CUBULATING SMALL CANCELLATION GROUPS

DANIEL T. WISE

ABSTRACT. We study the $B(6)$ and $B(4)$-$T(4)$ small cancellation groups. These classes include the usual $C'(\frac{1}{6})$ and $C'(\frac{1}{4})$-$T(4)$ metric small cancellation groups. We show that every finitely presented $B(4)$-$T(4)$ or word-hyperbolic $B(6)$ group acts properly discontinuously and cocompactly on a CAT(0) cube complex. We show that finitely generated infinite $B(6)$ and $B(4)$-$T(4)$ groups have codimension 1 subgroups and thus do not have property (T). We show that a finitely presented $B(6)$ group is word-hyperbolic if and only if it contains no $\mathbb{Z} \times \mathbb{Z}$ subgroup.

CONTENTS

1. INTRODUCTION

Let G be a finitely generated group and let H be a subgroup of G. The number of ends of (G, H) is the number of ends in the coset graph $H\backslash G$. This number does not depend upon the choice of generating set for G. We call H a *codimension* 1 subgroup if (G, H) has more than one end. A motivating example is the fundamental group of a closed orientable surface, which has two ends relative to any infinite cyclic subgroup, and similarly, a 3-manifold group has two ends with respect to the fundamental group of any incompressible two sided surface. A significant early result in the subject is Stalling's theorem that if G is finitely generated then the trivial subgroup of G is codimension 1 if and only if G splits over a finite subgroup [Sta68].

Date: October 22, 2003.
2000 *Mathematics Subject Classification.* 20F06, 20F67 .
Key words and phrases. CAT(0) cube complexes, small cancellation theory, property-(T).
Research supported by NSERC grant.

Scott observed that if G splits nontrivially over H then H is codimension 1, and we refer to [Sco78] for some further interesting results and examples. While Stallings' theorem gives a partial converse to this statement, the exact relationship between splittings over H and the number of ends of (G, H) has remained unclear. A major advance occurred in Sageev's thesis [Sag95] where he showed that a finitely generated group G acts nontrivially on a CAT(0) cube complex if and only if G contains a codimension 1 subgroup. This result successfully places the Bass-Serre theory in a more general context. Further work has refined the meaning of "nontrivially" so that it means "acting with no global fixed point". We refer the reader to [Ger97] and [NR98b].

One of the main classes of examples in geometric group theory are the small-cancellation groups, which are fundamental groups of 2-complexes with the property that attaching maps of 2-cells have small overlap in the 1-skeleton. In this paper we show that

Theorem 1.1. *Every infinite $C'(\frac{1}{6})$ or $C'(\frac{1}{4}) - T(4)$ group has a codimension* 1 *subgroup.*

In fact, we prove this theorem for a somewhat more general class of small-cancellation groups, and we actually obtain a significantly stronger conclusion when the groups are finitely presented. The primary objects of concern in this paper are simply-connected 2-complexes X, which satisfy either the $B(6)$ or $B(4)$-$T(4)$ conditions. Many such examples arise as the universal covers of standard 2-complexes of presentations. The $B(2p)$ condition on a presentation, states roughly that any subword of a relator which is more than half its length, cannot be expressed as the product of p or fewer pieces (see Definition 2.10). The $B(2p)$ condition is intermediate between the $C'(\frac{1}{2p})$ and $C(2p)$ small-cancellation conditions which are reviewed in Definition 2.4.

We utilize a variant of Sageev's construction to prove that certain finitely presented small-cancellation groups act properly discontinuously and cocompactly on CAT(0) cube complexes. Among the many previously known parallels between small-cancellation groups and groups acting on CAT(0) cube complexes, is that both are biautomatic [GS90, NR98a].

The following "geometrization theorem" for metric small-cancellation groups is a special case of our main results, Theorem 11.1 and Theorem 12.2, which we state at the end of the introduction.

Theorem 1.2. *Let G be a group with a finite $C'(\frac{1}{6})$ or $C'(\frac{1}{4})$-$T(4)$ presentation. Then G acts properly discontinuously and cocompactly on a CAT(0) cube complex C.*

It was shown in [NR98b] that if G has Kazhdan's property (T), then G cannot act on a CAT(0) cube complex without a global fixed point. We use this to obtain the following special case of Theorem 14.2.

