

Cubulating one-relator groups with torsion

Joseph Lauer

Department of Mathematics and Statistics,
McGill University, Montréal
Québec, Canada

April, 2007

A thesis submitted to McGill University in partial fulfillment of the requirements of
the degree of
Master of Science

Copyright © Joseph Lauer, 2007



Library and
Archives Canada

Published Heritage
Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et
Archives Canada

Direction du
Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence
ISBN: 978-0-494-38412-1
Our file Notre référence
ISBN: 978-0-494-38412-1

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

Acknowledgements

This thesis would have been impossible without the intense commitment and dedication of Dani Wise. In my experiences his work ethic and ambition are unparalleled, and I hope his influence on me will be apparent as I continue to develop mathematically.

I would also like to thank my fellow students in Montreal for their mathematical and social support. In particular Jitendra, Joe, Émilie, Nick, David, Ilya and Montserrat. You have all proved to be very useful.

Finally, thanks to Mom, Dad and Caleb for pretty much everything I have.

Abstract

Let $\langle a_1, \dots, a_m \mid w^n \rangle$ be a presentation of a group G , where w is freely and cyclically reduced and $n \geq 2$ is maximal. We define a system of codimension-1 subspaces in the universal cover, and invoke a construction essentially due to Sageev to define an action of G on a CAT(0) cube complex. By proving easily formulated geometric properties of the codimension-1 subspaces we show that when $n \geq 4$ the action is proper and cocompact, and that the cube complex is finite dimensional and locally finite. We also prove partial results when $n = 2$ or $n = 3$. It is also shown that the subgroups of G generated by non-empty proper subsets of $\{a_1, a_2, \dots, a_m\}$ embed by isometries into the whole group.

Résumé

Soit $\langle a_1, \dots, a_m \mid w^n \rangle$ une présentation d'un groupe G , où w est réduit librement et cycliquement, et $n \geq 2$ est maximale. Nous définissons une collection de codimension-1 sous-ensembles dans le revêtement universel, et utilisons le travail de Sageev pour définir une action de G sur un $CAT(0)$ cube complex. Nous montrons que l'action est propre et cocompacte quand $n \geq 4$. Nous montrons aussi que les sous-groupes de G engendrés par des sous-ensembles de $\{a_1, a_2, \dots, a_m\}$ sont paraboliques dans le groupe entier.

CONTENTS

| | |
|---|----|
| 1. Introduction | 5 |
| 2. Preliminaries | 8 |
| 3. Towers and staggered 2-complexes | 9 |
| 4. Exposed and extreme 2-cells | 11 |
| 5. Magnus subgroups | 16 |
| 6. Thin 2-complexes | 17 |
| 7. Hypergraphs and hypercarriers | 20 |
| 8. The hypergraphs are trees and embed in \bar{X} | 22 |
| 9. Constructing the CAT(0) cube complex | 24 |
| 10. Hypercarriers are quasiconvex | 26 |
| 11. Linear separation | 28 |
| 12. CAT(0) cubulation | 32 |
| References | 34 |

1. INTRODUCTION

In the 1970's Bass and Serre completely characterized groups that act essentially on trees. They proved that a group G acts freely on a tree if and only if G splits over a subgroup H as a non-trivial amalgamated free product or HNN extension. Their work was summarized by Serre in [26].

Trees are perhaps the simplest non-trivial examples of CAT(0) spaces. The CAT(0) condition, which is a global non-positive curvature condition, was named and given prominence by Gromov [7]. The condition itself is an old one, having been defined and developed by Alexandrov [1] and Toponogov beginning in the 1940's. Since trees are formed by identifying endpoints of a collection of edges isometric to an interval they are also cube complexes. In fact 1-dimensional CAT(0) cube complexes correspond exactly to trees.

In [25] Scott proved that if G splits over H , then H is an example of a codimension-1 subgroup, meaning that the coset graph $G \backslash H$ has more than one end. The reasoning behind this terminology is that the motivating examples are infinite cyclic subgroups of the fundamental group of a closed orientable surface, and the fundamental group of an incompressible two sided surface in a 3-manifold group.

The relationship between these two ideas was discovered in a fundamental work of Sageev [23]. He proved the following:

Sageev's Theorem *A group G acts essentially on a CAT(0) cube complex if and only if G contains a codimension-1 subgroup.*

By the work of Gerasimov [6] and Niblo-Roller [19] it was later understood that G acts *essentially* meant that the action of G had no global fixed point. The theory that has developed as a result of Sageev's Theorem can be considered a natural higher dimensional analogue of Bass-Serre theory. Indeed, if G splits over H , then Sageev's construction using H as the codimension-1 subgroup produces the Bass-Serre tree associated to the splitting.

An adaptation of Sageev's construction has emerged that allows one to proceed without explicit knowledge of a codimension-1 subgroup. This approach relies on the analysis of certain codimension-1 subspaces, which we call hypergraphs, in the universal cover of the standard 2-complex of a group G . Once these subspaces have been defined one obtains an

action on a CAT(0) cube complex, and the existence of desired properties of this action is reduced to answering easily formulated questions about the geometry of the codimension-1 subspaces.

This process has proved successful in *cubulating* many classes of groups. We say that a group has been cubulated if it acts properly and cocompactly on a CAT(0) cube complex: Finitely generated word-hyperbolic coxeter groups [18], a certain class of small cancellation groups [27], graphs of groups with free vertex groups and cyclic edge groups that do not contain Baumslag-Solitar groups [13], and random groups with density $< 1/6$ [20] have all been cubulated.

In this thesis we study the construction for one-relator groups with torsion.

The main tool in our analysis of the properties of hypergraphs is a new structure theorem for disc diagrams in one-relator groups with torsion. In the spirit of the B.B. Newman Spelling Theorem it places restrictions on the structure, more specifically on the external 2-cells, of disc diagrams that contain internal 2-cells. The usefulness of this result, as it is applied in this thesis, is born out of the fact that it applies not only to disc diagrams, which are planar, but to any simply-connected compact 2-complex.

We now give a brief outline of the sections in the thesis.

In Section 2 we review some basic results about one-relator groups with torsion and explain some notation used throughout the thesis.

In Sections 3 and 4 we provide the reader with the necessary background on towers and staggered complexes. We try to give the reader some sense of why Howie [10] calls the towers the standard method in one-relator group theory.

In Section 5 we use towers to study Magnus subgroups. These are subgroups generated by proper subsets of the generating set. Magnus himself proved in the Freiheitssatz [16] that such subgroups are free. We provide a simple proof that the elements of Magnus subgroups embed by isometries into the whole group, and so in particular Magnus subgroups are quasiconvex. This was first observed by I. Kapovich in [14].

The main result in Section 6 is that a certain class of compact 2-complexes cannot contain any internal 2-cells. It generalizes a theorem of Pride [22] that G satisfies the $C(2n)$ small cancellation condition.

In Section 7 we define the main objects of study: hypergraphs and hypercarriers. Each hypergraph consists of a connected graph Ω and a map from Ω to the universal cover of the standard 2-complex of G . The image of this map in the universal cover is the codimension-1 subspace we are interested in studying. Hypercarriers are the minimal closed subcomplexes that contain the image.

In order to construct a CAT(0) cube complex from a system of hypergraphs we must first verify certain basic properties. For example, the graphs cannot contain cycles and the maps must be embeddings. This is done in Section 8, and in Section 9 we explicitly construct the cube complex \mathcal{C} on which G will act, and prove that it is indeed CAT(0).

In Section 10 we prove that hypercarriers are quasi-convex. This is necessary in the proof of the main theorem to show that the action of G on \mathcal{C} is cocompact. When $n \geq 4$ the hypercarriers are actually convex, but we give examples where this fails in the $n = 2$ and $n = 3$ cases.

In Section 11 we restrict to the case where $n \geq 4$. In that case we prove that the number of hypergraphs that separate two points of the universal cover grows linearly compared to the distance between them. This allows us, in Section 12, to use a result in [11] that provides general conditions for the action to be cocompact. As a result we obtain our main theorem:

Theorem 12.4 *Let $\langle a_1, \dots, a_m \mid w^n \rangle$ be a presentation of a group G , where w is freely and cyclically reduced and $n \geq 4$. Then G acts properly and cocompactly on a locally finite and finite dimensional CAT(0) cube complex.*

This result has already been established since $n \geq 4$ implies that G falls into the class of small cancellation groups covered by [27]. However, the hypergraphs defined in this work depend more naturally on the group structure of G , and it is expected that it will be possible to use them to establish the theorem in the $n = 2$ and $n = 3$ cases when the small cancellation arguments do not apply.

2. PRELIMINARIES

Let G be a group defined by the presentation $\langle a_1, \dots, a_m \mid w^n \rangle$, where w is freely and cyclically reduced and $n \geq 2$. Throughout the thesis we use X to denote the standard 2-complex of such a presentation, and \tilde{X} for its universal cover.

