoy stronger properties still, including separability of quasiconvex subgroups (if Gromov hyperbolic) and linearity [14, 31].

Aside from hyperbolic 3-manifold groups, many classes of groups have been shown to be cubulable, including $C'\big(\frac{1}{6}\big)$ small cancellation groups [29]. One-relator groups with torsion of exponent $n \geq 4$, groups which admit a presentation of the form $\langle a_1, \ldots, a_m \mid w^n \rangle$ with $n \geq 4$, were cubulated by Lauer and Wise in 2013 [17]. These groups are $C'\big(\frac{1}{6}\big)$ when $n \geq 6$. An extension of Wise's result for $C'\big(\frac{1}{6}\big)$ groups was given by Martin and Steenbock in 2014 when they successfully cubulated $C'\big(\frac{1}{6}\big)$ small cancellation free products of cubulable groups [20]. In 2017, Jankiewicz and Wise gave an alternative proof of Martin and Steenbock's result relying on Wise's cubical small cancellation theory developed in [30], which they proved for $C'(\frac{1}{20})$ small cancellation free products [15]. In the present article, we generalize Lauer and Wise's cubulation results for one-relator groups with torsion to the free product setting.

A group is *locally indicable* if every finitely generated subgroup admits $\mathbb{Z}$ as a homomorphic image. For an element w of a group G, let $\langle\langle w \rangle\rangle$ denote the normal closure of w in G. The following is our main theorem.

Theorem 1.1. *Let A and B be locally indicable, cubulable groups, w a word in $A * B$ which is not conjugate into A or B, and $n \geq 4$. Then $G = A * B/\langle\langle w^n \rangle\rangle$ is cubulable.*

We remark that this is implied by the results of [20] when $n \geq 6$ and [15] when $n \geq 20$.

To prove Theorem 1.1, we are motivated to pass to a broader class of groups; namely, we consider "staggered" quotients of free products of finitely many locally indicable, cubulable groups. The topological models for these groups are *staggered generalized 2-complexes*. See Section 2 for the definition of such a complex X and its *minimal exponent* $n(X)$. Theorem 1.1 follows from the more general statement below by taking X to be a dumbbell space for the free product $A * B$ with a 2-cell corresponding to w^n glued to it.

Theorem 1.2. *Let X be a compact staggered generalized 2-complex. Suppose that X has locally indicable, cubulable vertex groups and that $n(X) \geq 4$. Then $\pi_1(X)$ is cubulable.*

Wise uses his theory of quasiconvex heirarchies to prove a strong generalization of the main result in [17], namely that all one-relator groups with torsion are virtually special [30, Corollary 18.2]. One-relator groups with torsion are (Gromov) hyperbolic, so when the exponent of the defining relator in a one-relator group is at least 4, this result also follows from [17] and Agol's theorem that a hyperbolic, cubulable group is virtually special [1, Theorem 1.1].

Local indicability of A and B also implies that $G = A * B/\langle\langle w^n \rangle\rangle$ is hyperbolic relative to $\{A, B\}$, a fact we will recover in Proposition 6.4. Thus if A and B are hyperbolic themselves, then so is G [21, Corollary 2.41], and [1, Theorem 1.1] gives the following as a corollary to Theorem 1.1:

Corollary 1.3. *Suppose that A and B are locally indicable, hyperbolic, and cubulable. Let w be a word in $A * B$ which is not conjugate into A or B, and $n \geq 4$. Then $G = A * B/\langle\langle w^n \rangle\rangle$ is virtually special.*

Though we suspect that Theorem 1.2 is true when $n(X) \geq 2$, we unfortunately find it necessary to impose the restriction that $n(X) \geq 4$, just as Lauer and Wise do, when seeking to prove properness of the action. In contrast to Lauer and Wise's setting, it also appears that the condition that $n(X) \geq 4$ is necessary for the cocompactness argument.

Question 1.4. Do Theorems 1.1 and 1.2 hold when $n(X) \in \{2, 3\}$?

In view of the fact that one-relator groups with torsion are virtually special, the following question is intriguing (but well beyond the scope of the present article).

Question 1.5. Let A and B be locally indicable, virtually special groups, w a word in $A * B$ which is not conjugate into A or B, and $n \geq 2$. Is $G = A * B/\langle\langle w^n \rangle\rangle$ virtually special?

1.1. Methods. Our methods are topological, and the following is what might be described as a naive approach to proving Theorem 1.1 that nonetheless captures many of the main ideas. First, build a model space X for $G = A * B/\langle\langle w^n \rangle\rangle$ by starting with a *dumbbell space* $X_A \vee X_B$ of non-positively curved cube complexes with $\pi_1(X_A) = A$ and $\pi_1(X_B) = B$. Then, attach a 2-cell to a path corresponding to the word w^n, so that $\pi_1(X) = G$. See Figures 1 and 2. The task, then, is to build a G-invariant collection of walls in the universal cover, invoke a construction of a dual cube complex with a G-action due to Sageev [24], and prove that the walls are geometrically nice enough to verify properness and cocompactness of the action.

A prerequisite for this method to work is to get good control over the geometry of X. It is in doing so that we are motivated to pass to the staggered generalized 2-complexes mentioned previously, and of which dumbbell spaces are a particular example.

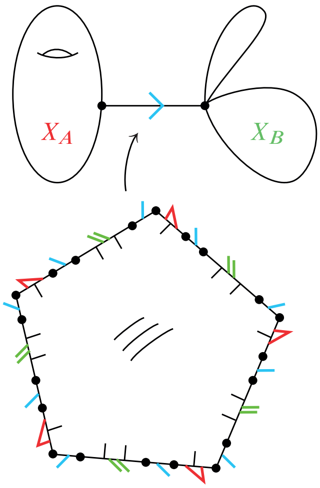

Figure 1. A presentation complex for G. The boundary path of the pentagonal cell corresponds to a word of the form w^5.

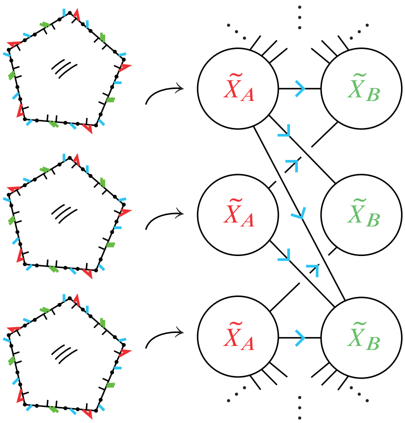

Figure 2. The universal cover of this presentation complex. We build our walls in this space by combining the Lauer–Wise walls considered in [17] (in the pentagonal cells) with the natural hyperplanes in the CAT(0) cube complex factors $\widetilde{X}_A$ and $\widetilde{X}_B$.

1.2. Outline. We follow the outline of [17] whenever possible. We define *staggered generalized* 2*-complexes* in Section 2. We also define the notion of a *tower* in this section, a fundamental tool for studying these complexes. Here we also establish results which illustrate the connections between staggerings, towers, and local indicability. The work in this section and the next is based heavily on work of James Howie [7, 8, 9].

Let G be the fundamental group of a staggered generalized 2-complex X with locally indicable, cubulable vertex groups and minimal exponent $n(X) \geq 2$. We prove geometric small cancellation results about *exposed* and *extreme* 2-cells in generalized van Kampen diagrams over G in Sections 3 and 4. These are strong statements about the local geometry of staggered generalized 2-complexes on which the rest of this work depends. These sections are direct generalizations of the work of [17].

In Section 5, we prove statements about the local geometry of a space $\bar{X}$ which is essentially the universal cover of X, and we develop a tool called *patchings* for producing the kinds of diagrams we can work with to prove results in later sections.

In Section 6, we recover relative hyperbolicity of G using Osin's idea of *linear relative Dehn functions* [21], which will be important for later arguments. The results up to this point in the outline do not depend on the fact that X has cubulable vertex groups.

We define the *walls* in $\bar{X}$ in Section 7, combining the Lauer–Wise walls of [17] with the natural walls in the portions of the universal cover which are already CAT(0) cube complexes. *Trellises* are defined as well – these are a convenient way to focus our study of the walls on the 2-skeleton of $\bar{X}$. We prove that walls embed and separate in Section 8.

At this point in the outline, we restrict to staggered generalized 2-complexes X with minimal exponent $n(X) \geq 4$.

We establish necessary conditions for the action on the dual cube complex to be cocompact in Section 9. Here the present work diverges from [17] significantly in order to deal with the fact that G is not a Gromov hyperbolic group, in general. We prove that wall stabilizers satisfy a property called *relative quasiconvexity*; this turns out to be the key to cocompactness of the action. Importantly, this argument involves attaching *combinatorial horoballs* (defined in [6]) to $\bar{X}$ to obtain a δ-hyperbolic space.

In Section 10, we show that the walls in $\bar{X}$ satisfy a criterion called *linear separation*, implying that the action on the dual cube complex is proper. This roughly means that the number of walls separating two points grows linearly with the distance between them.

We put everything together in Section 11. We use the Sageev construction to produce a dual cube complex with a G-action. Since our group is hyperbolic relative to the factors and our walls are relatively quasiconvex, a little more work allows us to apply a theorem of Hruska and Wise [13, Theorem 7.12] and prove cocompactness in this more general setting. Linear separation is used to show that the action is proper. Theorem 1.2 is proved in Theorem 11.5 and Theorem 1.1 is Corollary 11.6.

Note to the skimming reader. In an effort to make results as general as possible, additional assumptions are added as needed throughout the paper. We have tried to make this clear by making note of new standing assumptions at the beginning each section. The last place where new standing assumptions are added is Section 7.

Acknowledgments. The author wishes to thank Max Forester for his invaluable guidance throughout the duration of this project and without whom this work would not have been possible. Thanks as well to Paul Plummer and Jing Tao for helpful discussions, and to Noel Brady for helpful comments and questions during the post-production phase. He also wishes to thank the faculty and graduate students of Temple University for their hospitality and generosity in providing a place for him to work and discuss mathematics during the 2018–2019 academic year.

Finally, he wishes to thank the referee for providing extremely thorough feedback and productive suggestions which have greatly improved the exposition (including correcting many typos and several inaccuracies), and the editor for greatly improving the quality of the figures.

2. Preliminaries

2.1. Basic definitions and conventions. All maps are assumed to be continuous unless otherwise stated.

Definition 2.1 (graph of CW complexes). A *graph of CW complexes* Y is a connected CW complex which admits a construction as follows.

- Begin with a collection $\mathcal{V}$ of connected CW complexes called the *vertex spaces* of Y and an edge set $\mathcal{E}$. Let $\mathcal{V}^{(0)} = \bigcup_{V \in \mathcal{V}} V^{(0)}$ and let $i\colon \mathcal{E} \to \mathcal{V}^{(0)}$ and $t\colon \mathcal{E} \to \mathcal{V}^{(0)}$ be functions.
- For each $e \in \mathcal{E}$, attach a copy of $I = [0,1]$ to $\bigsqcup_{V \in \mathcal{V}} V$ by identifying 0 with $i(e)$ and 1 with $t(e)$. We call the copies of I coming from this step the *essential edges* of Y.

A graph of CW complexes is *finite* if $\mathcal{E}$ is finite.

Remark 2.2. A graph of CW complexes is similar to, but not quite the same as, a total space of a graph of spaces (with trivial edge spaces) in the sense of Scott and Wall [26]. An important distinction is that, in our case, the intersection of the essential edges with a given vertex space may consist of many distinct 0-cells of that vertex space. We will use the terminology "total space" in a slight abuse of notation in Definition 2.6.

Remark 2.3. The point above gives rise to a base point issue. There is no canonical graph of groups structure associated to a graph of CW complexes Y. However, we may associate a graph of groups to Y as follows.

1. To each vertex space V of Y, let T_V be an embedded path in $V^{(1)}$ which contains the intersection of the essential edges of Y with V.
2. The space obtained by collapsing each T_V to a point is homotopy equivalent to Y and is an honest total space of a graph of spaces in the sense of [26]. It thus gives rise to a natural graph of groups with trivial edge groups (using each collapsed T_V as a base point).
3. The essential edges in this total space form a connected graph. After choosing a maximal spanning tree T of essential edges in this graph, we may define the fundamental group of this graph of groups in the usual way.

Note that the fundamental group of this graph of groups is isomorphic to $\pi_1(Y)$, and the isomorphism is unique up to choice of T, each T_V, and base point in Y. We will casually refer to this algebraic structure as *the* graph of groups associated to Y when there is no danger of ambiguity.

Definition 2.4 (edge path/edge loop; reduced/cyclically reduced). Let Y be a CW complex. A map $I \to Y$ (resp. $S^1 \to Y$) is called an *edge path* (resp. *edge loop*) if there is a cell structure for I (resp. S^1) such that the map takes vertices to vertices and edges to edges. An edge path (resp. edge loop) is called *reduced* (resp. *cyclically reduced*) if it does not contain any backtracking.

Definition 2.5 (admissible cyclically reduced edge loop). Let Y be a graph of CW complexes. A cyclically reduced edge loop $S^1 \to Y$ is called *admissible* if it has the property that whenever a subpath of positive length maps to a loop in a single vertex space, then that loop is not nullhomotopic within that vertex space.

The following is a more topological definition of a staggered generalized 2-complex than that given in [10].

Definition 2.6 (staggered generalized 2-complex). A *staggered generalized 2-complex* X is a topological space which admits a construction as follows.

- Begin with a graph of CW complexes X_{tot} which we call the *total space*. Let $E(X)$ denote the set of essential edges of X_{tot}.
- Attach the elements of a (possibly empty) set of 2-cells to X_{tot} by their boundaries, with the property that each gluing is along an admissible cyclically reduced edge loop in X_{tot} and contains an edge of $E(X)$ in its image. Let $C(X)$ denote this set of 2-cells.
- We require that X admits a *staggering*:
    - a linear order on $C(X)$,
    - a linear order on $E(X)$,

- for $c, c' \in C(X)$, if $c < c'$ then $\max(c) < \max(c')$ and $\min(c) < \min(c')$, where $\min(c)$ is defined to be the least edge from $E(X)$ occurring in the attaching map for c, and similarly for $\max(c)$.

We call $C(X)$ the *essential* 2-cells of X and $E(X)$ the *essential* edges. When comparing cells of X we will sometimes use the notation $<_X$ to refer to the linear orders in the staggering. We will also sometimes write $\max_X(c)$ instead of $\max(c)$ to emphasize the staggering to which we are referring. We will reuse the phrase *vertex space of* X to refer to a subspace of X coming from a vertex space of X_{tot}. We will reuse the phrase *vertex group of* X to refer to the fundamental group of a vertex space. See Figure 3 for an example.

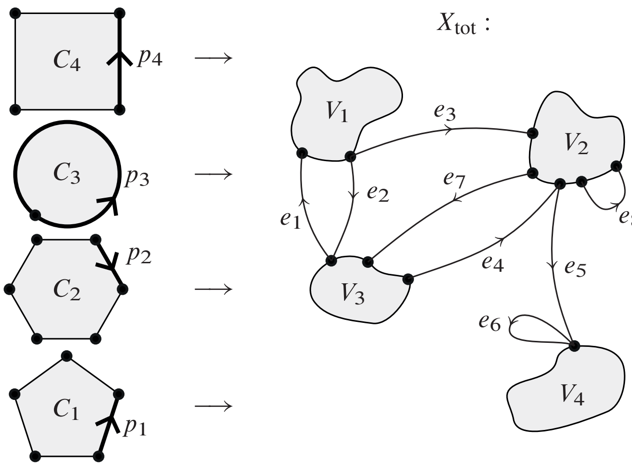

Figure 3. A staggered generalized 2-complex. The four elements of $C(X)$ are represented at the left of the figure. Each of the p_i specifies an admissible cyclically reduced edge loop and describes a portion of the attaching map of the corresponding c_i. Let the symbol v stand for an arbitrary path in any of the vertex spaces V_1, V_2, V_3, or V_4, and suppose that $p_1 = ve_1ve_2$, $p_2 = e_3ve_3^{-1}e_2ve_2^{-1}v$, $p_3 = e_4^{-1}ve_4ve_5ve_5^{-1}$, and $p_4 = ve_5e_6ve_5^{-1}ve^8$, and that the boundaries ∂c_1, ∂c_2, ∂c_3, ∂c_4 are labeled p_1^5, p_2^6, p_3^1, and p_4^4, respectively. Then the subscripts on $E(X)$ and $C(X)$ give a valid staggering ($e_1 < e_2 < e_3 < e_4 < e_5 < e_6 < e_7 < e_8$ and $c_1 < c_2 < c_3 < c_4$).

Definition 2.7 (exponent/proper power/minimal exponent $n(X)$). For an essential 2-cell α of $C(X)$, let R denote the admissible cyclically reduced edge loop in X_{tot} along which it is attached. Define the *exponent* of α to be $m(\alpha) = \max\{k \mid R = p^k \text{ for some loop } p\}$. If $m(\alpha) \geq 2$ we say that α is attached by a *proper power*. We define the *minimal exponent* $n(X) = \min\{m(\alpha) \mid \alpha \in C(X)\}$.

Definition 2.8 (tower/tower lift/maximal). A *tower* is a map $f: Y \to X$ between connected CW complexes such that $f = i_0 \circ p_1 \circ i_1 \circ \cdots \circ p_k \circ i_k$ where each i_j

is an inclusion of a finite subcomplex and each p_j is an infinite cyclic cover. Let K and X be connected CW complexes and $\psi: K \to X$ be a map. A *tower lift* is a map $\phi: K \to Y$ such that there is a tower $f: Y \to X$ and $\psi = f \circ \phi$. The map ϕ is called *maximal* if any tower lift $\phi': K \to Y'$ of ϕ has the property that the associated tower $f': Y' \to Y$ is a homeomorphism.

The following remark is straightforward, since it is easily verified for infinite cyclic covers (using the lifting criterion) and inclusions of finite subcomplexes.

Remark 2.9. If the attaching map of a 2-cell α in X is a proper power of exponent k, then for any 2-cell β in Y with $f(\beta) = \alpha$ under a tower $f: Y \to X$, the attaching map of β is a proper power of exponent k.

Definition 2.10 (indicable/locally indicable). A group is called *indicable* if it has $\mathbb{Z}$ as a quotient, and *locally indicable* if every nontrivial finitely generated subgroup is indicable.

Convention 2.11. Let Y be a CW complex.

1. Whenever we name an edge e in $Y^{(1)}$, we implicitly specify an orientation of e (or refer to one which has been previously defined). This orientation should not be taken to be absolute and may be modified freely when there is no risk of confusion or inconsistency.
2. In what follows, when we refer to a k-cell α in Y as a subspace of Y, it should be understood that this refers to all points x of that cell such that the preimage of x under the characteristic map has a neighborhood which is homeomorphic to $\mathbb{R}^k$, which we also call the *interior* of α in a slight abuse of terminology (depending on k). When we need to explicitly refer to the closure of α in Y, we will use the notation $\bar{\alpha}$.

Definition 2.12 (combinatorial map). A combinatorial map between CW complexes is one whose restriction to the interior of each cell is a homeomorphism.

2.2. The interplay between staggerings and local indicability. Let X be a staggered generalized 2-complex.

Let K be a compact and connected CW complex and let $\psi: K \to X$ be a combinatorial map. Howie shows [7, Lemma 3.1] that ψ has a maximal tower lift $\phi: K \to Y$. For us K will be an object similar to a van Kampen diagram, and we will use a maximal tower lift to study its boundary. By replacing Y with $\phi(K)$ and restricting the first inclusion, we may assume that maximal tower lifts are always surjective.

Note that a tower lift $\phi: K \to Y$ is not maximal if $\pi_1(K)$ is not indicable (e.g., if K is simply connected) and $\pi_1(Y)$ is. Indeed, for any nontrivial homomorphism

$g: \pi_1(Y) \to \mathbb{Z}$, Y admits an infinite cyclic cover $Y' \to Y$ corresponding to $\ker(g)$, and ϕ lifts since $\phi_*(\pi_1(K))$ lies in $\ker(g)$ by the fact that $\pi_1(K)$ is not indicable.

It is precisely this phenomenon which connects towers and local indicability. Informally, the map $K \to X$ may be hard to study because the image of K in X will be highly non-injective. In considering a maximal tower lift $K \to Y$, we will iteratively "unwind" the image of K through each successive infinite cyclic cover in the tower. Once at the top of the tower, we will use local indicability of vertex groups and other facts to draw conclusions about the boundary of K. This will in turn allow us to make conclusions about X.

The property of having a staggering is a flexible notion because it is preserved under towers. To see this, first note that staggerings are preserved under the maps a tower comprises:

Lemma 2.13. *If X is a staggered generalized 2-complex and $f: Y \to X$ is inclusion of a connected subcomplex or an infinite cyclic cover, then Y may also be expressed as a staggered generalized 2-complex.*

Proof. The essential cells of Y are exactly those which map to essential cells of X. In case f is inclusion of a connected subcomplex, note that the staggering of X restricts to a staggering of any subcomplex of X. In case f is an infinite cyclic cover, let ρ be a generator of the deck group of the cover, and define a "lexicographic" staggering on both the 1-cells and 2-cells of Y by the prescription that $\alpha < \beta$ if $f(\alpha) < f(\beta)$ or $\rho^k(\alpha) = \beta$ for some positive integer k. It is easy to check that this gives a staggering of Y. □

Lemma 2.14 (cf. [9, Lemma 2]). *If $f: Y \to X$ is a tower and X is a staggered generalized 2-complex, then Y may also be expressed as a staggered generalized 2-complex.*

Proof. Write $f = i_0 \circ p_1 \circ i_1 \circ \cdots \circ p_k \circ i_k$ where each i_j is an inclusion of a finite subcomplex and each p_j is an infinite cyclic cover. Apply Lemma 2.13 from left to right, starting with i_0. □

Remark 2.15. In general, there may be multiple ways to stagger Y. Whenever $Y \to X$ is a tower, we make the convention that the staggering on Y arises in the manner just described. This gives a unique staggering of Y up to choice of deck group generator of each cover.

Lemma 2.16 (cf. [9, Lemma 3] and [12, Lemma 2.6]). *Suppose that X is compact and has locally indicable vertex groups. Suppose additionally that X has no infinite cyclic cover and that α is the greatest essential 2-cell of X. If α is not attached along a proper power in $\pi_1(X_{tot})$, then X* ***collapses across α with free edge*** $\max \boldsymbol{\alpha}$, *i.e., X is homotopy equivalent to the complex obtained after removing α and* $\max \alpha$ *from X through a homotopy supported on $\bar{\alpha}$.*

Remark 2.17. Recall that the following are equivalent for any topological space Y: (1) $\pi_1(Y)$ is indicable; (2) $H^1(Y, \mathbb{Z}) \neq 0$; (3) Y has an infinite cyclic cover.

Proof. We follow Howie's proof in [9] – only minor changes are necessary. We say that X_{tot} is a *tree of spaces* if there is a unique choice of spanning tree T in the third step of Remark 2.3.

Now, if some essential 2-cell β is attached by a path of the form p^m in X_{tot} for some $m \geq 2$, then replacing β with the 2-cell β' attached by p will not affect $H^1(X)$, and giving β' the same position as β in the ordering of the 2-cells will not affect the staggering of X. So we may assume no essential 2-cell is attached by a proper power.

We induct on the number of essential 2-cells in X. If there is only one, then the rank of $H^1(X_{tot})$ is at most one, since $H^1(X) = 0$. If X_{tot} is a tree of spaces, then at most one vertex space can have nontrivial first cohomology by the Mayer-Vietoris Theorem. Also, since the attaching map of α is admissible and has positive length, there exists a closed subpath p' of the attaching map p of α which lies in a vertex space V of X_{tot} for which $H^1(V) = 0$. Since p is admissible, p' represents a nontrivial element g of $\pi_1(V)$. Since $\pi_1(V)$ is locally indicable and finitely generated since X is compact, we obtain a surjective map from $\pi_1(V)$ to $\mathbb{Z}$, giving us an infinite cyclic cover of V and contradicting that $H^1(V) = 0$. On the other hand, if X_{tot} is not a tree of spaces, then we must have $H^1(V) = 0$ for each vertex space, and there is a unique simple cycle in the underlying graph of X_{tot}. The attaching map of α must travel exactly once around this cycle, so that it uses $\max\alpha$ exactly once, and we can see that X collapses across α with free edge $\max\alpha$.