Corollary 1.3. *Let G be a finitely generated group with a $C'(\frac{1}{6})$ or $C'(\frac{1}{4})$-$T(4)$ presentation, then G does not have property (T) unless it is finite cyclic.*

Corollary 1.3 lies in contrast with a result of Zuk who gave a condition involving the combinatorial Laplace operator which he used to show that many fundamental groups of $C(3)$-$T(6)$ complexes do have property (T) [Żuk96]. Recently, Gromov gave another approach showing that most groups satisfying a condition related to $C'(\frac{1}{6})$ do not have property (T) [Gro03]. Corollary 1.3 is also interesting in view of Pride's examples of finitely presented $C'(\frac{1}{6})$ groups that do not split [Pri83]. The case of Corollary 1.3 where G acts on a CAT(0) 2-complex built out of Euclidean squares, or out of Euclidean regular hexagons was proven in [BŚ97].

Before describing the sections of this paper, we develop an analogy with the 3-manifold situation. A geometrically attractive case covered by Sageev's construction is a collection of immersed incompressible surfaces $S_1, \ldots, S_n$ in a 3-manifold M. In this case the action of $\pi_1 M$ on $\widetilde{M}$ permutes the translates of these surfaces. Sageev's work motivates an interesting (but hopelessly difficult) conjecture for a finite volume hyperbolic 3-manifold M. Does there exist a set of immersed incompressible surfaces $S_1, \ldots, S_n$ which 'fill' the 3-manifold, resulting in a free and cocompact action of $\pi_1 M$ on a CAT(0) cube complex?

We have solved the analogous problem for our class of small-cancellation complexes. Our hypergraphs in X, are analogous to the translates of the universal covers $\tilde{S}_1, \ldots, \tilde{S}_n$ in $\tilde{M}$. In both the 3-manifold and the small-cancellation scenarios, the cube complex corresponds to the pattern of intersections of these codimension 1 objects. The point of view adopted in this paper promises to be very fruitful for other natural classes of groups arising in geometric group theory.

We now provide a description of the sections of the paper.

In Section 2 we give a brief review of the definitions and theorems we shall need from small-cancellation theory. We then define the $B(2p)$ small-cancellation condition which is somewhat more general than the $C'(\frac{1}{p})$ metric small-cancellation conditions.

In Section 3 we define the 'hypergraphs' and 'hypercarriers' in X that carry the multi-ended subgroups of the group G acting properly discontinuously on X. We then prove various convexity properties of the hypercarriers that will be useful throughout the paper. We note that the hypergraphs in X are examples of what Dunwoody calls a *track* [Dun85].

In Section 4, we show that the family of hypergraphs in X is sufficiently rich in the sense that it "fills" X. This will later enable us to prove that the action of G on the CAT(0) cube complex C has finite stabilizers.

In Section 5, we describe the construction of the cube complex C, and the action of G on C. Our construction is a simplification of Sageev's original construction, that is catered to the small-cancellation situation. The vertices of C are certain choices of orientations on the set of all hypergraphs in X, and the n-cubes in C correspond to collections of pairwise crossing hypergraphs in X.

In order to understand the cocompactness properties of the G action on C, we are led to classify the cubes in C, which in turn leads us to an examination of families of pairwise crossing hypergraphs in X. This analysis is begun in Section 6, where we prove that any collection of pairwise crossing hypergraphs in a $B(4)$-$T(4)$ complex mutually meet at the center of a 2-cell. We also prove that a set of three pairwise crossing hypergraphs in a $B(6)$-complex meets along a "tricomb" (see Figure 10). This roughly means that there is a flat triangle whose sides are paths in the three hypergraphs.

In Section 7 we analyze the tricombs and flat planes occurring in a $B(6)$ complex X. A number of the results obtained will play an important role in Section 12. We also prove the following theorem which is of independent interest:

Theorem 7.10. *Suppose G acts properly discontinuously and cocompactly on a $B(6)$ complex X. The following are equivalent:*

(1) *G is not word-hyperbolic.*
(2) *X contains a flat plane.*
(3) *G contains a subgroup isomorphic to $\mathbb{Z} \times \mathbb{Z}$.*

In Section 8 we prove that in a $B(6)$ complex, any collection of pairwise crossing hypergraphs meets along a 2-cell, a vertex, or a tricomb. This is probably the single most important technical result that will allow us to understand the cube complex in Section 12.

In Section 9 we show that every cube in C is a subcube of an edge-cube, a cell-cube, a link-cube, or a tricomb-cube. These are the obvious cubes corresponding to families of pairwise crossing hypergraphs discussed earlier.

In Section 10 we prove that if X is locally finite then C is locally finite.

Finally, the various preliminary results are combined to prove our main results in Section 11 and Section 12.