In this section we will review some of the basic properties of such groups beginning with the following theorem announced by B.B.Newman [17] in 1968:

Theorem 2.1 (Spelling Theorem). *Let v be a non-empty freely reduced word in G . If v represents the trivial word in G , then there exists a subword u of v which is also a subword of w^n such that $|u| > |w^{n-1}|$.*

In particular, this implies that such a presentation is a Dehn presentation for G and so G is δ -hyperbolic.

Corollary 2.2. *If v is a non-empty proper subword of w^n , then v is non-trivial in G .*

Proof. Suppose otherwise. It is clear from the Spelling Theorem that $|v| > |w^{n-1}|$. Choose s so that that vs is a freely reduced cyclic conjugate of w^n . Then s is non-empty since $|v| < |w^n|$, and trivial in G . Thus $|s| > |w^{n-1}|$, and hence $|vs| > |w^n|$, a contradiction. \square

As a consequence of Corollary 2.2 each path in \tilde{X}^1 homotopic to w^n is a simple cycle, and each such cycle is the attaching map of exactly n 2-cells. In this thesis we will work with \bar{X} , the space obtained from \tilde{X} by identifying 2-cells which have the same boundary.

For any CW-complex, K , the set of i -cells in K is denoted by K^i . For $i \geq 1$, i -cells are always assumed to be open.

The importance of the next definition will become apparent in Section 3 when we will be concerned with the existence of covering spaces with infinite cyclic deck transformation group.

Definition 2.3. (indicable, locally indicable) A group G is *indicable* if there is a homomorphism of G onto \mathbb{Z} , and *locally indicable* if every finitely generated subgroup is indicable.

The following theorem was proved independently by Brodskiĭ [4] using algebraic methods and Howie [9] using topological methods.

Theorem 2.4. *Torsion-free one-relator groups are locally indicable.*

Let \bar{G} be the group defined by the presentation $\langle a_1, \dots, a_m \mid w \rangle$. Then there is a natural homomorphism $\phi : G \rightarrow \bar{G}$ sending a word in the generators of G to the same word considered as an element of \bar{G} . In 1972 Fischer, Karass and Solitar [5] proved

Theorem 2.5. *Every torsion element in G is contained in $\ker(\phi)$.*

In particular, this implies that the only torsion in G is the *obvious* torsion:

Corollary 2.6. *Every torsion element in G is conjugate to a power of w .*

3. TOWERS AND STAGGERED 2-COMPLEXES

The use of towers to study one-relator groups was suggested by Magnus' original proof of the Freiheitssatz [16]. However, the first time they appeared explicitly was nearly fifty years later in the realm of 3-manifolds in Papakyriakopoulos' proof [21] of Dehn's Lemma and the Sphere Theorem. More recently Hruska and Wise [12] have used towers to generalize the B.B. Newman Spelling Theorem to staggered 2-complexes, and to prove results about the structure of disc diagrams in one-relator groups with torsion.

The background presented in this section and more is given by Howie in [10].

A map $A \rightarrow B$ between CW-complexes is *combinatorial* if its restriction to each cell of A is a homeomorphism onto a cell of B . All maps in this section are assumed to be combinatorial.

Definition 3.1. (tower, tower lift) A map between CW-complexes $A \rightarrow B$ is a *tower* if it can be written as a composition

$$A = B_n \hookrightarrow \hat{B}_{n-1} \rightarrow B_{n-1} \hookrightarrow \dots \hookrightarrow \hat{B}_2 \rightarrow B_2 \hookrightarrow \hat{B}_1 \rightarrow B_1 \hookrightarrow B,$$

where the maps alternate between inclusions of subcomplexes and covering maps, where the covers are regular and connected with infinite cyclic covering transformation group.

Let $f : D \rightarrow B$ be a map of connected CW-complexes. A map $f' : D \rightarrow A$ is a *tower lift* of f provided that there is a tower $T : A \rightarrow B$ such that $T \circ f' = f$.

A tower lift $f' : D \rightarrow A$ is *maximal* if for any tower lift $D \rightarrow A'$ of f' , the map $A' \rightarrow A$ is a homeomorphism.

If $f' : D \rightarrow A$ is a maximal tower lift and D is finite, then A is also finite since f' is surjective.

Note that a tower lift $D \rightarrow A$ is not maximal if $\pi_1 A$ is indicable and $\pi_1 D$ is not. Indeed, if $\psi : \pi_1 A \rightarrow \mathbb{Z}$ is an epimorphism, then the covering space corresponding to $\ker(\psi) \triangleleft \pi_1 A$ is regular and has deck transformation group \mathbb{Z} .

We will consider maximal tower lifts of combinatorial maps $K \rightarrow B$ where K is a compact 2-complex. The existence of such lifts was proved by Howie for any compact CW-complex in [8].

Definition 3.2. (staggered) Let K be a 2-complex such that the attaching map of each 2-cell is an immersion. A *staggering* of K is a linear ordering on the 2-cells and a linear ordering on a subset \mathcal{O} of the 1-cells (called *ordered 1-cells*) such that:

- (1) each 2-cell contains at least one 1-cell from \mathcal{O} in the image of its attaching map, and
- (2) for 2-cells α and β , if $\alpha < \beta$ then $\min(\alpha) < \min(\beta)$ and $\max(\alpha) < \max(\beta)$, where $\min(\alpha)$ and $\max(\beta)$ are respectively the least and greatest ordered 1-cells in the attaching map of α .

A 2-complex together with a staggering is said to be *staggered*. A presentation is *staggered* if its standard 2-complex can be staggered.

Example 3.3. Any one-relator presentation is staggered since there is a unique 2-cell. Any non-empty subset of the 1-cells appearing in the relator can be chosen for the ordered 1-cells.

Example 3.4. The presentation $\langle a, b, c, d, e \mid bc^{-1}a^2, ca^2e^{-1}, ed^2c \rangle$ can be staggered by taking the ordered set of 1-cells to be $\{a, b, d, e\}$ ordered by $b < a < e < d$. Then $bc^{-1}a^2 < ca^2e^{-1} < ed^2c$ is an ordering on K^2 which produces a staggering.

The connection between towers and staggered 2-complexes is given by the following result:

Construction 3.5. [10, Lemma 2]. If $\phi : A \rightarrow B$ is a tower, and B is a staggered 2-complex then there exists a staggering for A .

Proof. It is clear that a subcomplex of a staggered 2-complex has an induced staggering so the proof of Lemma 3.5 reduces to the case where ϕ is an infinite cyclic cover. Let ρ be

a generator for the covering transformation group. Define the ordering on the 1-cells and 2-cells of A by $\alpha < \beta$ if either $\phi(\alpha) < \phi(\beta)$ or $\beta = \rho^n(\alpha)$ for some $n > 0$. It is elementary to check that this defines a staggering of A . \square

In general the staggering of A in Lemma 3.5 is not unique. In fact, given an arbitrary staggered 2-complex Y with at least two 2-cells there is a new staggering of Y which keeps the same set of ordered 1-cells and reverses the original ordering on both the ordered 1-cells and the 2-cells. This process has the property that the least 2-cells of Y in the original ordering become the greatest 2-cells in the new ordering. In particular, this implies that for a given staggering, statements about the greatest 2-cells also apply to the least 2-cells.

Whenever a tower map $A \rightarrow B$ is used in this thesis it will always be assumed that the staggering of A is one inherited from B using Construction 3.5. However, there is no canonical staggering since choosing ρ^{-1} instead of ρ has the affect of reversing the ordering on the 2-cells mapping to the same 2-cell via the covering map, but otherwise leaves the relation unchanged.

In the next section we will need the following lemma. Proofs can be found in [10] and [12]:

Lemma 3.6. *Let T be a finite staggered 2-complex which has no infinite cyclic cover. Then the greatest 2-cell of T , say α , is attached along a path w^n , and w is a closed path in T^1 that passes through $\max(\alpha)$ exactly once. Furthermore, no other 2-cell is attached along $\max(\alpha)$.*

4. EXPOSED AND EXTREME 2-CELLS

The main goal of this section is to prove Theorems 4.7 and 4.11. The former provides a geometric condition that can be used to find candidates for the greatest and least 2-cells in a staggered 2-complex, and the latter is an extremely useful structure theorem for reduced disc diagrams in one-relator groups with torsion.

A *cancelable pair* in a combinatorial map $K \rightarrow B$ is a pair of 2-cells in K whose boundary cycles share a common 1-cell, e , and such that the two boundary paths beginning at e are mapped to the same path in B . A map $K \rightarrow B$ is *reduced* if it does not contain any cancelable pairs.