For the inductive step, consider the Mayer-Vietoris sequence

$$\cdots \to H^1(X) \to H^1(X \setminus \alpha) \oplus H^1(D^2) \to H^1(S^1) \to \cdots$$

associated to attaching α to the rest of X. Exactness shows that the rank of $H^1(X \setminus \alpha)$ is at most one. Let X' be the subcomplex of X formed by removing α and $\max\alpha$ from X. If X' is connected, then $H^1(X \setminus \alpha) = H^1(X') \oplus \mathbb{Z}$, so $H^1(X') = 0$. Otherwise X' has two components X_1 and X_2 (say), and $H^1(X \setminus \alpha) = H^1(X_1) \oplus H^1(X_2)$; assume without loss of generality that $H^1(X_1) = 0$. In this case, note that X_1 must contain at least one essential 2-cell whose attaching map lies entirely inside it. If not, then $H^1(X_1) = 0$ implies that X_1 is a tree of spaces, with each vertex space having trivial first cohomology. Then since the attaching map p of α uses X_1 and is admissible, there exists a closed subpath p' of p lying in some vertex space V of X_1 such that p' represents a nontrivial element g of $\pi_1(V)$. As before (using compactness of X), indicability of $\pi_1(V)$ gives rise to an infinite cyclic cover of V, contradicting that $H^1(V) = 0$.

Thus we may apply the inductive hypothesis either to X' (in case X' is connected) or X_1 (in case X' is not connected), but using the staggering *opposite* to that inherited from X (i.e., the orderings of the 1-cells and 2-cells are reversed). By induction, the complex in question collapses across its least essential 2-cell β (in the original ordering) with free edge $\min\beta$. But the attaching map of α does not use $\min\beta$ since $\beta < \alpha$, so X also collapses across β with free edge $\min\beta$. Let $X'' = X \setminus \{\beta, \min\beta\}$ be the result of this collapse.

Now X'' has fewer essential 2-cells than X, so again apply the inductive hypothesis to X'' (using the original ordering) to see that X'' collapses across α with free edge $\max\alpha$. But the attaching map of β does not use $\max\alpha$ since $\beta < \alpha$. Thus $X = X'' \cup \{\beta, \min\beta\}$ also collapses across α with free edge $\max\alpha$. □

Lemma 2.18 (cf. [17, Lemma 3.10] and [12, Lemma 2.7]). *Suppose that X is compact and has locally indicable vertex groups. Suppose additionally that X has no infinite cyclic cover and that α is the greatest essential 2-cell of X. Then α is attached along a path p^m where p is a closed path in X_{tot} passing through* $\max(\alpha)$ *exactly once. Moreover, no other 2-cell has the edge* $\max(\alpha)$ *in the image of its attaching map.*

Proof. The proof is identical to the proof of [12, Lemma 2.7], except that we appeal to Lemma 2.16 rather than [12, Lemma 2.6]. The idea is to replace α by its "m^{th} root" and apply the previous lemma. □

3. Reduced diagrams and extreme 2-cells

Throughout this section, let X be a staggered generalized 2-complex.

3.1. The topology of van Kampen diagrams over X. We will now prove some helpful results about "van Kampen diagrams" over X. For our purposes it will be useful to allow diagrams which are not planar. In what follows, the *boundary* of a 2-complex E, denoted ∂E, is the closure of the union of the 1-cells of E which lie in the image of the attaching map of at most one 2-cell of E.

Let $E \to X$ be a combinatorial map. We refer to cells of E as essential or not according to whether or not their images in X are essential.

Definition 3.1 (cancelable pair/folding edge/reduced/diagram). Let Y be a CW complex and E a 2-complex. Let $\phi\colon E \to Y$ be a combinatorial map. Let α and β be a pair of 2-cells of E with attaching maps Φ_α and Φ_β. We say that α and β form a *cancelable pair* if there is a decomposition of $\partial\alpha$ as a loop $e_1\sigma_1$ for some edge e_1 and a decomposition of $\partial\beta$ as a loop $e_2\sigma_2$ for some edge e_2 such that $\Phi_\alpha(e_1) = \Phi_\beta(e_2)$ and $\phi \circ \Phi_\alpha(\sigma_1) = \phi \circ \Phi_\beta(\sigma_2)$. We call $\Phi_\alpha(e_1) = \Phi_\beta(e_2)$ a *folding edge*. The map ϕ is called *reduced* if E does not contain a cancelable pair. It is called a *diagram* if E is compact and simply connected.

The following remarks are straightforward.

Remark 3.2. Let Y be a CW complex, $\psi\colon D \to Y$ a diagram, and $\phi\colon D \to Z$ a lift of ψ to a cover $Z \to Y$. Then ϕ is reduced if and only if ψ is reduced.

Remark 3.3. Let Y be a CW complex, $\psi\colon D \to Y$ a diagram, and $\phi\colon D \to T$ a maximal tower lift. Then ϕ is reduced if and only if ψ is reduced.

The following fundamental result is due to van Kampen:

Theorem 3.4. *Let Y be a CW complex and let u be a closed path in $Y^{(1)}$. Then u is nullhomotopic if and only if there exists a reduced diagram $D \to Y$ with D a planar 2-complex such that there is a parametrization of ∂D mapping to u.*

3.1.1. Finding exposed essential 2-cells

Definition 3.5 (position). Let $\phi\colon E \to X$ be a combinatorial map. Let α be an essential 2-cell of E such that $\phi(\alpha)$ is of exponent m and attached by a path of the form p^m in X. Two consistently-oriented 1-cells e_1 and e_2 on the boundary of α are in the same *position* in α if a subpath γ of $\partial\alpha$ running from the terminal 0-cell of e_1 to the terminal 0-cell of e_2 has the property that $\phi(\gamma)$ is a cyclic conjugate of p^j for some $j \in \mathbb{Z}$. For a 1-cell e in $\partial\alpha$ we let $[e]_\alpha$ denote the collection of the m 1-cells in the same position as e in α.

Definition 3.6 (external/internal/exposed). Let $\phi\colon E \to X$ be a combinatorial map. An essential 2-cell α in E is *external* if there is an essential 1-cell in $\partial\alpha \cap \partial E$; otherwise it is called *internal*. An essential 2-cell α in E is *exposed* if there is an essential 1-cell e in $\partial\alpha$ such that every 1-cell in $[e]_\alpha$ lies in ∂E. In this case we also say e is an *exposed* edge.

We emphasize that only essential edges can be exposed. Note that if $\phi\colon E \to X$ is a combinatorial map, then any total order $<_X$ of a set of cells of X (such as those coming from the staggering) induces an order of the preimages of those cells of X in E, which we will also denote by $<_X$. Since two cells of E may map to the same cell of X, it may be the case that $\alpha =_X \beta$ for cells α and β of E. In this sense, $<_X$ is a *quasi-order* on the cells of E. By Remark 2.15, if $E \to T$ is a tower lift of ϕ and $\alpha <_X \beta$ for essential cells α and β of E, then $\alpha <_T \beta$.

Definition 3.7 (adjacent/adjacent along). Let E be a CW complex. We say that 2-cells α and β are *adjacent* (*along* e) if there is an edge e belonging to $\partial\alpha\cap\partial\beta$. For a path $\gamma\colon I \to E$ in $E^{(1)}$, we say α is *adjacent to* γ *along* e if e lies in $\operatorname{im}(\gamma) \cap \partial\alpha$.

Lemma 3.8 (cf. [17, Lemma 4.7] and [12, Lemma 4.1]). *Suppose X has locally indicable vertex groups. Let $\phi\colon D \to T$ be a maximal tower lift of a reduced diagram $\psi\colon D \to X$. If α is a greatest (resp. least) 2-cell of D (under $<_T$), then α is exposed with exposed edge* $\max_T \alpha$ *(resp.* $\min_T \alpha$*). In particular, every reduced diagram $D \to X$ with at least one essential 2-cell has an exposed essential 2-cell.*

Proof. Note that T is compact since D is. Let α' be the unique greatest 2-cell of T. By Lemma 2.18, α' is the unique 2-cell whose attaching map uses the edge $\max \alpha'$, and it uses it exactly m times if m is the exponent of α'. Since we may assume that ϕ is surjective, $\phi(\alpha) = \alpha'$. Let e be an essential 1-cell of α mapping to $\max \alpha'$ under ϕ. Assuming α is not exposed in D, there is a 2-cell β of D adjacent to α along some essential 1-cell e' belonging to $[e]_\alpha$ which also maps to $\max \alpha'$. Since α' is the unique 2-cell using $\max \alpha'$, we must have $\phi(\beta) = \alpha'$. Since the attaching map of α' uses $\max \alpha'$ exactly m times and is a proper power of exponent m, we must have that σ_α, the longer path from the terminal to the initial vertex of e' in $\partial\alpha$, and σ_β, the analogous path in $\partial\beta$, must map to the same path in T. This shows that α and β form a cancelable pair and contradicts that the map ϕ is reduced (by Remark 3.3). □

3.1.2. Essential 2-cells embed in diagrams. Let α be an essential 2-cell of X which is attached along a closed path p. Our next goal is to use known results to prove that no proper closed subpath of p is nullhomotopic in X. This fact is stated as Lemma 3.13 below.

Definition 3.9 (Magnus subcomplex, cf. [17, Definition 3.6]). A *Magnus subcomplex $Z \subset X$* is a subcomplex with the following properties:

i. the subcomplex Z contains the union of all vertex spaces;
ii. if α is an essential 2-cell of X with the property that all essential boundary 1-cells of α lie in Z, then α lies in Z;
iii. the essential 1-cells of X contained in Z form an interval.

Note that we do not require Z to be connected.

The following lemma is equivalent to Howie's "locally indicable" Freiheitssatz [7, Theorem 4.3]. We will reprove it for completeness.

Lemma 3.10 (cf. [12, Theorem 6.1]). *Suppose that X has locally indicable vertex groups. If Z is a Magnus subcomplex of X, then the inclusion $i\colon Z \to X$ is π_1-injective for any choice of base point in Z.*

Remark 3.11. This implies in particular that the vertex groups of X embed in $\pi_1(X)$.

Proof. We follow the proof in [12] – only minimal modifications are necessary.

Let $g \in \ker i_*$, and let u be a topological representative for g in Z for some choice of base point. Then u is nullhomotopic in X, so we may apply Theorem 3.4 to construct a reduced diagram $\psi\colon D \to X$ where D is a compact planar 2-complex and $\psi(\partial D) = u$. We will show that every 2-cell of D maps to Z; this will imply u is nullhomotopic in Z and so $g = 1$ in $\pi_1(Z)$.

If every essential 1-cell in D maps to Z (or no essential 1-cells appear in D), then conditions (i) and (ii) imply that every 2-cell in D maps to Z and we are done. So suppose there is an essential 1-cell in D not mapping to Z (for brevity, say *D has a 1-cell not in Z*). Reversing the staggering of X if necessary, we may assume by condition (iii) that D has a 1-cell not in Z which is greater than any essential 1-cell in Z. Let $\phi\colon D \to T$ be a maximal tower lift of ψ. Note that for any edge $e \in D$ with the property that e is greater (under $<_X$) than any essential 1-cell in Z, e is greater (under $<_T$) than any essential 1-cell of T mapping to Z by the tower $T \to X$. Surjectivity of ϕ implies that the greatest essential 1-cell of T, which we call e', does not map to Z. Therefore no edge in $\phi^{-1}(e')$ lies in ∂D.

Since e' is in the image of the surjective map ϕ, this last fact implies that e' must lie on the boundary of some essential 2-cell in T. Thus e' is $\max_T \alpha$ for the greatest essential 2-cell α of T. Applying Lemma 3.8, any essential 2-cell in D mapping to α under ϕ is exposed with some exposed edge e'' in $\phi^{-1}(e')$. This contradicts that no edge in $\phi^{-1}(e')$ lies in ∂D. □

Recall the following fact, the proof of which is technical but requires only Bass-Serre theory and Howie's Freiheitssatz (see [8]):

Lemma 3.12 ([8, Proposition 3.3]). *Let $(\mathcal{G}, Y)$ be a graph of groups with trivial edge groups and locally indicable vertex groups. Let $w_1, \dots, w_m$ ($m \geq 2$) be reduced closed words in $(\mathcal{G}, Y)$, not all of Bass-Serre length zero, such that $w = w_1 \dots w_m$ is defined and is cyclically reduced (in the algebraic sense). Let N denote the normal closure of w in $\pi(\mathcal{G}, Y)$. If $w_1 N = \cdots = w_m N$, then either*

1. *$w_1 = \cdots = w_m \notin N$ or*
2. *all but one of the w_i are the empty word.*

We may use the previous two results to prove the following:

Lemma 3.13 (cf. [17, Corollary 3.9]). *Suppose that X has locally indicable vertex groups. Let p be a nontrivial proper subpath of the attaching map of an essential 2-cell α, and suppose that p is a closed path in X. Then p is not nullhomotopic in X.*

Proof. Let Z be the Magnus subcomplex of X consisting of all vertex spaces and the 2-cell α. Let Z' be the component of Z containing α, and let q be the closed path in X such that α is attached along the path pq. After following the procedure of Remark 2.3, note that $\pi_1(Z' \setminus \alpha)$ is isomorphic to the fundamental group $\pi(\mathcal{G}, Y)$ of a graph of groups $(\mathcal{G}, Y)$ satisfying the hypotheses of Proposition 3.12. Let w_p and w_q be the images in $\pi(\mathcal{G}, Y)$ of $[p]$ and $[q]$, respectively, under this isomorphism, and let $w = w_p w_q$. Since the attaching map of α is along an admissible cyclically reduced edge loop and includes an essential edge in its image, we conclude that neither w_p nor w_q is the empty word, and that at least one of w_p or w_q has positive length in $\pi(\mathcal{G}, Y)$. Let N be the normal closure of w in $\pi(\mathcal{G}, Y)$. Now we claim that $w_p \notin N$. Indeed, if $w_p \in N$, then Lemma 3.12 implies that $w_p N \neq w_q N$. On the other hand, $w_q N = (w_p N)^{-1}(w_p w_q N) = N$, so $w_p N = w_q N$, a contradiction.

Thus p is not nullhomotopic in Z'. But $\pi_1(Z) = \pi_1(Z')$ for appropriate choice of base point, and $\pi_1(Z')$ injects into $\pi_1(X)$ by Lemma 3.10. Thus p is not nullhomotopic in X. □

Corollary 3.14. *Suppose X has locally indicable vertex groups. Let $D \to X$ be a reduced diagram, and α an essential 2-cell of D. Then $\partial\alpha$ is embedded in D. In particular, $\bar{\alpha}$ is a simply connected subset of D.*

Proof. If $\partial\alpha$ is not embedded in D, there is a closed loop in $\partial\alpha$ which maps to a nontrivial proper subpath p of the attaching map of an essential 2-cell of X. A nullhomotopy of this closed loop in D (which is simply connected) gives rise to a nullhomotopy of p in X, contradicting Lemma 3.13. □

3.1.3. Other simply connected subdiagrams. We now observe some consequences of Corollary 3.14.

Definition 3.15 (internally intersects and other notation for paths). Let Z be a subspace of a space Y and $\gamma\colon I \to Y$ a path. We say γ *internally intersects* Z if $\gamma((0,1)) \cap Z \neq \emptyset$. We will frequently abuse notation and refer to $\gamma(I)$ as γ and $\gamma((0,1))$ as $\mathrm{int}(\gamma)$ when there is no risk of confusion. In case γ is an edge path in a CW complex, we will also use $|\gamma|$ to mean the number of edges in γ.

The following basic topological fact will be quite useful throughout. The proof is straightforward.

Lemma 3.16 (snipping lemma). *Let E be a simply connected 2-complex. Let γ be an embedded, locally separating arc in E between two points x and y in ∂E, and suppose that γ does not internally intersect ∂E. We call γ a* ***snipping arc.*** *Then $E \setminus \gamma$ is disconnected (i.e, γ is separating). In particular, suppose int$(\gamma) \cap E$ is contained in a single 2-cell α, and fix a parametrization $p\colon S^1 \to \partial\alpha$. Let v and w be two points of S^1 which lie in distinct components of $S^1 \setminus p^{-1}(\gamma)$. Then there is no path from $p(v)$ to $p(w)$ in $E \setminus \gamma$.*

Lemma 3.17 (cf. [17, Lemma 4.9]). *Suppose that X has locally indicable vertex groups. Let $D \to X$ be a reduced diagram. Suppose an essential 2-cell α of D is external. Let B be a component of $\overline{D \setminus \bar{\alpha}}$. Then $B \cap \bar{\alpha}$, B, and $B \cup \bar{\alpha}$ are all simply connected.*

Proof. By van Kampen's Theorem and Corollary 3.14, it suffices to prove that B and $B \cap \bar{\alpha}$ are simply connected.

Observe that $B \cap \bar{\alpha}$ is connected. To see this, suppose that $B \cap \bar{\alpha}$ is disconnected and pick points v and w in distinct components therein. Also choose two points v' and w' in distinct components of $\partial\alpha \setminus B$. Connect v' and w' by a snipping arc γ through the interior of α. The fact that there is a path from v to w in B (thus avoiding γ) contradicts the snipping lemma. Thus $B \cap \bar{\alpha}$ is connected.

Since α is external, and by Corollary 3.14, $B \cap \bar{\alpha}$ is homeomorphic to an interval and is thus simply connected.

To prove that B is simply connected, note that D is the union of B and $\overline{D \setminus B}$, and that $B \cap \overline{D \setminus B} = B \cap \bar{\alpha}$. Since D and $B \cap \bar{\alpha}$ are simply connected, so is B by van Kampen's Theorem. □

3.2. Branches, extreme 2-cells, and a Spelling Theorem

Definition 3.18 (branch). Let $D \to X$ be a reduced diagram. If α is an exposed 2-cell of D with exposed edge e, then the components of $\overline{D \setminus \bar{\alpha}}$ which contain at least one essential 2-cell are called the *branches* of D at (α, e).

Lemma 3.17 implies the following:

Lemma 3.19. *Suppose that X has locally indicable vertex groups. Let $D \to X$ be a reduced diagram, and suppose α is an exposed 2-cell of D with exposed edge e. Let B be a branch of D at (α, e). Then $B \cup \bar{\alpha}$ is simply connected.*

Definition 3.20 (auxiliary diagram/extreme). Let $\phi\colon E \to X$ be a combinatorial map. The *auxiliary diagram* $\check{E}$ associated to E is obtained from E by collapsing all maximal (under inclusion) connected regions of E which map to vertex spaces of X to points. For any subset S of E, denote the image of S in $\check{E}$ by $\check{S}$. Let α be an essential 2-cell of E of exponent m. We say that α is *extreme* if there is a subpath γ of $\partial\alpha$ (called an *extreme* subpath) such that γ contains the union of all m elements of $[e]_\alpha$ for some exposed edge e in α, and $\check{\gamma}$ does not internally intersect $\check{\bar{\beta}}$ for all essential 2-cells $\beta \neq \alpha$ of E.

Remark 3.21. All extreme 2-cells are exposed. When $m = 1$ the definitions of exposed and extreme coincide. A generic (planar) reduced diagram is depicted in Figure 4.

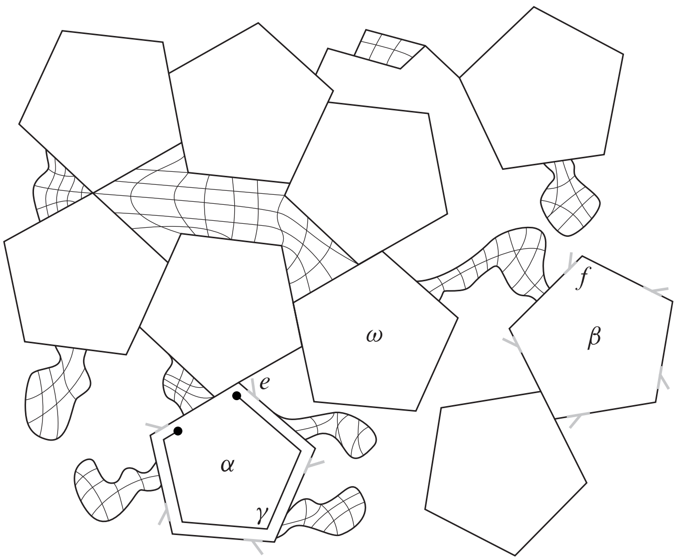

Figure 4. A generic planar reduced diagram. The regular bold pentagonal cells are exactly the essential 2-cells and persist in the auxiliary diagram. The essential 2-cell α is both exposed and extreme, as demonstrated by the path γ which contains all 5 elements of $[e]_\alpha$. The essential 2-cell β is exposed (consider $[f]_\beta$) but not extreme. The essential 2-cell ω is neither exposed nor extreme since at least $1/5^{\text{th}}$ of a contiguous portion of its boundary is internal.

Lemma 3.22. *Suppose that X has locally indicable vertex groups and let $\psi\colon D \to X$ be a reduced diagram. Let α be an exposed essential 2-cell in D with exposed edge e, and suppose that there is at most one branch of D at (α, e). Then α is extreme.*

Proof. This is obvious if there are *no* branches of D at (α, e), so assume there is exactly one and call it B. By Lemma 3.17, $B \cap \bar{\alpha}$ is contained in an arc of $\partial\alpha$ between two consecutive elements of $[e]_\alpha$, e_1 and e_2. Let γ be the arc of $\partial\alpha$ containing e_1 and e_2 which does not intersect B. Note that γ contains $[e]_\alpha$. Collapse D to the auxiliary diagram $\check{D}$. Let β be an essential 2-cell of B. Since γ does not internally intersect B, $\check{\gamma}$ does not internally intersect the closure of $\check{\beta}$. Thus α is extreme. ☐

We can now prove our first diagram result:

Proposition 3.23 (cf. [17, Theorem 4.11]). *Suppose that X has locally indicable vertex groups. Let $\psi\colon D \to X$ be a reduced diagram and suppose that D contains at least two essential 2-cells. Then D contains at least two extreme essential 2-cells.*

Proof. The proof is quite similar to that of [17, Theorem 4.11].

To prove the proposition, we induct on the number of essential 2-cells in D. Let $\phi\colon D \to T$ be a maximal tower lift of ψ with associated tower $f\colon T \to X$, and note that T is compact since D is.

First suppose there are exactly two essential 2-cells in D, α and β. Then α and β are both either greatest or least essential 2-cells (under $<_T$), and so Lemma 3.8 implies that they are both exposed. We claim that α and β are both extreme. To see α is extreme, let e be an exposed essential edge of α and note that there is a single branch B of D at (α, e). By Lemma 3.22, α is extreme. An identical argument shows that β is extreme.

For the inductive step, note first that we can find two exposed 2-cells α and β in D. Indeed, if T has only one essential 2-cell, then every essential 2-cell of D is a greatest 2-cell and so is exposed by Lemma 3.8, so choose α and β arbitrarily. On the other hand, if T has two or more essential 2-cells, and since ϕ is surjective, we can find a 2-cell in D (α, say) mapping to the greatest 2-cell of T, and a 2-cell in D (β, say) mapping to the least 2-cell of T; Lemma 3.8 implies that α and β are exposed. If α and β are extreme we are done, otherwise assume without loss of generality that α is not extreme. Then for an exposed edge e of α, there are at least two branches of D at (α, e) by Lemma 3.22. Call them B_1 and B_2. Now $B_1' = B_1 \cup \bar{\alpha}$ and $B_2' = B_2 \cup \bar{\alpha}$ are simply connected by Lemma 3.19, and thus $f \circ \phi_{|B_i'}$ is a reduced diagram for $i = 1, 2$ with fewer essential 2-cells than ψ. By induction, there is an extreme essential 2-cell $\alpha_1 \neq \alpha$ in B_1'. Observe that α_1 is also extreme in D since α separates B_1 from all other branches of D at (α, e). Similarly, we can find an extreme cell $\alpha_2 \neq \alpha$ in D which lies in B_2'. They are distinct since α_1 lies in B_1 and α_2 lies in B_2. □

Note. This generalizes part of the Spelling Theorem of Howie and Pride [10, Theorem 3.1(iii)], since the diagrams considered in that paper are planar.

