

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

CUBULATING ONE-RELATOR PRODUCTS WITH TORSION

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

BENJAMIN WALTNER STUCKY

Norman, Oklahoma

2019

CUBULATING ONE-RELATOR PRODUCTS WITH TORSION

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF MATHEMATICS

BY

Dr. Max Forester, Chair

Dr. Noel Brady

Dr. Scott Greene

Dr. Michael Jablonski

Dr. Jing Tao

DEDICATION

to

My Grandma, Marjorie Stucky, and

the memory of

my grandparents Marvin Stucky, LaNae Waltner, and LaVerne Waltner

all of whom showed me the importance of creativity

Acknowledgments

First and foremost, I wish to thank my adviser, Dr. Max Forester, for his invaluable guidance through the duration of my time as a graduate student and without whom this work would not have been possible. For introducing me to the beautiful world of geometric group theory, for his great generosity with his time, energy, and ideas, and for his thoughtful and lucid teaching style which has been an inspiration to me, I am extremely grateful.

I also wish to thank the rest of my advisory committee. Thanks to Dr. Noel Brady for the helpful feedback and probing questions about the thesis. Thanks to Dr. Michael Jablonski for serving on the committee in addition to his many other duties and for all of the professional advice he has given me. Thanks to Dr. Scott Greene for providing such thoughtful feedback about the dissertation and for suggesting the conclusion. Thanks to Dr. Jing Tao for helpful discussions about relative hyperbolicity and relative quasiconvexity as they pertain to the thesis.

I wish to thank Drs. Dave Futer, Keri Kornelson, Murad Özaydin, Ralf Schmidt, Krishnan Shankar, and Sepideh Stewart for their mentorship at various stages during my time as a graduate student. I wish to thank my Bethel College professors Dr. Karl Friesen and Dr. Tim Frye for their teaching and for their continued support. I wish to thank Dr. Lisa Thimm for yelling at me for sleeping during her linear algebra class and for pushing me to attend Budapest Semesters in Mathematics, where I first became interested in graduate school.

I cannot attempt to list all of the friends to whom I owe thanks for supporting and believing in me throughout graduate school. Whether or not they supported me directly in my studies, I feel strongly that I could not have completed my degree at OU without them. The following list is far from complete.

I wish to thank Paul Plummer, whose imaginative and inquisitive approach to research helped me to fully make the shift from a view of mathematics as something linear and static which needs

to be understood from the ground up to a view of the subject as something more dynamic, the frontiers of which can be probed at a high level in order to search for tractable projects to work on and new connections to make and exploit. Thank you as well for providing feedback about this project and for the TREE(3) – ϵ Hanabi games we played throughout graduate school.

I wish to thank Ignat Soroko, whose work ethic and fearless question-asking and professional connection-making has been an inspiration to me.

Thanks to all of my other former and current graduate student friends for your support. There are so many names I could list here, and I will get into trouble if I try to list them all, so I will just mention two: Thank you Long Tran and Shaoyun Yi for your friendship and constant support since Day 1 of this journey. We've been through so much and it's a pleasure to be graduating together.

I wish to thank the faculty and graduate students of Temple University for their hospitality and generosity in providing a place for me to work and discuss mathematics during the 2018 – 2019 academic year when I was ABD and living in Philadelphia – you made me feel so welcome in a new place. Thanks in particular to Dave Futer, Thomas Ng, and Kyle Rhoads. Thanks to my high school buddy Jon Hess for being a great friend to me in Philadelphia as well.

I wish to thank William Lonn for being particularly welcoming to me during my first year at OU, for generously introducing me to all of his friends and for being a great friend himself.

I wish to thank Ken and Bre Ward for their friendship and for serving as role models to me in my ongoing journey into adulthood. Through their example, they have paved the way for most of the milestones and intellectual shifts in my life since starting college – from questioning authority and finding my identity to getting married to adopting a pet to pursuing a career in academia – and I don't know where I would be without them.

I wish to thank Fred Frances for his friendship since high school, for teaching me to believe in myself, showing me how to better listen to and affirm others, and for sharing his sense of humor

with me and with the world.

I apologize to the friends I've omitted (including but not limited to my other friends from Kansas, Oklahoma, Pennsylvania, South Dakota, BSM, RTCS, PMC, my time in Kenya, math conferences, jazz bands, hockey and ultimate frisbee teams, and music festivals) – all of you are so important to me and I hope you know who you are.

I wish to thank my in-laws Pam Bracken, Buddy Johnson, and Nell Johnson for their love and hospitality. Thank you for letting me eat all of your food and also for the many rides to the bus station to get to Norman during trips to Oklahoma in my last year of graduate school.

I wish to thank my brothers Abe and Sam Stucky for their love and support. Seeing you both persevere to achieve your dreams has been a huge inspiration to me.

I wish to thank my parents Lynda and Max Stucky for everything that I have. In particular, thank you Mom for teaching me about the powers of positivity and creative problem-solving, and thank you Dad for teaching me to take responsibility for my actions and to communicate clearly and directly.

Finally, I wish to thank my wife Madeleine for being a constant source of support and inspiration, for her sense of humor, for her practical advice and organizational solutions for achieving my professional goals, and for continually reminding me how important it is to believe in other people. We're so alike in many ways and I guess I'm your beau, huh. Madeleine, I am so grateful to you for always pushing me to be the best person that I can be, for making sacrifices of your time so that I could finish graduate school, and for taking care of me in every way. You have enriched my life tremendously, and I am so glad that I met you. I can't wait to see what else we will accomplish together!

Contents

1	Introduction	1
1.1	Background	1
1.2	Summary of results	2
1.3	A naive approach	4
1.4	Outline	5
2	Preliminaries	7
2.1	The interplay between staggerings, towers, and local indicability	10
3	Van Kampen diagrams and extreme 2-cells	14
3.1	The topology of van Kampen diagrams over X	14
3.1.1	Finding exposed essential 2-cells	15
3.1.2	Essential 2-cells embed in diagrams	16
3.1.3	Other simply connected subdiagrams	19
3.2	Branches, extreme 2-cells, and a Spelling Theorem	20
4	Additional extreme 2-cells	23
5	Geometry of the universal cover	29
5.1	Admissible pseudometrics and relative geodesics	30
5.2	Local geometry of essential 2-cells	31
5.3	Patchings	34
5.4	More local geometry of essential 2-cells	35
5.5	Convexity of vertex spaces	38
6	Relative hyperbolicity	39

7 Walls and ladders	44
8 Walls embed and separate	49
9 Walls are relatively quasiconvex	54
9.1 Geometric relative quasiconvexity	54
9.2 Algebraic relative quasiconvexity	58
10 Walls satisfy linear separation	63
11 Existence of the action	74
12 Further directions	81
Bibliography	85

List of Figures

- 1.1 The presentation complex of a one-relator product with torsion 4
- 1.2 The universal cover of the presentation complex of a one-relator product with torsion 4

- 4.1 Proof of the main claim of Lemma 4.0.3 25

- 5.1 The form of a counterexample to Lemma 5.2.1 32
- 5.2 Proving Lemma 5.2.1 by pasting together short paths 32

- 7.1 Some portions of walls in \bar{X} 46

- 9.1 An illustration of the general case in Lemma 9.1.1 56

- 10.1 Some double-crossed ladders 64
- 10.2 Claims (i) and (ii) in the proof of Lemma 10.0.7 67
- 10.3 Case 1 of claim (iii) in the proof of Lemma 10.0.7 70
- 10.4 Case 2 of claim (iii) in the proof of Lemma 10.0.7 71
- 10.5 Case 3 of claim (iii) in the proof of Lemma 10.0.7 71

- 11.1 Proving Lemma 11.0.3 77

Abstract

Since the resolution of the virtual Haken conjecture in the theory of hyperbolic 3-manifolds, there has been much attention devoted to $CAT(0)$ cube complexes. These non-positively curved metric spaces are powerful tools for understanding infinite, finitely generated groups in part because of their “cubical” combinatorics. Simply knowing that a group is cubulable (acts geometrically – properly and cocompactly by isometries – on a $CAT(0)$ cube complex) is sufficient to unlock a good deal of structural information about it, and cubulating groups has become an important goal of modern geometric group theory.

In 2013, Lauer and Wise showed that a one-relator group with torsion whose defining relator has exponent at least 4 is cubulable. To achieve this, they build a system of nicely-behaved codimension-1 subspaces (“walls”) in the universal cover and invoke a construction due to Sageev.

In this thesis, we achieve a generalization of this result to one-relator products with torsion, namely, that a one-relator product of locally indicable groups whose defining relator has exponent at least 4 admits a geometric action on a $CAT(0)$ cube complex if the factors do. Our results are framed in the more general context of “staggered” quotients of free products of finitely many locally indicable and cubulable groups. The main tools are geometric small-cancellation results for van Kampen diagrams over these groups, which allow us to argue that walls are plentiful and geometrically well-behaved in the universal cover. Relative hyperbolicity of these one-relator products and relative quasiconvexity of wall stabilizers both play a central role.

Using Agol’s theorem that a hyperbolic, cubulable group is virtually special, we obtain as a corollary that the one-relator products we consider are virtually special provided that the factors are hyperbolic in addition to the other assumptions.

Chapter 1

Introduction

1.1 Background

The source of motivation for the questions which are asked and answered in this dissertation comes from 3-manifold topology. A classification of the possible geometries of closed 3-dimensional manifolds was boldly outlined by Thurston several decades ago [Thu82]. It took 30 years, but in 2012, Agol and Wise, building on the work of many others, finally proved the longstanding virtual Haken conjecture (VHC), thereby placing the last piece of the puzzle of a realization of a large portion of Thurston's vision [Ago13, BW12, KM12, Per03, Per02, Thu82].

Much of the theory that went into the proof of the VHC was developed by Wise, who had been studying objects called $CAT(0)$ *cube complexes*, simply-connected and non-positively curved topological spaces which are built by gluing cubes of various dimensions together along their faces (see Chapter 7 for the definition). The proof of the VHC involves replacing a given hyperbolic 3-manifold with a $CAT(0)$ cube complex on which the fundamental group of that manifold acts, using a construction of Sageev [Sag95]. It was by working cleverly with the combinatorics of this cube complex which allowed Agol to finish the proof of the VHC.

The resolution of the VHC brought cube complexes into the awareness of mathematicians all over the world. As it turns out, simply knowing that a group admits a proper, cocompact action on a $CAT(0)$ cube complex is sufficient to unlock a good deal of structural information about that group. For instance, these groups (henceforth referred to as *cubulable* groups) satisfy a Tits alternative [SW05], admit a quadratic-time solution to the word problem [Bri02], and satisfy

the Novikov and Baum-Connes conjectures [HP84, CCJ⁺01]. 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 and linearity [Wis12, HW99]. Thus, cubulating groups has become an important goal of modern geometric group theory.

1.2 Summary of results

We will be concerned with *one-relator products* in this dissertation, i.e., groups of the form $\frac{A*B}{\langle\langle R \rangle\rangle}$ where $\langle\langle R \rangle\rangle$ denotes the normal closure of an element R in $A * B$. These groups generalize one-relator groups, groups which admit a presentation of the form $\langle a_1, \dots, a_m \mid R \rangle$. One-relator groups with torsion of exponent $n \geq 4$ (i.e., $R = w^n$) were cubulated by Lauer and Wise in 2013 [LW13]. These groups satisfy the so-called $C'(\frac{1}{6})$ small-cancellation condition when $n \geq 6$, so this result is also covered in [Wis04]. An extension of Wise's result for $C'(\frac{1}{6})$ groups was pursued by Martin and Steenbock in 2014 when they successfully cubulated $C'(\frac{1}{6})$ small cancellation free products of cubulable groups [MS17] (see also [JW17]). In this thesis, 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. The following is our main theorem.

Theorem 1.2.1. *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 [MS17] when $n \geq 6$ and [JW17] when $n \geq 20$.

To prove Theorem 1.2.1, 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 Chapter 2 for the definition of such a complex X and its *minimal exponent* $n(X)$. We obtain the following:

Theorem 1.2.2. *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 hierarchies to prove a strong generalization of the main result in [LW13], namely that all one-relator groups with torsion are virtually special [Wis09, 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, Wise's result also follows from [LW13] and Agol's theorem that a hyperbolic, cubulable group is virtually special [Ago13, 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 $\{A, B\}$* (this can be deduced from [DH91, Theorem 3.3]). Thus if A and B are hyperbolic themselves, then so is G [Osi06, Corollary 2.41], and [Ago13, Theorem 1.1] gives the following as a corollary to Theorem 1.2.1:

Corollary 1.2.3. *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 1.2.2 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 1.2.4. Do Theorems 1.2.1 and 1.2.2 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 this thesis).