Theorem 11.1. *If G acts properly discontinuously and cocompactly on a $B(6)$-$C(7)$ or $B(4)$-$T(4)$ complex X, then G acts properly discontinuously and cocompactly on a cube complex C.*

Theorem 12.2. *If G acts properly discontinuously and cocompactly on a $B(6)$ complex X, then G acts properly discontinuously and cofinitely on the locally finite CAT(0) cube complex C. Furthermore, the conjugacy classes of maximal virtually $\mathbb{Z} \times \mathbb{Z}$ subgroups of G are represented by the cusp subgroups.*

The *cofinite* action and the *cusp subgroups* are explained at the beginning of Section 12. Roughly speaking, G acts cofinitely if the quotient $G\backslash C$ is quasi-isometric with the union of a finitely many copies of $[0, \infty)$ glued together along 0. The cusp subgroups are virtually $\mathbb{Z} \times \mathbb{Z}$ subgroups which stabilize copies of $\mathbb{E}^3$ in C. The quotients of these copies yield the infinite intervals above. The situation is in some ways similar to a finite volume cusped hyperbolic 3-manifold. The theorem is simplified when G is word-hyperbolic:

Corollary 12.3. *If X is $B(6)$ and G is word-hyperbolic then G acts properly discontinuously and cocompactly on the CAT(0) cube complex C.*

As the statement suggests, the $B(6)$ case is substantially more difficult than the $B(4)$-$T(4)$ case. This is perhaps surprising, since as mentioned above, many $C(3)$-$T(6)$ groups cannot act nontrivially on a CAT(0) cube complex. One interesting point is that as q gets large, a $C(p)$-$T(q)$ complex becomes closer to being nonpositively curved at the vertices. For instance, $C(4)$-$T(5)$ complexes are nonpositively curved and actually, automatically satisfy the $B(4)$-$T(4)$ condition since pieces have length 1. Similarly, $C(3)$-$T(6)$ complexes are already nonpositively curved. The reason why the $B(6)$ case is more difficult is because the cubes in the $B(6)$ case represent more global objects.

In solidarity with the widely held preference to avoid a thirteenth level in buildings and other edifices, we have chosen not to have a Section 13 in this paper.

In Section 14, we demonstrate the existence of codimension 1 subgroups and the failure of property (T) for groups acting on $B(4)$-$T(4)$ or $B(6)$ complexes, but we do this in a somewhat more general context than the properly-discontinuous and cocompact case that we focused on earlier.

Finally, the main theorems in this paper will often hold even when X fails to satisfy the $B(6)$ or $B(4)$-$T(4)$ conditions. This is especially true when X satisfies some combinatorial small-cancellation condition such as $C(6)$ or $C(4)$-$T(4)$. (Though we do provide examples where our exact construction does fail for $C(p)$-$T(4)$ complexes in Section 3.3.5.) In particular, this work necessarily raises the following:

Problem 1.8. Does every finitely presented $C(7)$ or $C(4)$-$T(4)$ group act properly discontinuously and cocompactly on a CAT(0) cube complex? Does every $C(6)$ group have a codimension 1 subgroup? Does every infinite $C(6)$ or $C(4)$-$T(4)$ group fail to satisfy Property (T)?

Finally, I wish to mention that after they learned of my results, but without knowing my method, N.Brady and J.McCammond have recovered the $B(4)$-$T(4)$ case of Theorem 11.1. Their method is more direct and promises other interesting applications.

2. Small-cancellation theory

In this section we give a brief review of the basic notions of small cancellation theory. The reader is referred to [MW02] for a rigorous development of these notions that is consistent with their use here.

We shall work in the category of combinatorial complexes:

Definition 2.1 (Combinatorial maps and complexes)**.** A map $Y \to X$ between CW complexes is *combinatorial* if its restriction to each open cell of Y is a homeomorphism onto an open cell of X. A CW complex X is *combinatorial* provided that the attaching map of each open cell of X is combinatorial for a suitable subdivision.

Definition 2.2 (Disc diagram)**.** A *disc diagram* D is a compact contractible 2-complex with a fixed embedding in the plane. A *boundary cycle* P of D is a closed path in ∂D which travels entirely around D (in a manner respecting the planar embedding of D).