The following is a simple topological criterion for identifying when an essential 2-cell in a diagram is *not* extreme which we record for reference. We will not use it until later.

Lemma 3.24. *Let $\phi\colon E \to X$ be a combinatorial map and let α be an essential 2-cell of E mapping to an essential 2-cell of X of exponent m. Suppose that there are two vertices x and y lying in $\partial\alpha$ with the following properties*:

i. *both paths from x to y in $\partial\alpha$ contain at least $|\partial\alpha|/m$ edges*;
ii. *each of the vertices $\check{x}$ and $\check{y}$ lies in the closure of at least two essential 2-cells in $\check{E}$.*

Then α is not extreme in E.

Proof. Let γ be a subpath of $\partial\alpha$ such that γ contains every 1-cell in $[e]_\alpha$ for an arbitrary essential edge e in α. Condition (i) implies that either x or y lies in $\operatorname{int}(\gamma)$, and condition (ii) implies that $\check{\gamma}$ internally intersects the closure of some 2-cell of $\check{E}$ other than the closure of $\check{\alpha}$. Thus α is not extreme. □

4. Additional extreme 2-cells

Standing assumptions from this point onward. Let X be a staggered generalized 2-complex with locally indicable vertex groups.

In this section, we will prove an additional criterion for the existence of extreme essential 2-cells in a reduced diagram $D \to X$.

Recall the main theorem from [8]:

Lemma 4.1 ([8, Theorem 4.2]). *Let A and B be locally indicable groups, and let G be the quotient of $A * B$ by the normal closure of a cyclically reduced word w (in the algebraic sense) of length at least 2 in the free product. Then the following are equivalent*:

i. *G is locally indicable*;
ii. *G is torsion free*;
iii. *w is not a proper power in $A * B$.*

Howie sketches the following corollary [8], which we prove here for completeness:

Corollary 4.2 (cf. [8, Corollary 4.5]). *Suppose X is such that the attaching map of each essential 2-cell is not a proper power. Then $\pi_1(X)$ is locally indicable.*

Proof. Consider the set of all staggered generalized 2-complexes X' which have all of the same data as X, except that $C(X')$ is a finite subset of $C(X)$. Then the set of the groups $\pi_1(X')$ forms a directed system for which $\pi_1(X)$ is the direct limit. Since a direct limit of locally indicable groups is locally indicable, it suffices to assume that $C(X)$ is finite.

Induct on the number of essential 2-cells in X.

For the base case, note that if there are no essential 2-cells in X, then $\pi_1(X)$ is locally indicable as the free product of locally indicable groups (by, e.g., the Kurosh Subgroup Theorem).

For the inductive step, let α be the greatest essential 2-cell of X and let $e = \max \alpha$. Then no other essential 2-cell uses e. If e separates $X \setminus \alpha$, then let X_A and X_B be the two components. Let $A = \pi_1(X_A)$, $B = \pi_1(X_B)$, and $w = [\partial\alpha]$. Now X_A and X_B are staggered generalized 2-complexes with locally indicable vertex groups and fewer essential 2-cells, and so A and B are locally indicable by induction. Now apply Lemma 4.1. On the other hand, if e does not separate $X \setminus \alpha$, we can see that $\pi_1(X \setminus \alpha)$ decomposes as a free product $A * \langle t \rangle$, where $A = \pi_1(X \setminus \{\alpha, e\})$ and t corresponds to a loop with winding number 1 over e, since no essential 2-cell uses e except α. Again observe that A is locally indicable by induction. Lemma 4.1 again applies with $B = \langle t \rangle$ and $w = [\partial\alpha]$ to give the result. □

We can use this fact to get a strong amplification of Remark 3.3:

Lemma 4.3 (cf. [17, Lemma 4.6]). *Let $\psi: D \to X$ be a reduced diagram. Let $\phi: D \to T$ be a maximal tower lift of ψ. If α and β are adjacent essential 2-cells of D then $\phi(\alpha) \neq \phi(\beta)$.*

Proof. The proof is in the same spirit as that of [17, Lemma 4.6].

Suppose that $\phi(\alpha) = \phi(\beta)$ and let e be a 1-cell in $\bar{\alpha} \cap \bar{\beta}$ (essential or not). Observe that $\psi(\alpha) = \psi(\beta)$. Let p be an edge loop of X such that p^m is the boundary path of $\psi(\alpha) = \psi(\beta)$ and p is not a proper power. By Remark 2.9, the boundary path of $\phi(\alpha) = \phi(\beta)$ is of the form $\hat{p}^m$ where $\hat{p}$ is a lift of p to T. In particular, $\hat{p}$ is a closed loop in T, and the map ϕ is periodic of period $|\hat{p}|$ on the edges of $\partial\alpha$. Let τ be the path of length $|\hat{p}|$ in $\partial\alpha$ which begins at the initial point of e and traverses e in the positive direction. Since α and β are identified under ϕ but do not form a cancelable pair, and all elements of $[e]_\alpha$ and $[e]_\beta$ map to $\phi(e)$, there an edge contained in τ which is not equal to e but which also maps to $\phi(e)$. This shows that there is a proper closed subpath of $\phi(\tau)$ in T. Thus there is a proper closed subpath γ of $\hat{p}$ in T. See Figure 5.

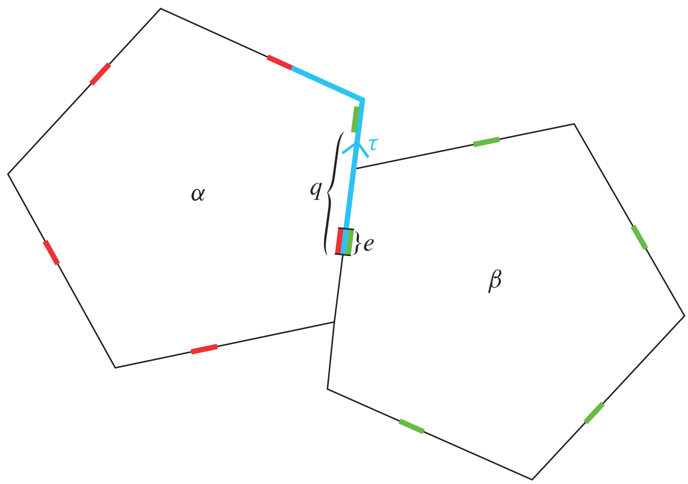

Figure 5. The path q (or possibly $q \setminus e$) maps to a closed loop in T.

Let X' be the staggered generalized 2-complex associated with X having nonperiodic attaching maps, built as follows. Let $X'_{\text{tot}} = X_{\text{tot}}$. Each element of $C(X)$ of exponent m gives rise to an exponent 1 element of $C(X')$ whose gluing map is the m^{th} root of the original, and the staggering on X' is the obvious one.

Consider the map $X \to X'$ which is the identity on X_{tot}, and an m-fold branched cover on each essential 2-cell if m is the exponent of that 2-cell. Let γ' be the image of γ in X'. By Lemma 3.13, γ' represents a nontrivial element of $\pi_1(X')$. Thus, via ϕ_*, $\pi_1(T)$ maps homomorphically to a nontrivial subgroup of $\pi_1(X')$, and that subgroup is finitely generated since T is compact. Since $\pi_1(X')$ is locally indicable by Corollary 4.2, there exists a surjective homomorphism $\pi_1(T) \to \mathbb{Z}$. Thus T has an infinite cylic cover and the tower lift $D \to T$ is not maximal, a contradiction. □

Now we can study connected subdiagrams of a reduced diagram:

Lemma 4.4 (cf. [17, Lemma 5.1]). *Let $D \to X$ be a reduced diagram. Let D' be a connected subcomplex of D, and let α be a greatest 2-cell of D'. Then α is exposed in D'.*

Note. The proof below is slightly more complicated than Lauer and Wise's proof of [17, Lemma 5.1]. There, the authors seem to assume that the analogue of the subcomplex B defined in the proof below is simply connected without justification.

Proof. Let $D \to T$ be a maximal tower lift of the diagram $D \to X$. By Lemma 4.3 applied to the map $D \to T$, each essential 2-cell adjacent to α in D' is strictly below α (under $<_T$). Let B be the smallest subcomplex of D' containing α and all 2-cells adjacent to α. Let B' be a minimal simply connected subcomplex of D containing B (under inclusion). Let $B' \to T'$ be a maximal tower lift of the composition $B' \hookrightarrow D \to T$, and let α' be a greatest essential 2-cell of B' under $<_{T'}$. Now Lemma 3.8 implies α' is exposed in B'. Note that since all essential 2-cells in $B \setminus \alpha$ are below α under $<_T$, they are also below α under $<_{T'}$. Thus $\alpha' \notin B \setminus \alpha$. If $\alpha' \neq \alpha$, then consider the component of $\overline{B' \setminus \overline{\alpha'}}$ containing α. This subcomplex of D contains B, is simply connected (by Lemma 3.17), and it is strictly contained in B'. This violates minimality of B'. Thus $\alpha' = \alpha$, so α is exposed in B'. But B' contains all 2-cells in D' adjacent to α, so α is also exposed in D'. □

Let $D \to X$ be a reduced diagram. Let V be the preimage in D of the disjoint union of the vertex spaces of X, and let α be an essential 2-cell of D. Define the following subcomplexes of D:

$$\widehat{G_\alpha} = \bigcup\{\bar{\beta} \in D \mid \beta \geq_X \alpha\} \cup V,$$

$$\widehat{L_\alpha} = \bigcup\{\bar{\beta} \in D \mid \beta <_X \alpha\} \cup \{\bar{\alpha}\} \cup V.$$

Let G_α and L_α be the components of $\widehat{G_\alpha}$ and $\widehat{L_\alpha}$, respectively, containing α.

Lemma 4.5 (cf. [17, Lemma 5.3]). *The components of $\widehat{G_\alpha}$ and $\widehat{L_\alpha}$ are simply connected.*

Proof. The proof is nearly identical to that of [17, Lemma 5.3]. We obtain $\widehat{G_\alpha}$ by successively removing the closure of a least essential 2-cell from D and passing to components of the closure of what remains. Reversing the staggering, Lemma 4.4 ensures that each successive essential 2-cell will be exposed, and Lemma 3.17 implies that removing each successive cell leaves simply connected components. In finitely many steps we obtain $\widehat{G_\alpha}$, and the argument is essentially the same for $\widehat{L_\alpha}$. □

We are ready to prove our second main criterion about extreme 2-cells in diagrams:

Proposition 4.6 (cf. [17, Theorem 5.4]). *Let $D \to X$ be a reduced diagram. If D has an internal essential 2-cell that maps to an exponent m 2-cell of X, then D contains at least $2m$ extreme 2-cells.*

Proof. The proof is essentially the same as that of [17, Theorem 5.4].

Let $D \to T$ be a maximal tower lift of $D \to X$, and let α be an internal essential 2-cell of D of exponent m. Define $\widehat{G_\alpha}$ and $\widehat{L_\alpha}$ with respect to $<_T$. Now Lemma 4.4 implies that α is exposed in both G_α and L_α, so there exist essential 1-cells e_G and e_L in α such that each 1-cell in $[e_G]_\alpha$ lies in ∂G_α and each 1-cell in $[e_L]_\alpha$ lies in ∂L_α. Since α is internal in D, $[e_G]_\alpha$ and $[e_L]_\alpha$ must be disjoint. By Lemma 4.5, $G_\alpha \to X$ is a reduced diagram, so Lemma 3.17 implies that each branch of G_α at (α, e_G) intersects $\partial\alpha$ in an arc. This fact, together with the fact that the m elements of $[e_L]_\alpha$ are internal in G_α (and the m elements of $[e_G]_\alpha$ lie on ∂G_α), implies that there are m branches of G_α at (α, e_G) adjacent to α along the m elements of $[e_L]_\alpha$. Call them $B_1, \dots, B_m$. Let G_i be the component of $\widehat{L_\alpha} \cup B_i$ containing α. Note that G_i contains at least one essential 2-cell strictly greater than α since B_i contains an essential 2-cell adjacent to α (applying Lemma 4.3 to $D \to T$). So any greatest 2-cell of G_i lies in B_i. Now Lemma 4.4 implies that there exists an essential 2-cell α' in B_i which is exposed in G_i. Note that α' is exposed in D since if β is a 2-cell of D adjacent to α', then β lies in $\widehat{L_\alpha}$ if $\beta < \alpha$ and lies in B_i if $\beta \geq \alpha$, so β lies in G_i already. Thus we obtain m distinct exposed 2-cells in D, one in each B_i, and all strictly greater than α.

We repeat almost the same argument for L_α to obtain m more distinct exposed 2-cells in D, all strictly less than α (in this case, the argument is actually simpler, as we don't need to apply Lemma 4.3). Thus we obtain $2m$ distinct exposed 2-cells in D. This completes the proof in the case $m = 1$, as the definitions of exposed and extreme coincide.

Thus assume $m \geq 2$, and let $\alpha_1, \dots, \alpha_{2m}$ be the $2m$ exposed 2-cells of D identified above. If α_i is not extreme, then D has at least two branches at (α_i, e_i)

for some e_i by Lemma 3.22. Let C_i be a branch not containing α, and note that $C_i \cup \overline{\alpha_i}$ is simply connected by Lemma 3.19. By Proposition 3.23, there are at least two extreme essential 2-cells in $C_i \cup \overline{\alpha_i}$; any one of these not equal to α_i is extreme in D. Call such a cell β_i and replace α_i by β_i in this case. Repeating for each i, we obtain $2m$ extreme 2-cells.

It remains to show that these $2m$ cells are distinct. Note first that α_i does not lie in C_j for any $i \neq j$, since α_i lies in the branch of D at (α_j, e_j) containing α. Thus $\alpha_i \neq \beta_j$ for all $i \neq j$ (when β_j is defined). On the other hand, if $\beta_i = \beta_j$ for some $i \neq j$, then $C_i \cap C_j \neq \emptyset$. Choose two elements of $[e_j]_{\alpha_j}$ such that there is no path in $\partial\alpha_j$ between C_j and the branch of D at (α_j, e_j) containing α, and connect these edges by a snipping arc running across the interior of α_j. Let B be the branch of D at (α_j, e_j) containing α (which contains α_i as remarked above). Now the complex $C_i \cup \overline{\alpha_i} \cup C_j \cup \overline{\alpha_j} \cup B$ shows that the snipping arc is non-separating, contradicting the snipping lemma (Lemma 3.16). Thus $\beta_i \neq \beta_j$ for all $i \neq j$ (when β_i and β_j are defined) and the $2m$ extreme 2-cells we have found are distinct. □

5. Geometry of the universal cover

Standing assumptions from this point onward. Let X be a staggered generalized 2-complex with locally indicable vertex groups. From now on, we assume that each essential 2-cell of X is attached by a proper power, that is, $n(X) \geq 2$.

We will soon be assuming that the vertex groups of X are cubulated. This section contains a collection of results about the geometry of X which do not depend on this assumption. In what follows, we will be working in a space $\overline{X}$ which is closely related to $\widetilde{X}$, the universal cover of X. Let Y denote the preimage of X_{tot} in $\widetilde{X}$.

Definition 5.1 ($\overline{X}$). By Lemma 3.10, $\pi_1(V)$ embeds naturally in $\pi_1(X)$ for each vertex space V of X, and thus Y may be viewed as a graph of simply connected CW complexes (each vertex space of which is $\widetilde{V}$ for some vertex space V of X) with preimages of essential edges running between them. Let $\overline{X}$ be the space obtained from $\widetilde{X}$ by identifying lifts of essential 2-cells of X which have the same attaching map up to cyclic permutation. The space $\overline{X}$ may be viewed as a subcomplex of $\widetilde{X}$ which contains Y, and there are thus combinatorial maps $\overline{X} \to \widetilde{X} \to X$.

Give $Y^{(1)}$ the combinatorial metric in which every edge has length 1. All of the metric statements in this section are really about $Y^{(1)} = \overline{X}^{(1)}$, and all paths of interest are edge paths.

5.1. Admissible pseudometrics and relative geodesics. We will work with paths in $\overline{X}$ which generalize geodesics. The idea of relative geodesics as defined below is that they allow for the possibility that paths may be "shorter than they look" in vertex spaces. At certain times in what follows, we will be strategically "augmenting" $\overline{X}$ in a manner which introduces this sort of behavior.

Definition 5.2 (admissible pseudometrics/relative length/relative geodesic). Let d denote the metric on $\overline{X}^{(1)}$ where every edge has length one. For each vertex space $\widetilde{V}$, choose a pseudometric $d_{\widetilde{V}}$ on $\widetilde{V}^{(0)}$. We require that this choice of pseudometrics is invariant with respect to the action of $G = \pi_1(X)$ on $\overline{X}$. If this holds we say the choice of pseudometrics is *admissible*.

Let $\gamma\colon I \to \overline{X}$ be an edge path whose endpoints are 0-cells x and y of $\overline{X}$. Decompose γ as a concatenation $\gamma_{v_1}e_1\dots\gamma_{v_k}e_k\gamma_{v_{k+1}}$, where each γ_{v_i} is a (possibly degenerate) maximal edge path mapping to a vertex space $\widetilde{V}_i$ of $\overline{X}$, and the e_i are essential edges. We define the *relative length* of γ, $\ell_r(\gamma)$, by the following formula:

$$\ell_r(\gamma) = k + \sum_{i=1}^{k+1} d_{\widetilde{V}_i}(i(\gamma_{v_i}), t(\gamma_{v_i})),$$

where $i(\lambda)$ and $t(\lambda)$ denote the initial and terminal vertices, respectively, of a path or edge λ. We say γ is a *relative geodesic* if $\ell_r(\gamma)$ is minimal among all paths from x to y. If we have not made an explicit choice of admissible pseudometrics on vertex spaces, the statement that γ is a relative geodesic should be taken to mean that there is a choice of admissible pseudometrics which makes γ a relative geodesic.

Some examples of admissible choices of pseudometrics are as follows (provided that the choices are made in a G-invariant manner):

- use the induced metric from $\overline{X}$. For some/all $\widetilde{V}$, define $d_{\widetilde{V}}(x, y) = d(x, y)$ for some/all $x, y \in \widetilde{V}^{(0)}$. Thus geodesics are relative geodesics;
- "Electrify" some/all $\widetilde{V}$ by defining $d_{\widetilde{V}}(x, y) = 0$ for all $x, y \in \widetilde{V}$;
- "cone off" some/all $\widetilde{V}$ by adding a new vertex and connecting all vertices of $\widetilde{V}$ to it by an edge of length 1/2, and define $d_{\widetilde{V}}$ by the metric this procedure induces, so that $d_{\widetilde{V}}(x, y) = 1$ for all distinct $x, y \in \widetilde{V}$;
- for some/all $\widetilde{V}$, choose $d_{\widetilde{V}}$ so that there is a constant C such that

  $$|d_{\widetilde{V}}(x, y) - 2\log(d(x, y) + 1)| < C$$

  for all $x, y \in \widetilde{V}$. This is the choice we will make later on when we attach so-called *combinatorial horoballs* to each $\widetilde{V}$.

5.2. Local geometry of essential 2-cells. The following fact is a crucially important statement about the boundaries of essential 2-cells in $\bar{X}$.

Lemma 5.3. *Suppose X is a staggered generalized 2-complex with locally indicable vertex groups and $n(X) \geq 2$. Let γ a relative geodesic in $\bar{X}$. Let e be an essential edge of an essential 2-cell α. Then there exists an element of $[e]_\alpha$ not contained in γ.*

Proof. Suppose that the lemma is false. Among all triples (α, e, γ) with the property that all members of $[e]_\alpha$ (for some essential edge e of some essential 2-cell α) lie in the relative geodesic γ, choose one for which $\ell_r(\gamma)$ is minimal. Note that γ contains at least two edges.

Label the elements of $[e]_\alpha$, $e_1, \ldots, e_m$ (where $m \geq 2$ is the exponent of α) in the order that they occur along γ, and orient them consistently with γ. Let $i(e_j)$ and $t(e_j)$ be the initial and terminal vertices, respectively, of e_j for $j \in \{1, \ldots, m\}$. We may also assume that the initial point of γ is $i(e_1)$ and the terminal point is $t(e_m)$ by removing edges from γ if necessary. Let σ_j be the subpath of γ between $t(e_j)$ and $i(e_{j+1})$, for $j \in \{1, \ldots, m-1\}$. Choose $\sigma \in \{\sigma_j\}$ such that $\ell_r(\sigma)$ is minimal. See Figure 6. Decompose the image of $\partial\alpha$ in X as a path p^m where p is not a proper power. Then p corresponds to an order m element w of $\pi_1(X)$ which acts on $\bar{X}$ by "rotation" through a point in the interior of α. Consider the paths $\{w^j\sigma\}$ for $j \in \{0, \ldots, m-1\}$. Each path connects two elements of $[e]_\alpha$, and the orbits chain together to form an m-pointed star shape with corners on members of $[e]_\alpha$ (there are two cases according to whether the $\{w^j\sigma\}$ meet at their endpoints or have endpoints separated by the elements of $[e]_\alpha$).

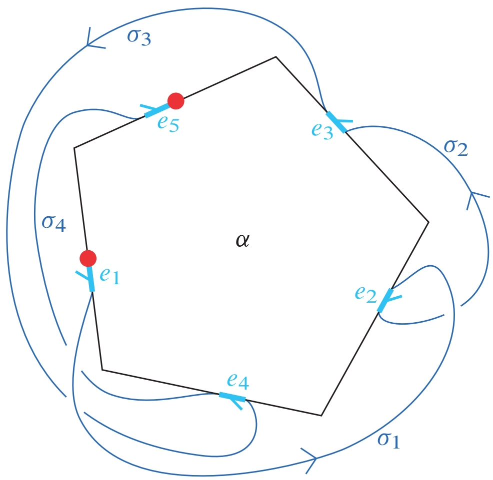

Figure 6. Decomposition of γ into the σ_j. Suppose that $\sigma = \sigma_4$.

Now, find a shortest relative path λ in $\bar{X}$ connecting $i(e_1)$ to $t(e_m)$ using only w-orbits of σ and members of $[e]_\alpha$. See Figure 7. It is clear that $\ell_r(\lambda) \leq \frac{m}{2}\ell_r(\sigma) + \frac{m}{2} + 1$. On the other hand, since γ is a relative geodesic with the same endpoints as λ, we have that $\ell_r(\lambda) \geq m\ell_r(\sigma) + m$. Unless $m = 2$, this contradicts the inequality

$$\frac{m}{2}\ell_r(\sigma) + \frac{m}{2} + 1 < m(\ell_r(\sigma) + 1),$$

which trivially holds whenever $\ell_r(\sigma) \geq 0$ and $m \geq 3$.

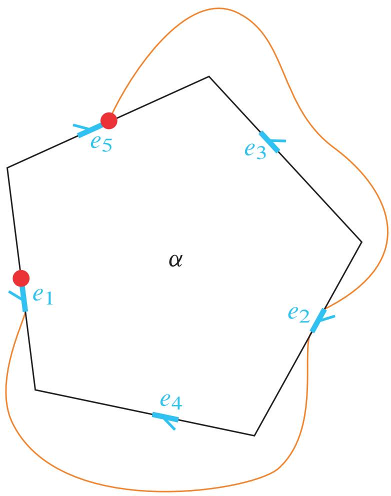

Figure 7. In this example, λ is made up of two orbits of σ and the edges e_1 and e_2.