Question 1.2.5. 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?

1.3 A naive approach

Our methods are topological, and the following is what might be described as a naive approach to proving Theorem 1.2.1 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 $\pi_1(X_B) = B$, and then attaching a 2-cell to a path corresponding to the word w^n , so that $\pi_1(X) = G$. See Figures 1.1 and 1.2. 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 [Sag95], and prove that the walls are geometrically nice enough to conclude 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, of which dumbbell spaces are a particular example.




Figure 1.1: A presentation complex for G . The boundary path of the pentagonal cell corresponds to a word of the form w^5 .




Figure 1.2: The universal cover of this presentation complex. We build our walls in this space by combining the Lauer-Wise walls considered in [LW13] (in the pentagonal cells) with the natural hyperplanes in the CAT(0) cube complex factors \tilde{X}_A and \tilde{X}_B .

1.4 Outline

We follow the outline of [LW13] whenever possible. We define *staggered generalized 2-complexes* in Chapter 2. We also define the notion of a *tower* in this chapter, 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 chapter and the next is based heavily on work of James Howie [How81, How82, How87].

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 Chapters 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 chapters are direct generalizations of the work of [LW13].

In Chapter 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 chapters.

In Chapter 6, we recover relative hyperbolicity of G using Osin's idea of *linear relative Dehn functions* [Osi06], 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 Chapter 7, combining the Lauer-Wise walls of [LW13] with the natural walls in the portions of the universal cover which are already CAT(0) cube complexes. *Ladders* 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 Chapter 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

Chapter 9. Here the present work diverges from [LW13] 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 [GM08]) to \bar{X} to obtain a δ -hyperbolic space.

In Chapter 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 in the distance between them.

We put everything together in Chapter 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 [HW14, Theorem 7.12] and prove cocompactness in this more general setting. Linear separation is used to show that the action is proper. Theorem 1.2.2 is proved in Theorem 11.0.5 and Theorem 1.2.1 is Corollary 11.0.6.

In Chapter 12, we provide concluding remarks and discuss some further directions.

Chapter 2

Preliminaries

We will be working extensively with graphs of spaces. We define them here following [SW79].

Definition 2.0.1. (Abstract graph/graph of spaces/total space/dumbbell space).

An *abstract graph* Γ consists of a vertex set $V(\Gamma)$, edge set $E(\Gamma)$, involution $\bar{\cdot} : E(\Gamma) \rightarrow E(\Gamma)$ sending each edge to its *inverse*, and boundary map $\delta_0 : E(\Gamma) \rightarrow V(\Gamma)$. We define $\delta_1(e) = \delta_0(\bar{e})$ and say that e joins $\delta_0(e)$ to $\delta_1(e)$. A *graph of spaces* with underlying graph Γ is a collection of based spaces (X_v, x_v) and (X_e, x_e) for each $v \in V(\Gamma)$ and $e \in E(\Gamma)$ (with $(X_e, x_e) = (X_{\bar{e}}, x_{\bar{e}})$) and continuous maps $f_e : (X_e, x_e) \rightarrow (X_{\delta_0(e)}, x_{\delta_0(e)})$. Let $I = [0, 1]$. The *total space* X_Γ of a graph of spaces with underlying graph Γ is the quotient of $\bigcup\{X_v \mid v \in V(\Gamma)\} \cup \bigcup\{X_e \times I \mid e \in E(\Gamma)\}$ by the identifications

$$X_e \times I \rightarrow X_{\bar{e}} \times I \text{ by } (x, t) \rightarrow (x, 1 - t)$$

$$X_e \times 0 \rightarrow X_{\delta_0(e)} \text{ by } (x, 0) \rightarrow f_e(x).$$

A *dumbbell space* is the total space of a graph of spaces whose underlying graph Γ is the graph with two vertices and single edge joining them.

Definition 2.0.2. (Regular map). Let X be a CW complex. A continuous map $S^1 \rightarrow X$ is called *regular* if there is a cell structure for S^1 such that the map takes vertices to vertices and edges to edges.

Definition 2.0.3. (Cyclically reduced edge path). Let X be the total space of a graph of spaces where each vertex space is a CW complex and each edge space is a point. A *cyclically reduced edge path* is a regular edge path in $X^{(1)}$ with no backtracking and with the property that

if it contains a path of the form $e\gamma e^{-1}$ or $e^{-1}\gamma e$ where e is an oriented edge not contained in a vertex space and γ maps to a single vertex space, then γ represents a nontrivial element of the fundamental group of that vertex space.

The following is a more topological definition of a staggered generalized 2-complex than that given in [HP84].

Definition 2.0.4. (Staggered generalized 2-complex). A *staggered generalized 2-complex* X is a topological space with some additional structure as specified by the following data:

- The *total space* X_{tot} : The total space of a graph of spaces where each vertex space is a CW complex and each edge space is a point. Let $E(X)$ denote the set of edges of X_{tot} corresponding to the edge spaces in the underlying graph.
- A set of 2-cells $C(X)$, each of whose boundaries is attached to a cyclically reduced edge path in X_{tot} and contains an edge of $E(X)$ in its image.
- 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.

Definition 2.0.5. (Exponent/proper power/minimal exponent $n(X)$). For an essential 2-cell α of $C(X)$, the assumptions on the attaching map of α imply that $R = \partial\alpha$, viewed as an

element of $\pi_1(X_{\text{tot}})$ for some choice of base point, is not conjugate into the fundamental group of any vertex space. This implies that R acts loxodromically on the Bass-Serre tree corresponding to X_{tot} , i.e., it has positive translation length. This implies that R is not infinitely divisible in $\pi_1(X_{\text{tot}})$. Thus there is a well-defined *exponent* $m = m(\alpha) = \max\{k \mid R = w^k \text{ for some } w \in \pi_1(X_{\text{tot}})\}$. If $m(\alpha) \geq 2$ we say that α is attached by a *proper power*. We define the *minimal exponent* $n(X) = \min\{m(\alpha) \mid \alpha \in C(X)\}$.

For any cell $\alpha \in C(X)$, we are free to adjust the attaching map by free homotopy in X without affecting $\pi_1(X)$. If the exponent of α is m , then the attaching map of α is freely homotopic to a cyclically reduced edge path of the form p^m . We thus adopt the convention that the attaching map of each $\alpha \in C(X)$ is periodic with period $m(\alpha)$.

Definition 2.0.6. (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.

Definition 2.0.7. (Tower/tower lift/maximal). A *tower* is a map $f : Y \rightarrow X$ 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_i is an inclusion of a finite subcomplex and each p_i is an infinite cyclic cover. Let K and X be connected CW complexes and $\psi : K \rightarrow X$ be a map. A *tower lift* is a map $\phi : K \rightarrow Y$ such that there is a tower $f : Y \rightarrow X$ and $\psi = f \circ \phi$. The map ϕ is called *maximal* if any tower lift $\phi' : K \rightarrow Y'$ of ψ has the property that the associated tower $f' : Y' \rightarrow X$ is a homeomorphism.

The following remark is straightforward, since it is easily verified for infinite cyclic covers and inclusions of finite subcomplexes (even with the free homotopy considerations following Definition 2.0.5).

Remark 2.0.8. 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 $f : Y \rightarrow X$, the attaching map of β is a proper power of exponent k .

2.1 The interplay between staggerings, towers, and local indicability

Convention 2.1.1. In what follows, when we refer to a k -cell α of a CW complex, it should be understood that α refers to the image of the interior of that k -cell under the characteristic map. When we need to explicitly refer to the closure of a cell α , we will use the notation $\bar{\alpha}$.

Let K be compact and $\psi : K \rightarrow X$ be a *combinatorial map* between connected CW complexes, that is, the restriction of ψ to the interior of each cell is a homeomorphism. Howie shows [How81, Lemma 3.1] that ψ has a maximal tower lift $\phi : K \rightarrow Y$. For us, K will be an object similar to a van Kampen diagram, and we will use maximal tower lifts to study its geometry.

Note that a tower lift $\phi : K \rightarrow Y$ 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 $g : \pi_1(Y) \rightarrow \mathbb{Z}$, Y admits an infinite cyclic cover $Y' \rightarrow Y$ 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 $K \rightarrow X$ may be hard to study because the image of K in X will be highly non-injective. By considering a maximal tower lift $K \rightarrow Y$, we will have effectively “unwound” the image of K in Y in an iterative manner, using local indicability of vertex spaces to produce infinite cyclic covers at each step. Once at the top of the tower, we can draw conclusions about the topology of K using maximality of the tower lift.

On the other hand, the property of having a staggering is a flexible notion because it is preserved under towers:

Lemma 2.1.2. (cf [How87, Lemma 2]). *If $f : Y \rightarrow X$ is a tower and X is a staggered generalized 2-complex, then so is Y .*

Proof. We induct on the number of maps f comprises, so it suffices to assume that f is an

inclusion of a connected subcomplex or an infinite cyclic cover. The essential cells of Y are exactly those which map to essential cells of X . In case f is an 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 . \square

In general, there are multiple ways to stagger Y . Whenever $Y \rightarrow X$ is a tower, we make the convention that the staggering on Y arises in the manner just described.

It may be useful to record here the basic observation that the following are equivalent for any topological space Y :

- $\pi_1(Y)$ is indicable.
- $H^1(Y, \mathbb{Z}) \neq 0$.
- Y has an infinite cyclic cover.

Lemma 2.1.3. (cf [How87, Lemma 3]; [HW01, Lemma 2.6]). *Suppose that X is a compact staggered generalized 2-complex with 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 $\max \alpha$, i.e., X is homotopy equivalent to the complex obtained after removing α and $\max \alpha$ from X through a homotopy supported on $\bar{\alpha}$.*

Proof. We follow Howie’s proof in [How87] – only minor changes are necessary.

Note that 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 $H^1(X_{\text{tot}})$ is at most one, since $H^1(X) = 0$. If the underlying graph of X_{tot} is a tree, then at most one vertex space can have nontrivial first cohomology by the Mayer-Vietoris theorem. Also, since the attaching map of α is cyclically reduced 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 reduced and cyclically reduced, 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 the underlying graph of X_{tot} is not a tree, 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

$$\dots \rightarrow H^1(X) \rightarrow H^1(X \setminus \alpha) \oplus H^1(D^2) \rightarrow H^1(S^1) \rightarrow \dots$$

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 cyclically reduced, 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 $\max \alpha$. \square

Lemma 2.1.4. *(cf [LW13, Lemma 3.10]; [HW01, Lemma 2.7]). Suppose that X is a compact staggered generalized 2-complex with 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 [HW01, Lemma 2.7], except that we appeal to Lemma 2.1.3 rather than [HW01, Lemma 2.6]. \square

Chapter 3

Van Kampen diagrams and extreme 2-cells

3.1 The topology of van Kampen diagrams over X

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

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 set of 1-cells in E which occur in the attaching map of at most one 2-cell of E .

Let $E \rightarrow X$ 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 3.1.1. (Cancelable pair/reduced/diagram). Let Y be a CW complex and E a 2-complex. Let $\phi : E \rightarrow Y$ 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 $\phi \circ \Phi_\alpha(\sigma_1) = \phi \circ \Phi_\beta(\sigma_2)$. 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 3.1.2. Let Y be a CW complex, $\psi : D \rightarrow Y$ a diagram, and $\phi : D \rightarrow Z$ a lift of ψ to a cover $Z \rightarrow Y$. Then ϕ is reduced if and only if ψ is reduced.

Remark 3.1.3. Let Y be a CW complex, $\psi : D \rightarrow Y$ a diagram, and $\phi : D \rightarrow T$ 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 3.1.4. *Let Y be a CW complex and let u be a closed path in $Y^{(1)}$. Then u is null-homotopic if and only if there exists a diagram $D \rightarrow Y$ with D a planar 2-complex such that there is a parametrization of ∂D mapping to u .*

In the above theorem, we may assume D is reduced if u is a cyclically reduced path, as there are standard moves which modify D to make it reduced *without* affecting ∂D .

3.1.1 Finding exposed essential 2-cells

Definition 3.1.5. (Position). Let $\phi : E \rightarrow X$ 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 3.1.6. (External/internal/exposed). Let $\phi : E \rightarrow X$ 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 $\phi : E \rightarrow X$ 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*. Note that by our convention for staggerings associated to towers, if $E \rightarrow T$ is a tower lift of ϕ and $\alpha <_X \beta$ for essential cells α and β of E , then $\alpha <_T \beta$.

Definition 3.1.7. (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 $\gamma : I \rightarrow E$ in $E^{(1)}$, we say α is *adjacent to γ along e* if e lies in $\text{im}(\gamma) \cap \partial\alpha$.