A *disc diagram* K is a 2-complex that is compact, simply-connected and planar. If K is a disc diagram then a combinatorial map $K \rightarrow B$ is called a *disc diagram in B* . If the map $K \rightarrow B$ is reduced then it is called a *reduced disc diagram in B* . Disc diagrams play a fundamental role in much of geometric group theory. For a complete introduction see [15].

The main property we will need is the following result due to van Kampen:

Theorem 4.1. *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 disc diagram $D \rightarrow Y$ with $\partial D = u$.*

If $\partial D = u$ then $D \rightarrow Y$ is called a *disc diagram for u* .

In what follows we will consider maps $D \rightarrow Y$ where D is a simply-connected compact 2-complex, but is not necessarily planar. In this case we introduce the following definition which extends the usual notion of the boundary of a disc diagram.

Definition 4.2. (isolated, boundary, external) Let C be a 2-complex. A 1-cell e in C is *isolated* if does not appear in the attaching map of some 2-cell.

Let E be the set of 1-cells in C that are either isolated or appear exactly once in the attaching map of a single 2-cell. Then the *boundary of D* , denoted by $\partial_E D$, is $\text{cl}(E)$.

A 2-cell α in D is *external* if $\partial\alpha \cap \partial_E D$ contains a 1-cell, and *internal* otherwise, where $\partial\alpha$ denotes the usual topological boundary of α .

Definition 4.3. (position) Let $\alpha \in \bar{X}^2$. Two 1-cells e_1 and e_2 in $\partial\alpha$ are in the *same position in α* if e_1 and e_2 map to the same 1-cell of X , and the path in $\partial\alpha$ from the terminal 0-cell of e_1 to the terminal 0-cell of e_2 is a cyclic conjugate of w^j for some $j \in \mathbb{Z}$. For a 1-cell $e \subset \partial\alpha$ we denote the n 1-cells in the same position as e in α by $[e]_\alpha$.

Definition 4.4. (extreme, exposed) Let $D \rightarrow X$ be a combinatorial map. A 2-cell α in D is *extreme* if there is a connected subpath γ of the attaching path of α such that:

- (1) γ is a subpath of $\partial_E D$,
- (2) $|\gamma| > (n - 1)|w|$.

A 2-cell α in D is *exposed* if there is a 1-cell e in $\partial\alpha$ such that every 1-cell in $[e]_\alpha$ lies in $\partial_E D$. In this case we say that the 1-cell e is exposed in α . Furthermore, if the 1-cell e maps to a 1-cell labeled a_i in X then we say that α has a 1-cell labeled a_i exposed.




FIGURE 1. In this picture $n = 3$. The 1-cells in $[e]_{\alpha_1}$ and $[e]_{\alpha_2}$ are labeled with a single arrow, and the 1-cells of $[f]_{\alpha_1}$ with a double arrow. Two cases are shown depending on whether or not e and f have the same orientation in $\partial\alpha_1$. The bold path is ρ .

Note that a 2-cell α is not extreme if there exists a 1-cell e such that two 1-cells in $[e]_\alpha$ are not in $\partial_E D$. However, it is possible for such a 2-cell to be exposed.

Definition 4.5. (adjacent) Two n -cells α and β in a CW-complex are *adjacent* if $\bar{\alpha} \cap \bar{\beta}$ contains an $(n - 1)$ -cell.

Lemma 4.6. Let $\psi : D \rightarrow X$ be a reduced map from a compact 2-complex, and let $\phi : D \rightarrow T$ be a maximal tower lift of ψ . If α_1 and α_2 are 2-cells of D that are adjacent, then $\phi(\alpha_1) \neq \phi(\alpha_2)$.

Proof. Suppose that $\phi(\alpha_1) = \phi(\alpha_2)$, and that e is a 1-cell in $\bar{\alpha}_1 \cap \bar{\alpha}_2$.

There exists a combinatorial homeomorphism $\Phi : \bar{\alpha}_1 \rightarrow \bar{\alpha}_2$ which preserves the labeling on the 1-cells. Since w is not a proper power this homeomorphism is unique up to rotating the image through a power of w .

Let f be a 1-cell in $\Phi^{-1}([e]_{\alpha_2})$ such that the distance between f and e in $\partial\alpha_1$ is minimal. Observe that $e \neq f$ since $\psi : D \rightarrow X$ is reduced.

Let ρ be the geodesic in $\partial\alpha_1$ between the terminal 0-cell of f and the terminal 0-cell of e . See Figure 1. Then ρ is a proper subword of w , and non-empty since w is freely and cyclically reduced. We claim that ρ gets mapped to a closed path ρ' in T . Indeed, since $\phi(\alpha_1) = \phi(\alpha_2)$ we have $\phi(f) = \phi \circ \Phi(f) = \phi(e_i) = \phi(e)$ for some $e_i \in [e]_{\alpha_2}$.

By Corollary 2.2 the tower map $T \rightarrow X$ sends ρ' to a path in X which is not null-homotopic. Thus $[\rho'] \neq 1$ in $\pi_1 T$. Let \bar{G} be the group defined by the presentation

$\langle a_1, \dots, a_m \mid w \rangle$. By Theorem 2.5 the image of $[\rho']$ under the map

$$\pi_1 T \rightarrow \pi_1 X \cong G \rightarrow \bar{G}$$

is non-trivial. Thus $\pi_1 T$ maps to a non-trivial subgroup of \bar{G} , and that subgroup is finitely generated since T is finite. Since \bar{G} is locally indicable this implies that $\pi_1 T$ is indicable, and so the tower lift $D \rightarrow T$ is not maximal, a contradiction. \square

Let $D \rightarrow X$ be a reduced combinatorial map into a staggered 2-complex, and let $\phi : D \rightarrow T$ be a tower lift. The ordering on the 2-cells of T given by Construction 3.5 induces a quasi-ordering on the 2-cells and ordered 1-cells of D in the following way: $(\phi(\alpha_1) < \phi(\alpha_2)) \Rightarrow (\alpha_1 < \alpha_2)$, and $(\phi(\alpha_1) = \phi(\alpha_2)) \Rightarrow (\alpha_1 = \alpha_2)$. Since this is only a quasi-ordering we can not expect there to be a unique greatest 2-cell in D . However, when the tower lift is maximal Lemma 4.6 implies that adjacent 2-cells in D are not equivalent.

The next result provides a connection between the exposed 2-cells of a reduced, compact, simply-connected 2-complex and its greatest [or least] 2-cells in the quasi-ordering induced by some maximal tower lift. We use Lemma 4.7 in Theorem 4.11 to prove that if the 2-complex has at least two 2-cells then there are at least two extreme 2-cells. The proofs are typical applications of towers to one-relator group theory. Both results are proved in [12] in the case where D is planar, and the same proofs work here.

Lemma 4.7. *Let D be a compact simply-connected 2-complex and let $\phi : D \rightarrow T$ be a maximal tower lift of a reduced map $D \rightarrow X$. If η is a greatest [or least] 2-cell of D , then η is exposed.*

Proof. There is no infinite cyclic cover of T since the tower lift is maximal, and T is finite since the map $D \rightarrow T$ is surjective. Let α be the unique greatest 2-cell of T . By Lemma 3.6 α is attached to T along w^n , and w is a closed path. Moreover, w passes through $\max(\alpha)$ exactly once, and no other 2-cell is attached along $\max(\alpha)$.

Let e be a 1-cell mapping to $\max(\alpha)$ and suppose $e_i \in [e]_\alpha$. If $e_i \notin \partial_E D$ then it must be contained in the attaching map of two distinct 2-cells. Each of these two cells must map to α since e_i maps to $\max(\alpha)$ and α is the only 2-cell attached along $\max(\alpha)$. Therefore, by Lemma 4.6, the map $D \rightarrow X$ is not reduced, a contradiction. \square

Notation 4.8. In what follows, we will often remove an exposed 2-cell, α , with an exposed 1-cell e in $\partial\alpha$, from a 2-complex D . This process also requires removing the (open) 1-cells in $[e]_\alpha$.

We write $D - (\alpha, e)$ to mean $D - \{\alpha \cup [e]_\alpha\}$.

Definition 4.9. (branch) Let $D \rightarrow X$ be a reduced map. If α is a 2-cell of D , and e is an exposed 1-cell in $\partial\alpha$ then the components of $D - (\alpha, e)$ are called the *branches* of D at (α, e) .

If α is an external 2-cell of D then there is a deformation retraction from D to $D - \alpha$. With this observation the following lemma is immediate:

Lemma 4.10. *Let D be a simply-connected 2-complex. If e is an exposed 1-cell in a 2-cell α then the branches of D at (α, e) are simply-connected.*

Theorem 4.11. *Let X be the standard 2-complex of a presentation $\langle a_1, \dots \mid w^n \rangle$, with $n \geq 2$. Let D be a compact, simply-connected 2-complex with at least two 2-cells and $D \rightarrow X$ be a reduced combinatorial map. Then D contains at least two extreme 2-cells.*

Proof. We prove the result by induction on the number of 2-cells in D . The result is true when D has exactly two 2-cells since if the common subpath is longer than w , then the 2-cells form a cancelable pair.