Thus, we have reduced to the case $m = 2$. We may assume that σ does not cross e_1 or e_2, as this would provide an obvious way to decrease the relative length of γ. We may also assume that σ connects antipodal points of $\partial\alpha$, for otherwise $w\sigma$ connects $i(e_1)$ to $t(e_2)$ and $\ell_r(w\sigma) < \ell_r(\gamma)$ since $w\sigma$ avoids e_1 and e_2.

Let ϕ be the characteristic map of α. Let $\alpha' = \phi^{-1}(\alpha)$ and $e_i' = \phi^{-1}(e_i)$ for $i \in \{1,2\}$. Consider the complex $E = \overline{\alpha'} \bigsqcup_{\{\phi^{-1}(i(\sigma)),\phi^{-1}(t(\sigma))\}} \sigma$, which has $\pi_1(E) = \mathbb{Z}$ and a natural map to $\bar{X}$. The assumption of the previous paragraph implies that σ joins the two distinct components of $\partial\alpha' \setminus (e_1' \cup e_2')$. Let q be a cyclically reduced path in E which represents a generator of $\pi_1(E)$, and $D' \to \bar{X}$ a reduced disk diagram with boundary filling the image of q. Let $D = E \bigsqcup_q D'$. If D is not reduced, then there is an essential 2-cell β of D' such that α' and β form a cancelable pair and share an edge f in their common boundary. If this happens, then "fold" β over α' by identifying the paths $\partial\beta \setminus \{f\}$ and $\partial\alpha' \setminus \{f\}$ and deleting β from D. This is a homotopy equivalence and has the effect of modifying q and deleting an essential 2-cell from D'. This process terminates after finitely many steps, so we may assume that D is reduced. Note that ∂D is contained in $\partial\alpha' \cup \sigma$.

Note also that at most one of e_1' and e_2' lies in ∂D. Otherwise, connect a point of e_1' to a point of e_2' by a snipping arc running across the interior of α', and observe that the path σ contradicts the snipping lemma (Lemma 3.16) since D is simply connected. Without loss of generality, assume that e_1' is internal in D. Thus e_1' lies in the boundary of at least two distinct essential 2-cells of D.

Thus there exist at least two essential 2-cells in D. Choose an inclusion $\overline{X} \to \widetilde{X}$ and consider the natural reduced map $D \to X$. By Proposition 3.23, there is an extreme essential 2-cell β of D distinct from α' with exposed edge f, say. Since ∂D is contained in $\partial\alpha' \cup \sigma$, all elements of $[f]_\beta$ are contained in this complex as well. In fact, all elements of $[f]_\beta$ are contained in σ since they lie on the boundary of D. Thus the triple (β, f, σ) yields a counterexample to the lemma. The fact that $\ell_r(\sigma) < \ell_r(\gamma)$ contradicts minimality of (α, e, γ), and the lemma is proved. □

An immediate consequence of this lemma is convexity of vertex spaces.

Lemma 5.4. *The vertex spaces of $\overline{X}$ are convex.*

Reminder. We are using the path metric on $\overline{X}^{(1)}$.

Proof. Let γ be a geodesic edge path between vertices x and y of a vertex space $\widetilde{V}$. By passing to an innermost subpath outside of $\widetilde{V}$, we may assume for contradiction that $\gamma \cap \widetilde{V} = \{x, y\}$. Let γ' be a shortest path from x to y in $\widetilde{V}$. Fill the loop $\gamma(\gamma')^{-1}$ with a reduced diagram D using Theorem 3.4. Since all edges of γ lie outside of $\widetilde{V}$, γ contains an essential edge. Since D is simply connected, this edge shows that D contains an essential 2-cell. Lemma 3.8 implies that there is an exposed essential 2-cell α with exposed edge e. Since γ' consists only of edges which are not essential, all elements of $[e]_\alpha$ lie on γ. This contradicts Lemma 5.3. □

5.3. Patchings. The following construction is of critical importance for later arguments. It is a technique we will use to turn certain reduced maps $\phi\colon E \to \overline{X}$ into reduced diagrams without introducing extra exposed or extreme 2-cells.

Definition 5.5 (patching). Let $\phi\colon E \to \overline{X}$ be reduced, where E is compact but not necessarily simply connected. A *patching* for ϕ is a simply connected 2-complex $E_\#$ and a reduced diagram $\phi_\#\colon E_\# \to \overline{X}$ such that $E_\#$ contains E as a subcomplex, $\phi_\#|_E = \phi$, and none of the essential 2-cells of $E_\# \setminus E$ are exposed in $E_\#$.

Remark 5.6. Fix an inclusion of $\overline{X}$ into $\widetilde{X}$. In view of the composition $\overline{X} \to \widetilde{X} \to X$, reduced diagrams $D \to \overline{X}$ give rise to reduced diagrams $D \to X$ and vice versa by Remark 3.2. Whenever we have a patching $E_\# \to \overline{X}$, we will casually confuse it with the corresponding reduced diagram $E_\# \to X$ in order to apply Propositions 3.23 and 4.6.

An *isolated edge* of a CW complex is one which is not in the boundary of any 2-cell. The following lemma formalizes the induction required in the key patching lemma, Lemma 5.8.

Lemma 5.7. *Let F be a compact connected 2-complex which is a subcomplex of $\overline{X}$. Let F_+ be a compact connected 2-complex with the following properties*:

1. *F is a subcomplex of F_+;*
2. *F_+ admits a reduced combinatorial map*

$$\phi_+ \colon F_+ \longrightarrow \overline{X}$$

(*with $\phi_+|_F$ equal to the inclusion map*).

Let p be a cyclically reduced edge loop in F_+. Then there is a simply connected planar 2-complex D (homeomorphic to a disk if p is embedded) and boundary gluing map p' (homotopy equivalent to p in F_+, and which we can arrange to be embedded if p is), such that $F_{++} = F_+ \bigsqcup_{p'} D$ has the following properties:

i. *F_+ is a subcomplex of F_{++};*

ii. *F_{++} admits a reduced combinatorial map*

$$\phi_{++} \colon F_{++} \longrightarrow \overline{X}$$

(*with $\phi_{++}|_{F_+} = \phi_+$*).

Moreover, we can arrange that the following hold:

a. *any isolated edge in F_{++} is also isolated in F_+;*

b. *any exposed essential 2-cell α belonging to $F_{++} \setminus F_+$ with exposed edge e has the property that all elements of $[e]_\alpha$ are isolated edges of F_+.*

Proof. Fill the image loop $\phi_+(p)$ with a reduced planar disk diagram $D \to \overline{X}$ using Lemma 3.4, and pull back the attaching map of D to F_+ along ϕ_+ so that D is attached to F_+ along p. If p is embedded, then D is homeomorphic to a disk. Define $p' = p$ and $F_{++} = F_+ \bigsqcup_{p'} D$, and use the defining map for D and ϕ_+ to define the combinatorial map $\phi_{++} \colon F_{++} \to \overline{X}$. If ϕ_{++} is reduced, (i) and (ii) are clear.

Otherwise, ϕ_{++} is not reduced, and there is a cancelable pair of 2-cells α_D and α_+ (essential or not) belonging to D and F_+, respectively. Let e denote a folding edge in p' for the cancelable pair (α_D, α_+). Let σ be a maximal subpath of p' contained in $\alpha_D \cup \alpha_+$, and orient it consistently with p'. Since D is planar, $\partial\alpha_D$ is embedded in D. Let δ be the closure of $\partial\alpha_D \setminus \sigma$ and orient it so that it has the same initial and terminal vertices as σ. Modify F_{++} and ϕ_{++} by replacing D by $\overline{D \setminus \overline{\alpha_D}}$ and replacing p' by $(p' \setminus \sigma) \cup \delta$. Note that this process preserves that F_+ is a subcomplex of F_{++}. It also preserves that p' is an embedded edge loop in $\bar{X}$ and that D is homeomorphic to a disk (assuming that p' was originally embedded and D was originally homeomorphic to a disk) and that D is a planar reduced disk diagram filling p'. Repeating as many times as necessary, we may assume that ϕ_{++} is reduced. This proves (i) and (ii).

Now (a) holds since the operation outlined in the previous paragraph cannot introduce isolated edges. The second observation (b) follows from (a): with α as described, every edge of $[e]_\alpha$ also lies in ∂D, and all such edges on ∂F_{++} must have been isolated in F_+ already. □

Lemma 5.8. *Let F be a compact connected 2-complex which is a subcomplex of $\bar{X}$, and let F_+ and ϕ_+ have properties (1) and (2) as in Lemma 5.7. Suppose that there is a path λ in F_+ with the property that λ contains every isolated edge of F_+ and maps to a relative geodesic in $\bar{X}$. Then a patching for ϕ_+ exists.*

Proof. Fix $k \geq 0$ elements $g_1, \ldots, g_k$ of $\pi_1(F_+)$ which generate this group. The list is finite since F_+ is compact.

Induct on k. If k=0, then F_+ is simply connected and ϕ_+ is a reduced diagram, so set $(\phi_+)_\# = \phi_+$ and we are done.

If $k > 0$, let p_k be a cyclically reduced edge loop in F_+ such that $[\gamma_k p_k \gamma_k^{-1}] = g_k$ after choice of base point and appropriate γ_k.

Apply Lemma 5.7 with p_k as defined to obtain the compact, connected subcomplex F_{++} and the reduced map ϕ_{++}. Observe that F is a subcomplex of F_+, since F is a subcomplex of F_+, and F_+ is a subcomplex of F_{++}. Also observe that $\phi_{++}|_F$ is equal to the inclusion map, since $\phi_{++}|_{F_+} = \phi_+$ and $\phi_+|_F$ is equal to the inclusion map. The first additional observation of Lemma 5.7 implies that all isolated edges of F_{++} belong to λ, since all isolated edges of F_+ belong to λ.

The disk D from the application of Lemma 5.7 shows that the group $\pi_1(F_{++})$ is generated by the elements corresponding to the paths $p_1, \ldots, p_{k-1}$. By the inductive hypothesis, a patching $\phi_\#\colon F_\# \to \bar{X}$ for ϕ_{++} exists, with the property that $\phi_\#|_{F_{++}} = \phi_{++}$, and none of the essential 2-cells of $F_\# \setminus F_{++}$ are exposed in $F_\#$. Note that $\phi_\#|_{F_+} = \phi_+$ since $\phi_{++}|_{F_+} = \phi_+$.

We claim that none of the essential 2-cells of $F_\# \setminus F_+$ are exposed in $F_\#$. Indeed, we already know that none of the essential 2-cells of $F_\# \setminus F_{++}$ are exposed in $F_\#$. On the other hand, no essential 2-cell α in $F_{++} \setminus F_+$ is exposed in F_{++} by the second additional observation of Lemma 5.7 and Lemma 5.3. Thus, no such α is exposed in $F_\#$, either.

This claim implies that $\phi_\#$ is a patching for ϕ_+ as well, and the proof is complete. □

Lemma 5.9. *Let $\phi\colon F \to \bar{X}$ be an inclusion of a compact connected 2-complex. Suppose that there is a path λ in F with the property that λ contains every isolated edge of F and maps to a relative geodesic in $\bar{X}$. Then a patching for ϕ exists.*

Proof. Apply Lemma 5.8 with $F_+ = F$ and ϕ_+ equal to inclusion. □

5.4. More local geometry of essential 2-cells. With patchings as the fundamental tool, we now prove some other useful statements about the local geometry of essential 2-cells.

Lemma 5.10. *Let α be an essential 2-cell of $\bar{X}$. Then $\bar{\alpha}$ embeds in $\bar{X}$.*

Proof. The subcomplex $E = \bar{\alpha}$ satisfies the hypotheses of Lemma 5.9, so let $E_\# \to X$ be a patching (writing X instead of $\bar{X}$ in the abuse of notation justified by Remark 5.6). Since $\bar{\alpha}$ is embedded in $E_\#$ (Corollary 3.14) and E is a subcomplex of $\bar{X}$, the result follows. □

Lemma 5.11. *Let α and β be distinct essential 2-cells of $\bar{X}$. Let e be an essential edge of α. Then at most one element of $[e]_\alpha$ lies in $\partial\beta$.*

Proof. Suppose that two elements e_1 and e_2 of $[e]_\alpha$ lie in $\partial\beta$. Then the complex $E = \bar{\alpha} \cup \bar{\beta}$ satisfies the hypotheses of Lemma 5.9, so let $E_\# \to X$ be a patching (writing X instead of $\bar{X}$ in the abuse of notation justified by Remark 5.6). By Proposition 3.23, α is extreme in $E_\#$ with exposed edge f. Note that $f \notin [e]_\alpha$ since e_1 and e_2 are internal in $E_\#$. Thus there are two elements of $[f]_\alpha$, f_1 and f_2, lying in distinct components of $\partial\alpha \setminus (e_1 \cup e_2)$. Connect midpoints of f_1 and f_2 by a snipping arc running through the interior of α, and observe that any path between e_1 and e_2 through the interior of β contradicts the snipping lemma (Lemma 3.16). □

The following strong statement rules out several more pathologies for a relative geodesic which intersects the boundary of an essential 2-cell in $\bar{X}$.

Lemma 5.12. *Let α be an essential 2-cell of $\bar{X}$, and let γ be a relative geodesic which uses at least 2 essential edges of $\partial\alpha$. Index the essential edges of γ from e_1 to e_k, where e_1 and e_k are the first and last essential edges in γ which lie in $\partial\alpha$, and the labels are with respect to an orientation of γ. The following statements hold*:

i. *each e_i lies in $\partial\alpha$*;

ii. *for $i \in \{1, \dots, k-1\}$, there is a path λ_i in $\partial\alpha$ connecting e_i to e_{i+1} which does not use any essential edges*;

iii. *the orientations of the e_i are consistent with either orientation of $\partial\alpha$.*

Proof. Let $E = \bar{\alpha} \cup \gamma$. Then E satisfies the hypothesis of Lemma 5.9, so let $E_\#$ be a patching for E. By Proposition 3.23, there is only one essential 2-cell in $E_\#$.

(i) Assume that some e_i does not lie in $\partial\alpha$. In particular, $i \notin \{1, k\}$. The fact that $E_\#$ is simply connected implies e_i is contained in an essential 2-cell of $E_\#$ distinct from α, but this is a contradiction.

(ii) For fixed $i \in \{1, \dots, k-1\}$, let λ_1 and λ_2 be the two subpaths of $\partial\alpha$ connecting e_i to e_{i+1} which do not internally intersect e_i or e_{i+1}, and assume for contradiction that they both use at least one essential edge. Note that at least one of λ_1 or λ_2 has the property that all essential edges therein lie in the interior of $E_\#$: otherwise, we may join two boundary essential edges of λ_1 and λ_2 by a snipping arc running across the interior of α, and observe that the portion of γ between e_i and e_{i+1} contradicts the snipping lemma (Lemma 3.16). Without loss of generality, we may assume λ_1 has this property. By the initial assumption, λ_1 contains an essential edge e. Now Lemma 5.10 implies that there is an essential 2-cell of $E_\#$ distinct from α which is adjacent to α along e. This is a contradiction.

(iii) If this statement is false, then there is a pair of edges e_i and e_{i+1} which have opposite orientations in $\partial\alpha$. Now, observe that at least one of e_i or e_{i+1} is internal in $E_\#$. Indeed, if this is not the case, then connect e_i and e_{i+1} together by a snipping arc running across the interior of α. Because of the opposite orientation of e_i and e_{i+1} in $\partial\alpha$, the portion of γ between e_i and e_{i+1} now contradicts the snipping lemma (Lemma 3.16). Thus at least one of e_i or e_{i+1} is internal. As in (ii), there is an essential 2-cell in the diagram distinct from α, a contradiction. □

Let $\lceil x \rceil$ be the smallest integer greater than or equal to x. The following is also useful:

Lemma 5.13. *Let α be an essential 2-cell in $\overline{X}$ of exponent m and boundary path p^m in X, and let γ be a relative geodesic. Let e be an essential edge of $\partial\alpha$. Then γ contains at most $\lceil \frac{m}{2} \rceil$ elements of $[e]_\alpha$.*

Proof. The path p is a loop in X which corresponds to an order m element w of $\pi_1(X)$ which acts by "rotation" of $\overline{X}$ through a point in the interior of α. Assume for contradiction that γ contains k elements of $[e]_\alpha$, where $k \geq \lceil \frac{m}{2} \rceil + 1$. After possibly replacing γ by a path with fewer edges, we may assume that the first and last edges of γ are elements of $[e]_\alpha$. Let $e_1, \dots, e_k$ be the elements of $[e]_\alpha$ lying in γ. By Lemma 5.12, there is an orientation of γ such that γ traverses each of e_1 through e_k in the positive direction, in turn, and $we_i = e_{i+1}$ for $i \in \{1, \dots, k-1\}$ (after possibly replacing w by w^{-1}).

Now γ runs from $i(e_1)$ to $t(e_k)$, and since $k \geq \lceil \frac{m}{2} \rceil + 1$, $w^{k-1}\gamma$ runs from $i(e_k)$ to $t(e_{k'})$ for some $k' \in \{1, \dots, k-1\}$. The observations of the previous paragraph imply that $w^{k-1}\gamma$ contains the points $t(e_k)$ and $i(e_1)$ in its interior. Let γ' be the subpath of $w^{k-1}\gamma$ running from $t(e_k)$ to $i(e_1)$. Note that $\ell_r(\gamma') < \ell_r(w^{k-1}\gamma)$ since $w^{k-1}\gamma$ uses e_k and e_1 but γ' does not. Since $\ell_r(w^{k-1}\gamma) = \ell_r(\gamma)$ by G-invariance of ℓ_r, the path γ' is an "ℓ_r-shortcut" between $i(e_1)$ and $t(e_k)$. This contradicts that γ is a relative geodesic. □

6. Relative hyperbolicity

Standing assumptions from this point onward. Let X be a staggered generalized 2-complex with locally indicable vertex groups and $n(X) \geq 2$. From this point onward, assume that X_{tot} is a finite graph of CW complexes. In other words, the graph obtained by collapsing each vertex space of X_{tot} to a point is finite. Note that this does not imply that X_{tot} is compact as vertex spaces may not be. However, the staggering does imply that $C(X)$ is finite.

A result of crucial importance later on is that $\pi_1(X)$ is relatively hyperbolic with these assumptions. We prove this now.

Definition 6.1 (finite relative presentation/finite relative generating set). Suppose $\mathbb{P}$ is a finite collection of infinite subgroups of a countable group G (called *peripheral subgroups*) and let $\mathcal{P}$ be the union of all $P \in \mathbb{P}$. We say that $(G, \mathbb{P})$ has a *finite relative presentation* with *finite relative generating set* $\mathcal{S}$ if $\mathcal{S}$ is finite and symmetrized ($\mathcal{S} = S \sqcup \overline{S}$), $\mathcal{S} \cup \mathcal{P}$ is a generating set for G, and the kernel of the natural map from $F(S) * (*_{P\in\mathbb{P}} P) \to G$ is finitely normally generated, where $F(S)$ denotes the free group on the set S.

Definition 6.2 (linear relative Dehn function). Suppose $(G, \mathbb{P})$ has a finite relative presentation with finite relative generating set $\mathcal{S} = S \sqcup \overline{S}$. Let $\mathcal{P}$ be the union of all $P \in \mathbb{P}$. Let $H = F(S) * (*_{P\in\mathbb{P}} P)$ and $\mathcal{R}$ be a finite and symmetrized normal generating set for the kernel of the natural map $H \to G$. For any word W over $\mathcal{S} \cup \mathcal{P}$ representing the identity of G (called a *trivial word*), we have an equation in H of the form $W = \Pi_{i=1}^{k} h_i^{-1} R_i h_i$ where $R_i \in \mathcal{R}$ and $h_i \in H$ for each i. The smallest such k (ranging over equations of this form) is called the *area* of W and denoted by $A(W)$. We say $(G, \mathbb{P})$ has a *linear relative Dehn function* for this relative presentation if there is a linear function $f: \mathbb{N} \to \mathbb{N}$ such that for each trivial word W of length at most m in $\mathcal{S} \cup \mathcal{P}$, $A(W) \leq f(m)$.

Definition 6.3 (relatively hyperbolic [11, Definition 3.7]). Suppose $(G, \mathbb{P})$ has a finite relative presentation. If $(G, \mathbb{P})$ has a linear relative Dehn function for some finite relative presentation of $(G, \mathbb{P})$, then we say $(G, \mathbb{P})$ is *relatively hyperbolic* (or *G is hyperbolic relative to* $\mathbb{P}$).

Note. The definition above was introduced in a more general form by Osin in [21]. Hruska shows it is equivalent to no fewer than five others in the case that the set of peripheral subgroups is finite [11].

Proposition 6.4. *Suppose X is a staggered generalized 2-complex with locally indicable vertex groups, $n(X) \geq 2$, and X_{tot} is a finite graph of CW complexes. Let $\mathbb{P}$ be the collection of vertex groups of X. Then $(\pi_1(X), \mathbb{P})$ is relatively hyperbolic.*

Remark 6.5. This result seems to be known, though we were unable to find a suitable reference in the literature. However, the isoperimetric inequality proved in [5, Theorem 3.3] implies that the pair $\left(\frac{A*B}{\langle\langle w^m\rangle\rangle}, \{A, B\}\right)$ (for A and B locally indicable, w not conjugate into A or B, and $m \geq 2$) is relatively hyperbolic, which covers the case that X_{tot} is a dumbbell space.

Proof. We first construct a finite relative generating set for $G = \pi_1(X)$. Within X_{tot}, choose T_V for each vertex space V of X as in Remark 2.3, and let $q\colon X_{\text{tot}} \to \widetilde{X_{\text{tot}}}$ be the quotient map and homotopy equivalence which is the result of collapsing each T_V to a point. Choose a spanning tree T' of essential edges in $\widetilde{X_{\text{tot}}}$ (including endpoints), and let $T = q^{-1}(T')$. Note that T contains the union of the T_V and is a finite tree. Define S as follows. Fix a base vertex x in T and orient the essential edges of $X_{\text{tot}} \setminus T$. Each essential edge e therein contributes an element to S corresponding to a reduced path which starts and ends at x, traverses e exactly once in the positive direction, and otherwise does not leave the tree T. Let $\mathcal{S} = S \sqcup \overline{S}$ (where $\overline{S}$ is the collection of inverses of elements of S).

For each $P \in \mathbb{P}$, let t_P be the shortest path in T from x to the vertex space V_P corresponding to P, and let x_P be the other endpoint of this path. For each nontrivial group element $p \in P$, choose a representative loop of the form $t_P \rho t_P^{-1}$ where ρ is a cyclically reduced edge loop of minimal length in V_P such that $[\rho] = p$ under the isomorphism between $\pi_1(V_P, x_P)$ and P. These representatives give rise to a natural surjective homomorphism $H \to \pi_1(X, x)$, where $H = F(S) * (*_{P\in\mathbb{P}} P)$. Note also that this homomorphism factors through an isomorphism $H \to \pi_1(X_{\text{tot}}, x)$.

Let $\mathcal{P}$ be the union of all $P \in \mathbb{P}$. For each essential 2-cell α of X, choose a reduced edge path in X_{tot} from x to a point of $\partial\alpha$. We may thus view $\partial\alpha$ with a particular orientation as an element of $\pi_1(X, x)$ expressed as word R_α over $\mathcal{S} \cup \mathcal{P}$ (using the surjection above). Note that R_α gives rise to a particular loop based at x (expressed as a product of the based loops previously defined) which is freely homotopic to $\partial\alpha$ in X_{tot}. Fix such a homotopy f_α and let x_α be the image of the base point x in $\partial\alpha$. Let $\mathcal{R}$ be the set of the R_α and their inverses. Applying van Kampen's Theorem, we see that $\mathcal{R}$ normally generates the kernel of the map from the previous paragraph and is in one-to-one correspondence with the set of oriented boundary paths of essential 2-cells of X (and their inverses). Thus $\mathcal{S}$ is a finite relative generating set for $(G, \mathbb{P})$.