Lemma 3.1.8. (cf [LW13, Lemma 4.7]; [HW01, Lemma 4.1]). *Suppose X has locally indicable vertex groups. Let $\phi : D \rightarrow T$ be a maximal tower lift of a reduced diagram $\psi : D \rightarrow X$. 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 $D \rightarrow X$ 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 2.1.4, 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' . 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 3.1.3). □

3.1.2 Essential 2-cells embed in diagrams

Let α be an essential 2-cell of X which is attached to 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 3.1.12 below.

Definition 3.1.9. (Magnus subcomplex) (cf [LW13, Definition 3.6]). A *Magnus subcomplex* $Z \subset X$ is a subcomplex with the following properties:

- (i) The subcomplex Z contains the disjoint 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.

The following lemma is equivalent to Howie's "locally indicable" Freiheitssatz [How81, Theorem 4.3]. We will reprove it for completeness.

Lemma 3.1.10. (cf [HW01, Theorem 6.1]). *Suppose that X has locally indicable vertex groups. If Z is a Magnus subcomplex of X , then the inclusion $i : Z \rightarrow X$ is π_1 -injective for any choice of base point in Z .*

Proof. We follow the proof in [HW01] – minimal modifications are necessary.

Let $g \in \ker i_*$. Then any loop u representing $i_*(g)$ is nullhomotopic in X , so we may apply Theorem 3.1.4 to construct a reduced diagram $\psi : D \rightarrow X$ where D is a disk 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 $\phi : D \rightarrow T$ 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 $T \rightarrow X$.

Thus 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 3.1.8, 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 . \square

Recall the following fact, the proof of which is technical but requires only Bass-Serre theory and Howie's Freiheitssatz (see [How82]):

Lemma 3.1.11. [How82, Corollary 3.4] *Let (\mathcal{G}, Y) be a graph of groups with trivial edge groups and locally indicable vertex groups. Let w be a cyclically reduced closed word of positive length in (\mathcal{G}, Y) , and let N be the normal closure of the subgroup generated by w . Then no proper closed subword of w represents an element of N .*

A topological interpretation of this gives the following:

Lemma 3.1.12. (cf [LW13, 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 . Then $\pi_1(Z' \setminus \alpha)$ decomposes as a graph of groups satisfying the hypotheses of Lemma 3.1.11. Let $w = [\partial\alpha]$. Since $\partial\alpha$ is cyclically reduced, we realize $[p]$ as a proper closed subword of w . Applying Lemma 3.1.11, 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 3.1.10. Thus p is not nullhomotopic in X . \square

Corollary 3.1.13. *Suppose X has locally indicable vertex groups. Let $D \rightarrow X$ 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 .

3.1.3 Other simply connected subdiagrams

We now observe some consequences of Corollary 3.1.13.

Definition 3.1.14. (*Internally intersects and other notation for paths*). Let Z be a subspace of a space Y and $\gamma : I \rightarrow Y$ a path. We say γ *internally intersects* Z if $\gamma(\text{int}(I)) \cap Z \neq \emptyset$. We will frequently abuse notation and refer to $\gamma(I)$ as γ and $\gamma(\text{int}(I))$ as $\text{int}(\gamma)$. 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 3.1.15. (*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 $\text{int}(\gamma) \cap E$ is contained in a single 2-cell α , and fix a parametrization $p : S^1 \rightarrow \partial\alpha$. 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 3.1.16. (*cf [LW13, Lemma 4.9]*). *Suppose that X has locally indicable vertex groups. Let $D \rightarrow X$ be a reduced diagram. Suppose an essential 2-cell α of D is external. Let B be a component of $\overline{D \setminus \alpha}$. Then $B \cap \overline{\alpha}$, B , and $B \cup \overline{\alpha}$ are all simply connected.*

Proof. By van Kampen's Theorem and Corollary 3.1.13, it suffices to prove that B and $B \cap \overline{\alpha}$ are simply connected.

Observe that $B \cap \overline{\alpha}$ is connected. To see this, suppose that $B \cap \overline{\alpha}$ 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 $B \cap \overline{\alpha}$ is connected.

Since α is external, and by Corollary 3.1.13, $B \cap \bar{\alpha}$ 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. \square

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

Definition 3.2.1. (Branch). Let $D \rightarrow X$ be a reduced diagram. If α is an exposed 2-cell of D with exposed edge e , then the components of $\overline{D \setminus \alpha}$ which contain at least one essential 2-cell are called the *branches* of D at (α, e) .

Lemma 3.1.16 implies the following:

Lemma 3.2.2. *Let $D \rightarrow X$ 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 3.2.3. (Auxiliary diagram/extreme). Let $\phi : E \rightarrow X$ be a combinatorial map. The *auxiliary diagram* \check{E} associated to E is obtained from E by collapsing all 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 $\check{\beta}$ for all essential 2-cells $\beta \neq \alpha$ of E .

Remark 3.2.4. All extreme 2-cells are exposed. When $m = 1$ the definitions of exposed and extreme coincide.

Lemma 3.2.5. *Suppose that X has locally indicable vertex groups and let $\psi : D \rightarrow X$ 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 3.1.16, $B \cap \bar{\alpha}$ 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. \square

We can now prove our first diagram result:

Proposition 3.2.6. (cf [LW13, Theorem 4.11]). *Suppose that X has locally indicable vertex groups. Let $\psi : D \rightarrow X$ 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 [LW13, Theorem 4.11].

To prove the Proposition, we induct on the number of essential 2-cells in D . Let $\phi : D \rightarrow T$ be a maximal tower lift of ψ with associated tower $f : T \rightarrow X$, 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 3.1.8 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 3.2.5, 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 3.1.8, 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 3.1.8 implies that α and β are exposed. If α and β are extreme we are done, otherwise assume without loss 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 3.2.5. 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 3.2.2, 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 [HP84, Theorem 3.1(iii)], since the diagrams considered in that paper are planar.

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

Lemma 3.2.7. *Let $\phi : E \rightarrow X$ 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 with boundary path p^m , where the loop p is not a proper power. 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 as many edges as p .*
- (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 some essential edge e in α . Condition (i) implies that either x or y lies in $\text{int}(\gamma)$, 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. □

Chapter 4

Additional extreme 2-cells

In this chapter, let X be a staggered generalized 2-complex with locally indicable vertex groups. In this chapter, we will prove additional statements about extreme essential cells in reduced diagrams $D \rightarrow X$.

Recall the main theorem from [How82]:

Lemma 4.0.1. *[How82, 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 of positive length. 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 [How82], which we prove here for completeness:

Corollary 4.0.2. *(cf [How82, 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 $C(X)$ 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 4.0.1. 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 4.0.1 again applies with $B = \langle t \rangle$ and $w = [\partial \alpha]$ to give the result. \square

We can use this fact to get a strong amplification of Remark 3.1.3:

Lemma 4.0.3. (cf [LW13, Lemma 4.6]). *Let $\psi : D \rightarrow X$ be a reduced diagram. Let $\phi : D \rightarrow T$ 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 [LW13, 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^m be the boundary path of $\psi(\alpha) = \psi(\beta)$, where p is not a proper power. By Remark 2.0.8, 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 . 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. The path $\phi(\tau)$ is a closed loop, and we claim that there is a proper closed subpath of $\phi(\tau)$ in T . To this end, assume that τ is embedded in D except possibly at its endpoints. Consider the set S of edges in $\phi^{-1}(\phi(e)) \cap \partial \alpha$ which belong to τ , which is nonempty since it contains e . If this set has exactly one element, then $[e]_\alpha$ is the only orbit of edges in $\partial \alpha$ mapping to the edge $\psi(e)$. Since $\psi(\alpha) = \psi(\beta)$, this implies that $\psi^{-1}(\psi([e]_\alpha)) \cap \partial \beta = [e]_\beta$ so that α and β form a cancelable pair, which contradicts that D is reduced. Thus S contains

at least two distinct elements, and so there are two distinct edges of τ which become identified under ϕ . This proves the claim. Thus there is a proper closed subpath γ of \hat{p} in T . See Figure 4.1.

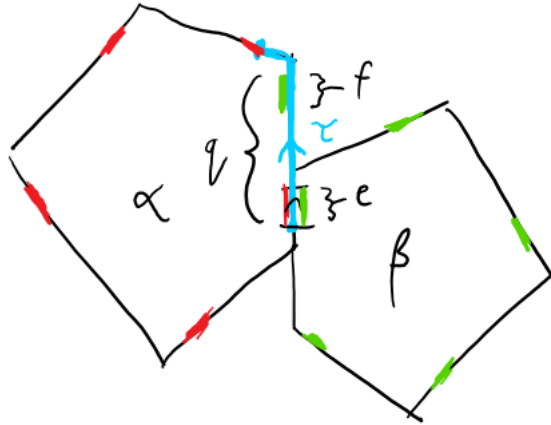


Figure 4.1: Proving the claim: The fact that S contains two distinct edges e and f implies that the path q contains the desired path γ , since e and f (in fact, all red and green edges) become identified under ϕ .

Let X' be the 2-complex associated with X having nonperiodic attaching maps, and consider the map $X \rightarrow X'$ which is the identity on the 1-skeleton of X , 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 3.1.12, 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 4.0.2, there exists a surjective homomorphism $\pi_1(T) \rightarrow \mathbb{Z}$. Thus T has an infinite cyclic cover and the tower lift $D \rightarrow T$ is not maximal, a contradiction. \square

Now we can study connected subdiagrams of a reduced diagram:

Lemma 4.0.4. (cf [LW13, Lemma 5.1]). *Let $D \rightarrow X$ 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 [LW13, Lemma 5.1]. There, the authors seem to assume that the subcomplex B defined in the proof

below is simply connected without justification.

Proof. Let $D \rightarrow T$ be a maximal tower lift of the diagram $D \rightarrow X$. By Lemma 4.0.3 applied to the map $D \rightarrow T$, 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 $B' \rightarrow T'$ be a maximal tower lift of the composition $B' \hookrightarrow D \rightarrow T$, and let α' be a greatest essential 2-cell of B' under $<_{T'}$. Now Lemma 3.1.8 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 $\overline{B' \setminus \alpha'}$ containing α . This subcomplex of D contains B , is simply connected (by Lemma 3.1.16), 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' . \square

Let $D \rightarrow X$ 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}_α and \widehat{L}_α , respectively, containing α .

Lemma 4.0.5. (cf [LW13, Lemma 5.3]). *The components of \widehat{G}_α and \widehat{L}_α are simply connected.*

Proof. The proof is nearly identical to that of [LW13, Lemma 5.3]. We obtain \widehat{G}_α 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 4.0.4 ensures that each successive essential 2-cell will be exposed, and Lemma 3.1.16 implies that removing each successive cell leaves simply

connected components. In finitely many steps we obtain \widehat{G}_α , and the argument is essentially the same for \widehat{L}_α . \square

We are ready to prove our second main diagram theorem:

Proposition 4.0.6. (cf [LW13, Theorem 5.4]). *Let $D \rightarrow X$ 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 [LW13, Theorem 5.4].

Let $D \rightarrow T$ be a maximal tower lift of $D \rightarrow X$, and let α be an internal essential 2-cell of D of exponent m . Define \widehat{G}_α and \widehat{L}_α with respect to $<_T$. Now Lemma 4.0.4 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, this implies that $[e_G]_\alpha$ and $[e_L]_\alpha$ must be distinct. Since the m elements of $[e_L]_\alpha$ are internal in G_α , and because each branch of G_α at (α, e_G) intersects $\partial\alpha$ in an arc (Lemma 3.1.16), there are exactly m branches of G_α at (α, e_G) . 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 4.0.3 to $D \rightarrow T$). So any greatest 2-cell of G_i lies in B_i . Now Lemma 4.0.4 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' and β doesn't lie in \widehat{L}_α , then β is essential and $\beta \geq \alpha$, so β lies in G_i . 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 4.0.3). Thus we obtain $2m$ 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 3.2.5. Let B be a branch not containing α , and note that $B \cup \overline{\alpha_i}$ is simply connected by Lemma 3.2.2. By Proposition 3.2.6, there are at least two extreme essential 2-cells in $B \cup \overline{\alpha_i}$; any one of these not equal to α_i is extreme in D . Repeating for each i , we obtain $2m$ extreme 2-cells. They are distinct since for $j \neq i$, α_j lies in the branch of D at (α_i, e_i) containing α . \square

Chapter 5

Geometry of the universal cover

Throughout this chapter, let X be a staggered generalized 2-complex with locally indicable vertex groups. From now on, we also 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 chapter 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 \bar{X} which is closely related to \tilde{X} , the universal cover of X . Let Y denote the preimage of X_{tot} in \tilde{X} .