Let $D \rightarrow T$ be a maximal tower lift of $D \rightarrow X$. If T has a unique 2-cell, say α , then every 2-cell in D maps to α . Let η_1 and η_2 be distinct 2-cells of D . Then by Lemma 4.7 each η_i is exposed so the branches of each η_i with respect to some collection of 1-cells are defined. If each η_i has only one branch then they are both extreme and we are done. If one of them, say η_1 has at least two branches B_1 and B_2 . Then the two disc diagrams $D_1 = B_1 \cup \eta_1$ and $D_2 = B_2 \cup \eta_1$ are strictly smaller than D . Thus there exists a 2-cell $\alpha_1 \neq \eta_1$ in D_1 which is extreme in D_1 and a 2-cell $\alpha_2 \neq \eta_2$ in D_2 which is extreme in D_2 by the inductive hypothesis. Then since B_1 and B_2 are branches α_1 and α_2 are also extreme in D .

Now suppose Y has two distinct 2-cells. Choose 2-cells σ and τ in D which map to the greatest and least 2-cells of Y respectively. Since the ordering can be reversed τ is exposed. We then proceed as above with σ and τ replacing η_1 and η_2 . \square

5. MAGNUS SUBGROUPS

We will now demonstrate the usefulness of the methods developed in the last two sections by using them to give an easy proof that magnus subgroups are quasi-convex. This fact was first observed by I. Kapovich in [14].

Definition 5.1. (K-quasiconvex, convex) For a metric space (X, d) and $K > 0$, a subset $C \subseteq X$ is *K-quasiconvex* if $\gamma \subseteq \mathcal{N}_K(C)$ for every geodesic γ with endpoints in C .

A subset $C \subseteq X$ is said to be *convex* if it is 0-quasiconvex.

For a group G generated by a set A , we denote the cayley graph of G with respect to A by $\Gamma_A(G)$. If H is a subgroup of G then $V(H)$ denotes the vertices of $\Gamma_A(G)$ which correspond to elements of H .

Definition 5.2. (quasiconvex subgroup) Let G be group with generating set A . Then $H \leq G$ is *quasi-convex with respect to A* if $V(H)$ is a K-quasiconvex subset of $\Gamma_A(G)$ for some $K > 0$.

It is well known that the quasi-convexity of a subgroup of a δ -hyperbolic group does not depend on the choice of generating set. See [2] and [3] for details.

Definition 5.3. (Magnus subgroup) Let G be the group defined by the presentation $\langle a_1, a_2, \dots, a_m \mid w^n \rangle$. A *Magnus subgroup* of G is a subgroup of the form $\langle M \rangle \leq G$, where $M \subsetneq \{a_1, a_2, \dots, a_m\}$.

Magnus subgroups are free with basis M by the Freiheitsatz [16].

The following lemma is useful in the case, as with Magnus subgroups, when a particular generator is distinguished.

Lemma 5.4. *Let G be the group given by presentation $\langle a_1, a_2, \dots, a_m \mid w^n \rangle$, where $n \geq 2$, and let $D \rightarrow X$ be a reduced disc diagram. Then for each a_i there exists a 2-cell α in D such that $\partial\alpha$ contains an exposed 1-cell labeled a_i .*

Proof. As noted in Example 3.3 a staggering of X is determined by a choice of the ordered 1-cells \mathcal{O}_X and an ordering on that set. Given a generator a_i , choose the staggering of X induced by setting $\mathcal{O}_X = \{a_i\}$.

Let $D \rightarrow Y$ be a maximal tower lift of the disc diagram $D \rightarrow X$. Every 1-cell in \mathcal{O}_Y , and in particular the greatest 1-cell of Y , must map to a_i in X . Thus by Lemma 4.7 any 2-cell of D mapping to the greatest 2-cell of Y has an exposed 1-cell labeled a_i . \square

Theorem 5.5. *Let G be the group defined by a presentation $\langle A \mid w^n \rangle$. Let $H = \langle M \rangle$ be a Magnus subgroup. Then the map induced by the inclusion $H \hookrightarrow G$ is an isometric embedding of cayley graphs.*

Proof. Suppose otherwise. Then there exists a shortest possible subword in H , say γ , which is not a geodesic in G . Let δ be a geodesic in $\Gamma(G)$ between the endpoints of γ . Since γ is the shortest such word the closed path $\gamma\delta^{-1}$ is a simple loop. Let $D \rightarrow X$ be a reduced disc diagram for this closed path.

Let $a_i \in A$ be a generator of G such that $a_i \notin M$. By Lemma 5.4 there exists a 2-cell α of D whose boundary contains a 1-cell e , which is labeled a_i , and such that each 1-cell in $[e]_\alpha$ lies in $\partial_E D$.

By assumption there are no 1-cells labeled a_i in γ and so each e_i is contained in δ . Since δ is a geodesic $\partial\alpha \cap \delta$ is connected. Thus δ contains a subword of w^n longer than $|w^{n-1}|$ contradicting δ as a geodesic. \square

As an immediate corollary we obtain:

Corollary 5.6. *Magnus subgroups are quasiconvex.*

6. THIN 2-COMPLEXES

In this section we investigate the structure of reduced 2-complexes that contain internal 2-cells. In Theorem 6.3 we show that such a complex contains at least $2n$ extreme 2-cells.

Let $D \rightarrow T$ be a maximal tower lift of a reduced map $D \rightarrow X$. An important part of our analysis is understanding the behaviour of the (not necessarily unique) greatest and least 2-cells in subcomplexes of D . If $D' \subset D$ is a subcomplex, then the restriction of the map $D \rightarrow T$ to D' may not be a maximal tower lift of the composition $D' \hookrightarrow D \rightarrow X$. However, by looking locally at each of the greatest 2-cells of D' we obtain the same conclusion as in Lemma 4.7:

Lemma 6.1. *Suppose $D \rightarrow T$ is a maximal tower lift of a reduced map $D \rightarrow X$. Let D' be a connected subcomplex of D , and suppose that η is a greatest 2-cell of D' . Then η is exposed in D' .*

The quasi-ordering referred to in the following proof is the one induced by the original maximal tower lift $D \rightarrow T$.

Proof. Let $D' \rightarrow T'$ be a maximal tower lift of the map $D' \hookrightarrow D \rightarrow T$.

Let B be the smallest closed subcomplex of D' containing η and all the 2-cells adjacent to η . By Lemma 4.6 no 2-cell adjacent to η is equivalent to η in the quasi-ordering and so η is the unique greatest 2-cell in B .

Thus if $B \rightarrow T''$ is a maximal tower lift of $B \hookrightarrow D \rightarrow X$, then η maps to the greatest 2-cell of T'' . Therefore η is exposed in B by Lemma 4.7, and hence in D' since B includes all the 2-cells of D' adjacent to η . \square

Again, let $D \rightarrow T$ be a maximal tower lift of a reduced map $D \rightarrow X$. Let μ be an internal 2-cell of D . Using the quasi-ordering on 2-cells of D induced by the map $D \rightarrow T$ we define

$$\hat{G}_\mu = cl(\{\alpha \in D^2 \mid \alpha \geq \mu\})$$

and

$$\hat{L}_\mu = cl(\{\alpha \in D^2 \mid \alpha < \mu\}) \cup \bar{\mu}$$

and let G_μ and L_μ be the connected components of \hat{G}_μ and \hat{L}_μ respectively that contain μ . Note that μ is exposed in each of L_μ and G_μ by Lemma 6.1.

Lemma 6.2. *Each component of \hat{G}_μ or \hat{L}_μ is simply-connected.*

Proof. The components of \hat{G}_μ can be obtained from D by successively removing the least 2-cells. Let α_0 be any least 2-cell of D . By Lemma 6.1, α_0 will be exposed, and by Lemma 4.10 the components of $D - (\alpha_0, e)$ are simply-connected, where e is an exposed 1-cell in α . Repeat this process for a least 2-cell (in the original quasi-ordering) of $D - (\alpha_0, e)$. After finitely many steps we obtain \hat{G}_μ .