Let $P \to X$ be a closed null-homotopic path. A *disc diagram in X for P* is a disc diagram D together with a map $D \to X$ such that the closed path $P \to X$ factors as $P \to D \to X$ where $P \to D$ is the boundary cycle of D. The van Kampen Lemma [vK33] essentially states that every null-homotopic path $P \to X$ is the boundary cycle of a disc diagram. We define $\mathsf{Area}(D)$ to be the number of 2-cells in D. For a null-homotopic path $P \to X$, we define $\mathsf{Area}(P)$ to equal the minimal number of 2-cells in a disc diagram $D \to X$ that has boundary cycle P. The disc diagram $D \to X$ will then be referred to as a *minimal area disc diagram* for P.

Let R_1 and R_2 be 2-cells that meet along a 1-cell e in the disc diagram $D \to X$. We say R_1 and R_2 form a *cancellable pair* provided that the boundary paths of R_1 and R_2 starting at e map to the same closed path in X. A disc diagram is *reduced* if it has no cancellable pair of 2-cells. It is a fact that every minimal area disc diagram is reduced.

A disc diagram is *nonsingular* if it is homeomorphic to a closed 2-dimensional ball. Note that it is *singular* if and only if it has a cutpoint, or it is *trivial* which means that it consists of a single 0-cell.

Definition 2.3 (Piece)**.** Let X be a combinatorial 2-complex. Intuitively, a piece of X is a path which is contained in the boundaries of the 2-cells of X in at least two distinct ways. More precisely, a nontrivial path $P \to X$ is a *piece* of X if there are 2-cells R_1 and R_2 such that $P \to X$ factors as $P \to R_1 \to X$ and as $P \to R_2 \to X$ but there does not exist a homeomorphism $\partial R_1 \to \partial R_2$ such that there is a commutative diagram

$$\begin{array}{ccc} P & \to & \partial R_2 \\ \downarrow & \nearrow & \downarrow \\ \partial R_1 & \to & X \end{array}$$

Excluding commutative diagrams of this form ensures that P occurs in ∂R_1 and ∂R_2 in essentially distinct ways.

Definition 2.4 ($C(p)$-$T(q)$-complex)**.** An *arc* in a diagram is a path whose internal vertices have valence 2 and whose initial and terminal vertices have valence ≥ 3. The arc is *internal* if its interior lies in the interior of D, and it is a *boundary arc* if it lies entirely in ∂D.

A 2-complex X satisfies the $T(q)$ condition if for every reduced disc diagram $D \to X$, each internal 0-cell of D has either valence 2 or valence $\geq q$. Similarly, X satisfies the $C(p)$ condition if the boundary path of each 2-cell in D either contains a nontrivial boundary arc, or is the concatenation of at least p nontrivial internal arcs. Finally, for a fixed positive real number α, the complex X satisfies $C'(\alpha)$ provided that for each 2-cell $R \to X$, and each piece $P \to X$ which factors as $P \to R \to X$, we have $|P| < \alpha|\partial R|$. Note that if X satisfies $C'(\alpha)$ and $n \leq \frac{1}{\alpha} + 1$ then X satisfies $C(n)$. Note that if $D \to X$ is reduced then each nontrivial arc in the interior of D is a piece in the sense of Definition 2.3.

The main idea of small-cancellation theory is that when p and q are sufficiently large, nontrivial reduced $C(p)$-$T(q)$ disc diagrams always contain either spurs or i-shells.

Definition 2.5 (i-shells and spurs)**.** Let D be a diagram. An *i-shell* of D is a 2-cell $R \hookrightarrow D$ whose boundary cycle ∂R is the concatenation $P_0 P_1 \cdots P_i$ where $P_0 \to D$ is a boundary arc, the interior of $P_1 \cdots P_i$ maps to the interior of D, and $P_j \to D$ is a nontrivial interior arc of D for all $j > 0$. The path P_0 is the *outer path* of the i-shell.

A 1-cell e in ∂D which is incident with a valence 1 0-cell v is a *spur*. In analogy with the outer path of an i-shell, we will refer to the length 2 path (either ee^{-1} or $e^{-1}e$) that passes through v as the *outer path* of the spur.

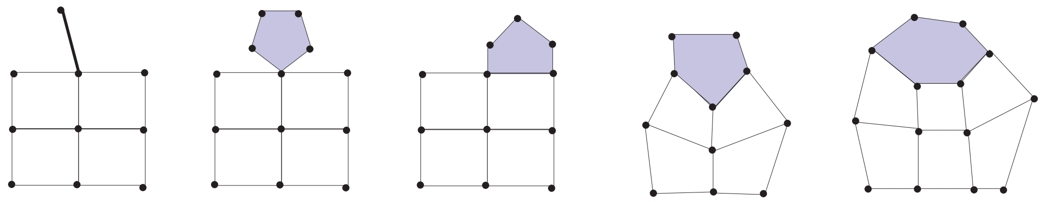

FIGURE 1. Spurs and i-shells

Illustrated from left to right in Figure 1 are disc diagrams containing a spur, a 0-shell, a 1-shell, a 2-shell, and a 3-shell. In each case, the 2-cell R is shaded, and the boundary arc P_0 is $\partial R \cap \partial D$.