By Lemma 3.1.10, $\pi_1(V)$ embeds naturally in $\pi_1(X)$ for each vertex space V of X , and thus Y may be viewed as a collection of CAT(0) cube complexes (each of which is \tilde{V} for some vertex space V of X , and to which we also refer as a *vertex space* by slight abuse of notation) with essential edges running between them. Let \bar{X} be the space obtained from \tilde{X} by identifying elevations of essential 2-cells of X which have the same boundary; it may be viewed as a subcomplex of \tilde{X} which contains Y . Give $Y^{(1)}$ the combinatorial metric in which every edge has length 1. All of the metric statements in this chapter are really about $Y^{(1)} = \bar{X}^{(1)}$, and all paths of interest are edge paths. From now on, let d be the graph metric on $\bar{X}^{(1)}$.

We may assume that X has the property that for each essential 2-cell α of X , any lift of a maximal subpath of $\partial\alpha$ mapping to a vertex space V is a shortest path in the copy of $\tilde{V}^{(1)}$ to which it lifts. To achieve this, we argue as follows: Suppose that the exponent of α is m , so the boundary $\partial\alpha$ is a path of the form p^m , where p is a loop in $X_{\text{tot}}^{(1)}$. For each maximal subpath p_V of p mapping entirely to a vertex space V of X , note that p_V is a loop. We modify p by

replacing p_V by a loop p'_V in $V^{(1)}$ with the properties that p'_V has the same basepoint as p_V , p'_V and p_V represent the same element of $\pi_1(X)$, and p'_V uses a minimal number of edges. Let p' be the result of modifying p in this way. Replace α by a 2-cell α' with attaching map $(p')^m$. Doing this for all essential 2-cells does not affect $\pi_1(X)$ or the staggering, and the resulting staggered generalized 2-complex has the desired property.

5.1 Admissible pseudometrics and relative geodesics

We will work with paths in \bar{X} which generalize geodesics. The idea of relative geodesics as defined below is that they allow for the possibility that paths can be “shorter than they look,” but only in vertex spaces. At certain times in what follows, we will be “augmenting” \bar{X} and allowing for this sort of behavior.

Definition 5.1.1. (Admissible pseudometrics/relative length/relative geodesic). Let d denote the metric on $\bar{X}^{(1)}$ where every edge has length one. For each vertex space \tilde{V} , choose a pseudometric $d_{\tilde{V}}$ on $\tilde{V}^{(0)}$. We require that this choice of pseudometrics is invariant with respect to the action of G on \bar{X} . If this holds we say the choice of pseudometrics is *admissible*.

Let $\gamma : I \rightarrow \bar{X}$ be an edge path whose endpoints are 0-cells x and y of \bar{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 \tilde{V}_i of \bar{X} , and the e_i are essential edges. We define the *relative length* of γ , $\ell_r(\gamma)$, by the following formula:

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

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 the restriction of γ to each vertex space is a geodesic in the one-skeleton of that vertex space, and $\ell_r(\gamma)$ 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):

- Make no change: For some/all \tilde{V} , define $d_{\tilde{V}}(x, y) = d(x, y)$ for some/all $x, y \in \tilde{V}^{(0)}$. Thus geodesics are relative geodesics.
- “Electrify” some/all \tilde{V} by defining $d_{\tilde{V}}(x, y) = 0$ for all $x, y \in \tilde{V}$.
- “Cone off” some/all \tilde{V} by adding a new vertex and connecting all vertices of \tilde{V} to it by an edge of length $1/2$, and define $d_{\tilde{V}}$ by the metric this procedure induces, so that $d_{\tilde{V}}(x, y) = 1$ for all distinct $x, y \in \tilde{V}$.
- For some/all \tilde{V} , choose $d_{\tilde{V}}$ so that there is a constant C such that

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

for all $x, y \in \tilde{V}$. This is the choice we will make later on when we attach so-called *combinatorial horoballs* to each \tilde{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 5.2.1. *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 $[e]_\alpha$ lie in the relative geodesic γ , choose one for which the number of edges in γ is minimal. Note that γ will contain at least two edges.

Label the elements of $[e]_\alpha$, e_1, \dots, 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 $j \in \{1, \dots, m\}$. By minimality, the initial point of γ is $i(e_1)$ and the terminal point is $t(e_m)$. Let σ_j be the subpath of γ between $t(e_j)$ and $i(e_{j+1})$, for $j \in \{1, \dots, m-1\}$. Choose $\sigma \in \{\sigma_j\}$ such that $\ell_r(\sigma)$ is minimal. See Figure 5.1. Decompose the image of $\partial\alpha$ in X as a path p^m where p is not a proper power. The closed path 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 $j \in \{0, \dots, m-1\}$. Each path will connect two elements of $[e]_\alpha$ and the orbits will 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 5.1: Decomposition of γ into the σ_j . Suppose that $\sigma = \sigma_4$.




Figure 5.2: In this example, λ is made up of two orbits of σ and the edges e_1 and e_2 .

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 5.2. 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

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

which holds when $L \geq 0$ and $m \geq 3$.

Thus we have reduced to the case $m = 2$. 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 .

Corollary 3.1.13 tells us that $\partial\alpha$ embeds in \bar{X} , so the two paths λ_1 and λ_2 of $\partial\alpha \setminus \{e_1, e_2\}$ do not intersect in \bar{X} (labeled so that $t(e_1) \in \lambda_1$). Since σ starts in λ_1 and ends in λ_2 , we can find an innermost subpath σ' of σ whose endpoints lie in λ_1 and λ_2 , respectively, and which does not internally intersect $\partial\alpha \setminus \{e_1, e_2\}$. Note that σ' does not cross e_1 or e_2 , as this would provide an obvious way to decrease the relative length of γ .

Consider the compact subcomplex $E = \alpha \cup \sigma'$ of \bar{X} . By choice of σ' , $\pi_1(E) = \mathbb{Z}$. Let q be a reduced path in \bar{X} which represents a generator of $\pi_1(E)$, and $D' \rightarrow \bar{X}$ a reduced disk diagram with boundary q . Let $D = E \cup D'$. 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. We may also assume that ∂D is contained in $\partial\alpha \cup \sigma'$, since any 2-cell contributing an edge g to ∂D not in $\partial\alpha \cup \sigma'$ may simply be removed from D along with g without affecting that D is simply connected. Note 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 Lemma 3.1.15. 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 . Consider the natural reduced map $D \rightarrow X$. By

Proposition 3.2.6, 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 subcomplex of \bar{X} as well. In fact, all elements of $[f]_\beta$ are contained in σ' since they lie on the boundary of D . Thus (β, f, σ') is 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. \square

5.3 Patchings

The following construction is of critical importance for later arguments. It shows that certain non-simply connected subcomplexes of \bar{X} can be made simply connected without introducing extra exposed or extreme 2-cells, as follows.

Definition 5.3.1. (Patching). Let $\phi : E \rightarrow \bar{X}$ 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 $\phi_\# : E_\# \rightarrow \bar{X}$ 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 5.3.2. In view of the composition $\bar{X} \rightarrow \tilde{X} \rightarrow X$, where the first map is any inclusion of \bar{X} into \tilde{X} , reduced diagrams $D \rightarrow \bar{X}$ give rise to reduced diagrams $D \rightarrow X$ and vice versa by Remark 3.1.2. Whenever we have a patching $E_\# \rightarrow \bar{X}$, we will casually confuse it with the corresponding diagram $E_\# \rightarrow X$ in order to apply Propositions 3.2.6 and 4.0.6.

An *isolated edge* of a CW complex is one which is not in the boundary of any 2-cell.

Lemma 5.3.3. *Let $\phi : E \rightarrow \bar{X}$ be an inclusion of a compact connected 2-complex. Suppose that there is a path λ in E with the property that λ contains every isolated edge of E and maps to a relative geodesic in \bar{X} . Then a patching for ϕ exists.*

Proof. If E is simply connected, then ϕ is a reduced diagram, so set $\phi_\# = \phi$ and we are done. Otherwise let g_1, \dots, g_k be generators of $\pi_1(E)$. Let $E_0 = E$ and $\phi_0 = \phi$. For each i , let p_i be a reduced path in $E^{(1)}$ such that $[p_i] = g_i$. Let $\rho_i : D_i \rightarrow \bar{X}$ be a reduced disk diagram such

that $\rho_i(\partial D_i) = \phi(\rho_i)$. Inductively define $E_i = E_{i-1} \sqcup_{\rho_i} D_i$, and observe that there is a natural combinatorial map $\phi_i : E_i \rightarrow \bar{X}$. The map ϕ_0 is reduced, and we can make ϕ_i reduced by the following inductive procedure: If ϕ_i is not reduced, then by induction and the fact that D_i is reduced, there is a cancelable pair of 2-cells α and β in E_{i-1} and D_i , respectively. Let e denote the shared edge between α and β , and let σ_α and σ_β be the paths in $\partial\alpha \setminus e$ and $\partial\beta \setminus e$, respectively, from the terminal to the initial vertex of e , which are identified under ϕ_i . Modify E_i and ϕ_i by replacing D_i with $D_i \setminus \{\beta, e\}$ and identifying σ_α with σ_β . Note that this process preserves E as a subcomplex of E_i , and that, although we are modifying ∂D_i , $\rho_i(\partial(D_i \setminus \{\beta, e\}))$ is homotopic to ρ_i in $\phi_{i-1}(E_{i-1})$. It preserves homotopy type of E_i because it is a homotopy equivalence. Repeating as many times as necessary, we may assume that there is no cancelable pair between E_{i-1} and D_i , and thus that ϕ_i is reduced. Now $E_\# = E_k$ contains E , and since $E_\#$ is simply connected, $\phi_\# = \phi_k$ is a reduced diagram. By construction, it is also clear that $\phi_\#|_E = \phi$.

Observe that for each i , all isolated edges of E_i belong to λ .

It remains to prove that any essential 2-cell α belonging to $E_\# \setminus E$ is not exposed in $E_\#$. To that end, let α be an essential 2-cell belonging to $E_\# \setminus E$. Then α belongs to the complex D_i for some $i \geq 1$. Now, if α is exposed in $E_\#$, then there is some exposed edge e in $\partial\alpha$ such that $[e]_\alpha$ lies in ∂E_i . Since each edge of $[e]_\alpha$ also lies in ∂D_i , it must be the case that every edge of $[e]_\alpha$ is an isolated edge of E_{i-1} , and thus belongs to λ by the observation above. This contradicts Lemma 5.2.1. □

5.4 More local geometry of essential 2-cells

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

Lemma 5.4.1. *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 5.3.3, so let $E_\# \rightarrow X$ be a patching (writing X instead of \bar{X} in the abuse of notation justified by Remark 5.3.2). By Proposition 3.2.6, α is extreme in $E_\#$ with exposed edge f . Note that $f \notin [e]_\alpha$ since e_1 and e_2 are internal in $K_\#$. Thus there are two elements of $[f]_\alpha$, f_1 and f_2 , lying in distinct components of $\partial\alpha \setminus \{e_1, e_2\}$. 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 Lemma 3.1.15. \square

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 5.4.2. *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 $E = \alpha \cup \gamma$. Then E satisfies the hypothesis of Lemma 5.3.3, so let $E_\#$ be a patching for E . By Proposition 3.2.6, 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): Assume that every path in $\partial\alpha$ connecting e_i to e_{i+1} uses at least one essential edge. 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} . 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 Lemma 3.1.15. Without loss of generality, we may assume λ_1 has this property. By the initial assumption, λ_1 contains an essential edge. By Corollary 3.1.13, there is an essential 2-cell of $E_\#$ distinct from α , 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 Lemma 3.1.15. 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. \square

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

Lemma 5.4.3. *Let α be an essential 2-cell in \bar{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 \bar{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 5.4.2, there is an orientation of γ such that γ traverses each of e_1 through e_k in the positive direction, in turn, and $w e_i = e_{i+1}$ 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. \square

5.5 Convexity of vertex spaces

The following fact will also be useful.

Lemma 5.5.1. *The vertex spaces of \bar{X} are convex.*

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

Proof. Let γ be a geodesic edge path between vertices x and y of a vertex space \tilde{V} . By passing to an innermost subpath outside of \tilde{V} , we may assume that $\gamma \cap \tilde{V} = \{x, y\}$. Let γ' be a shortest path from x to y in \tilde{V} . Note that neither γ nor γ' backtrack. Also, the first edges of γ and γ' are not identified by the innermost subpath assumption; neither are the last edges. Thus the loop $\gamma(\gamma')^{-1}$ is cyclically reduced, so we may fill it with a reduced diagram D by Theorem 3.1.4. If D contains an essential 2-cell, then by Lemma 3.1.8, 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 5.2.1. Thus D contains no essential 2-cells and so γ also maps to \tilde{V} , which is also a contradiction. \square

Chapter 6

Relative hyperbolicity

Let X be a staggered generalized 2-complex with locally indicable vertex groups and $n(X) \geq 2$. From this point onward, assume that the underlying graph of the total space X_{tot} is finite. Note that this does not imply that X_{tot} is compact as vertex spaces may not be. However, it 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 fact in this chapter.