We now make and prove the following claim, which amounts to one direction of a "relative" van Kampen Lemma (cf. [18, Lemma 1.2, p. 239]). Let $D \to X$ be a planar reduced diagram with k essential 2-cells. Then there is a trivial word W over $\mathcal{S} \cup \mathcal{P}$ such that $W = \Pi_{i=1}^{k} h_i^{-1} R_i h_i$ where $R_i \in \mathcal{R}$ and $h_i \in H$ for each i, and the topological representation of W is freely homotopic to ∂D in X_{tot}.

We prove the claim above by induction on k. When $k = 0$, ∂D is nullhomotopic in X_{tot}. By the isomorphism $H \to \pi_1(X_{\text{tot}}, x)$, we see that $W = 1$.

For the inductive step, Lemma 3.8 implies that there is an essential edge on ∂D. At least one endpoint of this edge maps to T. Connect this endpoint to x within T and glue the resulting path to the diagram D; the result is a planar reduced diagram which is naturally homotopic to D. Thus, we may assume without loss of generality that there is a point of ∂D mapping to the base point x. Fix such a point on ∂D and call it x_D. Now, traverse ∂D in the counterclockwise direction starting at x_D. Again by Lemma 3.8, we will encounter a point mapping to an essential 2-cell α. Let x_1 be the first such point, and call the path traversed from x_D to x_1 in D, p_1. This path gives rise to a diagram D_1 which consists solely of the path p_1 glued to the 2-cell α at the point x_1. Let q_1 denote the maximal initial segment of $\partial\alpha$, starting from x_1 and proceeding counterclockwise, which lies on ∂D. Now, let $D' = \overline{D \setminus (\alpha \cup q_1)}$ and observe that D is naturally homotopic to $D' \bigsqcup_{x_D} D_1$ by a homotopy which "peels away" α from D.

Now define a path extending p_1 as follows. Traverse $\partial\alpha$ in the counterclockwise direction from x_1 until reaching the unique point which maps to x_α. Following the homotopy f_α (or its "mirror image" depending on the orientation of α), we end at another point which maps to x. The image of this extended path in X_{tot} defines a loop in X_{tot} based at x. Let h_1 be the corresponding group element of H. If reading $\partial\alpha$ counterclockwise from x_α agrees with the orientation chosen to define R_α, then the diagram D_1 gives rise to a based homotopy between ∂D_1 and the topological realization of the word $h_1 R_\alpha h_1^{-1}$. Otherwise, the same is true with R_α replaced by R_α^{-1}. Let $R_1 = R_\alpha^{\pm 1}$ as the case may be.

Now apply the inductive hypothesis to D' to obtain a trivial word W' over $\mathcal{S} \cup \mathcal{P}$ such that $W' = \Pi_{i=2}^{k} h_i^{-1} R_i h_i$ where $R_i \in \mathcal{R}$ and $h_i \in H$ for each $2 \leq i \leq k$, and the topological representation of W' is freely homotopic to $\partial D'$ in X_{tot}. Let $W = h_1 R_1 h_1^{-1} W'$ and note that gluing this homotopy and the homotopy from the previous paragraph along x_D gives rise to a free homotopy between the topological representation of W and $D' \bigsqcup_{x_D} D_1$ in X_{tot}. A caveat is that x_D may not be fixed through the homotopy of $\partial D'$; this may necessitate a modification of the homotopy in the previous paragraph (so that it is no longer based) and/or h_1. Finally, ∂D is freely homotopic to $D' \bigsqcup_{x_D} D_1$. This proves the claim.

Now let W be a trivial word over $\mathcal{S} \cup \mathcal{P}$ of length m. Viewing the elements of $\mathcal{S} \cup \mathcal{P}$ as loops in X_{tot} based at x, we may associate a topological representative p_W to W which is an edge loop in X_{tot} (with backtracking) based at x. Let $L(p_W)$ denote the number of essential edges of p_W in $X_{\mathrm{tot}} \setminus T$ plus the number of nontrivial maximal subpaths of p_W which lie entirely in $V \setminus T_V$ for a single vertex space V. By the choice of the topological representatives of $\mathcal{S} \cup \mathcal{P}$, it is clear that $L(p_W) = m$.

Now p_W is nullhomotopic in X. Let $D \to X$ be a planar reduced diagram for p_W which uses a minimal number of essential 2-cells, and call the number of essential 2-cells in such a diagram $\mathcal{A}(p_W)$. By the claim proved above, $A(W) \leq \mathcal{A}(p_W)$.

For an arbitrary edge path p in X_{tot}, define $\ell(p)$ to be the number of essential edges in p. Since $A(W) \leq \mathcal{A}(p_W)$ and $L(p_W) = m$, the result follows from the following claim:

1. $\ell(p_W)$ is bounded above by a linear function of $L(p_W)$ and
2. $\mathcal{A}(p_W)$ is bounded above by a linear function of $\ell(p_W)$.

It remains to prove the claim.

Let $p = p_W$. To see the first claim, note that since T is a finite tree and all of the topological representatives of the elements of $\mathcal{S} \cup \mathcal{P}$ exit T, there is a constant c such that any subpath of p which stays entirely inside T uses at most c essential edges. Thus if p' is a subpath of p with $\ell(p') = c+1$, then p' contains an essential edge of $X_{\text{tot}} \setminus T$ and contributes at least one unit of length to $L(p)$. This shows that

$$\frac{\ell(p)}{c+1} - 1 \leq L(p),$$

i.e.,

$$\ell(p) \leq (c+1)L(p) + (c+1).$$

For the second claim, use Dehn's algorithm. Let $D \to X$ be a reduced diagram for p which uses a minimal number of essential 2-cells. Suppose first that D contains at least two essential 2-cells. Then D contains an extreme essential 2-cell α by Proposition 3.23. Since $n(X) \geq 2$, α has exponent at least 2, and thus strictly more than half of the essential edges of $\partial\alpha$ lie on ∂D. Let D' be the unique component of $\overline{D \setminus \bar{\alpha}}$ which contains essential 2-cells (it is unique since α is extreme). The path $p' = \operatorname{im}(\partial D')$ has the property that $\ell(p') \leq \ell(p) - 1$. Also, D' uses a minimal number of essential 2-cells since D does. By induction on $\ell(p)$, we may assume that there exist positive constants a' and b' such that $\mathcal{A}(p') \leq a'\ell(p') + b'$. Assume without loss of generality that $a', b' \geq 1$. We have that

$$\mathcal{A}(p) = \mathcal{A}(p') + 1 \leq a'\ell(p') + b' + 1 \leq a'\ell(p) - a' + b' + 1 \leq a'\ell(p) + b'$$

as well. On the other hand, if D contains one or fewer essential 2-cells, then $\mathcal{A}(p) \leq 1$. In particular, we again have that $\mathcal{A}(p) \leq a'\ell(p) + b'$. □

7. Walls and trellises

Standing assumptions from this point onward. From now on, assume that the staggered generalized 2-complex X with $n(X) \geq 2$ and locally indicable vertex groups has the additional property that each of the vertex groups of X admits a proper and cocompact action on a CAT(0) cube complex. We also continue to assume that X_{tot} is a finite graph of CW complexes, so that $C(X)$ is finite. We will not be adding any additional standing assumptions for the rest of the paper.

Remark 7.1. Our standing assumptions imply that X is compact. Indeed, since locally indicable groups are necessarily torsion-free, our assumption that the vertex groups are cubulable implies that each vertex group acts freely on its associated cube complex. We may thus assume that each vertex space V is a compact non-positively curved (NPC) cube complex, and the universal cover $\widetilde{V}$ is a CAT(0) cube complex. Note that this implies in particular that each vertex group is finitely presented, since V is a finite $K(G,1)$ for its vertex group. Since $C(X)$ is finite, this also implies that the complex $\overline{X}$ is locally finite and X is compact.

Remark 7.2. Since the non-essential 2-cells of X ($\overline{X}$) are precisely 2-dimensional cubes in vertex spaces, we will use the term *square* to refer to a 2-cell which is not essential in the sequel. Thus the 2-cells in diagrams mapping to X and $\overline{X}$ obey the inheritance relationships depicted in Figure 8.

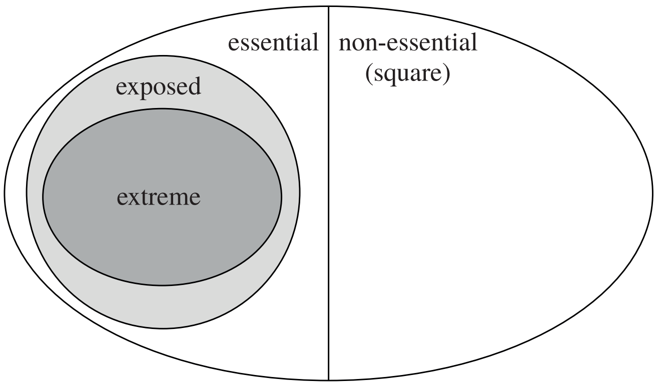

Figure 8. Inclusion relationships of 2-cells in diagrams mapping to X ($\overline{X}$).

In this section, we will define walls as codimension-1 immersed hyperspaces in $\overline{X}$. The construction of [24] will be used to obtain an action of $G = \pi_1(X)$ on an associated *dual cube complex*.

For metric statements in what follows, we will use the ℓ_1 metric in the 1-skeleton of $\widetilde{V}$ unless otherwise specified.

Definition 7.3 (walls). Similarly to the description in [19], we define *walls* as components of a "midcube complex," $M(\overline{X})$, which comes equipped with a natural map to $\overline{X}$:

- We first describe the disjoint union of the cubes of $M(\overline{X})$. Fix $\frac{1}{2} > \epsilon > 0$. Each cell of $\overline{X}$ is either a cube of some dimension or an essential 2-cell. Each k-dimensional cube C of $\overline{X}$ contains k midcubes of codimension 1 obtained by setting exactly one coordinate equal to $\frac{1}{2}$. For us, each of these midcubes C' will give rise to exactly two $(k-1)$-dimensional cubes of $M(\overline{X})$

equipped with homeomorphisms to two parallel copies of C' distance ϵ from C' on opposite sides of C'. On the other hand, each essential 2-cell α of $\bar{X}$ contributes edges to $M(\bar{X})$ as follows. Suppose that α is of exponent m. Each edge e in $\partial\alpha$ is either an essential edge or a 1-dimensional cube in some $\widetilde{V}$. In either case, consider two points in the interior of e which are distance ϵ from the midpoint of e. After choosing an orientation of $\partial\alpha$ we may label them v_e^- and v_e^+. There are an analogous pair of points in each edge of $[e]_\alpha$, and we add m edges (1-dimensional cubes) to $M(\bar{X})$ where each edge maps to a path in $\bar{\alpha}$ running from the v_e^+ in each edge of $[e]_\alpha$ to the v_e^- in the next edge of $[e]_\alpha$ through $\mathrm{int}(\alpha)$, and such that the images of these n edges are disjoint. Moreover, we arrange that the image of edges of $M(\bar{X})$ mapping to essential 2-cells is invariant with respect to the action of $\pi_1(X)$ on $\bar{X}$.

- Now identify faces of cubes of $M(\bar{X})$ as follows. Whenever one of the face identifications of $\bar{X}$ identifies the images of two faces of cubes of $M(\bar{X})$, we identify those faces in $M(\bar{X})$. The *walls* of $\bar{X}$ are defined as the components of $M(\bar{X})$. Figure 9 shows an illustration of some portions of walls in $\bar{X}$.

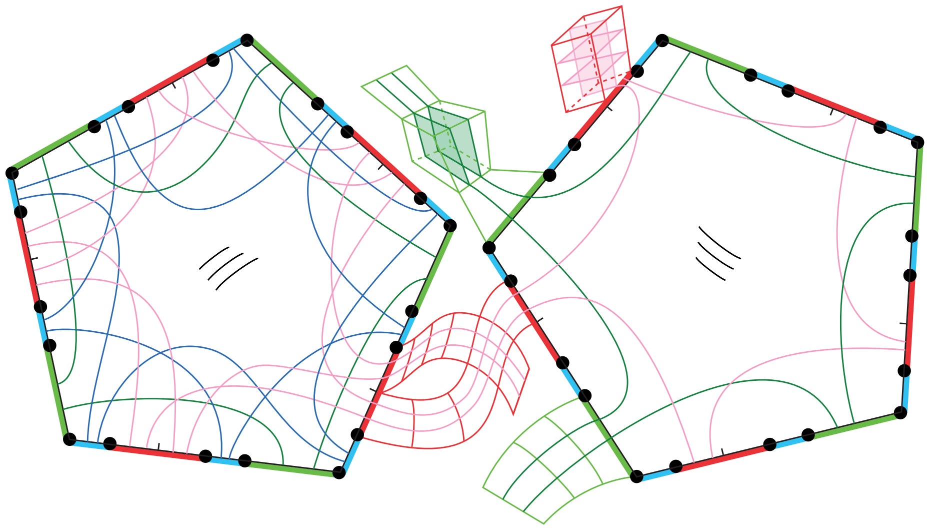

Figure 9. Some portions of walls in $\bar{X}$. Not all walls are shown.

Each wall comes equipped with a natural map to $\bar{X}$ which is the restriction of the map $M(\bar{X}) \to \bar{X}$. Note that the action of $\pi_1(X)$ on $\bar{X}$ preserves the system of walls just defined.

Definition 7.4 (types of walls). There are two types of walls in $\bar{X}$.

- The *type* 1 walls are those which are dual to essential edges and do not intersect any $\widetilde{V}$; these walls are graphs.

- The *type* 2 walls are those which nontrivially intersect some $\widetilde{V}$. These walls may be higher dimensional. More precisely, these walls are *graphs of hyperplanes*, i.e., they consist of hyperplanes of vertex spaces which are joined to each other by edges crossing essential 2-cells, with the property that the endpoints of each edge are connected to vertices of hyperplanes.

A straightforward observation about walls is that they are locally determined:

Lemma 7.5. *For any cell ω and walls Λ and Λ' of $\overline{X}$, if $im(\Lambda) \cap \omega$ is nonempty and $im(\Lambda) \cap \omega = im(\Lambda') \cap \omega$, then $\Lambda = \Lambda'$.*

It is not clear that the walls we have just defined are well-behaved in $\overline{X}$. For one, walls may not map to $\overline{X}$ injectively. A priori, a wall could travel in some vertex space $\widetilde{V}$, leave the space through some essential 2-cell α, and later come back to that same vertex space so that its image in $\overline{X}$ intersects itself.

However, we can make some basic observations about walls, vertex spaces of $\overline{X}$, and how walls behave therein. These facts follow directly from the definition of a CAT(0) cube complex and the well-known behavior of the hyperplanes therein, and the proofs are omitted. See [27], for example.

Lemma 7.6. *Let $\widetilde{V}$ be a vertex space of $\overline{X}$. Let $\Lambda \looparrowright \overline{X}$ be a wall and let Λ_V be a maximal connected component of the preimage of $\widetilde{V}$ in Λ. Let γ be a geodesic edge path in $\widetilde{V}$ and let s be a square of $\widetilde{V}$. Then*

- *Λ is an NPC cube complex*;
- *∂s embeds in $\overline{X}$*;
- *Λ_V embeds in $\overline{X}$ (since it is a hyperplane of a* CAT(0) *cube complex*);
- *$s \cap \Lambda_V$ is either empty or a single edge of Λ_V*;
- *$\gamma \cap \Lambda_V$ is either empty or a single point.*

Since each wall is an NPC cube complex, it makes sense to speak of a local geodesic in the 1-skeleton of a wall.

Definition 7.7 (carrier/cable/trellis). For a wall $\Lambda \looparrowright \overline{X}$, the *carrier* of Λ is the smallest subcomplex of $\overline{X}$ containing the image of Λ. A *cable* λ in a wall Λ is a local geodesic in $\Lambda^{(1)}$, embedded except possibly at its endpoints. The *trellis* associated to λ is the smallest subcomplex of $\overline{X}$ containing the image of λ.

Remark 7.8. This is a modification of the terminology in [17]. There, the authors use the term *wall segment* to refer to a cable, and *ladder* to refer to a trellis.

Note that trellises are necessarily at most 2-dimensional subcomplexes of $\overline{X}$.

8. Walls embed and separate

In [17], trellises turn out to be simply connected. This is not necessarily true in our case, but they can be patched:

Lemma 8.1. *Let H be the trellis associated to a cable. Then H contains at most two extreme essential 2-cells, and there is a patching $H_\# \to \overline{X}$ for H.*

Proof. Consider the inclusion of H into $\overline{X}$, which is a reduced map. Note that the first and last essential 2-cells of H are the only candidates for extreme 2-cells. Indeed, let λ be the cable for which H is the associated trellis, and observe that Lemma 3.24 may be applied to any essential 2-cell α of H which is not the first or last (taking the points x and y to be respective endpoints of the two edges of $\partial\alpha$ dual to λ and on opposite sides of λ in α). Note also that H has no isolated 1-cells unless H is a single edge, so the hypotheses of Lemma 5.9 are satisfied and $H_\# \to \overline{X}$ exists. □

The fact that walls embed and separate is a consequence of the following lemma.

Lemma 8.2. *Let α be a 2-cell of $\overline{X}$ (essential or not). If λ is a cable with both endpoints mapping to α, then $\mathrm{im}(\lambda)$ is contained in α.*

Proof. Let H be the trellis associated to λ and let $K = \alpha \cup H$. Note that $\partial\alpha$ embeds in $\overline{X}$ by either Corollary 5.10 or Lemma 7.6. We will show that K contains no 2-cells besides α, which proves the lemma.

If K contains a 2-cell besides α then we may choose distinct points u and v in $\partial\alpha \cap \mathrm{im}(\lambda)$ such that the portion of λ (of positive length) whose image is a path from u to v (which we denote by λ') does not have image internally intersecting α. Let H' be the trellis associated to λ', and note that $K' = \alpha \cup H'$ is itself a trellis (by possibly extending λ' across α if necessary). By Lemma 8.1, K' has a patching $K'_\# \to \overline{X}$.

First suppose that α is a square. Then the image of λ' passes through an essential 2-cell by Lemma 7.6. Let u' and v' be the first points along $\mathrm{im}(\lambda')$ from u and v, respectively, which lie in the boundary of some essential 2-cells α_u and α_v, which may or may not be distinct. Note that α_u and α_v are the only candidates for extreme essential 2-cells of $K'_\#$ by Lemma 8.1. On the other hand, u' and v' become identified in the auxiliary diagram, so in fact neither α_u nor α_v can be extreme by Lemma 3.24. The complex $K'_\#$ contradicts Proposition 3.23.

Now suppose α is an essential 2-cell. By extending λ' through α if necessary, we see that α is both the first and last essential 2-cell through which λ passes. Since α is the only candidate for an extreme 2-cell of $K'_\#$ by Lemma 8.1, Proposition 3.23 implies that α is the only essential 2-cell of $K'_\#$. Thus H' is made entirely of squares. Let e_u and e_v be the edges of $\partial\alpha$ containing u and v. Let σ and σ'

be the two arcs of $\partial\alpha \setminus \{u, v\}$. Suppose one of these arcs, say σ, contains no essential edges. The arc $e_u \cup \sigma \cup e_v$ is a combinatorial geodesic in a CAT(0) cube complex, and the cable λ' shows that some cable (lying entirely in that CAT(0) cube complex) crosses it twice. This contradicts Lemma 7.6. Thus there are essential edges e and e' in σ and σ' respectively. On the other hand, e and e' lie on $\partial K'_\#$ by the fact that α is the only essential 2-cell of $K'_\#$ and Corollary 3.14. Connect midpoints of e and e' by an arc running through the interior of α. The cable λ' shows that this arc is non-separating. This contradicts the snipping lemma (Lemma 3.16).

It follows that K contains no 2-cells besides α, and the lemma is proved. □

Lemma 8.3 (cf. [17, Theorem 7.4]). *Each wall is a tree of hyperplanes and embeds in $\overline{X}$.*

Proof. If some wall Λ is not simply connected, then there exists a cable λ of positive length in $\Lambda^{(1)}$ which is a loop. Let H be the trellis associated to λ. Note that H contains at least two 2-cells since the boundaries of 2-cells of $\overline{X}$ embed and by Lemma 7.6 and Corollary 5.10. Pick a 2-cell α in H. Note that we may find an arc of λ with endpoints mapping to $\partial\alpha$ whose image does not internally intersect α. This contradicts Lemma 8.2.

Thus Λ is simply connected. Since it is an NPC cube complex, it is in fact a CAT(0) cube complex. We thus see that Λ is a tree (a tree of trivial hyperplanes) if it is a wall of type 1, and a tree of hyperplanes if it is a wall of type 2.

Now suppose that a wall Λ does not embed in $\overline{X}$. Then Λ intersects itself in some essential 2-cell α or some cube c. In the latter case, there is some 2-dimensional face of c which witnesses the intersection of Λ with itself. Thus we may choose a cable λ which intersects itself exactly once in a 2-cell α (essential or not) and let H be the trellis associated to λ. Note that H contains at least two 2-cells since the boundaries of 2-cells of $\overline{X}$ embed by Lemma 7.6 and Corollary 5.10. Thus, we may find an arc of λ with endpoints mapping to $\partial\alpha$ whose image does not internally intersect α. This contradicts Lemma 8.2. □

This result permits us to casually confuse a wall Λ with its image in $\overline{X}$, a liberty we will take freely in what follows.

Corollary 8.4. *Each wall in $\overline{X}$ is separating.*

Proof. For any point p in a wall Λ, Λ separates a neighborhood of p into exactly two components, by Lemma 8.3 and construction. Thus each wall is locally separating and has an I-bundle neighborhood. And since each wall is a tree of hyperplanes (and by Lemma 8.3), each wall is a contractible subspace of $\overline{X}$. Thus each I-bundle neighborhood is actually a product $\Lambda \times I$. Thus for each wall,

$\overline{X}$ decomposes as a graph of spaces with a single simply connected edge space. Since $H^1(\overline{X}) = 0$, this graph of spaces is a dumbbell space (not a loop), and each wall is separating. $\square$

Here are some miscellaneous convenient lemmas about the geometry of walls.

Lemma 8.5. *Let γ be a relative geodesic edge path in a vertex space $\widetilde{V}$ of $\overline{X}$. Let Λ be a wall. Then $\Lambda \cap \gamma$ is either empty or a single point.*

Proof. Since γ lies in a vertex space, it is a combinatorial geodesic. Suppose Λ intersects γ in two distinct points x and y. Let λ be a cable connecting x to y and let H be the associated trellis. The subcomplex $K = H \cup \gamma$ satisfies the hypotheses of Lemma 5.9, so let $K_\#$ be a patching. Note that $K_\#$ has a maximum of two extreme 2-cells by Lemma 8.1 applied to H. If $K_\#$ has an essential 2-cell, then H contains essential 2-cells and the first one α through which λ passes is extreme in $K_\#$ by Proposition 3.23. Let e be an exposed essential edge lying in the boundary of α, and choose two elements e_1 and e_2 of $[e]_\alpha$ which lie on opposite sides of $\lambda \cap \alpha$. Connect e_1 and e_2 by a snipping arc across the interior of α, and observe that this snipping arc is non-separating, contradicting the snipping lemma (Lemma 3.16). Indeed we can get from one side to the other by following λ to γ, traversing γ from x to y (or y to x), and then going through the other portion of λ until reaching the snipping arc. This works because there are no essential edges in γ. Thus there are no essential 2-cells in $K_\#$. But this means that a connected component of $\Lambda \cap \widetilde{V}$ (which is a hyperplane in $\widetilde{V}$ by Lemma 7.6) crosses the geodesic γ twice, which contradicts the behavior of hyperplanes in CAT(0) cube complexes. $\square$

We record the following immediate corollary.