Before stating the small-cancellation theorem we shall use, we define ladders which will play an important role in this paper.

Definition 2.6 (Ladder)**.** A *ladder* L is a disc diagram which is the union of a sequence of closed 1-cells and 2-cells $c_1, \ldots, c_n$, such that for $1 < j < n$, there are exactly two components in $L - c_j$, and exactly one component in $L - c_1$ and $L - c_n$. Finally, any c_i which is a 1-cell is not contained in any other closed c_j. Note that L is nonsingular if each c_j is a 2-cell, and each $c_j \cap c_{j+1}$ is a nontrivial arc called a *rung*.

The classical result which forms the basis of small cancellation theory is called Greendlinger's Lemma (see [LS77, Thm V.4.5]). The following strengthening of Greendlinger's Lemma was proven in [MW02, Thm 9.4].

Theorem 2.7. *If D is a $C(4)$-$T(4)$ [$C(6)$-$T(3)$] disc diagram, then one of the following holds:*

1. *D contains at least three spurs and/or i-shells with $i \leq 2$ [$i \leq 3$].*
2. *D is a ladder and hence has a spur, 0-shell or 1-shell at each end.*
3. *D consists of a single 0-cell or a single 2-cell.*

Moreover, if D is nontrivial and v is a 0-cell in ∂D, then D contains a spur or an i-shell with $i \leq 2$ [$i \leq 3$] whose outer path does not contain v in its interior.

We shall also employ the following immediate corollary.

Corollary 2.8. *Let D be a $C(4)$-$T(4)$ [$C(6)$-$T(3)$] disc diagram and let P and Q be immersed paths such that PQ^{-1} is the boundary cycle of D. If neither path contains the outer path of an i-shell in D with $i \leq 2$ [$i \leq 3$], then D is a ladder. Furthermore, P and Q are "opposite rails" of this ladder in the sense that following the notation of Definition 2.6, each closed c_j cell in D intersects both P and Q.*

Finally, we shall use the following fact which follows readily from Theorem 2.7 (see [LS77], or [MW02, Lemma 13.2]).

Corollary 2.9. *Every 2-cell in a simply-connected $C(4)$-$T(4)$ or $C(6)$ complex embeds.*

2.1. **The $B(2p)$ condition.** We now introduce the $B(2p)$ condition, which is the most natural small-cancellation condition for the results presented in this paper.

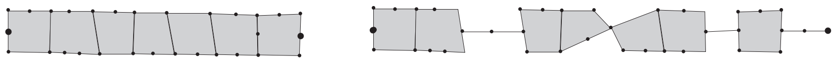

FIGURE 2. The ladder on the left is nonsingular, and the one on the right is singular.

Definition 2.10. The 2-complex X satisfies the $B(2p)$ condition if for each 2-cell R, and for each path $S \to \partial R$ which is the concatenation of at most p pieces in X, we have $|S| \leq \frac{1}{2}|\partial R|$.

For example, if for each piece V in the boundary of a 2-cell R, we have $|V| \leq \frac{1}{2p}|\partial R|$ then X satisfies the $B(2p)$ condition. In particular, $C'(\frac{1}{2p})$ complexes strictly satisfy this inequality so they are $B(2p)$.

Proposition 2.11. *$B(2n) \Rightarrow C(2n)$.*

Proof. Let R be a 2-cell, and let ∂R be expressed as the concatenation of pieces $\partial R = P_1 P_2 \cdots P_k$. Choose j so that $|P_1 \cdots P_j| < \frac{1}{2}|\partial R|$, but $|P_1 \cdots P_j P_{j+1}| \geq \frac{1}{2}|\partial R|$. Since $|P_{j+1} \cdots P_k| > \frac{1}{2}|\partial R|$, the $B(2n)$ condition implies that $k - j \geq n + 1$. Similarly, since $|P_1 \cdots P_j P_{j+1} P_{j+2}| > \frac{1}{2}|\partial R|$, we see that $j + 2 \geq n + 1$. Adding these two inequalities, we find that $k + 2 \geq 2n + 2$, so $k \geq 2n$ as claimed. □