By way of motivation, recall that δ -hyperbolic groups are a class of groups defined by Gromov in terms of a δ -thin triangle condition. This is a property defined for a general geodesic metric space Y which can be concisely stated as follows: There is a $\delta > 0$ such that for any geodesic triangle in Y , the δ -neighborhood of any two sides contains the third side. It is a coarse negative curvature property which, when it appears in the Cayley graph of a given finitely generated group, is independent of generating set (after possibly modifying δ). See [Bow06] or [Sis14] for a synopsis.

For a group G with a fixed finite presentation P , it is well-known (see [Bri02] for example) that the Cayley graph of G with respect to P is δ -hyperbolic for some δ if and only if G has a linear Dehn function for P , which means, roughly, that the maximum area of a van Kampen diagram in the Cayley 2-complex for G with respect to P grows linearly in its perimeter.

In this spirit, we will use a definition of relative hyperbolicity in terms of linear *relative* Dehn functions, as defined below (following [Hru10]).

Definition 6.0.1. (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 ($\mathcal{S} = \mathcal{S} \sqcup \overline{\mathcal{S}}$), $\mathcal{S} \cup \mathcal{P}$ is a generating set for G , and the kernel of the natural map from $F(\mathcal{S}) * (*_{P \in \mathbb{P}} P) \rightarrow G$ is finitely normally generated, where $F(\mathcal{S})$ denotes the free group on the set \mathcal{S} .

Definition 6.0.2. (Linear relative Dehn function). Suppose (G, \mathbb{P}) has a finite relative presentation with finite relative generating set $\mathcal{S} = \mathcal{S} \sqcup \overline{\mathcal{S}}$. Let \mathcal{P} be the union of all $P \in \mathbb{P}$. Let $H = F(\mathcal{S}) * (*_{P \in \mathbb{P}} P)$ and \mathcal{R} be a finite normal generating set for the kernel of the natural map $H \rightarrow G$. 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 $W = \prod_{i=1}^k h_i^{-1} R_i h_i$ 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 $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for each trivial word W of length at most m in $\mathcal{S} \cup \mathcal{P}$, $A(W) \leq f(m)$.

Definition 6.0.3. (Relatively hyperbolic) [Hru10, 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 [Osi06]. Hruska shows it is equivalent to no fewer than five others in the case that the set of peripheral subgroups is finite [Hru10].

Lemma 6.0.4. *Suppose X is a staggered generalized 2-complex with locally indicable vertex groups, $n(X) \geq 2$, and the underlying graph of the total space X_{tot} is finite. Let \mathbb{P} be the collection of vertex groups of X . Then $(\pi_1(X), \mathbb{P})$ is relatively hyperbolic.*

Remark: This result seems to be known, though we were unable to find a suitable reference in the literature. However, the isoperimetric inequality proved in [DH91, Theorem 3.3] implies that

the pair $(\langle\langle \frac{A*B}{w^m} \rangle\rangle, \{A, B\})$ (for A and B locally indicable, w not conjugate into A or B , and $m \geq 2$) is relatively hyperbolic, which certainly covers the case that $G(X)$ is a dumbbell space.

Proof. We first construct a finite relative generating set for $G = \pi_1(X)$. Choose a maximal spanning tree T of essential edges in X_{tot} . Define S as follows. Fix a base point in T and orient the essential edges of $X_{\text{tot}} \setminus T$. Each edge e_i therein contributes an element to S corresponding to a reduced path which starts and ends at the base point, traverses e_i exactly once in the positive direction, and otherwise does not leave the tree T . Then $\mathcal{S} = S \sqcup \bar{S}$ is a finite relative generating set (where \bar{S} is the collection of inverses of elements of S). Moreover, by van Kampen's Theorem, a normal generating set \mathcal{R} for the kernel of the natural map from $F(\mathcal{S}) * (*_{P \in \mathbb{P}} P) \rightarrow G$ may be identified with the set of boundary paths of essential 2-cells of X .

Let \mathcal{P} be the union of all $P \in \mathbb{P}$. Let p be a reduced, cyclically reduced edge path in X_{tot} such that $[p]$ represents the trivial element of G . We may also view p as a trivial word over $\mathcal{S} \cup \mathcal{P}$. Let $L(p)$ denote the word length of p in $\mathcal{S} \cup \mathcal{P}$, and note that we can compute $L(p)$ by counting the number of essential edges of p in $X_{\text{tot}} \setminus T$ plus the number of nontrivial maximal subloops of p which lie entirely in a single vertex space. Let $D \rightarrow X$ be a planar reduced diagram for p which uses a minimal number of essential 2-cells, and call the number of essential 2-cells in such a diagram $\mathcal{A}(p)$. There is standard one-to-one correspondence (see [Bri02, Theorem 4.2.2], for example) between planar diagrams with k essential 2-cells and equations of the form $W = \prod_{i=1}^k h_i^{-1} R_i h_i$ for $R_i \in \mathcal{R}$ and $h_i \in H$, so $\mathcal{A}(p) = A(p)$. Thus, having a linear relative Dehn function with respect to the finite relative generating set above is equivalent to requiring that there exist constants a, b such that $\mathcal{A}(p) \leq am + b$ for each trivial word p with $L(p) \leq m$.

To find such constants, we will also need to consider the *Bass-Serre length* of p , denoted by $\ell(p)$ and defined for arbitrary paths in $\bar{X}^{(1)}$, which is just the number of essential edges occurring in p . We claim that:

- (1) $\ell(p)$ is bounded above by a linear function of $L(p)$, and

(2) $\mathcal{A}(p)$ is bounded above by a linear function of $\ell(p)$.

To see the first claim, note that since T is finite, there is a constant d such that any reduced path which stays entirely inside it (using only essential edges) can use at most d essential edges. In particular, any reduced path p' in X_{tot} with $\ell(p') > d$ will either use an essential edge of $X_{\text{tot}} \setminus T$ or contain a subloop representing a nontrivial element of some vertex space. Thus if p' is a subpath of p with $\ell(p') = d + 1$, p' contributes at least one unit of length to $L(p)$. This shows that

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

i.e.

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

For the second claim, use Dehn's algorithm: Let $D \rightarrow X$ 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 3.2.6. Since $n(X) \geq 2$, α has exponent at least two, and thus strictly more than half of the essential edges of $\partial\alpha$ lie on ∂D . Let D' be the unique component of $\overline{D \setminus \alpha}$ which contains essential 2-cells (it is unique since α is extreme). The path $p' = \text{im}(\partial D')$ 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 $\mathcal{A}(p) \leq a'\ell(p) + b'$.

Stacking the inequalities from claims (1) and (2) gives the required linear relative Dehn function.

□

Chapter 7

Walls and ladders

In this chapter, we will assume our vertex groups are cubulated and define walls as codimension-1 immersed hyperspaces in \bar{X} . The construction of [Sag95] will be used to obtain an action of $G = \pi_1(X)$ on an associated *dual cube complex*.

We first define the notion of a CAT(0) cube complex, following [Man16].

Definition 7.0.1. (Cube complex/midcube/NPC cube complex/CAT(0) cube complex/hyperplane). Let $I = [0, 1]$. An n -cube is a copy of I^n metrized as a subset of Euclidean space. A k -dimensional *face* of I^n is a subset in which all but k of the coordinates are held constant at either 0 or 1, and a *midcube* of a cube is obtained by setting exactly one coordinate equal to $\frac{1}{2}$. A *cube complex* is a metric space built from a disjoint union of cubes of various dimensions, glued together by isometries of faces. The 0-cubes will also be referred to as vertices; the 1-cubes as edges. The *link* of a vertex is the boundary of an ε -neighborhood for some $0 < \varepsilon < \frac{1}{2}$, viewed as a Δ -complex. A simplicial complex is called *flag* if, whenever the one-skeleton of a k -simplex is present in the complex, that k -simplex is present. A cube complex is called *non-positively curved (NPC)* if the link of each vertex is a flag simplicial complex. It is called CAT(0) if it is NPC and simply connected. A *hyperplane* of a CAT(0)-cube complex is a nonempty, closed, connected subspace whose intersection with each cube is either empty or a midcube. Hyperplanes are simply connected, 2-sided, divide the CAT(0) cube complex into two components, and are themselves CAT(0) cube complexes (see [Sag14]).

There is another meaning of CAT(0) for a general geodesic metric space Y (originally defined by Cartan, Alexandrov and Toponogov whose names form the acronym) which is stated in terms

of *comparison triangles*, and says roughly that geodesic triangles in Y are no “fatter” than their Euclidean counterparts. See [BH99]. As outlined by Sageev in [Sag14], CAT(0) cube complexes are CAT(0) in this more general sense by work of Bridson, Gromov, and the Cartan-Hadamard Theorem. The name is therefore justified since they possess this notion of non-positive curvature.

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 the underlying graph of X_{tot} is finite. In this setting, this is equivalent to the assumption that X is compact, as outlined below.

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 each copy of the universal cover \tilde{V} in \bar{X} 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 \bar{X} is locally finite and X is compact.

In what follows, for metric statements about copies of \tilde{V} in \bar{X} for a vertex space V of X (to which we also refer as *vertex spaces* by slight abuse of notation), we will use the ℓ_1 metric in the 1-skeleton of \tilde{V} unless otherwise specified.

Similarly to the description in [Man16], we define walls as components of a “midcube complex,” $M(\bar{X})$. The cube complex $M(\bar{X})$ and its natural map to \bar{X} are defined as follows.

We first describe the disjoint union of the cubes of $M(\bar{X})$. Fix $\frac{1}{2} > \varepsilon > 0$. Each cell of \bar{X} is either a cube of some dimension or an essential 2-cell. Each k -dimensional cube C of \bar{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 $M(\bar{X})$ 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 \tilde{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 $\text{int}(\alpha)$, 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 7.1 shows an illustration of some portions of walls in \bar{X} .




Figure 7.1: Some portions of walls in \bar{X} .

Note that the action of $\pi_1(X)$ on \bar{X} preserves the system of walls just defined. Also note that there are two types of walls in \bar{X} :

- (i) The walls which are dual to essential edges and do not intersect any \tilde{V} ; these walls are

graphs.

- (ii) The walls which nontrivially intersect some \tilde{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 7.0.2. *For any cell ω and walls Λ and Λ' of \bar{X} , if $im(\Lambda) \cap \omega$ is nonempty and $im(\Lambda) \cap \omega = im(\Lambda') \cap \omega$, then $\Lambda = \Lambda'$.*

However, it is not clear that the walls we have just defined are well-behaved in \bar{X} . For example, a priori, a wall could travel in some vertex space \tilde{V} , leave the space through some essential 2-cell α , and later come back to that same vertex space so that its image in \bar{X} intersects itself.

On the other hand, it is clear that portions of walls behave well in vertex spaces. We make the following basic observations about walls, vertex spaces of \bar{X} , and how walls behave therein. A *square* is a 2-cell which is not essential. These facts follow directly from the definition of a CAT(0) cube complex and the well-known behavior of the standard hyperplanes therein, and the proofs are omitted. See [Sis14], for example.

Lemma 7.0.3. *Let \tilde{V} be a vertex space of \bar{X} . Let $\Lambda \looparrowright \bar{X}$ be a wall and let Λ_V be a maximal connected component of the preimage of \tilde{V} in Λ . Let γ be a geodesic edge path in \tilde{V} and let s be a square of \tilde{V} . Then*

- Λ is an NPC cube complex.
- ∂s embeds in \bar{X} .
- Λ_V embeds in \bar{X} .
- $s \cap \Lambda_V$ 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 7.0.4. (Carrier/wall segment/ladder). For a wall $\Lambda \looparrowright \bar{X}$, the *carrier* of Λ is the smallest subcomplex of \bar{X} containing the image of Λ . A *wall segment* λ in a wall Λ is a local geodesic in $\Lambda^{(1)}$, embedded except possibly at its endpoints. The *ladder* associated to λ is the smallest subcomplex of \bar{X} containing the image of λ .

Note that ladders are necessarily at most 2-dimensional subcomplexes of \bar{X} .