Corollary 8.6. *For each wall Λ and each vertex space $\widetilde{V}$, $\Lambda \cap \widetilde{V}$ is either empty or consists of a single hyperplane in $\widetilde{V}$.*

Lemma 8.7. *Let γ be a geodesic in $\overline{X}$ and suppose $\Lambda \cap \gamma$ consists of at least two distinct points x and y. If λ is a cable in Λ connecting x to y, then λ passes through at least one essential 2-cell.*

Proof. Let H be the trellis associated to λ, and let $K = H \cup \gamma$. Then K satisfies the hypotheses of Lemma 5.9, so let $K_\# \to \overline{X}$ be a patching. If λ does not pass through an essential 2-cell, then H is made entirely of squares, and thus so is $K_\#$ by Lemma 3.8. This implies that there are no essential edges in γ, because any such edge is isolated and non-separating in $K_\#$. Thus $K_\#$ maps to a single vertex space $\widetilde{V}$ of $\overline{X}$. Thus γ is a combinatorial geodesic in that vertex space. The fact that $\Lambda \cap \widetilde{V}$ crosses γ twice is a contradiction to Lemma 7.6. $\square$

9. Walls are relatively quasiconvex

In [17], walls turn out to be quasi-convex. This is used in conjunction with the fact that one-relator groups with torsion are hyperbolic to apply a theorem of Sageev and conclude that the action of such a group on its associated dual cube complex is cocompact.

We will use a relative version of this argument. As we argued in Proposition 6.4, $G = \pi_1(X)$ is hyperbolic relative to the vertex groups. In this secton, this will be an ingredient in a proof that each wall stabilizer is *quasiconvex relative to the vertex groups*, a notion to be made precise in what follows. This result will be used in Section 11 when we apply a generalization of Sageev's theorem by Hruska–Wise to conclude that the action on the dual cube complex is cocompact.

9.1. Geometric relative quasiconvexity. We will first prove the following geometric relative quasiconvexity statement about wall carriers and then translate it to the algebraic relative quasiconvexity of wall stabilizers. In this lemma, we only use the metric on $\bar{X}^{(1)}$.

Lemma 9.1 (cf. [17, Theorem 8.4]). *Let X be a compact staggered generalized 2-complex with locally indicable, cubulable vertex groups. Suppose that $n(X) \geq 4$. Let Λ be a wall in $\bar{X}$. There is a constant $W = W(X)$ such that if γ is a relative geodesic in $\bar{X}^{(1)}$ between vertices in the carrier C of Λ, then every vertex of γ which lies in an essential edge is within distance W of C.*

Proof. First note that since $C(X)$ is finite by our standing assumptions, there is an upper bound W_X on the number of edges (essential or not) in the attaching map of any element of $C(X)$. We will show that $W = W_X$ satisfies the conclusion of the lemma.

Let γ be a relative geodesic in $\bar{X}^{(1)}$ whose endpoints x and y are vertices in C. If γ is contained in C, then we are done. By passing to an innermost subpath of γ which lies outside of C, we may assume that $\gamma \cap C = \{x, y\}$. Since x and y lie in C, there is a trellis H in C containing x and y with associated cable λ, and γ does not internally intersect H.

Let σ be an embedded edge path in H between x and y. Since σ and γ are embedded in $\bar{X}$ and γ does not internally intersect H, the loop $\sigma \cup \gamma$ is also embedded. With $p = \sigma \cup \gamma$, the complex $H \cup \gamma$ satisfies the hypothesis of Lemma 5.7 with $F = F_+ = H \cup \gamma$. Applying Lemma 5.7, there exists a planar reduced disk diagram D (homeomorphic to a disk) such that $D \bigsqcup_{p'} (H \cup \gamma)$ has a reduced map to $\bar{X}$, where p' embeds in $H \cup \gamma$ and is homeomorphic to $\sigma \cup \gamma$. Redefine σ to be the embedded edge path $\overline{p' \setminus \gamma}$.

In fact, the complex $K = D \bigsqcup_{\sigma \cup \gamma} (H \cup \gamma)$ satisfies the hypotheses of Lemma 5.8 (with $F = (H \cup \gamma)$ and $F_+ = K$), so there is a patching $K_\# \to \bar{X}$.

By construction, D is a planar subcomplex of $K_\#$ (homeomorphic to a disk), γ is one arc of ∂D, and the other arc (σ) lies in H. Note also that σ has no edges on $\partial K_\#$ since H has no isolated edges.

By construction, K has no more extreme 2-cells than H, and the patching property implies that $K_\#$ has no more extreme 2-cells than K. By Lemma 8.1, H has a maximum of two extreme 2-cells. Thus, $K_\#$ has a maximum of two extreme 2-cells.

Now Proposition 4.6 implies that every essential 2-cell of $K_\#$ is external (since $n(X) \geq 2$). In particular, this holds for every essential 2-cell of D, and in fact every essential 2-cell of D has an essential edge lying along γ since H has no isolated edges.

Let A be the union of essential 2-cells of D whose closures intersect H (i.e., their boundaries intersect σ). Let z be a point in an essential edge e of γ. We will show all such z are uniformly close to H. If $z \in \overline{A}$, then $d(z, H) \leq \frac{W_X}{2}$. If $z \notin \overline{A}$, let δ be the maximal connected subpath of γ containing z such that $\operatorname{int}(\delta) \cap \overline{A}$ is empty. Since every 2-cell of A has an edge on γ, the complex $\overline{D \setminus \overline{A}}$ is a tree of disks. Let D' be the maximal subcomplex of $\overline{D \setminus \overline{A}}$ which contains z and is homeomorphic to a disk. Let δ' be the path $\partial D' \setminus \operatorname{int}(\delta)$ (the other boundary arc of D'), and label the endpoints of δ', x' and y' in such a way that x' lies on the subpath of γ between y' and x. See Figure 10.

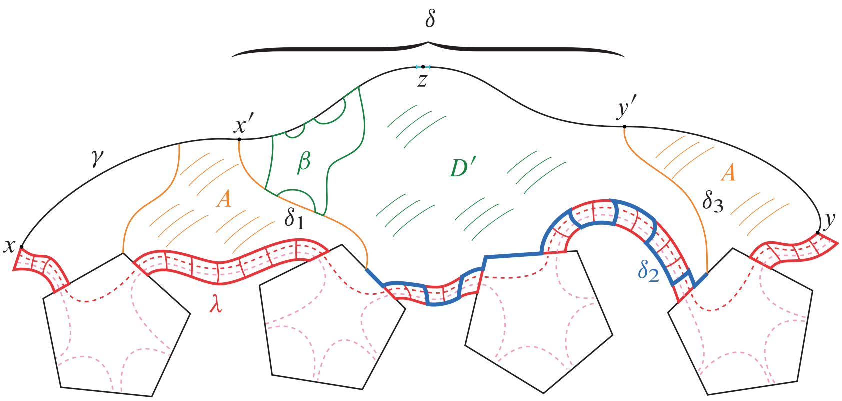

Figure 10. An illustration of the general case. Because δ_1 and δ_3 are so short, δ is a relative geodesic, δ_2 contains no essential edges, and $n(X) \geq 4$, any candidate β for an extreme essential 2-cell of D' must have exposed edges on all of δ_1, δ, and δ_3. This shows that D' contains a single essential 2-cell which contains z and intersects $\delta_1 \cup \delta_3$, so that z is close to A.

We claim that at most two essential 2-cells in A are adjacent to δ' along essential edges. Indeed, if there are three or more let α be one which is not the first, α_1, or the last, α_2, encountered while traversing δ' in the positive direction (for a chosen orientation). Since α is external in $K_\#$ and lies in D, there is an essential edge f of α on $\partial K_\#$, and f lies on γ. Without loss of generality, suppose that f lies in the portion of γ between z and x. Because D is planar, whichever of α_1 or α_2 intersects the subpath of δ' between $\bar{\alpha} \cap \delta'$ and x' cannot also intersect σ, contradicting that it lies in A. This proves the claim.

The above claim shows that δ' decomposes as a path $\delta_1\delta_2\delta_3$, where δ_1 and δ_3 are (possibly degenerate) paths, each of which lies along the boundary of an essential 2-cell of A, and δ_2 is a (possibly degenerate) subpath of σ which does not use any essential edges and maps to a single vertex space.

Next, we claim that D' contains at most one essential 2-cell. To see this claim, suppose that D' contains two or more essential 2-cells. Then D' contains at least two extreme 2-cells α and β by Proposition 3.23, with, say, exposed edges f and g, respectively. Note that all elements of $[f]_\alpha$ and $[g]_\beta$ lie along $\delta_1 \cup \delta \cup \delta_3$ since δ_2 contains no essential edges. In fact, it must be the case that at least two elements f_1 and f_2 of $[f]_\alpha$ lie along $\delta_1 \cup \delta_3$. Indeed, otherwise $m-1$ elements of $[f]_\alpha$ lie along δ, where m is the exponent of α. Since $m \geq n(X) \geq 4$, $m-1 \geq \lceil \frac{m}{2} \rceil + 1$, but this contradicts Lemma 5.13 since δ is a relative geodesic. Similarly, at least two elements g_1 and g_2 of $[g]_\beta$ lie along $\delta_1 \cup \delta_3$. Now consider the following statements:

- f_1 and f_2 lie along δ_1;
- f_1 and f_2 lie along δ_3;
- g_1 and g_2 lie along δ_1;
- g_1 and g_2 lie along δ_3.

If *none* of these statements hold, then both α and β have boundary intersecting both δ_1 and δ_3, so either α or β is internal in $K_\#$ by planarity of D'. This contradicts Proposition 4.6. On the other hand, if *any* of these statements hold, we immediately obtain a contradiction to Lemma 5.11, since δ_1 and δ_3 both lie in the boundary of a single essential 2-cell. This contradiction proves the claim.

Since $z \notin A$, D' contains a single essential 2-cell α, and $z \in \partial\alpha$. By Lemma 3.8, α is exposed in D' with exposed edge e, say. By Lemma 5.3, some element of $[e]_\alpha$ lies in $\delta_1 \cup \delta_3$. This shows that $d(z, A) \leq \frac{W_X}{2}$ and $d(z, H) \leq W_X$, so setting $W = W_X$ proves the lemma. □

Remark 9.2. We wonder if Lemma 9.1 holds when $n(X) \in \{2, 3\}$. One seems to run into trouble when trying to rule out the case where D' contains a "fat" region of squares in its interior. Lauer and Wise do not experience this difficulty in [17].

To apply the Hruska–Wise cocompactness criterion, we also need to know that wall stabilizers act cocompactly on their associated walls:

Lemma 9.3. *Let Λ be a wall of $\overline{X}$. Then $H = stab(\Lambda)$ acts cocompactly on the carrier of Λ, and thus on Λ.*

Proof. Let C be the carrier of Λ in $\overline{X}$. We claim that there are finitely many H-orbits of cells of C, which implies the result. To see this, let $\phi\colon \overline{X} \to X$ be the natural map and let β be any 2-cell of X which intersects $\phi(C)$. Now $(\phi|_\Lambda)^{-1}(\beta)$ consists of a collection of cables of Λ. Each such segment λ has the property that $\phi(\lambda)$ separates β into two components, and $\phi(\lambda)$ is one of finitely many possibly images. Enumerate these images $\lambda_1, \cdots, \lambda_k$. By Lemma 8.2, any 2-cell α of C which maps to β has a well-defined type $i \in \{1, \cdots, k\}$, defined to be the unique index for which $\phi^{-1}(\lambda_i) \cap \alpha$ lies in Λ. Fix i and suppose α and α' are cells of type i. Since the action of $G = \pi_1(X)$ on $\overline{X}$ is essentially a covering space action, there is an element $g \in G$ which takes α to α'. Moreover, because these cells are both of type i, $\phi^{-1}(\lambda_i) \cap \alpha'$ lies in both $g\Lambda$ and Λ. Now, since walls are locally determined (Lemma 7.5), this shows that g in fact stabilizes Λ, i.e. $g \in H$. Thus the number of H-orbits of $\phi^{-1}(\beta) \cap C$ is bounded above by k. This proves the claim and the lemma. □

9.2. Algebraic relative quasiconvexity. To show wall stabilizers are relatively quasiconvex, we will use the following definition of relative quasiconvexity, which we quote from [11]. In that paper, Hruska shows that this notion of relative quasiconvexity is well-defined and equivalent to no fewer than four others, at least in the case that the peripheral subgroups are finitely generated and there are finitely many of them. See [11] for the definitions of cusp-uniform action and truncated space.

Definition 9.4. [relatively quasiconvex [11, Definition 6.6] ("QC-3")] Suppose G is countable, $\mathbb{P} = \{P_1, \ldots, P_m\}$ is a finite collection of subgroups, and that $(G, \mathbb{P})$ is relatively hyperbolic. A subgroup $H \leq G$ is *relatively quasiconvex* (with respect to $\mathbb{P}$) if the following holds. Let (Y, ρ) be a proper δ-hyperbolic metric space on which $(G, \mathbb{P})$ has a cusp-uniform action. Let $Y \setminus U$ be a truncated space for G acting on Y. For some base point $x \in Y \setminus U$, there is a constant $\mu \geq 0$ such that whenever γ is a geodesic in Y with endpoints in the orbit Hx, we have

$$\gamma \cap (Y \setminus U) \subset N_\mu(Hx),$$

where the μ-neighborhood $N_\mu(Hx)$ of Hx is taken with respect to the metric ρ on Y.

We will proceed by "augmenting" the space $\overline{X}^{(1)}$, which is decidedly not δ-hyperbolic, in general, by attaching "combinatorial horoballs" to form a proper δ-hyperbolic metric space $A(\overline{X}^{(1)})$ on which G acts in a cusp uniform manner. The space $A(\overline{X}^{(1)})$ will play the role of Y in the definition above, and the disjoint union of essential edges of $\overline{X}^{(1)}$ will play the role of $Y \setminus U$.

Proposition 9.5. *Let X be a compact staggered generalized 2-complex with locally indicable, cubulable vertex groups and $n(X) \geq 4$. Then the stabilizer of each wall in $\overline{X}$ is quasiconvex relative to the collection of vertex groups of X.*

Proof. Let $G = \pi_1(X)$. As in the first paragraph of the proof of Proposition 6.4, construct a tree T in X_{tot} which is the union of paths T_V in each vertex space V (chosen as in Remark 2.3) with a spanning tree T' of essential edges in the space formed from X_{tot} by collapsing each T_V. Let $\mathbb{P}$ be the set of vertex groups of X. Let $\mathcal{S} = S \sqcup \overline{S}$ be the set of oriented essential edges of X not in T and their formal inverses. As argued in the first paragraph of the proof of Proposition 6.4, $\mathcal{S}$ is a finite relative generating set for $(G, \mathbb{P})$. The Cayley graph Γ of G with respect to $\mathcal{S}$ is disconnected, in general.

Now, attach Groves-Manning *combinatorial horoballs* to Γ to form the *augmented space* $A(\Gamma)$ associated to the data $(G, \mathbb{P}, \mathcal{S})$. See [11, Definitions 4.1 and 4.3] for the precise construction. To each $P \in \mathbb{P}$ is associated a CAT(0) cube complex which induces a natural left-invariant metric d_P on it. The rough idea is that for each coset gP, we begin with a set of copies of the coset gP indexed by the naturals (called *levels*). Form a graph as follows. For each $j \geq 0$ and each element of gP at level j, attach a *vertical edge* to the corresponding element in level $j + 1$. For each pair of elements of gP at level j whose d_P-distance is less than or equal to 2^j, attach a *horizontal edge* connecting the pair. Now glue this graph to Γ by identifying the vertices at level 0 with the subset of vertices of Γ corresponding to the coset gP. Identify any duplicate edges at level 0. Let $\mathcal{H}_\Gamma(g, P)$ be the combinatorial horoball above the coset gP, which by convention includes the original gP at level 0, as well as any edges added there. By [11, Theorem 4.4] (originally proved by Groves and Manning) and relative hyperbolicity of $(G, \mathbb{P})$, the augmented space $A(\Gamma)$ is connected and δ-hyperbolic.

Next we build the augmented space $A(\overline{X}^{(1)})$. Each vertex space of $\overline{X}^{(1)}$ is stabilized by gPg^{-1} for some $g \in G$ and $P \in \mathbb{P}$; this is a one-to-one correspondence. We label this vertex space $\widetilde{V}_g^P$. The space $A(\overline{X}^{(1)})$ is built by attaching a combinatorial horoball $\mathcal{H}_X(g, P)$ above the zero-skeleton of $\widetilde{V}_g^P$, again with respect to the cube complex metric, for each (g, P) (as before, identify any duplicate edges at level 0, and adopt the convention that $\mathcal{H}_X(g, P)$ includes the one-skeleton of $\widetilde{V}_g^P$). The space $A(\overline{X}^{(1)})$ is proper since it is a locally finite graph.

Let p denote the natural map from $\overline{X}^{(1)}$ to $X^{(1)}$. By covering space theory, $p^{-1}(T)$ consists of disjoint, homeomorphic copies of T which become identified under the action of G on $\overline{X}^{(1)}$. Pick any 0-cell $x \in p^{-1}(T)$, and consider the orbit map $\phi: G \to \overline{X}^{(1)}$. Via ϕ, embed the vertices of Γ into $\overline{X}^{(1)}$. Each vertex of Γ belongs to a unique component of $p^{-1}(T)$, and all components of $p^{-1}(T)$ contain a vertex of Γ. By construction of the finite relative generating set $\mathcal{S}$ and general covering space theory, there is a one-to-one correspondence between edges of Γ connecting group elements g and h with reduced edge paths of $\overline{X}^{(1)}$ connecting $\phi(g)$ to $\phi(h)$ which stay entirely inside of $p^{-1}(T)$ except to traverse exactly one essential edge of $\overline{X}^{(1)} \setminus p^{-1}(T)$. This observation gives rise to a G-equivariant inclusion of Γ into $\overline{X}^{(1)}$.

We may identify the group elements of gPg^{-1} with vertices of $\widetilde{V}_g^P$ via the orbit map with respect to x. This shows that $\mathcal{H}_\Gamma(g,P)$ is a full subgraph of $\mathcal{H}_X(g,P)$. The inclusion described in the previous paragraph thus extends to a G-equivariant inclusion $A(\Gamma) \hookrightarrow A(\overline{X}^{(1)})$, which we now claim is a quasi-isometry. Assuming this claim, we have that $A(\overline{X}^{(1)})$ is δ-hyperbolic (after possibly modifying δ).

To see the claim, let d be the graph metric on $X^{(1)}$, and choose $K > \max_{x \in X^{(1)}} d(x, p(\Gamma))$, which is defined since X is compact. It is clear that $A(\Gamma)$ is K-cobounded in $A(\overline{X}^{(1)})$. It remains to show that $A(\Gamma)$ is quasi-isometrically embedded. For points x and y of $A(\Gamma)^{(0)}$, it is also clear that $d_{A(\overline{X}^{(1)})}(x,y) \leq d_{A(\Gamma)}(x,y)$. It thus remains to find a constant K' such that $d_{A(\Gamma)}(x,y) \leq K' d_{A(\overline{X}^{(1)})}(x,y) + K'$. Let γ be a geodesic in $A(\overline{X}^{(1)})$ between x and y. Then γ decomposes as a path of the form $\gamma_0 e_1 \gamma_1 e_2 \ldots e_k \gamma_k$ where each e_j is an essential edge and each γ_j is a (possibly degenerate) edge path in some $\mathcal{H}_X(g,P)$. By [6, Lemma 3.10], we may assume that each γ_j consists of at most two vertical segments and a single (possibly degenerate) horizontal segment of length at most 3. This implies that each such vertical segment must contain an endpoint of γ_j. Since each vertex of $\mathcal{H}_\Gamma(g,P)$ at level 0 lies in the image of the orbit map, each such vertical segment also lie in $\mathcal{H}_\Gamma(g,P)$. Now, the horizontal segment h_j may not belong to $\mathcal{H}_\Gamma(g,P)$, but because its endpoints are connected by a path of length at most 3, there is a path h'_j of length 5 in $\mathcal{H}_\Gamma(g,P)$ between its endpoints, where h'_j consists of two vertical segments of length 2 and a single horizontal edge two levels above h_j. Replacing each h_j by h'_j, we obtain a path γ' between x and y in $A(\Gamma)$, and since $|h'_j| \leq |h_j| + 4$, we have that $|\gamma'| \leq |\gamma| + 4(k+1)$. But also $d_{A(\Gamma)}(x,y) \leq |\gamma'|$ and $k \leq |\gamma| = d_{A(\overline{X}^{(1)})}(x,y)$, so $d_{A(\Gamma)}(x,y) \leq 5 d_{A(\overline{X}^{(1)})}(x,y) + 4$. Thus, the claim is true with $K' = 5$.

Now, we claim that G has a cusp-uniform action on $A(\overline{X}^{(1)})$ with truncated space the disconnected union of all essential edges of $\overline{X}^{(1)}$. In other words, the vertex spaces of $\overline{X}^{(1)}$, along with their combinatorial horoballs, form a collection of disjoint G-equivariant horoballs (in the cusp-uniform sense) centered at the parabolic points of G. It is clear that G acts coboundedly on this truncated space with quotient the essential edges of X.

To see the claim, one can construct explicit horofunctions on these horoballs. For each vertex space $\widetilde{V}$ of $\overline{X}^{(1)}$, let $\mathcal{H}_{\widetilde{V}}$ be the combinatorial horoball above it. Let d_A be the graph metric on $A(\overline{X}^{(1)})$. Define a function $\widetilde{v}\colon A(\overline{X}^{(1)}) \to \mathbb{R}$ by

$$\widetilde{v}(x) = \begin{cases} d_A(x, \widetilde{V}) & \text{if } x \in \mathcal{H}_{\widetilde{V}}, \\ -d_A(x, \widetilde{V}) & \text{otherwise.} \end{cases}$$

It is easy to check using elementary hyperbolic geometry that $\widetilde{v}$ is a horofunction centered at the parabolic point in the Gromov boundary of $A(\overline{X}^{(1)})$ which can be identified with any geodesic ray starting in $\widetilde{V}^{(0)}$ and using only vertical edges. This proves the claim.

For each vertex space $\widetilde{V}$ of $\overline{X}$, define $d_{\widetilde{V}}(x, y) = d_A(x, y)$ for all $x, y \in \widetilde{V}^{(0)}$. The property of G-invariance is clear, so this is an admissible choice of pseudometrics.

To complete the proof, pick a base point vertex x in the carrier C of Λ and let $H = \mathrm{stab}(\Lambda)$, so that Hx lies in C. Let x', y' in Hx, and let γ be a geodesic in $A(\overline{X}^{(1)})$ between x' and y'. Note that the intersection of γ with the truncated space is precisely the set of essential edges of γ. Form a relative geodesic γ' in $\overline{X}^{(1)}$ (with respect to the admissible choice of pseudometrics above) which agrees with γ on essential edges by deleting the maximal subpaths of γ which map to horoballs and replacing them by arbitrary edge paths in the associated vertex spaces with the same endpoints. Applying Lemma 9.1 to γ', we see that every essential edge of γ' lies uniformly close to C, and thus to H_X. Thus the same is true for γ, and the proposition is proved. □

10. Bridges and linear separation

In this section, we continue to assume that X is a compact staggered generalized 2-complex with locally indicable, cubulable vertex groups and $n(X) \geq 2$.

In order to conclude that the action of $G = \pi_1(X)$ on its associated dual cube complex is proper, we will argue that the walls in $\overline{X}$ satisfy the *linear separation property*, which roughly means that the number of walls separating pairs of points in $\overline{X}$ grows at least linearly with their distance. Hruska and Wise describe how the linear separation property leads to properness of the dual cube complex action in [13, Theorem 5.2].

The precise statement we will prove is as follows.