By successively removing the greatest 2-cells of D the same argument shows that the components of $L_0 = cl(\{\alpha \in D^2 \mid \alpha \leq \mu\})$ are simply-connected. Now let α be a 2-cell




FIGURE 2. In this example $n=3$, and the unshaded 2-cell is μ . This picture shows how a disc diagram might look when it is decomposed into \hat{G}_μ and \hat{L}_μ . We can see the 6 disjoint branches which we prove must each contain at least one extreme 2-cell of D .

different from μ such that $\alpha = \mu$ in the quasi-ordering. By Lemma 6.1, α has an exposed 1-cell e_1 and so the components of $L_0 - (\alpha, e_1)$ are simply-connected. Repeating this process for each such 2-cell we obtain \hat{L}_μ in finitely many steps. \square

The following theorem is used heavily in subsequent sections to restrict certain combinatorial configurations in \bar{X} .

Theorem 6.3. *Let $D \rightarrow X$ be a reduced map, where D is a simply-connected compact 2-complex and X is the standard 2-complex of the presentation $\langle a_1, \dots, a_m \mid w^n \rangle$ with $n \geq 2$. If D has an internal 2-cell then D contains at least $2n$ extreme 2-cells.*

Proof. Let $D \rightarrow T$ be a maximal tower lift of the map $D \rightarrow X$.

Let μ be an internal 2-cell of D , and let \hat{G}_μ , \hat{L}_μ , G_μ and L_μ be defined as above using the quasi-ordering induced by the map $D \rightarrow T$.

Since μ is exposed in G_μ there exists a 1-cell e in $\partial\mu$ such that each 1-cell in $[e]_\mu$ lies in ∂G_μ . Let B_1, \dots, B_n be the branches of G_μ at (μ, e) . See Figure 2. Each branch contains a 2-cell since μ is internal, and branches are disjoint since G_μ is simply-connected by Lemma 6.2. For $1 \leq i \leq n$ let G_i be the component of $\hat{L}_\mu \cup B_i$ containing μ .

The complex G_i contains at least one 2-cell which is strictly greater than μ since B_i contains a 2-cell adjacent to μ . This implies that any greatest 2-cell of G_i lies in B_i . By Lemma 6.1 there exists a 2-cell α in B_i such that α is exposed in G_i . But we claim that $\bar{\alpha} \cap \partial G_i = \bar{\alpha} \cap \partial D$. Indeed, if β is a 2-cell of D adjacent to α and β does not lie in \hat{L}_μ then $\beta \geq \mu$ and so β lies in G_i . Therefore α is exposed in D .

Thus there is a 2-cell in each of the n disjoint branches that is exposed in D . By the same (slightly simpler) argument applied to L_μ we obtain n more exposed 2-cells. Thus D contains at least $2n$ exposed 2-cells. To finish off the proof we need to translate this into information about extreme 2-cells.

Let R be the maximal subcomplex of D containing μ so that all exposed 2-cells in R are extreme. Clearly μ is internal in R and so by the work above R contains $2n$ extreme 2-cells $\alpha_1, \dots, \alpha_{2n}$. Suppose α_i is not extreme in D . Then α_i has at least two non-trivial branches. Let B be a non-trivial branch of α_i that does not contain μ . By Lemma 5.4 there exists $\beta_i \neq \alpha_i$, which is extreme in $B \cup \alpha_i$, and hence extreme in D . Repeating this for each i we obtain $2n$ extreme 2-cells. They are distinct since for $j \neq i$ each α_j lies in the branch of α_i containing μ . This completes the proof. \square

7. HYPERGRAPHS AND HYPERCARRIERS

We now define the main object of study in the sequel: hypergraphs and hypercarriers. Hypercarriers themselves play no part in the construction of the cube complex, but are introduced to facilitate easier proofs of the properties of hypergraphs.

An n -star is the unique graph with n vertices of degree one, and one vertex of degree n . Let α be a 2-cell in \bar{X} and let e be a 1-cell in $\partial\alpha$. Define $T_{[e]_\alpha} \subset \bar{\alpha}$ to be an embedded n -star whose degree one vertices map to the midpoints of the n 1-cells in $[e]_\alpha$. See Figure 3[a].

Let $T'_{[e]_\alpha}$ be an ϵ -thickening of $T_{[e]_\alpha}$. Then $\partial T'_{[e]_\alpha} \cap \alpha$ consists of n connected components homeomorphic to a 1-cell. See Figure 3[b].

Construct a graph Γ as follows:

- (1) $V(\Gamma) = \{v_x \mid x \in \partial(\bigcup T'_{[e]_\alpha}) \cap \bar{X}^1\}$, where the union is taken over all possible $[e]_\alpha$.




FIGURE 3. In this example $w = a^2b$ and $n = 4$. [a] The three subspaces $T_{[e_i]_\alpha}$. [b] The corresponding spaces $\partial T'_{[e_i]_\alpha}$. [c] The three images from the second row have been superimposed. This shows the image $\varphi(\Gamma)$ in a particular 2-cell of \bar{X} . Each such 2-cell contains the image of $n|w|$ 1-cells of Γ and vertices of Γ are mapped into the open 1-cells of \bar{X} . It is clear from this picture that φ is far from an embedding. However, we will prove in the next section that φ restricted to each component of Γ is an embedding.




FIGURE 4. Five 2-cells in \bar{X} , each containing an embedded 3-star. The bold lines represent the image of a subgraph of a hypergraph.

(2) $(v_x, v_y) \in E(\Gamma)$ if there exists a pair (α, e) such that x and y lie in the same connected component of $\partial T'_{[e]_\alpha} \cap \bar{\alpha}$.

There are exactly two vertices of Γ corresponding to each 1-cell in \bar{X} , and there is a natural injective map $\varphi_v : V(\Gamma) \rightarrow \bar{X}$, sending v_x to x . This map can be extended to an immersion $\varphi : \Gamma \rightarrow \bar{X}$ from the whole graph. If the edge (v_x, v_y) is witnessed by the pair (α, e) then φ sends the edge (v_x, v_y) to the path between x and y in $\partial T'_{[e]_\alpha} \cap \alpha$. See Figure 3[c].

Definition 7.1. (hypergraph) Let $\Lambda_i, i \in I$, be the connected components of Γ . A *hypergraph* in \bar{X} is any of the maps $\varphi : \Lambda_i \rightarrow \bar{X}$. Usually we will suppress the functional notation and only refer to the hypergraph Λ_i , and consider it a subset of \bar{X} . The 1-cells of \bar{X} which cross a hypergraph are said to be *dual* to it.

Remark 7.2. When $n = 2$ the ϵ -thickening is unnecessary since the construction results in paired hypergraphs running parallel through \bar{X} . Without the thickening, a single hypergraph replaces each pair, and φ then sends each vertex of Γ to the midpoint of a 1-cell. When we refer to the $n=2$ case in subsequent sections it will always be with the unthickened construction in mind.

Definition 7.3. (hypercarrier) Let Λ be a hypergraph. Then the *hypercarrier* of Λ is the smallest closed subcomplex of \bar{X} containing $\varphi(\Lambda)$.

A *hypergraph segment* is a finite path in a hypergraph. The *hypercarrier chain* associated to a hypergraph segment is the smallest closed subcomplex of \bar{X} containing that hypergraph segment. The *end 2-cells* of a hypercarrier chain are the 2-cells that contain the first and last edges of the associated hypergraph segment.

8. THE HYPERGRAPHS ARE TREES AND EMBED IN \bar{X}

We will now prove that the hypergraphs defined in the last section have the basic properties necessary to define the cube complex on which the group G will act.

Theorem 8.1. *Let H be the hypercarrier chain associated to an arbitrary hypergraph segment. Then H is simply-connected.*




FIGURE 5. By gluing a reduced disc diagram to each of these complexes we create simply-connected 2-complexes with internal 2-cells that have at most 2, 0 and 1 extreme 2-cells respectively (from left to right).

Proof. Suppose H is not simply-connected. By considering a counterexample so that the hypergraph segment is as short as possible we may assume that $\pi_1 H \cong \mathbb{Z}$. See Figure 5 for the different possibilities, depending on how the hypergraph segment sits inside \bar{X} .

Let \mathcal{P} be the set of closed paths in H^1 that represent a generator of $\pi_1 H$. Let $D \rightarrow X$ be a minimal area reduced disc diagram for paths in \mathcal{P} , and set $\rho = \partial D$. The 2-complex $L = H \cup_\rho D$ is simply-connected since $[\rho]$ is a generator of $\pi_1 H$.

Observe that since H has no isolated 1-cells every 2-cell of D is internal in L . In particular, since D is non-empty this implies that L contains an internal 2-cell.

Now, only the end 2-cell(s) of H can be extreme in L . Indeed, any other 2-cells contain two dual 1-cells in its boundary which are internal in H , and hence in L . Therefore L contains at most two extreme 2-cells and an internal 2-cell and so by Theorem 6.3 there exists a cancelable pair in L .