Chapter 8

Walls embed and separate

In [LW13], ladders turn out to be simply connected. This is not necessarily true in our case, but they can be patched:

Lemma 8.0.1. *Let H be the ladder associated to a wall segment. Then H contains at most two extreme essential 2-cells, and there is a patching $H_{\#} \rightarrow \bar{X}$ for H .*

Proof. Consider the inclusion of H into \bar{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 wall segment for which H is the associated ladder, and observe that Lemma 3.2.7 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 5.3.3 are satisfied and $H_{\#} \rightarrow \bar{X}$ exists. \square

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

Lemma 8.0.2. *Let α be a 2-cell of \bar{X} (essential or not). If λ is a wall segment with both endpoints mapping to α , then $\text{im}(\lambda)$ is contained in α .*

Proof. Let H be the ladder associated to λ and let $K = \alpha \cup H$. Note that $\partial\alpha$ embeds in \bar{X} by either Corollary 3.1.13 or Lemma 7.0.3. 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 $\partial\alpha \cap \text{im}(\lambda)$ 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 ladder associated to λ' , and note that $K' = \alpha \cup H'$ is itself a ladder (by possibly extending λ' across α if necessary). By Lemma 8.0.1, K' has a patching $K'_{\#} \rightarrow \bar{X}$.

First suppose that α is a square. Then the image of λ' passes through an essential 2-cell by Lemma 7.0.3. Let u' and v' be the first points along $\text{im}(\lambda')$ 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 8.0.1. 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 3.2.7. The complex $K'_{\#}$ contradicts Proposition 3.2.6.

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 8.0.1, Proposition 3.2.6 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 geodesic in a CAT(0) cube complex, and the wall segment λ' shows that some wall segment (lying entirely in that CAT(0) cube complex) crosses it twice. This contradicts Lemma 7.0.3. 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 3.1.13. Connect midpoints of e and e' by a snipping arc running through the interior of α and observe that the wall segment λ' furnishes a contradiction to Lemma 3.1.15.

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

Proposition 8.0.3. (cf [LW13, Theorem 7.4]). *Each wall is a tree of hyperplanes and embeds in \bar{X} .*

Proof. If some wall Λ is not simply connected, then there exists a wall segment λ of positive

length in $\Lambda^{(1)}$ which is a loop. Let H be the ladder associated to λ . Note that H contains at least two 2-cells by since the boundaries of 2-cells of \bar{X} embed by Lemma 7.0.3 and Corollary 3.1.13. 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 8.0.2.

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 (i), and a tree of hyperplanes if it is a wall of type (ii).

Now suppose that a wall Λ does not embed in \bar{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 in which we will witness the intersection of Λ with itself. Thus we may choose a wall segment λ which intersects itself exactly once in a 2-cell α (essential or not) and let H be the ladder associated to λ . Note that H contains at least two 2-cells since the boundaries of 2-cells of \bar{X} embed by Lemma 7.0.3 and Corollary 3.1.13. Thus, we may find an arc of λ with endpoints mapping to $\partial\alpha$ whose image does not internally intersect α . This contradicts Lemma 8.0.2. \square

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

Corollary 8.0.4. *Each wall in \bar{X} is separating.*

Proof. For any point p in a wall Λ , Λ separates a neighborhood of p into exactly two components, by Lemma 8.0.3 and construction. Thus each wall is locally separating and has an I -bundle neighborhood. And since each wall is a tree of hyperplanes (also Lemma 8.0.3), each wall is contractible. Thus each I -bundle neighborhood is actually a product $\Lambda \times I$. Thus for each wall, \bar{X} decomposes as a graph of spaces with a single simply connected edge space. Since $H^1(\bar{X}) = 0$, 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 8.0.5. *Let γ be a relative geodesic edge path in a vertex space \tilde{V} of \bar{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 geodesic. Suppose Λ intersects γ in two distinct points x and y . Let λ be a wall segment connecting x to y and let H be the associated ladder. The subcomplex $K = H \cup \gamma$ satisfies the hypotheses of Lemma 5.3.3, so let $K_{\#}$ be a patching. Note $K_{\#}$ has a maximum of two extreme 2-cells by Lemma 8.0.1 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 3.2.6. 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 Lemma 3.1.15. 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 $\Lambda \cap \tilde{V}$ (which is a hyperplane in \tilde{V} by Lemma 7.0.3) crosses the geodesic γ twice, which contradicts the behavior of hyperplanes in CAT(0) cube complexes. \square

We record the following immediate corollary.

Corollary 8.0.6. *For each wall Λ and each vertex space \tilde{V} , $\Lambda \cap \tilde{V}$ is either empty or consists of a single hyperplane in \tilde{V} .*

Lemma 8.0.7. *Let γ be a relative geodesic in \bar{X} and suppose $\Lambda \cap \gamma$ consists of at least two distinct points x and y . If λ is a wall segment in Λ connecting x to y , then λ passes through at least one essential 2-cell.*

Proof. Let H be the ladder associated to λ , and let $K = H \cup \gamma$. Then K satisfies the hypotheses of Lemma 5.3.3, so let $K_{\#} \rightarrow \bar{X}$ 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 3.1.8. This implies that there

are no essential edges in γ , because any such edge would be isolated and nonseparating in $K_{\#}$. Thus $K_{\#}$ maps to a single vertex space \tilde{V} of \bar{X} . Since γ is a relative geodesic mapping to a single vertex space, it is a geodesic in that vertex space. The fact that $\Lambda \cap \tilde{V}$ crosses γ twice is a contradiction to Lemma 7.0.3. □

Chapter 9

Walls are relatively quasiconvex

In [LW13], 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 Lemma 6.0.4, $G = \pi_1(X)$ is hyperbolic relative to the vertex groups. In this chapter, 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 Chapter 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 9.1.1. *(cf [LW13, 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 the underlying graph of X_{tot} is finite, the set $C(X)$ is finite, and

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 ladder H in C containing x and y with associated wall segment λ , and γ does not internally intersect H . The subcomplex $K = \gamma \cup H$ satisfies Lemma 5.3.3, so let $K_{\#} \rightarrow \bar{X}$ be a patching. When choosing generators of $\pi_1(K)$ to perform the patching, choose them so that there is exactly one generator which uses the path γ . Call the disk associated to this generator D and make the choice that this is D_1 , the first disk, in the patching construction. With this choice we may assume there is a planar subcomplex D of $K_{\#}$, homeomorphic to a disk, such that 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.

Note that $K_{\#}$ has a maximum of two extreme 2-cells since H does (by Lemma 8.0.1). Thus Proposition 4.0.6 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 . These are the points we will show are uniformly close to H . If $z \in \bar{A}$, then $d(z, H) \leq \frac{W_X}{2}$. If $z \notin \bar{A}$, let δ be the maximal connected subpath of γ containing z such that $\text{int}(\delta) \cap \bar{A}$ is empty. Since every 2-cell of A has an edge on γ , the complex $\overline{D \setminus \bar{A}}$ is a tree of disks. Let D' be the maximal subcomplex of $\overline{D \setminus \bar{A}}$ which contains z and is homeomorphic to a disk. Let δ' be the path $\partial D' \setminus \text{int}(\delta)$ (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 9.1.

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.




Figure 9.1: 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 .

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 3.2.6, 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 5.4.3 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 4.0.6. On the other hand, if *any* of these statements hold, we immediately obtain a contradiction to Lemma 5.4.1, 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 3.1.8, α is exposed in D' with exposed edge e , say. By Lemma 5.2.1, 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. \square

Problem: Does Lemma 9.1.1 hold 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 [LW13].

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

Lemma 9.1.2. *Let Λ be a wall of \bar{X} . Then $H = \text{stab}(\Lambda)$ acts cocompactly on the carrier of Λ , and thus on Λ .*

Proof. Let C be the carrier of Λ in \bar{X} . We claim that there are finitely many H -orbits of cells

of C , which implies the result. To see this, let $\phi : \bar{X} \rightarrow X$ 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 wall segments 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 possible images. Enumerate these images $\lambda_1, \dots, \lambda_k$. By Lemma 8.0.2, any 2-cell α of C which maps to β has a well-defined type $i \in \{1, \dots, k\}$, 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 \bar{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 7.0.2), 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. \square

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 [Hru10]. 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 [Hru10] for the definitions of cusp-uniform action and truncated space.

Definition 9.2.1. (Relatively quasiconvex) [Hru10, Definition 6.6] (“QC-3”) Suppose G is countable, $\mathbb{P} = \{P_1, \dots, P_m\}$ 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 $\bar{X}^{(1)}$, which is decidedly not δ -hyperbolic, in general, by attaching “combinatorial horoballs” to form a space $A(\bar{X}^{(1)})$ which is δ -hyperbolic and on which G acts in a cusp uniform manner. The space $A(\bar{X}^{(1)})$ will play the role of Y in the definition above, and the disjoint union of essential edges of $\bar{X}^{(1)}$ will play the role of $Y \setminus U$.

Proposition 9.2.2. *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 \bar{X} is quasiconvex relative to the collection of vertex groups of X .*

Proof. As in Chapter 6, let $\mathbb{P} = P_1, \dots, P_m$ be the vertex groups of X and choose a maximal spanning tree T of essential edges of X_{tot} . Let $\mathcal{S} = S \sqcup \bar{S}$ be the set of oriented essential edges of X not in T and their formal inverses. Then \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 [Hru10, Definitions 4.1 and 4.3] for the precise construction. To each P_i is associated a CAT(0) cube complex which induces a natural left-invariant metric d_i on it. The rough idea is that for each coset gP_i , we take countably many copies of gP_i indexed by the naturals, attach “vertical edges” between each element of gP_i in every level and the corresponding element above and below it, and “horizontal edges” between elements of gP_i in the same level of d_i -distance less than or equal to 2^j , where j is the level. The original coset gP_i sits at level 0. Let $\mathcal{H}_\Gamma(g, i)$ be the combinatorial horoball above the coset gP_i , which by convention includes the original gP_i at level 0, as well as any edges added there. By [Hru10, 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.

On the other hand, let X_c be the space obtained by collapsing T to a point. This collapse lifts to a G -equivariant quotient map $f : \bar{X}^{(1)} \rightarrow \bar{X}_c^{(1)}$, where the target is obtained by collapsing each

copy of T in $\bar{X}^{(1)}$; this map is a quasi-isometry which is the identity when restricted to vertex spaces of $\bar{X}^{(1)}$.

Now, G acts naturally on $\bar{X}_c^{(1)}$, and each vertex space of $\bar{X}_c^{(1)}$ is stabilized by some gPg^{-1} . We label this vertex space \tilde{V}_g^i . We now form the augmented space $A(\bar{X}_c^{(1)})$ by building a combinatorial horoball $\mathcal{H}_X(g, i)$ above the zero-skeleton of \tilde{V}_g^i , again with respect to the cube complex metric, for each (g, i) (as before, $\mathcal{H}_X(g, i)$ includes the one-skeleton of \tilde{V}_g^i by convention). We can identify the group elements of gPg^{-1} with vertices of \tilde{V}_g^i via the orbit map (choosing the image of T in $\bar{X}_c^{(1)}$ as a base point). Thus, $\mathcal{H}_\Gamma(g, i)$ is a full subgraph of $\mathcal{H}_X(g, i)$ for each (g, i) .

Observe now that the Cayley graph Γ includes naturally inside of $\bar{X}_c^{(1)}$ via the orbit map, with edges of Γ mapping bijectively to essential edges of $\bar{X}_c^{(1)}$. By the observation of the previous paragraph, there is also a natural inclusion $A(\Gamma) \hookrightarrow A(\bar{X}_c^{(1)})$, which we now claim is a quasi-isometry. Assuming this claim, we have that $A(\bar{X}_c^{(1)})$ is δ -hyperbolic (after possibly modifying δ).

To see the claim, first choose $K > \max_i(\text{diam}_{d_i}(P_i))$. It is clear that $A(\Gamma)$ is K -cobounded in $A(\bar{X}_c^{(1)})$. 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 $d_{A(\bar{X}_c^{(1)})}(x, y) \leq d_{A(\Gamma)}(x, y)$. It remains to find a constant K' such that $d_{A(\Gamma)}(x, y) \leq K' d_{A(\bar{X}_c^{(1)})}(x, y) + K'$. Let γ be a geodesic in $A(\bar{X}_c^{(1)})$ between x and y . Then γ decomposes as a path of the form $\gamma_0 e_1 \gamma_1 e_2 \dots e_k \gamma_k$ where each e_j is an essential edge and each γ_j is a (possibly empty) edge path in some $\mathcal{H}_X(g, i)$. By [GM08, Lemma 3.10], we may assume that each γ_j consists of at most two vertical segments and a single horizontal segment of length at most 3. Moreover, since the endpoints of γ_j lie in the image of the orbit map, these vertical segments also lie in $\mathcal{H}_\Gamma(g, i)$. Now, the horizontal segment h_j may not belong to $\mathcal{H}_\Gamma(g, i)$, but because its endpoints are connected by a path of length at most 3, there is a path h'_j of length 5 in $\mathcal{H}_\Gamma(g, i)$ 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 $k \leq |\gamma| = d_{A(\bar{X}_c^{(1)})}(x, y)$, so $d_{A(\Gamma)}(x, y) \leq 5d_{A(\bar{X}_c^{(1)})}(x, y) + 4$. Setting $K' = 5$ proves the claim.