3. Hypergraphs and the convexity of hypercarriers

3.1. The hypergraphs and hypercarriers.

We begin by adopting the following conventions which we shall follow for the remainder of the paper:

Convention 3.1. Let X denote a 2-complex. We suppose that the attaching map of every 2-cell of X has even length. Note that this can always be arranged by possibly subdividing all 1-cells of X. We shall assume that no two 2-cells of X have the same attaching map. This can be arranged by identifying 2-cells with the same attaching map. Later we shall focus on the situation where a group G acts on X, and we note that this action induces an action on the quotient space. Furthermore, if the original action has finite cell stabilizers, then so does the action on the quotient. We will primarily be interested in the case where X is simply-connected. However, the definitions and constructions below are meaningful and carry useful information in general.

Finally, in a few places we have assumed that there are no *free faces* which are 1-cells occurring on the boundary of a unique 2-cell. We do however allow *isolated edges* which are 1-cells that do not occur on any 2-cell. Because there are no redundant 2-cells, the convention that there are no free faces implies that each 1-cell is a piece. (Many of the small-cancellation theorems in this paper continue to hold if instead of assuming that there are no free faces, we regard a single 1-cell as a piece in the $B(2p)$ definition.)

Definition 3.2 (Hypergraphs)**.** Let Λ_X be the graph defined as follows: The 0-cells of Λ_X are in correspondence with 1-cells of X. The 1-cells of Λ_X are in correspondence with pairs of antipodal edges in the attaching maps of 2-cells of X. Each component of Λ_X is a *hypergraph* in X.

See Figure 3 for a 2-complex with eight distinct hypergraphs.

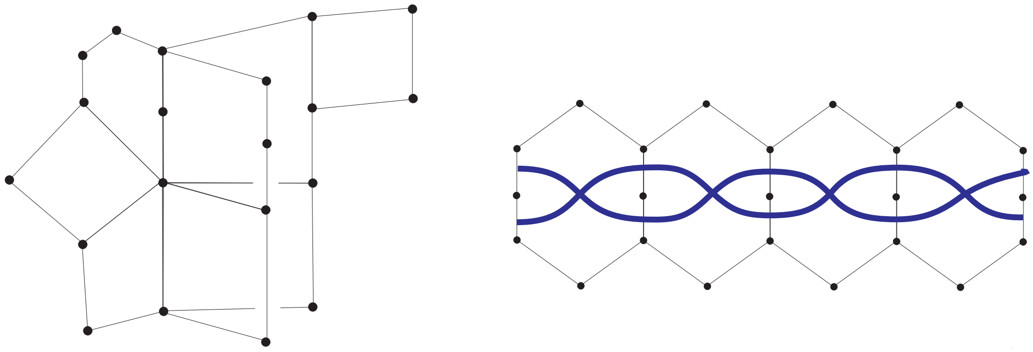

FIGURE 3. There are eight 2-cells and eight unillustrated hypergraphs in the left 2-complex. The right 2-complex consisting of four octagonal 2-cells, has a pair of illustrated hypergraphs that give rise to the same hypercarrier. Each 2-cell contains two additional unillustrated hypergraphs.

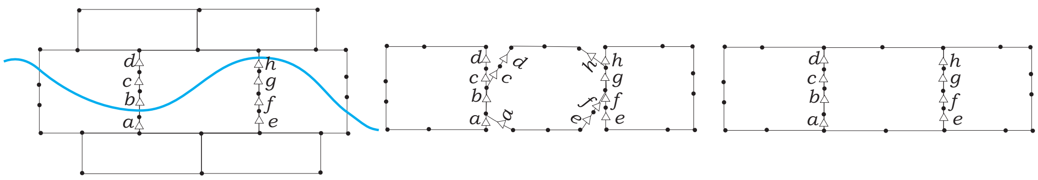

FIGURE 4. The hypercarrier Y is obtained by folding edges of H.

Definition 3.3 (Hypercarriers)**.** To each hypergraph Λ, we associate a 2-complex Y called a *hypercarrier* that we shall now define. The term "hypercarrier" is short for "hypergraph carrier".

We first form a 2-complex denoted by H which is defined as follows: First take the disjoint union of closed 1-cells of X corresponding to the vertices of Λ. These 1-cells will be called *dual* 1-cells. Next, for each edge e in Λ that is attached to vertices $\iota(e)$ and $\tau(e)$, we take a copy of the corresponding closed 2-cell R of X, and attach it along its antipodal 1-cells to the corresponding dual 1-cells. (Note that a priori, it is possible for a single 2-cell to be attached along both its antipodal 1-cells to a single dual 1-cell.)