Since H and D are reduced any cancelable pair in L must contain one 2-cell from D , say α , and one 2-cell from H , say β . Let $\gamma = \bar{\alpha} \cap \bar{\beta}$, and define δ to be the path with the same endpoints as γ such that $\gamma^{-1}\delta = \partial\beta$. Let ρ' be the path obtained from ρ by replacing the subpath γ by δ . Then $[\rho'] = [\rho]$ in $\pi_1 H$, and $D - \alpha$ is a disc diagram for ρ' which has strictly smaller area than D , a contradiction. \square

Theorem 8.2. *Each hypergraph is a tree and embeds in \bar{X} .*

Proof. Suppose that there exists a cycle, C_k , in a hypergraph Λ . By Theorem 8.1 the hypercarrier chain, K , associated to C_k is simply-connected, and clearly K contains at least two 2-cells. Thus it must contain at least two extreme 2-cells by Theorem 4.11. But

no 2-cell of K is extreme since the 1-cells dual to Λ are internal in K . Thus each hypergraph is a tree.

Now suppose e_1 and e_2 are two edges of a hypergraph Λ which get mapped into the same 2-cell, say α , of \bar{X} . Let K be hypercarrier chain containing the image of the unique path between e_1 and e_2 in Λ . The same argument as above applies to K except that the 2-cell α could potentially be extreme in K . Since α is the only such 2-cell we arrive at the same contradiction. \square

The last result of this section is an easy consequence of Theorem 8.1.

Corollary 8.3. *For each $i \in I$, $\bar{X} - \Lambda_i$ consists of two infinite connected components.*

9. CONSTRUCTING THE CAT(0) CUBE COMPLEX

We now demonstrate how to use a hypergraph system satisfying the properties proved in Section 8 to construct a CAT(0) cube complex on which the group G will act. This is essentially Sageev's construction, but appears in the current context in [27], an account we follow closely here. Superficial differences arise simply because, when $n \geq 3$, our construction produces two hypergraphs dual to each edge.

Before beginning the construction we will quickly review the notion of a CAT(0) cube complex.

Definition 9.1. (CAT(0) cube complex) A *cube complex* is a metric polyhedral complex where each cell is isometric to a Euclidean n -cube $[-1, 1]^n$, and such that the gluing maps are isometries.

It is *CAT(0)* if the following conditions are satisfied:

- (1) C is simply-connected
- (2) For each $n \geq 0$ if c_1, c_2, c_3 are three $(n+2)$ -cubes such that each pairwise intersection is an $(n+1)$ -cube and $c_1 \cap c_2 \cap c_3$ is an n -cube, then there exists an $(n+3)$ -cube that contains each c_i as a face.

Definition 9.2. (halfspace) For each hypergraph Λ the set $\bar{X} - \Lambda$ contains exactly two components. A *halfspace* is the closure of one such component. If A is a halfspace then we use A^C to denote the other halfspace associated to the same hypergraph.

We now define a graph Ω that will serve as the 1-skeleton of a cube complex \mathcal{C}_0 . The CAT(0) cube complex \mathcal{C} on which G will act will be a connected component of \mathcal{C}_0 .

Let S be the set of all halfspaces. A vertex of Γ is a subset $\mathcal{V} \subset S$ such that for any halfspaces A and B the following hold:

- (1) $A \subset B$ and $A \in \mathcal{V}$ implies $B \in \mathcal{V}$, and
- (2) $|\mathcal{V} \cap \{A, A^C\}| = 1$.

Two vertices $\mathcal{V}_1, \mathcal{V}_2 \subset S$ are joined by an edge in Ω if and only if $\mathcal{V}_1 = \{\mathcal{V}_2 - A\} \cup A^C$ for some halfspace A . In this case we say that \mathcal{V}_1 is obtained from \mathcal{V}_2 by *switching* the halfspace A .

The construction of \mathcal{C}_0 is completed inductively by gluing an n -cube to Ω anywhere its $(n-1)$ -skeleton exists.

An n -cube in \mathcal{C}_0 arises from a collection, say Σ , of n pairwise intersecting hypergraphs. For any hypergraph $\Lambda \notin \Sigma$, there cannot be two hypergraphs in Σ which are contained in distinct halfspaces of Λ , since then those two hypergraphs would not intersect. Let \mathcal{V}_0 be the set of halfspaces which contains for each $\Lambda \notin \Sigma$ a halfspace A so that any hypergraph in Σ intersects A non-trivially. This ensures that condition (1) in Definition 9.1 will be satisfied for any of the 2^n possible collections of halfspaces which extend \mathcal{V}_0 and satisfy condition (2). These 2^n vertices will be the 0-cells of an n -cube.

For any 0-cell $x \in \bar{X}$ let $\mathcal{V}_x = \{A \in S \mid x \in A\}$. It is easily verified that such a set of halfspaces satisfies the conditions above and so correspond to a vertex in \mathcal{C}_0 .

Lemma 9.3. *For any $x, y \in \bar{X}^0$ the vertices \mathcal{V}_x and \mathcal{V}_y lie in the same connected component of \mathcal{C}_0 .*

Proof. Since \bar{X} is connected it is enough to prove the result in the case where x and y are connected by a 1-cell e in \bar{X} . Suppose $n \geq 3$. Let Λ_x and Λ_y be the hypergraphs dual to e . Clearly the sets \mathcal{V}_x and \mathcal{V}_y are identical except for the halfspaces corresponding to Λ_x and Λ_y . There is a length two path from \mathcal{V}_x to \mathcal{V}_y found by first switching the halfspace associated to Λ_x and then switching the halfspace associated to Λ_y .

If $n = 2$ then there is only one hypergraph between x and y , and so \mathcal{V}_x and \mathcal{V}_y are adjacent in \mathcal{C}_0 . □

Remark 9.4. When $n \geq 3$ switching halfspaces in the opposite order does not define a path in \mathcal{C}_0^1 since at the intermediate step the set of halfspaces would not satisfy condition (1) in the definition of the vertices of Γ .

Definition 9.5. Let \mathcal{C} be the connected component of \mathcal{C}_0 containing \mathcal{V}_x for each $x \in \bar{X}^0$.

Since the hypergraph system is G -equivariant the action of G on \bar{X} induces an action on the set of halfspaces. This in turn induces an action of G on the cube complex \mathcal{C}_0 . For instance, for a vertex $\mathcal{V} = \{A_\sigma\}_{\sigma \in \Sigma}$ we have $g\mathcal{V} = \{gA_\sigma\}_{\sigma \in \Sigma}$. But since $g\mathcal{V}_x = \mathcal{V}_{gx}$ the action on \mathcal{C}_0 actually stabilizes \mathcal{C} giving an action of G on \mathcal{C} .

Theorem 9.6. \mathcal{C} is a $CAT(0)$ cube complex.

We will not prove that \mathcal{C} is simply-connected. A proof of this fact that transfers exactly to our situation can be found in [27]. We will however prove that \mathcal{C} satisfies condition (2) in Definition 9.1.

Proof. Let c_1, c_2, c_3 be three $(n+2)$ -cubes such that $c_i \cap c_j$ is an $(n+1)$ -cube when $i \neq j$ and $c_1 \cap c_2 \cap c_3$ is an n -cube.

Each c_i is determined by a set of $n+2$ independent hypergraphs, say S_i . Then $|S_1 \cap S_2 \cap S_3| = n$ by assumption and $|S_i \cap S_j| = n+1$ when $i \neq j$. Recalling the formula for the number of elements in a Venn diagram we get that $|S_1 \cup S_2 \cup S_3| = 3(n+2) - 3(n+1) + n = n+3$.

Moreover, each pair of hypergraphs in $S_1 \cup S_2 \cup S_3$ cross. Let Λ_1 be a hypergraph and assume without loss of generality that it lies in S_1 . By definition Λ_1 crosses each of the other hypergraphs in S_1 . There is only one hypergraph Λ_2 not in S_1 , but it lies in $S_2 \cap S_3$. Since Λ_1 lies in either S_2 or S_3 this implies that Λ_1 and Λ_2 intersect non-trivially.

Thus the $n+3$ hypergraphs in $S_1 \cup S_2 \cup S_3$ correspond to an $(n+3)$ -cube of which each c_i is face. \square

10. HYPERCARRIERS ARE QUASICONVEX

In this section we apply Theorem 6.3 to prove that hypercarriers are quasiconvex subspaces of \bar{X} . This fact will be used in Theorem 12.1 to help conclude that the group action on \mathcal{C} is cocompact.




FIGURE 6. The disc diagram D on the left is glued to $H' \cup \gamma$ along $\partial_E D$. Since the resulting complex cannot contain any internal 2-cells every 2-cell of D must lie along the subpath of $\partial_E D$ that gets mapped to γ . In the figure this subpath of $\partial_E D$ is bold. The shaded 2-cells of D are the 2-cells that lie along H' in the resultin complex.