Proposition 10.1. *Suppose that $n(X) \geq 4$. Let d be the graph metric on $\overline{X}^{(1)}$ as before. There are constants $\kappa > 0$ and ϵ such that for any vertices $x, y \in \overline{X}$, the number of walls separating x and y is at least $\kappa d(x, y) - \epsilon$.*

We will be assuming for contradiction that walls repeatedly cross geodesics in the sense of the following definition.

Definition 10.2 (bridges/bridge). Let γ be a geodesic in $\overline{X}^{(1)}$ between two 0-cells x and y of $\overline{X}$. For every edge e of γ, there are two walls dual to e which intersect e in the points v_e^x and v_e^y, labeled so that $d(x, v_e^x) < d(x, v_e^y)$. Call the wall which passes through v_e^x, Λ_e^x, and the wall passing through v_e^y, Λ_e^y. We say that Λ_e^x *bridges* γ if there is a cable λ in Λ_e^x between v_e^x and another distinct point along γ. Starting from v_e^x and traversing λ, give the label u_e^x to the first such point encountered. Define λ_e^x to be the portion of λ between v_e^x and u_e^x. There is a unique trellis H_e^x associated to λ_e^x. Let γ_e^x be the subsegment of γ connecting the edges containing v_e^x and u_e^x. Let $Y = Y_e^x = \gamma_e^x \cup H_e^x$. We call the subcomplex Y_e^x a *bridge* of γ at (e, x), if it exists. See Figure 11 for an illustration.

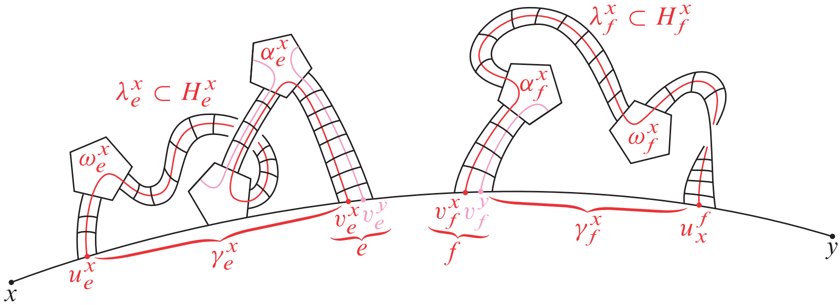

Figure 11. Some bridges. The trellis H^x_e bends in the direction of x, and H^x_f bends in the direction of y. Here the rank of $\pi_1(Y^x_e)$ is 2. Some pathologies for bridges may be ruled out immediately. For example, the depicted half-twist in H^x_f is ruled out by Corollary 8.4.

Definition 10.3 (returns). Let Y^x_e be a bridge of γ at (e, x), with associated trellis H^x_e. We say that Y^x_e (or H^x_e) *returns* through an essential 2-cell if that 2-cell is the first or last essential 2-cell of H^x_e through which the cable λ^x_e passes, as we traverse λ^x_e starting from v^x_e. We use the notation α^x_e for the first 2-cell through which Y^x_e returns, and ω^x_e for the last.

Lemma 8.7 implies that whenever Y^x_e is a bridge, α^x_e and ω^x_e always exist, and they are clearly unique. It is possible that $\alpha^x_e = \omega^x_e$.

Definition 10.4 (bends in the direction of). Let $z \in \{x, y\}$. Let Y^z_e be a bridge of γ at (e, z) with associated trellis H^z_e. We say that Y^z_e (or H^z_e) *bends in the direction of* x if $d(u^z_e, x) < d(v^z_e, x)$. Otherwise we say that Y^z_e (or H^z_e) bends in the direction of y.

The following lemma allows us to determine the direction in which walls bend, but only when $n(X) \geq 4$. The lemma is false for $n(X) \in \{2, 3\}$.

Lemma 10.5. *Suppose that $n(X) \geq 4$. Let γ be a geodesic in $\bar{X}^{(1)}$ between two 0-cells x and y of $\bar{X}$. For some edge e of γ, suppose that the wall Λ^x_e bridges γ. Then there exists a bridge Y^x_e of γ at (e, x) with associated trellis H^x_e which bends in the direction of x.*

Proof. Let $Y = Y^x_e$ be a bridge with the property that Λ^x_e does not cross γ between $v = v^x_e$ and $u = u^x_e$. We will show that this bridge bends in the direction of x. For contradiction, assume that it bends in the direction of y.

By Corollary 8.4, $\bar{X} \backslash \Lambda^x_e$ decomposes into two components $\bar{X}_{\text{in}}$ and $\bar{X}_{\text{out}}$. Label these so that the $\gamma' = \gamma^x_e$ maps to $\bar{X}_{\text{in}}$.

Let $\alpha = \alpha^x_e$ and let e_1 and e_2 be the edges of $\partial\alpha$ which are dual to $\lambda = \lambda^x_e$ (they may be essential or not), labeled so that there is a path from e_1 to v inside

λ which does not internally intersect α. Orient e_1 so that it crosses λ in the same direction that e crosses it, and extend this orientation to $\partial\alpha$. Let σ_{in} and σ_{out} be the two subpaths of $\partial\alpha \setminus \{e_1, e_2\}$, oriented consistently with $\partial\alpha$, and labeled so that σ_{in} maps to $\bar{X}_{\text{in}}$ and σ_{out} maps to $\bar{X}_{\text{out}}$ (we may do this since $\alpha \cap \Lambda_e^x$ consists only of the arc $\alpha \cap \lambda$ by Lemma 8.2). Thus no point of σ_{out} lies along γ'. Note that all elements $[e_1]_\alpha$ distinct from e_1 and e_2 lie along σ_{in}. Thus there are no elements of $[e_1]_\alpha$ properly contained within σ_{out}. See Figure 12.

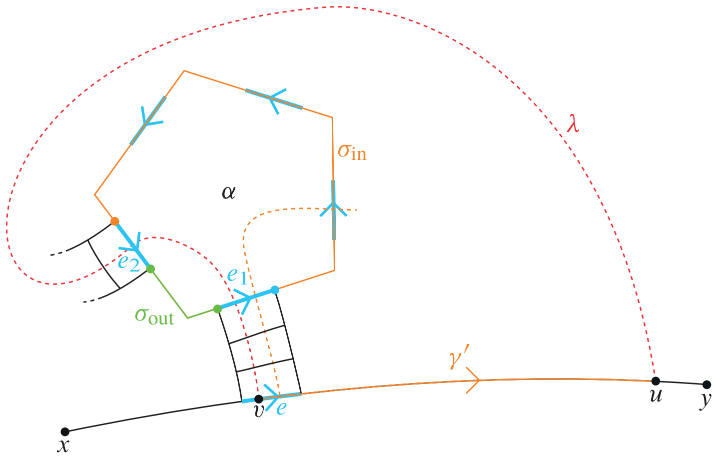

Figure 12. Proving Lemma 10.5.

Note that Y satisfies the hypotheses of Lemma 5.9 and let $Y_\#$ be a patching for Y. By Lemma 8.1, α and ω_e^x are the only essential 2-cells of $Y_\#$ which can be extreme, and in fact α is exposed by Lemma 3.8 or Proposition 3.23. We claim that σ_{out} is not internal in $Y_\#$. To see this, let f be an exposed essential edge of α. Suppose $\text{im}(\partial\alpha) = p^m$ in X, where p is not a proper power. Since σ_{out} has length $|p| - 1$, either some element of $[f]_\alpha$ lies along σ_{out}, in which case we are done, or e_1 and e_2 belong to $[f]_\alpha$. In the latter case, $\alpha = \omega_e^x$ and both e_1 and e_2 lie along γ'. Lemma 5.12 implies that every element of $[f]_\alpha$ lies along γ', which contradicts Lemma 5.3. This proves the claim.

Since e_1 and e_2 do not lie in $[f]_\alpha$, we may choose f to be the element of $[f]_\alpha$ which lies in σ_{out}. The other $m - 1$ elements of $[f]_\alpha$ lie in σ_{in}. Note that every such element must lie along γ'. Indeed, if this is not the case then given an element $f' \in [f]_\alpha$ which lies in σ_{in} but not along γ', we may join f and f' by a snipping arc running through the interior of α. The graph $Y \cap (\gamma \cup \lambda)$ shows that this arc is non-separating, contradicting the snipping lemma (Lemma 3.16). Thus the geodesic γ' visits $m - 1$ elements of $[f]_\alpha$. Since $m \geq 4$, $m - 1 \geq \lceil \frac{m}{2} \rceil + 1$. This contradicts Lemma 5.13. □

The following definition describes an impossible configuration of a pair of bridges in $\overline{X}$. We will show that if linear separation fails we can find such a configuration.

Definition 10.6 (double bridge). Let γ be a geodesic in $\overline{X}^{(1)}$ with endpoints 0-cells x and y. Let e_a and e_b be adjacent edges along γ. Suppose that Y_a and Y_b are bridges at (e_a, z_a) and (e_b, z_b), respectively, where $z_a, z_b \in \{x, y\}$. Suppose further that Y_a and Y_b bend in the same direction and that $\alpha_a = \alpha_{e_a}^{z_a}$ and $\alpha_b = \alpha_{e_b}^{z_b}$ are distinct. In this case we call the subcomplex $Y = Y_a \cup Y_b$ of $\overline{X}$ a *double bridge*. We denote by ω_a the last essential 2-cell through which Y_a returns, λ_a the cable associated to Y_a, and H_a its associated trellis. Similarly define ω_b, λ_b, and H_b.

Lemma 10.7. *There does not exist a double bridge in* $\overline{X}$.

Remark 10.8. This lemma is true when $n(X) \in \{2, 3\}$. This is what makes the following proof so technical.

Proof. Let $Y = Y_a \cup Y_b$ be a double bridge. Suppose without loss of generality that Y_a and Y_b bend in the direction of x. Note that Y satisfies the hypotheses of Lemma 5.9, and let $Y_\#$ be a patching. By Lemma 8.1, the only candidates for extreme 2-cells of $Y_\#$ are α_a, ω_a, α_b, and ω_b. We also know that $Y_\#$ contains at least two essential 2-cells since α_a and α_b are distinct. Observe that H_a and H_b embed in $Y_\#$, but they may overlap with each other.

We will prove the following statements:

i. if $\alpha_a \neq \omega_a$, then α_a is not extreme;

ii. if $\alpha_b \neq \omega_b$, then α_b is not extreme;

iii. if $\omega_a \neq \omega_b$, then at most one of ω_a and ω_b can be extreme.

Taken together, these statements imply that $Y_\#$ contains at most one extreme essential 2-cell. This contradicts Proposition 3.23.

To see statement (i), temporarily orient e_a and e_b so that their terminal points coincide. Let f_a and g_a be the edges of $\partial\alpha_a$ which are dual to λ_a (they may be essential or not), labeled so that there is a path from f_a to e_a inside λ_a which does not internally intersect α_a. Suppose $\text{im}(\partial\alpha) = p^m$ in X, where p is not a proper power. Orient f_a so that it crosses λ_a in the same direction that e_a crosses it, and extend this orientation to $\partial\alpha_a$. Now the terminal points $t(f_a)$ and $t(g_a)$ of f_a and g_a are the length of p apart in $\partial\alpha_a$. Moreover, in the auxiliary diagram $\check{Y}$, $\widetilde{t(f_a)}$ lies in $\breve{\alpha_b}$ and $\widetilde{t(g_a)}$ lies in $\check{\beta}$ for some essential 2-cell of Y_a distinct from α_a, since $\alpha_a \neq \omega_a$. Lemma 3.24 proves the claim. Note that this argument does not depend on the direction in which λ_a bends. Switching the symbols a and b, an identical argument shows that α_b is not extreme if $\alpha_b \neq \omega_b$, and statement (ii) is proved. See Figure 13.

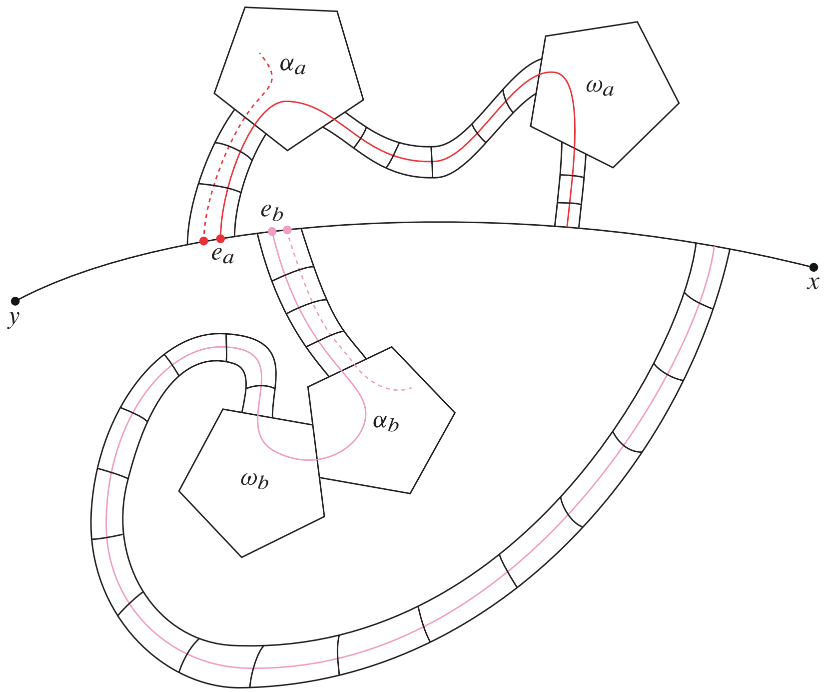

Figure 13. Proving statements (i) and (ii). The point is that α_a and α_b prevent each other from being extreme, provided that H_a and H_b both contain at least two essential 2-cells.

The following claim will be useful in proving statement (iii). Suppose ω_a is extreme with exposed essential edge f_a. Then some element of $[f_a]_{\omega_a}$ lies along γ. To see this, first note that the claim is obvious if some element of $[f_a]_{\omega_a}$ contains the terminal point of λ_a along γ. Otherwise, we may pick two elements from $[f_a]_{\omega_a}$ on opposite sides of λ_a, neither of which lies along γ, for contradiction. Connect these two edges by a snipping arc running across ω_a. This arc is non-separating in $Y_\#$, since there is a path from one side to the other in the graph $(\gamma \cup \lambda_a) \cap Y_a$; this contradicts Lemma 3.16. Similarly, if ω_b is extreme with exposed essential edge f_b, then some element of $[f_b]_{\omega_b}$ lies along γ.

We now prove statement (iii). Suppose for contradiction that $\omega_a \neq \omega_b$, but both are extreme. Among all exposed essential edges e' of ω_a (meaning that all members of $[e']_{\omega_a}$ lie on the boundary of $Y_\#$), choose the one which is on γ and closest to x along γ, and call it f_a. Define f_b similarly. Note $f_a \neq f_b$ since all elements of both $[f_a]_{\omega_a}$ and $[f_b]_{\omega_b}$ lie in $\partial Y_\#$. There are two cases according to whether f_b is closer to x than f_a or vice-versa.

Suppose first that f_b is closer to x than f_a. In this case we will show that there are two edges in $\partial\omega_a \cap \partial Y_\#$ which can be connected together by a non-separating snipping arc through ω_a, contradicting Lemma 3.16. Orient f_a so that it points towards x along γ and extend this orientation to $\partial\omega_a$. Let g_a be the next element of $[f_a]_{\omega_a}$ after f_a. Note that g_a does not lie along γ. Indeed, if it does, then by choice of f_a, g_a lies closer to y along γ than f_a by Lemma 5.12. Lemma 5.12 also implies that every element of $[f_a]_{\omega_a}$ lies along γ, which contradicts Lemma 5.3.

Connect midpoints of f_a and g_a together by a snipping arc that runs across ω_a and let S be a closed neighborhood of this arc which includes the vertices $i(f_a)$, $t(f_a)$, $i(g_a)$, and $t(g_a)$ but is small enough so that $\partial S \cap \partial\omega_a = f_a \cup g_a$. Orient S by declaring that the edge of S running from $t(f_a)$ to $i(g_a)$ is the front edge of S, and the edge running from $i(f_a)$ to $t(g_a)$ is the back edge. Let v_a denote the first point (with respect to the orientation of λ_a) in $\omega_a \cap \lambda_a$. Note that v_a is the first point of $\omega_a \cap \lambda_a$ in $\partial\omega_a$ after f_a with respect to the given orientation of $\partial\omega_a$. Note that v_a does not lie in S, for otherwise λ_a runs through the center of S connecting g_a to f_a, but because g_a lies on the boundary of $Y_\#$ this would mean $g_a = e_a$, contradicting that g_a does not lie on γ. Note also that $e_a \neq f_a$, as this scenario implies $\alpha_a = \omega_a$ and forces g_a to lie on γ (after possibly applying Lemma 5.12), which we have already ruled out.

There are now two cases to consider.

- Case 1. The vertices v_a and $t(f_a)$ lie in different components of $\overline{\omega_a \setminus S}$. This case is illustrated in Figure 14. In this case we find a path from $t(f_a)$ to the back edge of S in $\overline{Y_\# \setminus S}$ as follows.

  Starting from $t(f_a)$, travel along γ until reaching f_b. From $i(f_b)$, travel inside the interior of ω_b to reach λ_b. Next, travel backwards along λ_b all the way through H_b until reaching e_b. If at any point we cross S, then it means that ω_a is identified with an essential 2-cell in the trellis H_b *distinct* from ω_b, but this cannot happen since we already know that none of these 2-cells are extreme. Once arriving at e_b, travel within $e_b \cup e_a$ to λ_a. Here, we will not touch S because $e_a \neq g_a$ and $e_b \neq g_a$ since g_a does not lie on γ, $e_b \neq f_a$ since $\alpha_b \neq \omega_a$ but f_a lies on the boundary of $Y_\#$, and $e_a \neq f_a$ as previously observed. Finally, continue along λ_a all the way through H_a until entering ω_a through v_a and reaching the back edge of S in ω_a (we will not touch S in any other essential 2-cell since H_a is a subcomplex of $\bar{X}$). The path we have found connects the front and back edges of S in $\overline{Y_\# \setminus S}$ and contradicts Lemma 3.16.

- Case 2. The vertices v_a and $t(f_a)$ lie in the same component of $\overline{\omega_a \setminus S}$. This case further breaks into two subcases. Note that $e_a \neq f_a$ as previously observed.

  - Subcase 1. The edge e_a is strictly closer to y along γ than f_a is. This subcase is illustrated in Figure 15. In this case we find a path from $t(f_a)$ to the back edge of S in $\overline{Y_\# \setminus S}$ as follows.

    Starting from $t(f_a)$, travel along γ to $i(f_b)$, and then through the interior of ω_b to reach λ_b. Travel backwards through λ_b to reach e_b

(for the same reasons as the previous case, this path does not touch the interior of S). Since e_b is adjacent to e_a and $e_b \neq f_a$ (as in the previous case), it is the case that e_b is strictly closer to y along γ than f_a is. Thus there is a path in γ from the initial point of λ_b to $i(f_a)$ which avoids S. We have again contradicted Lemma 3.16.

– SUBCASE 2. The edge e_a is strictly closer to x along γ than f_a is. This subcase is illustrated in Figure 16. Let e'_a be the edge of γ which is dual to the terminal edge of λ_a, and oriented so that it points in the direction of x. Note that $e_a \neq e'_a$ by Lemma 8.2, and e'_a is strictly closer to x along γ than e_a. Let w_a^{front} and w_a^{back} be the vertices of $S \cap \lambda_a$, labeled according to whether they are on the front or back edge of S. In this case we find a path from w_a^{back} to w_a^{front} in $\overline{Y_\# \setminus S}$ as follows.

  Travel from w_a^{back} to e'_a along λ_a in the forward direction, and travel backwards along γ from e'_a to e_a. Then simply travel forward along λ_a through H_a until reaching w_a^{front}. This again contradicts Lemma 3.16.

For the case in which f_a is closer to x than f_b, the argument is identical, except that we exchange the roles of a and b in the above argument. Note that the above argument does not depend on the order in which e_a and e_b occur along γ, but only uses that these edges are adjacent in γ. □

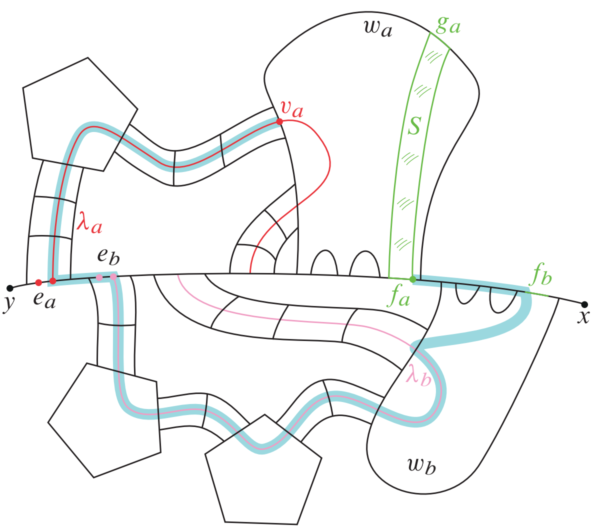

Figure 14. An example of what could happen in Case 1. The highlighted path gives the contradiction to Lemma 3.16.

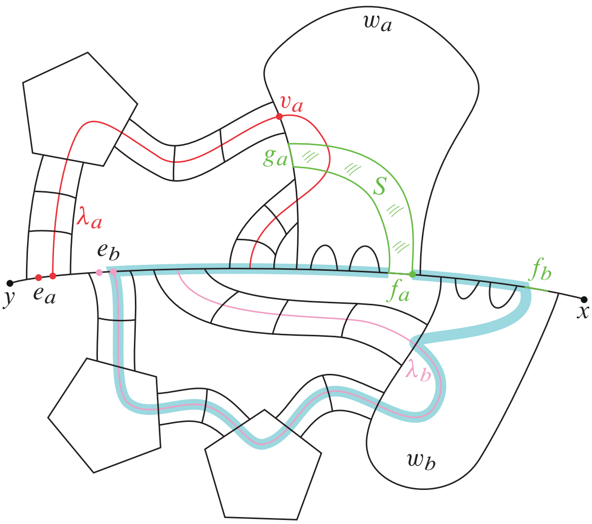

Figure 15. An example of Subcase 1. The highlighted path gives the contradiction to Lemma 3.16.

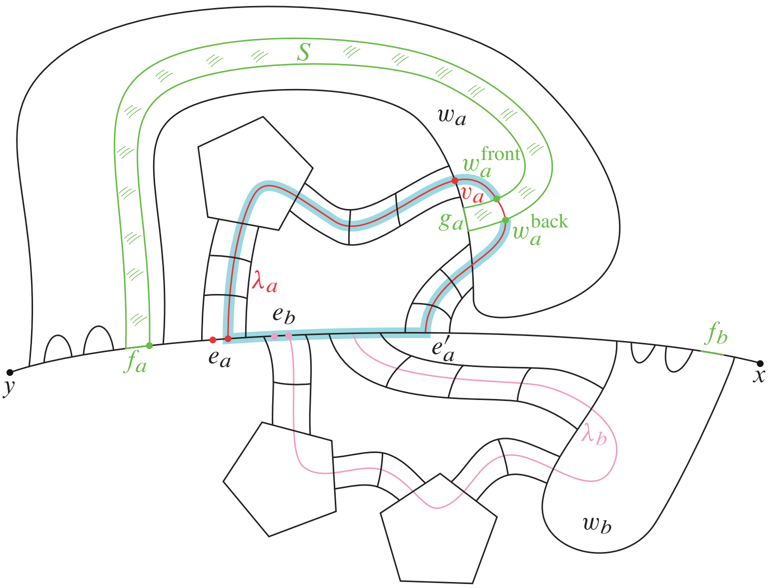

Figure 16. The general picture in Subcase 2. The highlighted path gives the contradiction to Lemma 3.16.

The following lemma now easily implies linear separation.