Finally, build the augmented space $A(\bar{X}^{(1)})$. For each vertex space \tilde{V}_g^i of \bar{X} which is stabilized by $gP_i g^{-1}$, build a combinatorial horoball above its zero-skeleton using the cube complex metric as in the case of X_c . In fact, since the map f is the identity on the one skeleton of \tilde{V}_g^i , the horoball just added will be an isometric copy of $\mathcal{H}_X(g, i)$. The map f thus extends to a quasi-isometry $\tilde{f} : A(\bar{X}^{(1)}) \rightarrow A(\bar{X}_c^{(1)})$ which is the identity on combinatorial horoballs, so that $A(\bar{X}^{(1)})$ is δ -hyperbolic (after possibly modifying δ).

Now, we claim that G has a cusp-uniform action on $A(\bar{X}^{(1)})$ with truncated space the disconnected union of all essential edges of $\bar{X}^{(1)}$. In other words, the vertex spaces of $\bar{X}^{(1)}$, 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 edges of S .

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

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

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

For each vertex space \tilde{V} of \bar{X} , define $d_{\tilde{V}}(x, y) = d_A(x, y)$ for all $x, y \in \tilde{V}^{(0)}$. 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 $H = \text{stab}(\Lambda)$, so

that H_X lies in C . Let x', y' in H_X , and let γ' be a relative geodesic in $\bar{X}^{(1)}$ between x' and y' (with respect to the admissible choice of pseudometrics above). Let γ be a geodesic in $A(\bar{X}^{(1)})$ which agrees with γ' on essential edges (it is clear by the construction of the pseudometrics that such a geodesic exists). Note that the intersection of γ with the truncated space is precisely the set of essential edges of γ . Applying Lemma 9.1.1 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. □

Chapter 10

Walls satisfy linear separation

In this chapter, 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 \bar{X} satisfy the *linear separation property*, which roughly means that the number of walls separating pairs of points in \bar{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 [HW14, Theorem 5.2].

The precise statement we will prove is as follows:

Proposition 10.0.1. *Suppose that $n(X) \geq 4$. Let d be the graph metric on $\bar{X}^{(1)}$ as before. There are constants $\kappa > 0$ and ε such that for any vertices $x, y \in \bar{X}$, the number of walls separating x and y is at least $\kappa d(x, y) - \varepsilon$.*

We will be assuming for contradiction that walls frequently *double-cross* geodesics. We will use the following definition.

Definition 10.0.2. (Double-crosses/double-crossed ladder). Let γ be a geodesic in $\bar{X}^{(1)}$ between two 0-cells x and y of \bar{X} . For every edge e of γ , there are two dual walls 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 *double-crosses* γ if there is a wall segment λ_e^x in Λ_e^x between v_e^x and another distinct point u_e^x along γ . If this behavior occurs we will pass to an initial such wall segment emanating from v_e^x and assume that

Λ_e^x does not cross γ between v_e^x and u_e^x . There is a unique ladder 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 *double-crossed ladder* of γ at (e, x) , if it exists. See Figure 10.1 for an illustration.




Figure 10.1: Some double-crossed ladders. The ladder H_e^x bends in the direction of x , and H_f^x bends in the direction of y . Here the rank of $\pi_1(Y_e^x)$ is 2. Some pathologies for double-crossed ladders may be ruled out immediately. For example, the depicted half-twist in H_f^x is ruled out by Corollary 8.0.4.

Definition 10.0.3. (Returns). Let Y_e^x be a double-crossed ladder of γ at (e, x) , with associated ladder H_e^x . We say that Y_e^x (or H_e^x) *returns* through an essential 2-cell if that 2-cell is the first or last essential 2-cell of H_e^x through which the wall segment λ_e^x passes, as we traverse λ_e^x starting from v_e^x . We use the notation α_e^x for the first 2-cell through which Y_e^x returns, and ω_e^x for the last.

Lemma 8.0.7 implies that whenever Y_e^x is a double-crossed ladder, α_e^x and ω_e^x always exist, and they are clearly unique. It is possible that $\alpha_e^x = \omega_e^x$.

Definition 10.0.4. (Bends in the direction of). Let $z \in \{x, y\}$. Let Y_e^z be a double-crossed ladder of γ at (e, z) with associated ladder H_e^z . We say that Y_e^z (or H_e^z) *bends in the direction of x* if $d(u_e^z, x) < d(v_e^z, x)$. Otherwise we say that Y_e^z (or H_e^z) 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 10.0.5. *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_e^x double-crosses γ . Then there is a double-crossed ladder Y_e^x of γ at (e, x) with associated ladder H_e^x which bends in the direction of x .*

Proof. Suppose that every double-crossed ladder Y_e^x bends in the direction of y . Let $Y = Y_e^x$ be a double-crossed ladder with the property that Λ_e^x does not cross γ between $v = v_e^x$ and $u = u_e^x$. By Corollary 8.0.4, $\bar{X} \setminus \Lambda_e^x$ decomposes into two components \bar{X}_{in} and \bar{X}_{out} , labeled so that $\gamma' = \gamma_e^x$ maps to \bar{X}_{in} .

Let $\alpha = \alpha_e^x$ and let e_1 and e_2 be the edges of $\partial\alpha$ which are dual to $\lambda = \lambda_e^x$ (they may be essential or not), labeled so that there is a path from e_1 to e 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}_{in} and σ_{out} maps to \bar{X}_{out} (we may do this since $\alpha \cap \Lambda_e^x$ consists only of the arc $\alpha \cap \lambda$ by Lemma 8.0.2). Thus no point of σ_{out} lies along γ' .

Note that Y satisfies the hypotheses of Lemma 5.3.3 and let $Y_{\#}$ be a patching for Y . By Lemma 8.0.1, 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 3.1.8 or Proposition 3.2.6. 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 5.4.2 implies that every element of $[f]_{\alpha}$ lies along γ' , which contradicts Lemma 5.2.1. 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 $Y \cap (\gamma \cup \lambda)$ now furnishes a contradiction to Lemma 3.1.15. 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 5.4.3. \square

The following definition describes an impossible configuration of a pair of double-crossed ladders in \bar{X} . We will show that if linear separation fails we can find such a configuration.

Definition 10.0.6. (Double-crossed pair of ladders). Let γ be a geodesic in $\bar{X}^{(1)}$ 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 double-crossed ladders 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 \bar{X} a *double-crossed pair of ladders*. We denote by ω_a the last essential 2-cell through which Y_a returns, λ_a the wall segment associated to Y_a , and H_a its associated ladder. Similarly define ω_b , λ_b , and H_b .

Lemma 10.0.7. *There does not exist a double-crossed pair of ladders in \bar{X} .*

Remark: 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-crossed pair of ladders. 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 5.3.3, and let $Y_\#$ be a patching. By Lemma 8.0.1, 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 3.2.6.

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 $\check{\alpha}_b$ 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 3.2.7 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 10.2.




Figure 10.2: 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 3.1.15. 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 3.1.15. 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 5.4.2. Lemma 5.4.2 also implies that every element of $[f_a]_{\omega_a}$ lies along γ , which contradicts Lemma 5.2.1.

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 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 5.4.2), which we have already ruled out.

There are now some 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 10.3. 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 ladder 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 3.1.15.

- 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 10.4. 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




Figure 10.3: An example of what could happen in case 1. The highlighted blue path gives the contradiction to Lemma 3.1.15.

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 3.1.15.

- Subcase 2: The edge e_a is strictly closer to x along γ than f_a is. This subcase is illustrated in Figure 10.5. 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 8.0.2, 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 3.1.15.

For the case in which f_a is closer to x than f_b , the argument is identical, except that we exchange




Figure 10.4: An example of subcase 1. The highlighted blue path gives the contradiction to Lemma 3.1.15.




Figure 10.5: The general picture in subcase 2. The highlighted blue path gives the contradiction to Lemma 3.1.15.

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 . \square

The following lemma now easily implies linear separation.

Lemma 10.0.8. *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 9.1.1, 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 double-cross γ , then we are done. Thus, assume that Λ_e^x double-crosses γ . Fix a wall segment λ_e^x associated to this double crossing and let Y_e^x be the associated double-crossed ladder. By Lemma 10.0.5, 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 , $\{e = e_1, e_2, e_3, \dots\}$. Let k be the largest integer with the property that $\Lambda_{e_k}^x$ double crosses γ and such that α_a is the first essential 2-cell through which some wall segment $\lambda_{e_k}^x$ in $\Lambda_{e_k}^x$ returns. Since there are at most W_X wall segments passing through α_a , $k \leq W_X$. Define Y_a to be the double-crossed ladder associated to $\lambda_{e_k}^x$. By Lemma 10.0.5, we may assume Y_a bends in the direction of x . In particular, e_{k+1} exists.

Now, observe that the wall $\Lambda_{e_{k+1}}^x$ crosses γ exactly once. Indeed, if not, then there is a double-crossed ladder $Y_b = Y_{e_{k+1}}^x$ at (e_{k+1}, x) which bends in the direction of x by Lemma 10.0.5, and $\alpha_a \neq \alpha_b$ by definition of k . Thus $Y_a \cup Y_b$ is a pair of double-crossed ladders. This contradicts Lemma 10.0.7. Thus $W = W_X + 1$ satisfies the conclusion of the lemma. \square

Proof of Proposition 10.0.1. By Lemma 10.0.8, $\kappa = \frac{1}{W_X+1}$ and $\varepsilon = 1$ do the trick. \square

Problem: Just as Lauer and Wise ask in [LW13], 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 pair of double-crossed ladders in this situation, since one has less control over the direction in which double-crossed ladders bend.

Chapter 11

Existence of the action

In this chapter 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 construction” to obtain an action of $\pi_1(X)$ on a CAT(0) cube complex.

Definition 11.0.1. (Wallspace/dual cube complex). Let Y be a metric space and let \mathcal{W} be a collection of closed, connected subspaces of Y (called *walls*), 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* [Sag95]. The hyperplanes of this cube complex are in one-to-one correspondence with the original walls. Very roughly, the k -dimensional cubes of $\mathcal{C}(Y)$ correspond to k -element pairwise-transverse subsets of \mathcal{W} in Y . A summary can be found in [HW14, Construction 3.2, Theorem 3.7, Remark 3.11].

Properness of this action in our setting will follow immediately from what we proved in Chapter 10. Cocompactness will follow by an application of [HW14, Theorem 7.12]. We state a simplified version of this theorem below.

Theorem 11.0.2. (cf [JW17, 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 $\text{stab}(\Lambda)$ 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 $\text{diam}(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ for some $d = d(\Lambda)$.

Then there exists a compact subcomplex K such that $\mathcal{C}(Y) = GK \cup \bigcup_{P \in \mathbb{P}} GC_*(Y_P)$. In particular, G acts cocompactly on $\mathcal{C}(Y)$ provided that each $\mathcal{C}_*(Y_P)$ is P -cocompact.

For us, $G = \pi_1(X)$, $Y = \bar{X}$, \mathcal{W} is the collection of walls we defined in \bar{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 \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 $\text{stab}(Y_P) = P$.

In order to apply this theorem, it remains to show that each $\mathcal{C}_*(Y_P)$ is P -cocompact, as we will see. The following key lemma says, roughly, that a geodesic with large projection to Y_P comes very close to Y_P .

Lemma 11.0.3. *Fix Y_P . Suppose γ is a geodesic in $\bar{X}^{(1)}$ 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 from x to π_x (resp. y to π_y), and let γ' be a geodesic edge path from π_x to π_y . Orient everything so that $\sigma = \gamma\gamma_y\gamma'\gamma_x$ is a closed loop. Note that γ' lies in Y_P by Lemma 5.5.1. 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. We make σ cyclically reduced as follows. First 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 . Now, there may be backtracking at x , so 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 $\gamma_0 = \gamma|_{[x', y']}$, $\gamma_{x'} = \gamma_x|_{[\pi_x, x']}$, and $\gamma_{y'} = \gamma_y|_{[y', \pi_y]}$. Redefine $\sigma = \gamma_0 \gamma_{y'} \gamma' \gamma_{x'}$. It is clear that there is no folding of σ at x' or y' , so σ is cyclically reduced. See Figure 11.1.

Fill σ with a planar reduced disk diagram $D \rightarrow \bar{X}$ using Lemma 3.1.4. 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 5.2.1).
- No element of $[e]_\alpha$ lies along γ' (since by Lemma 5.5.1 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 3.2.7, α is not extreme.