Theorem 10.1. *Hypercarriers are quasiconvex.*

Proof. Let H be the hypercarrier associated to a hypergraph Λ . Suppose that there exists a geodesic γ with endpoints $x, y \in H$ such that $\gamma \cap H = \{x, y\}$. Let H' be the minimal connected hypercarrier chain containing both x and y . We will prove that any point of γ lies within a $\frac{3}{2}|w^n|$ -neighbourhood of H' .

Let $D \rightarrow X$ be a minimal area reduced disc diagram whose boundary is a generator of $\pi_1(H' \cup \gamma) \cong \mathbb{Z}$. Let K be the compact 2-complex formed by gluing D to $H' \cup \gamma$ along $\partial_E D$.

Observe that K is reduced since D was chosen to be minimal area. This follows from the same argument as in the proof of Theorem 8.1. K is also simply-connected since D is glued to $H' \cup \gamma$ along a generator of $\pi_1(H \cup \gamma)$.

Let α be a 2-cell in D . Then α is not extreme in K . Indeed, any subpath of $\partial\alpha$ which is external in K lies along the geodesic and so has length less than $\frac{|w^n|}{2}$. The only 2-cells of H' that could possibly be extreme in K are the end 2-cells of H' . The other 2-cells are not even extreme in H' . Thus K contains at most two extreme 2-cells, and so every 2-cell of K is external by Theorem 6.3. In particular, for each 2-cell α of D we have that $\partial\alpha \cap \gamma \subset K$ contains an edge. See Figure 6.

Let A be the set of 2-cells of D whose closures intersect H' in K , and let $x \in \gamma$. If $x \in \bar{A}$ then clearly we have $d(x, H') \leq \frac{1}{2}|w^n|$. If $x \notin \bar{A}$ let δ be the maximal connected subpath of γ containing x such that $\text{int}(\delta) \cap \bar{A} = \emptyset$. Let $u_1, u_2 \in \bar{A}$ be the endpoints of δ , and let δ' be the path between u_1 and u_2 in $\partial\bar{A} - H'$.

We claim that at most two 2-cells of A lie along $\text{int}(\delta')$. If not, then there exists a 2-cell α in A , and a 0-cell v such that $v \in \partial\alpha \cap \delta' \subset \text{int}(\delta')$. Note that $\partial\alpha \cap \delta'$ is connected since




FIGURE 7. Let X be the standard 2-complex of the presentation $\langle a, b \mid ((abab^{-1})^2a)^2 \rangle$. The complex shown is a subcomplex of \bar{X} , and the dotted line is a hypergraph segment. The corresponding hypercarrier, which is shaded, is not convex since the path outside it is a geodesic. The unshaded 2-cell does not lie in the hypercarrier since the hypergraphs embed. Convexity also fails in $\langle a, b \mid ((abab^{-1})^5a)^3 \rangle$ using an analogous picture.

otherwise D contains an internal 2-cell. Also, since every 2-cell of D lies along γ there exists a 0-cell $w \in \bar{\alpha} \cap (\gamma - \delta)$. We may assume that w lies in the connected component of $\gamma - \delta$ containing u_1 . Let ρ_1 be the path between w and v in $\bar{\alpha} - H'$ and ρ_2 be the path between v and u_2 contained in δ' . Then $\rho_1 \cup \rho_2$ separates u_1 from H' in K , contradicting the fact that $u_1 \in \partial A$.

Hence $|\delta| \leq |\delta'| < 2|w^n|$ since δ is a geodesic. Thus $d(u_i, x) < |w^n|$ for $i = 1$ or $i = 2$, and since $u_1, u_2 \in \bar{A}$ this yields $d(x, H) < \frac{3}{2}|w^n|$ as required. \square

When $n \geq 4$ a stronger result is possible. In this case the hypercarriers are actually convex. However, when $n = 2$ or $n = 3$ convexity fails, as can be seen in Figure 7.

11. LINEAR SEPARATION

In this section we restrict to the case where $n \geq 4$. Each hypergraph splits \bar{X} into two infinite components. For $x, y \in \bar{X}^0$ let $\#(x, y)$ be the number of hypergraphs, Λ , for which x and y lie in distinct components of $\bar{X} - \Lambda$. Clearly $\#(x, y) \leq 2d(x, y)$, since any separating hypergraph must be dual to some edge in each geodesic between x and y .

In order to satisfy the hypothesis of Theorem 12.2 we need to show that $\#(x, y)$ grows proportionally to $d(x, y)$. This will allow us to conclude that the cube complex constructed from the hypergraphs is locally finite and that the action is proper.

We begin by introducing some terminology. Let γ be a geodesic with endpoints $x, y \in \bar{X}^0$, and let e be a 1-cell in γ . There are two hypergraphs, Λ_{e_x} and Λ_{e_y} , dual to e , labeled so




FIGURE 8. The path between x and y is a geodesic γ and the two hypergraphs are dual to the same edge e . The hypergraph on the left is denoted Λ_{e_x} and the hypergraph on the right is denoted Λ_{e_y} .

that $d(\Lambda_{e_x} \cap e, x) < d(\Lambda_{e_y} \cap e, x)$. Let C_x and C_y be the components of $\gamma - e$ which contain x and y respectively. See Figure 8. With this notation we have the following lemma, which says that if a hypergraph crosses a geodesic twice it is turning in the *expected* direction.

Lemma 11.1. *Let γ be a geodesic in \bar{X}^1 between two 0-cells x and y , and let e be a 1-cell in γ . Then $\gamma \cap \Lambda_{e_x} \subset C_x$.*

Proof. Suppose otherwise. Define $v_x = e \cap \Lambda_x$, and let $v_{x'}$ be the point in $\Lambda_x \cap C_y$ which is closest to v_x . Let H be the hypercarrier chain associated to the unique hypergraph segment between v_x and $v_{x'}$. Note that H contains at least two 2-cells since otherwise the geodesic contains a subword of w^n longer than w^{n-1} .

If γ is not contained in H then we proceed as in the proof of Theorem 10.1 by attaching a minimal area reduced disc diagram to the complex $H \cup \gamma$ to form a new complex L . If $\gamma \subset H$ then let $L = H$.

Let α be the 2-cell in H such that $v_x \in \partial\alpha$, and let δ be the subpath of $\partial\alpha$ which is internal in L . Since the only 2-cells of L that can be extreme are the end 2-cells of H α must be extreme in L by Theorem 4.11. Thus $|\delta| < |w|$.

Let $\Delta = \bar{\alpha} \cap \gamma$. Then $|\Delta| + |\delta| > (n-1)|w|$ since $\Delta \cup \delta$ contains the longer subpath of $\partial\alpha$ containing two 1-cells dual to the same hypergraph. Combining this with $|\delta| < |w|$ we have that

$$|\Delta| > (n-1)|w| - |\delta| > (n-2)|w| \geq \frac{n}{2}|w|$$

since $n \geq 4$, contradicting the fact that γ is a geodesic. □




FIGURE 9. The hypercarrier chains H' and K with H' lying above the geodesic γ and K below. Once a reduced disc diagram is glued to this complex the shaded 2-cell is the only 2-cell that could be extreme.

Lemma 11.1 is not true when $n = 3$, and cannot even be stated properly when $n = 2$ since there is only one hypergraph dual to each 1-cell.

Definition 11.2. (comes back) Let δ be any simple path in \bar{X}^1 . For a hypergraph Λ and a 2-cell α we say that Λ *comes back to δ through α* if there is a hypergraph segment between two distinct points in $\delta \cap \Lambda$ such that α is extreme in the associated hypercarrier chain.

Theorem 11.3. Let γ be a geodesic in \bar{X}^1 with endpoints x and y . For any 1-cell e in γ , there exists a hypergraph that intersects γ exactly once, and the point in the intersection is within $\frac{|w^n|}{2}$ edges of e .

Proof. If Λ_{e_y} intersects γ once then we are done, so suppose d is a second 1-cell in γ dual to Λ_{e_y} chosen so that the distance between e and d is minimal.

By Lemma 11.1 we know that d is contained in the component of $\gamma - e$ that contains y .

Let H be the hypercarrier chain for the unique hypergraph segment between e and d , and let α_1 be the 2-cell in H such that $e \subset \partial\alpha_1$.

Let e_1, e_2, \dots, e_m be the 1-cells in the path $\partial\alpha_1 \cap C_y$ (where C_y is defined as for Lemma 11.1). Choose $0 \leq s \leq m$ to be maximal so that $\Lambda_{e_{s_y}}$ comes back through α_1 , and let H' be the hypercarrier chain for this new hypergraph segment between points of γ . Note that $s < m$ since α_1 must be extreme in H' . Label the 2-cells of H' , $\alpha_1, \alpha_2, \dots, \alpha_m$, in the order that the hypergraph passes through them.