Lemma 10.9. *There is a constant $W = W(X)$ so that the following holds. Let γ be a geodesic in $\bar{X}^{(1)}$ with endpoints 0-cells x and y. Suppose that $n(X) \geq 4$. For any 1-cell e of γ, there exists a wall that intersects γ exactly once, and the point of intersection is within W edges of e.*

Proof. As in the proof of Lemma 9.1, let W_X be an upper bound on the number of edges (essential or not) in the attaching map of any element of $C(X)$. We will show that $W = W_X + 1$ satisfies the conclusion of the lemma.

If either wall dual to e does not bridge γ, then we are done. Thus, assume that Λ_e^x bridges γ. Fix a cable λ_e^x associated to this bridging and let Y_e^x be the associated bridge. By Lemma 10.5, we may assume that Y_e^x bends in the direction of x. Let $\alpha_a = \alpha_e^x$. Let γ_x be the subsegment of γ between e and x, including e. Consider the sequence of successive edges of γ_x starting with e and moving towards x, $\{e = e_1, e_2, e_3, \ldots\}$. Let k be the largest integer with the property that $\Lambda_{e_k}^x$ bridges γ and such that α_a is the first essential 2-cell through which some cable $\lambda_{e_k}^x$ in Λ_e^x returns. Since there are at most W_X cables passing through α_a, $k \leq W_X$. Define Y_a to be the bridge associated to $\lambda_{e_k}^x$. By Lemma 10.5, we may assume Y_a bends in the direction of x. In particular, e_{k+1} exists.

Now, observe that the wall Λ_{k+1}^x crosses γ exactly once. Indeed, if not, then there is a bridge $Y_b = Y_{k+1}^x$ at (e_{k+1}, x) which bends in the direction of x by Lemma 10.5, and $\alpha_a \neq \alpha_b$ by definition of k. Thus $Y_a \cup Y_b$ is a double bridge. This contradicts Lemma 10.7. Thus $W = W_X + 1$ satisfies the conclusion of the lemma. □

Proof of Proposition 10.1. By Lemma 10.9, $\kappa = \frac{1}{W_X+1}$ and $\epsilon = 1$ do the trick. □

Remark 10.10. Just as Lauer and Wise ask in [17], we wonder – Does $\bar{X}$ satisfy the linear separation property relative to its walls when $n(X) \in \{2, 3\}$? It appears difficult to produce a double bridge in this situation, since one has less control over the direction in which bridges bend.

11. Existence of the action

In this section we will prove the main theorem, that is that $\pi_1(X)$ acts properly and cocompactly on a CAT(0) cube complex. We first invoke the so-called "Sageev contruction" to obtain an action of $\pi_1(X)$ on a CAT(0) cube complex.

Definition 11.1 (wallspace/dual cube complex). Let Y be a metric space and let $\mathcal{W}$ be a collection of closed, connected subspaces of Y, each of which separates Y into two components. We call $(Y, \mathcal{W})$ a (geometric) *wallspace*. If a group G

acts properly and cocompactly on Y preserving both its metric and wallspace structures, then Sageev shows that G acts on a CAT(0) cube complex $\mathcal{C}(Y)$, called the *dual cube complex* [24].

We strongly suggest that the reader become familiar with the construction above before proceeding. A summary can be found in [13, Construction 3.2, Theorem 3.7, Remark 3.11].

Properness of this action in our setting will follow immediately from what we proved in Section 10. Cocompactness will follow by an application of [13, Theorem 7.12]. We state a simplified version of this theorem below.

Theorem 11.2 (cf. [15, Theorem 3.1]). *Let $(Y, \mathcal{W})$ be a wallspace. Suppose G acts properly and cocompactly on Y preserving both its metric and wallspace structures, and the action on $\mathcal{W}$ has only finitely many G-orbits of walls. Suppose G is hyperbolic relative to $\mathbb{P}$ with $\mathbb{P}$ finite. Suppose stab(Λ) acts cocompactly on Λ and is relatively quasiconvex for each wall $\Lambda \in \mathcal{W}$. For each $P \in \mathbb{P}$ let $Y_P \subset Y$ be a nonempty P-invariant P-cocompact subspace. Let $\mathcal{C}(Y)$ be the cube complex dual to $(Y, \mathcal{W})$ and for each $P \in \mathbb{P}$ let $\mathcal{C}_*(Y_P)$ be the cube complex dual to $(Y_P, \mathcal{W}_P)$, where $\mathcal{W}_P$ consists of all walls Λ with the property that diam$(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ for some $d = d(\Lambda)$.*

Then there exists a compact subcomplex K such that

$$\mathcal{C}(Y) = GK \cup \bigcup_{P \in \mathbb{P}} G\mathcal{C}_*(Y_P).$$

In particular, G acts cocompactly on $\mathcal{C}(Y)$ provided that each $\mathcal{C}_*(Y_P)$ is P-cocompact.

For us, $G = \pi_1(X)$, $Y = \overline{X}$, $\mathcal{W}$ is the collection of walls we defined in $\overline{X}$, and $\mathbb{P}$ is the finite collection of vertex groups of X. Each vertex group P has an associated vertex space V_P in X (a compact NPC cube complex). Fix a base point in $\overline{X}$ and let Y_P to be the copy of the universal cover of V_P in $\overline{X}$ (a CAT(0) cube complex) with stab$(Y_P) = P$.

The bulk of the remaining work needed to apply this theorem is to show that each $\mathcal{C}_*(Y_P)$ is P-cocompact. The following key lemma says, roughly, that a geodesic with large projection to Y_P comes very close to Y_P.

Lemma 11.3. *Fix Y_P. Suppose γ is a geodesic in $\overline{X}^{(1)}$ with endpoints 0-cells x and y, at least one of which does not belong to Y_P. Let π_x and π_y be nearest-point projections of x and y to the vertex set of Y_P. For all $d \geq 0$, there exists $R \geq 0$ such that if $d(x, \pi_x) \leq d$, $d(y, \pi_y) \leq d$, and $d(\pi_x, \pi_y) > R$, then there is an essential edge e of γ within $W_X/2$ edges of Y_P (where W_X is an upper bound on the lengths of attaching maps of essential 2-cells in X).*

Proof. Let d be given and assume $d(x, \pi_x) \leq d$ and $d(y, \pi_y) \leq d$. We claim that the conclusion of the lemma is satisfied with $R = W_X + 4d + 2$. Assume that $d(\pi_x, \pi_y) > R$. By the triangle inequality, this implies that $d(x, y) > 2d$.

Since either x or y does not belong to Y_P, note that if any edge of γ maps to Y_P, then γ contains at least one essential edge. In that case, the closest essential edge along γ to this edge has distance 0 to Y_P, and we are done.

Form a quadrilateral as follows. Let γ_x (resp. γ_y) be a geodesic edge path between x and π_x (resp. y and π_y), and let γ' be a geodesic edge path from π_y to π_x. Orient the paths so that $\sigma = \gamma\gamma_y\gamma'\gamma_x$ is a closed loop. Note that γ' lies in Y_P by Lemma 5.4. Also note that there is no backtracking in any of γ, γ_y, γ_x, or γ', so there can only be backtracking where these paths meet at their endpoints. Note that there is no backtracking of σ at π_x or π_y by the fact that these points are nearest-point projections of x and y to Y_P and γ' lies in Y_P. There may be backtracking at x. Let x' be the last vertex along γ (from x) in the image of γ_x, and similarly define y' to be the last vertex along γ (from y) in the image of γ_y. The fact that $d(x, y) > 2d$ ensures that there will remain at least one edge of γ running from x' to y'. Note also that if $x' = \pi_x$ or $y' = \pi_y$, then $\gamma \cap Y_P$ is nonempty and we are done. Let $\gamma_0 = \gamma|_{[x',y']}$, $\gamma_{y'} = \gamma_y|_{[y',\pi_y]}$, and $\gamma_{x'} = \gamma_x|_{[\pi_x,x']}$. See Figure 17.

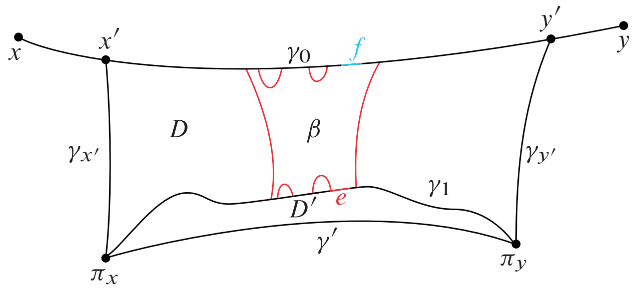

Figure 17. The general case in this lemma. The subdiagram D' maps entirely to Y_P. By choosing π_x and π_y sufficiently far apart, we can find the essential 2-cell β which does not intersect $\gamma_{x'}$ or $\gamma_{y'}$. Since β is external in D', we can find the essential edge f on γ, showing that γ passes close to Y_P.

Fill the loop $\gamma_0\gamma_{y'}\gamma'\gamma_{x'}$ with a planar reduced disk diagram $D \to \overline{X}$ using Lemma 3.4. If D has no essential 2-cells then all of D maps to Y_P. In particular γ_0 maps to Y_P and we are done. Otherwise, Suppose α is an exposed 2-cell of D with exposed edge e. We make the following observations.

- It cannot happen that there exist $e, f \in [e]_\alpha$ with e along $\gamma_{x'}$ and f along $\gamma_{y'}$. Indeed, if this happens, then $\partial\alpha$ offers a shortcut between $\gamma_{x'}$ and $\gamma_{y'}$ so that $d(\pi_x, \pi_y) \leq W_X/2 + 2d < R$, a contradiction.

- For each of $\gamma_{x'}$, $\gamma_{y'}$, and γ_0, there is an element of $[e]_\alpha$ not belonging to it, since all of these paths are relative geodesics (by Lemma 5.3).
- No element of $[e]_\alpha$ lies along γ' (since by Lemma 5.4 no edge of γ' is essential).

It may be the case that α *straddles* x' in the following sense: at least one element of $[e]_\alpha$ lies in γ_0 and at least one in $\gamma_{x'}$, and all elements of $[e]_\alpha$ lie in $\gamma_{x'} \cup \gamma_0$. Alternatively, α could straddle y'. However, these are the only possibilities allowed by the observations above.

Now we claim that D contains at most 2 extreme 2-cells. To see this, first note that there is a natural linear order on the extreme two cells of D induced by the order in which their boundaries are encountered while traversing γ_0 from x' to y'. If there are three or more extreme essential 2-cells, then we may choose one which is not the first or last with respect to this order. Call this 2-cell α and suppose that α is exposed with exposed edge e. Without loss of generality, we may assume that α straddles x'. Let e_1 be an element of $[e]_\alpha$ along γ_0 and e_2 an element of $[e]_\alpha$ along $\gamma_{x'}$. Let γ_1 and γ_2 be the two minimal paths in $\partial\alpha$ containing e_1 and e_2, and labeled so that the component of $D \setminus \gamma_2$ which contains x' also contains α. If p^m is the boundary path of the image of $\partial\alpha$ in X for p not a proper power, then $|\gamma_1|, |\gamma_2| \geq |p| + 1$. Also note that the image of γ_1 in the auxiliary diagram $\check{D}$ internally intersects an essential 2-cell of $\check{D}$ which lies before α in the order determined by γ_0. Similarly, the image of γ_2 in $\check{D}$ internally intersects an essential 2-cell of $\check{D}$ which lies after α in the order determined by γ_0. By Lemma 3.24, α is not extreme.

Using this claim and applying Proposition 4.6 and Lemma 3.8, we see that every essential 2-cell of D is external.

Now, let D' be the maximal connected subdiagram of D containing γ' and mapping to Y_P. Call the other arc of $\partial D'$ from π_y to π_x, γ_1. Note that no edge of γ_1 lies in $\gamma_{x'}$ or $\gamma_{y'}$ since π_y and π_x are nearest-point projections. If any edge of γ_1 belongs to γ_0, then some edge of γ maps Y_P and we are done. Thus we may assume that every edge of γ_1 belongs to an essential 2-cell of D lying in $D \setminus D'$.

Since $|\gamma_1| \geq |\gamma'| > R \geq W_X + 2d + 2$, we may choose an edge e of γ_1 with the property that $d(e, \pi_x) > W_X/2 + d$ and $d(e, \pi_y) > W_X/2 + d$. Let β be the essential 2-cell of D with e in its boundary. The observation above implies β is external with essential edge f (say) along ∂D. Observe that f does not lie along $\gamma_{x'}$, as this would offer a shortcut through $\partial\beta$ from e to π_x of length less than or equal to $W_X/2 + d$, contradicting the triangle inequality. Similarly, f does not lie along $\gamma_{y'}$. Thus f lies along γ_0. Now the shorter path along $\partial\beta$ from e to f maps to a path in $\bar{X}$ from Y_P to an essential edge of γ of length less than or equal to $W_X/2$, and we see that R satisfies the conclusion of the lemma. □

Lemma 11.4. *Each $\mathcal{C}_*(Y_P)$ is P-cocompact.*

Proof. Suppose that Λ is a wall of $\bar{X}$ with the property that $\operatorname{diam}(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ for some d. We claim that Λ passes within distance $d' = 3W_X/2$ of Y_P, where W_X is an upper bound on the lengths of attaching maps of essential 2-cells in X. To see this, note that we may choose vertices x and y of $\Lambda \cap \mathcal{N}_d(Y_P)$ with $d(x, y)$ arbitrarily large by assumption. By the triangle inequality, $d(\pi_x, \pi_y)$ grows with $d(x, y)$, so we may assume that $d(x, y)$ is large enough that $d(\pi_x, \pi_y) > R$, where $R(d)$ is chosen according to Lemma 11.3. Moreover, we may assume that x does not belong to Y_P, for otherwise the claim is obvious. Let π_x and π_y be the projections of x and y to Y_P, and let γ be a geodesic edge path between them. By Lemma 11.3, there is a point z in Y_P within distance $W_X/2$ of an essential edge e of γ. By geometric relative quasiconvexity of wall carriers (Lemma 9.1), the distance from e to the carrier of Λ is bounded by W_X, which means the distance from e to Λ is bounded by $3W_X/2$ since any point in the carrier is within $W_X/2$ of Λ. This proves the claim.

Now, since $P = \operatorname{stab}(Y_P)$ acts cocompactly on Y_P (its action is a covering space action and the vertex space for P is a compact NPC cube complex), P also acts cocompactly on $\mathcal{N}_{d'}(Y_P)$ by local finiteness of $\bar{X}$. Since every wall Λ with $\operatorname{diam}(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ for some d meets $\mathcal{N}_{d'}(Y_P)$ as shown above, there are finitely many P-orbits of such walls. This is exactly what it means for $\mathcal{C}_*(Y_P)$ to be P-cocompact. ☐

Putting everything together, we have the main theorem for staggered generalized 2-complexes with locally indicable vertex groups and $n(X) \geq 4$.

Theorem 11.5. *Let X be a compact staggered generalized 2-complex. Suppose that X has locally indicable vertex groups and that $n(X) \geq 4$. Suppose that for each vertex space V of X, $\pi_1(V)$ acts properly and cocompactly on a* CAT(0) *cube complex. Then $\pi_1(X)$ acts properly and cocompactly on a* CAT(0) *cube complex.*

Proof. As before, let $G = \pi_1(X)$. Let $\mathcal{W}$ be the collection of walls in $\bar{X}$ coming from the construction of Section 7. Let $\mathcal{C}$ be the cube complex dual to the action of G on the wallspace $(\bar{X}, \mathcal{W})$.

By Proposition 10.1, the wallspace $(\bar{X}, \mathcal{W})$ satisfies linear separation. By [13, Theorem 5.2], the action of G on $\mathcal{C}$ is proper.

Let $\mathbb{P}$ be the finite collection of vertex groups of X. Each vertex group P has an associated vertex space V_P in X (a compact NPC cube complex). Fix a base point in $\bar{X}$ and let Y_P to be the copy of the universal cover of V_P in $\bar{X}$ (a CAT(0) cube complex) with $\operatorname{stab}(Y_P) = P$.

Observe that all hypotheses of Theorem 11.2 are satisfied. Indeed, it is clear that G acts properly and cocompactly on $\bar{X}$ preserving both its metric and wallspace structures, and the action on $\mathcal{W}$ has only finitely many G-orbits of walls. Relative hyperbolicity of $(G, \mathbb{P})$ was shown in Proposition 6.4. For each wall Λ, Lemma 9.3 implies stab(Λ) acts cocompactly on it, and we showed stab(Λ) is relatively quasiconvex in Proposition 9.5. Finally, each $\mathcal{C}_*(Y_P)$ is P-cocompact by Lemma 11.4.

Applying Theorem 11.2, the action of G on $\mathcal{C}$ is cocompact and the theorem is proved. □

Corollary 11.6. *Let A and B be locally indicable, cubulable groups, w a word in $A * B$ which is not conjugate into A or B, and $n \geq 4$. Then $G = A * B/\langle\langle w^n \rangle\rangle$ is cubulable.*

Proof. Build a model space X for $G = A * B/\langle\langle w^n \rangle\rangle$ by starting with a dumbbell space $X_A \vee X_B$ of non-positively curved cube complexes with $\pi_1(X_A) = A$ and $\pi_1(X_B) = B$. Choose a base point and topological representative of w^n which is of period at least n, and attach a 2-cell to a path corresponding to the word w^n, so that $\pi_1(X) = G$. By performing free reductions and deletions of trivial loops in vertex spaces (while preserving periodicity), we may assume that the attaching map is along an admissible cyclically reduced edge loop and the essential 2-cell has exponent at least n. Note also that the attaching path uses the single essential edge of the dumbbell since w is not conjugate into A or B. Observe that X is trivially staggered generalized and Theorem 11.5 applies. □

References

[1] I. Agol, The virtual Haken conjecture. *Doc. Math.* **18** (2013), 1045–1087. With an appendix by I. Agol, D. Groves, and J. Manning. Zbl 1286.57019 MR 3104553

[2] N. Bergeron and D. T. Wise, A boundary criterion for cubulation. *Amer. J. Math.* **134** (2012), no. 3, 843–859. Zbl 1279.20051 MR 2931226

[3] M. R. Bridson, The geometry of the word problem. In M. R. Bridson and S. M. Salamon (eds.), *Invitations to geometry and topology.* Dedicated to Brian Steer to mark his 60th birthday. Oxford Graduate Texts in Mathematics, 7. Oxford University Press, Oxford, 2002, 29–91. Zbl 0996.54507 MR 1967746

[4] P.-A. Cherix, M. Cowling, P. Jolissaint, P. Julg, and A. Valette, *Groups with the Haagerup property.* Gromov's a-T-menability. Progress in Mathematics, 197. Birkhäuser Verlag, Basel, 2001. Zbl 1030.43002 MR 1852148

[5] A. J. Duncan and J. Howie, The genus problem for one-relator products of locally indicable groups. *Math. Z.* **208** (1991), no. 2, 225–237. Zbl 0724.20024 MR 1128707

[6] D. Groves and J. F. Manning, Dehn filling in relatively hyperbolic groups. *Israel J. Math.* **168** (2008), 317–429. Zbl 0724.20024 MR 2448064

[7] J. Howie, On pairs of 2-complexes and systems of equations over groups. *J. Reine Angew. Math.* **324** (1981), 165–174. Zbl 0447.20032 MR 614523

[8] J. Howie, On locally indicable groups. *Math. Z.* **180** (1982), no. 4, 445–461. Zbl 0471.20017 MR 667000

[9] J. Howie, How to generalize one-relator group theory. In S. M. Gersten and J. R. Stallings (eds.), *Combinatorial group theory and topology.* (Alta, UT, 1984.) Annals of Mathematics Studies, 111. Princeton University Press, Princeton, N.J., 1987, 53–78. Zbl 0632.20021 MR 895609

[10] J. Howie and S. J. Pride, A spelling theorem for staggered generalized 2-complexes, with applications. *Invent. Math.* **76** (1984), no. 1, 55–74. Zbl 0544.20033 MR 739624

[11] G. C. Hruska, Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.* **10** (2010), no. 3, 1807–1856. Zbl 1202.20046 MR 2684983

[12] G. C. Hruska and D. T. Wise, Towers, ladders and the B. B. Newman spelling theorem. *J. Aust. Math. Soc.* **71** (2001), no. 1, 53–69. Zbl 0995.20018 MR 1840493

[13] G. C. Hruska and D. T. Wise, Finiteness properties of cubulated groups. *Compos. Math.* **150** (2014), no. 3, 453–506. Zbl 1335.20043 MR 3187627

[14] T. Hsu and D. T. Wise, On linear and residual properties of graph products. *Michigan Math. J.* **46** (1999), no. 2, 251–259. Zbl 0962.20016 MR 1704150

[15] K. Jankiewizc and D. Wise, Cubulating small cancellation free products. Preprint, 2017. https://www.math.uchicago.edu/~kasia/freeprodsmallcanc.pdf

[16] J. Kahn and V. Markovic, Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math.* (2) **175** (2012), no. 3, 1127–1190. Zbl 1254.57014 MR 2912704

[17] J. Lauer and D. T. Wise, Cubulating one-relator groups with torsion. *Math. Proc. Cambridge Philos. Soc.* **155** (2013), no. 3, 411–429. Zbl 1330.20064 MR 3118410

[18] R. Lyndon and P. Schupp, *Combinatorial group theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete, 89. Springer-Verlag, Berlin etc., 1977. Zbl 0368.20023 MR 0577064

[19] J. F. Manning, Cubulating spaces and groups. Lecture notes, 2016–2020. https://pi.math.cornell.edu/~jfmanning/teaching/notes/cubulating20200303.pdf

[20] A. Martin and M. Steenbock, A combination theorem for cubulation in small cancellation theory over free products. *Ann. Inst. Fourier* (*Grenoble*) **67** (2017), no. 4, 1613–1670. Zbl 06984821 MR 3711135

[21] D. V. Osin, Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.* **179** (2006), no. 843, vi+100 pp. Zbl 1093.20025 MR 2182268

[22] G. Perelman, The entropy formula for the Ricci flow and its geometric applications. Preprint, 2002. arXiv:math/0211159 [math.DG] Zbl 1130.53001

[23] G. Perelman, Ricci flow with surgery on three-manifolds. Preprint, 2003. arXiv:math/0303109 Zbl 1130.53002

[24] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3) **71** (1995), no. 3, 585–617. Zbl 0861.20041 MR 1347406

[25] M. Sageev and D. T. Wise, The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.* **37** (2005), no. 5, 706–710. Zbl 1081.20051 MR 2164832

[26] P. Scott and T. Wall, Topological methods in group theory. In C. T. C. Wall (ed.), *Homological group theory*. (Durham, 1977.) London Mathematical Society Lecture Note Series, 36. Cambridge University Press, Cambridge and New York, 1979. 137–203. Zbl 0423.20023 MR 564422

[27] A. Sisto, *Lecture notes on geometric group theory*. Preprint, 2014. https://people.math.ethz.ch/~alsisto/LectureNotesGGT.pdf

[28] W. P. Thurston Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc.* (*N.S.*) **6** (1982), no. 3, 357–381. Zbl 0496.57005 MR 648524

[29] D. T. Wise, Cubulating small cancellation groups. *Geom. Funct. Anal.* **14** (2004), no. 1, 150–214. Zbl 1071.20038 MR 2053602

[30] D. T. Wise, Research announcement: the structure of groups with a quasiconvex hierarchy. *Electron. Res. Announc. Math. Sci.* **16** (2009), 44–55. Zbl 1183.20043 MR 2558631

[31] D. T. Wise, *From riches to raags*: 3*-manifolds, right-angled Artin groups, and cubical geometry*. CBMS Regional Conference Series in Mathematics, 117. Published for the Conference Board of the Mathematical Sciences, Washington, D.C., by the American Mathematical Society, Providence, R.I., 2012. Zbl 1278.20055 MR 2986461

Received March 11, 2019

Ben Stucky, Department of Mathematics & Computer Science, Beloit College,
700 College St, Beloit, WI 53511, USA

home page: https://www.beloit.edu/live/profiles/3214-ben-stucky

e-mail: stuckybw@beloit.edu