Using this claim and applying Proposition 4.0.6 and Lemma 3.1.8, 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 $D \setminus D'$.




Figure 11.1: 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 blue essential edge f on γ , showing that γ passes close to Y_p .

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. \square

Lemma 11.0.4. *Each $\mathcal{C}_*(Y_P)$ is P -cocompact.*

Proof. Suppose that Λ is a wall of \bar{X} with the property that $\text{diam}(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ 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 11.0.3. 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 11.0.3, 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 9.1.1), 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 $P = \text{stab}(Y_P)$ 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 $\text{diam}(\Lambda \cap \mathcal{N}_d(Y_P)) = \infty$ 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. \square

Putting everything together, we have the main theorem for staggered generalized 2-complexes

with locally indicable vertex groups and $n(X) \geq 4$.

Theorem 11.0.5. *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 Chapter 7. Let \mathcal{C} be the cube complex dual to the action of G on the wallspace (\bar{X}, \mathcal{W}) .

By Proposition 10.0.1, the wallspace (\bar{X}, \mathcal{W}) satisfies linear separation. By [HW14, 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 $\text{stab}(Y_P) = P$.

Observe that all hypotheses of Theorem 11.0.2 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 Lemma 6.0.4. For each wall Λ , Lemma 9.1.2 implies $\text{stab}(\Lambda)$ acts cocompactly on it, and we showed $\text{stab}(\Lambda)$ is relatively quasiconvex in Proposition 9.2.2. Finally, each $\mathcal{C}_*(Y_P)$ is P -cocompact by Lemma 11.0.4.

Applying Theorem 11.0.2, the action of G on \mathcal{C} is cocompact and the theorem is proved. \square

Corollary 11.0.6. *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. We may assume that w is cyclically reduced. 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 $\pi_1(X_B) = B$, and then attaching a 2-cell to a path corresponding to the word w^n , so that $\pi_1(X) = G$. Observe that X is trivially a staggered generalized 2-complex and Theorem 11.0.5 applies. □

Chapter 12

Further directions

As we have hinted throughout, the most pressing question concerns the smallest exponent $n(X)$ for which Theorem 11.0.5 holds. We suspect that Theorem 11.0.5 is true when $n(X) \geq 2$, mainly because the main theorem of [LW13] is true in the exponent 2 case in a strong way (see Theorem 12.0.6 below). For our argument, we unfortunately found it necessary to impose the restriction that $n(X) \geq 4$, just as Lauer and Wise did, when seeking to prove properness of the action. In contrast to Lauer and Wise's setting, it also appeared that the condition that $n(X) \geq 4$ was necessary for the cocompactness argument.

Question 12.0.1. (Question 1.2.4 in Chapter 1). Does Theorem 11.0.5 hold when $n(X) \in \{2, 3\}$?

It is well known that one-relator groups with and without torsion (corresponding to groups where the relator is or is not a proper power, respectively, by Lemma 4.0.1) have very different behavior. We have focused on generalizations of the proper power case in this dissertation, but one may also ask what other hypotheses are needed to make Theorem 11.0.5 hold if $n(X) = 1$. It certainly holds for some groups. For example the fundamental groups of hyperbolic surfaces (which are cubulable), like all torsion free one-relator groups, decompose as groups of the form $A * B / \langle\langle w \rangle\rangle$ with w not a proper power and A and B free (and thus locally indicable and cubulable). But there are many one-relator groups for which it does not hold. The Baumslag-Solitar group $\langle a, b \mid b^{-1}ab = a^2 \rangle$, for example, (like many torsion free one-relator groups, see [GW19]) has Dehn function which is not linear or quadratic and thus cannot be cubulable. Thus the $n(X) = 1$ case is very interesting, but would require a more subtle statement and probably very different

techniques.

In the hyperbolic 3-manifold setting, a significant portion of the story to prove the virtual Haken conjecture (VHC) involved not only showing that a hyperbolic 3-manifold group G is cubulable, but that the cube complex \mathcal{C} on which it acts has the property of being *virtually special*. This condition was originally defined by Wise in terms of certain hyperplane pathologies avoided in some finite cover of the quotient of \mathcal{C} by the G -action, and it is not immediately clear that it is a property of *groups*. However, we state the following equivalent definition which was also mentioned in Chapter 1.

Definition 12.0.2. (Right-angled Artin group/virtually special). A *right-angled Artin group* (RAAG) is a group with a presentation of the form $\langle a_1, \dots, a_m \mid * \rangle$, where $*$ represents some collection of commutators of the a_i 's. A group G is *virtually special* if it has a finite index subgroup which embeds in a RAAG.

Agol's important contribution to proving the VHC was his proof that any cubulable group which is hyperbolic is virtually special. Beyond the properties enjoyed by cubulable groups listed in Chapter 1, a virtually special group G enjoys the following strong properties (see, e.g., [Wis12]):

- G is residually finite.
- If G is hyperbolic, the quasiconvex subgroups of G are separable.
- G is linear (also see [HW99]).

It is thus quite desirable to show that our favorite groups are virtually special, and the following is a natural question:

Question 12.0.3. Let w be a word in $A * B$ which is not conjugate into A or B , and $n \geq 2$. Under what circumstances is the group $G = A * B / \langle\langle w^n \rangle\rangle$ virtually special?

To avoid pathological examples, we would like to work in a setting in which the natural maps $A \rightarrow G$ and $B \rightarrow G$ are injective. Thus it seems necessary to require that A and B are locally

indicable, or at least torsion free. This forces us to assume that A and B are virtually special.

In case A and B are locally indicable, G is hyperbolic relative to $\{A, B\}$ (Lemma 6.0.4). Thus if A and B are hyperbolic themselves, then so is G [Osi06, Corollary 2.41], and Agol's theorem gives the following as a corollary to Theorem 11.0.6 (which we have already mentioned as Corollary 1.2.3):

Theorem 12.0.4. *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.*

But what if the factors are not hyperbolic? Namely:

Question 12.0.5. (Question 1.2.5 in Chapter 1). 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?

One strategy of proof would be to generalize Wise's theory of *quasiconvex hierarchies* to the relatively hyperbolic setting. In fact, this program is already carried out under the assumption that the factors A and B are virtually abelian [Wis09]. These hierarchies can be used to directly show the following strong generalization of cubulability of one-relator groups with torsion:

Theorem 12.0.6. [Wis09, Corollary 18.2] *Let $H = \langle a_1, \dots, a_m \mid w^n \rangle$ where w is cyclically reduced, and $n \geq 2$. Then H is virtually special.*

There is also a question in a slightly different direction. One could also consider small cancellation quotients of free products of virtually special groups. As remarked in Chapter 1, Martin and Steenbock show that a $C'(\frac{1}{6})$ quotient of $A * B$, where A and B are cubulable, is itself cubulable. The natural question is the following:

Question 12.0.7. Let G be a $C'(\frac{1}{6})$ quotient of $A * B$, where A and B are virtually special. Is G virtually special?

These groups are also hyperbolic relative to $\{A, B\}$, and perhaps it is conceivable that a suitable theory of quasiconvex hierarchies in the relatively hyperbolic setting could be used to tackle this question as well.

Bibliography

- [Ago13] Ian Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning. MR 3104553
- [BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486
- [Bow06] Brian H. Bowditch, *A course on geometric group theory*, MSJ Memoirs, vol. 16, Mathematical Society of Japan, Tokyo, 2006. MR 2243589
- [Bri02] Martin R. Bridson, *The geometry of the word problem*, Invitations to geometry and topology, Oxf. Grad. Texts Math., vol. 7, Oxford Univ. Press, Oxford, 2002, pp. 29–91. MR 1967746
- [BW12] Nicolas Bergeron and Daniel T. Wise, *A boundary criterion for cubulation*, Amer. J. Math. **134** (2012), no. 3, 843–859. MR 2931226
- [CCJ⁺01] Pierre-Alain Cherix, Michael Cowling, Paul Jolissaint, Pierre Julg, and Alain Valette, *Groups with the Haagerup property*, Progress in Mathematics, vol. 197, Birkhäuser Verlag, Basel, 2001, Gromov’s α -T-menability. MR 1852148
- [DH91] Andrew J. Duncan and James Howie, *The genus problem for one-relator products of locally indicable groups*, Math. Z. **208** (1991), no. 2, 225–237. MR 1128707
- [GM08] Daniel Groves and Jason Fox Manning, *Dehn filling in relatively hyperbolic groups*, Israel J. Math. **168** (2008), 317–429. MR 2448064
- [GW19] Giles Gardam and Daniel J. Woodhouse, *The geometry of one-relator groups satisfying a polynomial isoperimetric inequality*, Proc. Amer. Math. Soc. **147** (2019), no. 1, 125–

129. MR 3876736

- [How81] James Howie, *On pairs of 2-complexes and systems of equations over groups*, J. Reine Angew. Math. **324** (1981), 165–174. MR 614523
- [How82] ———, *On locally indicable groups*, Math. Z. **180** (1982), no. 4, 445–461. MR 667000
- [How87] ———, *How to generalize one-relator group theory*, Combinatorial group theory and topology (Alta, Utah, 1984), Ann. of Math. Stud., vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 53–78. MR 895609
- [HP84] J. Howie and S. J. Pride, *A spelling theorem for staggered generalized 2-complexes, with applications*, Invent. Math. **76** (1984), no. 1, 55–74. MR 739624
- [Hru10] G. Christopher Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. **10** (2010), no. 3, 1807–1856. MR 2684983
- [HW99] Tim Hsu and Daniel T. Wise, *On linear and residual properties of graph products*, Michigan Math. J. **46** (1999), no. 2, 251–259. MR 1704150
- [HW01] G. Christopher Hruska and Daniel T. Wise, *Towers, ladders and the B. B. Newman spelling theorem*, J. Aust. Math. Soc. **71** (2001), no. 1, 53–69. MR 1840493
- [HW14] G. C. Hruska and Daniel T. Wise, *Finiteness properties of cubulated groups*, Compos. Math. **150** (2014), no. 3, 453–506. MR 3187627
- [JW17] K. Jankiewicz and D. Wise, *Cubulating small cancellation free products*, 2017.
- [KM12] Jeremy Kahn and Vladimir Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. (2) **175** (2012), no. 3, 1127–1190. MR 2912704
- [LW13] Joseph Lauer and Daniel T. Wise, *Cubulating one-relator groups with torsion*, Math. Proc. Cambridge Philos. Soc. **155** (2013), no. 3, 411–429. MR 3118410

- [Man16] Jason F. Manning, *Cubulating spaces and groups, lecture notes (working draft)*, 2016.
- [MS17] Alexandre Martin and Markus Steenbock, *A combination theorem for cubulation in small cancellation theory over free products*, *Ann. Inst. Fourier (Grenoble)* **67** (2017), no. 4, 1613–1670. MR 3711135
- [Osi06] Denis V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, *Mem. Amer. Math. Soc.* **179** (2006), no. 843, vi+100. MR 2182268
- [Per02] G. Perelman, *The entropy formula for the Ricci flow and its geometric applications*, *ArXiv Mathematics e-prints* (2002).
- [Per03] ———, *Ricci flow with surgery on three-manifolds*, *ArXiv Mathematics e-prints* (2003).
- [Sag95] Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, *Proc. London Math. Soc. (3)* **71** (1995), no. 3, 585–617. MR 1347406
- [Sag14] ———, *CAT(0) cube complexes and groups*, *Geometric group theory, IAS/Park City Math. Ser.*, vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 7–54. MR 3329724
- [Sis14] Alessandro Sisto, *Lecture notes on geometric group theory*, 2014.
- [SW79] Peter Scott and Terry Wall, *Topological methods in group theory*, *Homological group theory (Proc. Sympos., Durham, 1977)*, *London Math. Soc. Lecture Note Ser.*, vol. 36, Cambridge Univ. Press, Cambridge-New York, 1979, pp. 137–203. MR 564422
- [SW05] Michah Sageev and Daniel T. Wise, *The Tits alternative for CAT(0) cubical complexes*, *Bull. London Math. Soc.* **37** (2005), no. 5, 706–710. MR 2164832
- [Thu82] William P. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, *Bull. Amer. Math. Soc. (N.S.)* **6** (1982), no. 3, 357–381. MR 648524

- [Wis04] D. T. Wise, *Cubulating small cancellation groups*, *Geom. Funct. Anal.* **14** (2004), no. 1, 150–214. MR 2053602
- [Wis09] Daniel T. Wise, *Research announcement: the structure of groups with a quasiconvex hierarchy*, *Electron. Res. Announc. Math. Sci.* **16** (2009), 44–55. MR 2558631
- [Wis12] _____, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012. MR 2986461