Now we claim that $\Lambda_{e_{s+1_y}}$ intersects γ exactly once. By Lemma 11.1 this hypergraph does not intersect γ in the component of $\gamma - e_{s+1}$ containing x . By assumption it does

not come back to γ through α_1 . Thus if it intersects γ more than once it comes back to γ through a 2-cell, β_1 , giving rise to a second hypercarrier chain K . We label the 2-cells of K in order $\beta_1, \beta_2, \dots, \beta_\ell$. See Figure 9.

Let δ be the subpath of γ between the two endpoints of the hypergraph segments that are not contained in $\bar{\alpha}_1$ and $\bar{\beta}_1$. This is the bold path in Figure 9. If $H' \cup K \cup \delta$ is simply-connected let $L = H' \cup K$. Otherwise, let D be a minimal area reduced disc diagram for a generator of $\pi_1(H' \cup K \cup \delta)$. Gluing D to $H' \cup K \cup \delta$ we obtain a simply-connected compact 2-complex, which we call L .

No 2-cell of D is extreme in L since any 2-cell of D that is external in L lies along δ in L so its external subpath has length at most $\frac{|w^n|}{2}$. No 2-cell α_j , for $1 < j < m$, or β_i , for $1 \leq i < \ell$ is extreme in $H' \cup K$ since for each such 2-cell there are two 1-cells in its boundary that are internal in $H' \cup K$ and dual to the same hypergraph. If $e_s \subset \partial\beta_1$ then the same is true for α_1 , but certainly $e_{s+1} \subset \partial\beta_1$. Thus, although the internal subpath of $\partial\alpha_1$ may not contain two 1-cells dual to the same hypergraph it still has length at least $|w|$, and so α_1 is also not extreme in L .

We finish the proof by noting that at most one of α_m and β_ℓ can be extreme in L . Indeed, assume without loss of generality that β_ℓ is closer to y than α_m (this is the situation shown in Figure 9). Then both dual 1-cells in α_m are internal in L and so α_m cannot be extreme in L . Therefore L contains at most one extreme 2-cell contradicting Theorem 4.11, and so $\Lambda_{e_{s+1}y}$ intersects γ exactly once.

Finally, note that $s + 1 \leq \frac{|w^n|}{2}$ since both e_0 and e_{s+1} lie on a geodesic contained in boundary of a single 2-cell. \square

Corollary 11.4. $\#(x, y) \geq \frac{d(x, y)}{|w^n|}$.

Corollary 11.5. For any $r \geq 1$, if $d(x, y) \geq r|w^n|$ apart can be separated by at least r hypergraphs. In particular, $\#(x, y) \geq \frac{d(x, y)}{|w^n|}$.

Proof. Let γ be a geodesic between x and y . Then γ contains at least r disjoint subpaths of length $|w^n|$. By Theorem 11.3 each such subpath contains at least one hypergraph which separates x and y . \square

12. CAT(0) CUBULATION

The first three theorems listed in this section connect what we have proved so far about the geometry of hypergraphs and the action of G on C defined in Section 9. Theorem 12.3 is due to Sageev [24], but all three theorems can be found as stated here in [11] where Hruska and Wise present a systematic approach to cubulating a class of groups.

Theorem 12.1. *Suppose G acts cocompactly on K . Then G acts cocompactly on the cube complex C if the following hold:*

- (1) K is δ -hyperbolic.
- (2) The hypergraphs are quasiconvex.
- (3) The hypergraph system is locally finite.

Theorem 12.2. *Let K be a 2-complex equipped with a collection of hypergraphs satisfying the following properties. Then the cube complex C associated to K is locally finite.*

- (1) K is locally finite.
- (2) The hypergraph system is uniformly locally finite.
- (3) There is a constant M so that for each $r \geq 1$, every pair of points at a distance at least rM apart are separated by at least r distinct hypergraphs.
- (4) There is a constant δ such that every hypergraph triangle is δ -thin.

Theorem 12.3. *Let K be locally finite 2-complex with a collection of hypergraphs such that the associated cube complex C is locally finite. Then if G acts properly on K , G acts properly on C .*

Using these theorems we are now in position to prove our main theorem:

Theorem 12.4. *Let $\langle a_1, \dots, a_m \mid w^n \rangle$ be a presentation of a group G , where w is freely and cyclically reduced and $n \geq 4$. Then G acts properly and cocompactly on a locally finite CAT(0) cube complex C .*

Proof. It is clear that \bar{X} is locally finite, and that the hypergraph system is uniformly locally finite. Condition (3) of Theorem 12.2 is satisfied by Theorem 11.5. By the B.B. Newman Spelling Theorem, the presentation $\langle a_1, \dots, a_m \mid w^n \rangle$ is a Dehn presentation for G and so

G is word-hyperbolic. By Theorem 10.1 the hypergraphs embed by quasi-isometries and so the hypergraph triangles in \bar{X} are δ -thin, with δ depending on the hyperbolicity constant for G and the quasi-isometry constants for the hypergraphs. Therefore, the cube complex C associated to the hypergraph system is locally finite.

Since G acts properly on \tilde{X} and hence on \bar{X} , Theorem 12.3 implies that the action of G on C is proper.

We have shown that all the conditions of Theorem 12.1 hold so the action of G on C is cocompact. \square

REFERENCES

- [1] A. Aleksandrov. A theorem on triangles in a metric space and some applications of it. *Tr. Mat. Inst. Steklov*, 38:5–23, 1951.
- [2] J. M. Alonso and et al. Notes on word hyperbolic groups. In É. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 3–63. World Sci. Publishing, River Edge, NJ, 1991. Edited by H. Short.
- [3] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.
- [4] S. D. Brodskii. Equations over groups, and groups with one defining relation. *Sibirsk. Mat. Zh.*, 25(2):84–103, 1984.
- [5] J. Fischer, A. Karrass, and D. Solitar. On one-relator groups having elements of finite order. *Proc. Amer. Math. Soc.*, 33:297–301, 1972.
- [6] V. N. Gerasimov. Semi-splittings of groups and actions on cubings. In *Algebra, geometry, analysis and mathematical physics (Russian) (Novosibirsk, 1996)*, pages 91–109, 190. Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk, 1997.
- [7] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [8] J. Howie. On pairs of 2-complexes and systems of equations over groups. *J. Reine Angew. Math.*, 324:165–174, 1981.
- [9] J. Howie. On locally indicable groups. *Math. Z.*, 180(4):445–461, 1982.
- [10] J. Howie. How to generalize one-relator group theory. In S. M. Gersten and J. R. Stallings, editors, *Combinatorial group theory and topology*, pages 53–78, Princeton, N.J., 1987. Princeton Univ. Press.
- [11] C. Hruska and D. T. Wise. Axioms for finiteness of cubulations. Preprint, 2004.
- [12] G. C. Hruska and D. T. Wise. Towers, ladders and the B. B. Newman spelling theorem. *J. Aust. Math. Soc.*, 71(1):53–69, 2001.
- [13] T. Hsu and D. T. Wise. Cubulating graphs of free groups with cyclic edge groups. Preprint, 2005.
- [14] I. Kapovich. The combination theorem and quasiconvexity. *Intern. J. Algebra Comput.*, 11(2).
- [15] R. C. Lyndon and P. E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977. *Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89*.
- [16] W. Magnus. Über diskontinuierliche Gruppen mit einer definierenden Relation (Der Freiheitssatz). *J. Reine Angew. Math.*, 163:141–165, 1930.
- [17] B. B. Newman. Some results on one-relator groups. *Bull. Amer. Math. Soc.*, 74:568–571, 1968.
- [18] G. A. Niblo and L. D. Reeves. Coxeter groups act on $CAT(0)$ cube complexes. *J. Group Theory*, 6(3):399–413, 2003.
- [19] G. A. Niblo and M. A. Roller. Groups acting on cubes and Kazhdan's property (T). *Proc. Amer. Math. Soc.*, 126(3):693–699, 1998.

- [20] Y. Ollivier and D. T. Wise. Cubulating groups at density $< \frac{1}{6}$. Preprint, 2004.
- [21] C. D. Papakyriakopoulos. On Dehn's lemma and the asphericity of knots. *Ann. of Math. (2)*, 66:1–26, 1957.
- [22] S. J. Pride. On the generation of one-relator groups. *Trans. Amer. Math. Soc.*, 210:331–384, 1975.
- [23] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [24] M. Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
- [25] P. Scott. Ends of pairs of groups. *J. Pure Appl. Algebra*, 11(1-3):179–198, 1977/78.
- [26] J.-P. Serre. *Trees*. Springer-Verlag, Berlin, 1980. Translated from the French by John Stillwell.
- [27] D. T. Wise. Cubulating small cancellation groups. *GAFA, Geom. Funct. Anal.*, 14(1):150–214, 2004.