Observe that there is a combinatorial map $H \to X$, and we define Y to be the coarsest quotient of H such that $Y^1 \to X$ is an immersion, and $H \to X$ factors as $H \to Y \to X$. In practice, Y can be constructed from H, by successively folding edges in H^1 that map to the same edge of X [Sta83]. For example, in Figure 4, the planar complex X on the left has seven 2-cells, the hypergraph Λ is indicated by a bold path, the complex H is in the center, and after finitely many edges are folded together we obtain the hypercarrier Y on the right.

Remark 3.4. We note that there is an obvious embedding $\Lambda \hookrightarrow H$ so that each vertex of Λ maps to the center of its corresponding dual 1-cell, and each edge of Λ passes through the center of its corresponding 2-cell. We shall soon see that under favorable conditions, Λ projects to an embedding in its associated hypercarrier Y.

While we have associated a hypercarrier Y to each hypergraph Λ, we shall often suppress mentioning the family of dual 1-cells in Y from which the hypergraph Λ can be reconstructed. It is possible for two distinct hypergraphs Λ_1 and Λ_2 to yield hypercarriers that are identical as 2-complexes mapping to X, and are only distinguished by having different families of dual 1-cells. This occurs in a scenario where "braided" hypergraphs $\Lambda_1 \hookrightarrow Y$ and $\Lambda_2 \hookrightarrow Y$ repeatedly cross each other in successive 2-cells of Y. See the two hypergraphs on the right in Figure 3.

We shall soon give conditions under which the hypercarriers and hypergraphs embed in X. The reader who prefers a one-to-one correspondence between hypercarriers and hypergraphs can then achieve this by removing the sequence of vertices in X that lie between the two hypergraphs with the same hypercarrier. This can be done throughout X and effectively identifies the braided hypergraphs.

Definition 3.5. The *interior*, $\text{Int}(Y)$, of a hypercarrier Y, is the union of its open 2-cells, and its open dual 1-cells. The *boundary*, ∂Y, is $Y - \text{Int}(Y)$. Note that distinct hypergraphs Λ_1 and Λ_2 might have the same hypercarrier $Y_1 = Y_2$ but $\text{Int}(Y_1) \neq \text{Int}(Y_2)$ and $\partial Y_1 \neq \partial Y_2$.

3.2. **The hypercarriers embed in X.** The following generalizes an idea first explored in [Wis01] and [Wis].

Lemma 3.6. *Let Λ be a hypergraph in the $B(4)$ complex X and let Y be its hypercarrier. Then*

(1) *The boundary of each 2-cell of Y contains a pair of disjoint nontrivial arcs called* isolated rails*, and the interior of each isolated rail is disjoint from any other closed cell of Y.*
(2) *For each dual edge there is a tree called a* rung tree *such that distinct dual edges have disjoint rung trees.*
(3) *Y^1 is the union of rung trees and isolated rails such that rung trees corresponding to adjacent vertices of Λ are attached together by a pair of isolated rails corresponding to the edge of Λ connecting these vertices.*
(4) *the map $Y \to X$ is an immersion.*

Proof. Let R be a 2-cell in the complex H associated with Λ. We first show that ∂R is the concatenation of paths $e_1C_1A_1B_1e_2C_2A_2B_2$ with the following property: The edges e_1 and e_2 are dual 1-cells of H, the B_i and C_i paths are possibly trivial, and A_1 and A_2 are nontrivial arcs called *isolated rails* of R such that in the quotient map $H \to Y$, no other edge of H is identified with edges in either A_1 or A_2.

We prove the statement for A_1 and note that the proof for A_2 is similar. Let $\partial R = e_1D_1e_2D_2$. Let C_1 be the longest initial subpath of D_1 which can fold with another path in H with the same initial 0-cell as D_1, and similarly let B_1 be the longest terminal folding subpath. Then D_i is certainly not equal to either C_1 or B_1, for this would violate $B(2)$. If C_1 and B_1 intersect in the interior of A_1, then $A_1 = C_1'B_1'$ for some subpaths C_1' and B_1' of C_1 and B_1, and this would contradict $B(4)$. Thus $A_1 = C_1A_1B_1$ where A_1 is a nontrivial subpath of D_1, and no edge of A_1 folds with an edge on the boundary of another 2-cell of H incident with the dual edge e_1.

Thus, for each dual edge e, there is a collection of arcs BeC that are subpaths of boundaries of 2-cells incident with e. These arcs can be identified along e, and

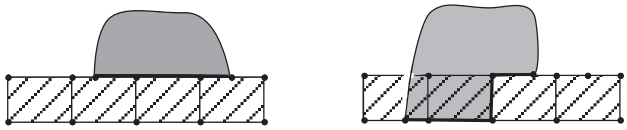

FIGURE 5. In each $Y \cup_P R$ above, P is indicated by the very bold path, Y is indicated by the striped ladder, and R is indicated by the shaded 2-cell.

folded together to form a *rung tree* of Y. Y^1 is then the union of a collection of disjoint rung trees corresponding to the dual edges of Λ, and a disjoint collection of pairs of isolated rails corresponding to the edges of Λ, such that consecutive rung trees are connected by pairs of isolated rails along their endpoints.

Finally, we show that the map $Y \to X$ is an immersion. Indeed, $Y^1 \to X$ is an immersion by construction, and since the isolated rails are nontrivial, distinct 2-cells of Y that meet along a 0-cell, are distinct 2-cells of Y meeting along a dual 1-cell, and therefore map to distinct 2-cells of X. □

Definition 3.7 (Missing)**.** Consider the commutative diagram on the left

$$\begin{array}{ccc} P & \to & Y \\ \downarrow & & \downarrow \\ R & \to & X \end{array} \qquad\qquad \begin{array}{ccc} P & \to & Y \\ \downarrow & \nearrow & \downarrow \\ R & \to & X \end{array}$$

where Y is a 2-complex, R is a closed 2-cell, and $P \to X$ is a path which factors through both Y and R. We say that R is *missing in* Y *along* P if the map $P \to Y$ does not extend to a map $R \to Y$ so that the diagram on the right commutes.

Lemma 3.8. *Let Y be a hypercarrier in X, and let R be a 2-cell which is missing in Y along a path $P \to Y$. If X is $B(6)$ [or $B(4)$-$T(4)$] then P can be expressed as the concatenation of two [one] or fewer pieces in $Y \cup_P R$.*

See Figure 5 for illustrations of $Y \cup_P R$ that contradict the $B(6)$ or $B(4)$-$T(4)$ hypotheses.

Proof. From Lemma 3.6, Y is the union of rung trees centered at dual 1-cells such that adjacent trees are disjoint but joined by pairs of isolated rails. If P traversed an entire isolated rail then the $B(6)$ condition would be violated for the 2-cell containing this isolated rail. Thus P is the concatenation of part of an isolated rail followed by a path in a rung tree followed by part of another isolated rail.

In the $B(4)$-$T(4)$ case we note that if P transitioned from the boundary of one 2-cell to the boundary of another, then the $T(4)$ condition would be violated since the 2-cells also meet along an edge (in the rung tree) at the transition. Thus P lies entirely along the boundary of a single 2-cell. □

Definition 3.9 (i-closed)**.** Let $Z \to X$ be an immersion. We say Z is *i-closed* if for any j-shell F with $j \leq i$, and with outer path $Q \to F \to X$, any lift of Q to Z extends to a lift of F to Z. That is, no j-shell is missing in Z along its outer path.

Using the terminology of Definition 3.9 we have the following consequence of Lemma 3.8:

Corollary 3.10. *If X satisfies $B(6)$ [or $B(4)$-$T(4)$] then each hypercarrier Y is 3-closed [2-closed].*

Lemma 3.11. *Let X be $B(6)$ [or $B(4)$-$T(4)$] then each hypercarrier Y embeds in X. Furthermore Y is simply-connected.*

Proof. Consider an essential immersed path $P \to Y$ which projects to a closed path in X, where *essential* means that either P is not nullhomotopic in Y, or P has distinct endpoints in Y. Assume $P \to Y$ has been chosen so that $\mathsf{Area}(P \to X)$ is minimal among all such choices, and let $D \to X$ be a minimal area disc diagram for $P \to X$.

Since $P \to Y$ is essential, it cannot be a trivial path. By Theorem 2.7, D contains an i-shell F with $i \leq 3$ [or $i \leq 2$] whose outer path Q is a subpath of P. But then F lifts to Y by Corollary 3.10, so we can push P across F to obtain a new path P' with the same endpoints as P, but with a disc diagram D' obtained from D by removing Q and F, and so D' has smaller area. This contradicts our choice of P. □

Corollary 3.12. *If X is $C(6)$ or $B(4)$-$T(4)$ then each hypergraph is a tree and embeds in X.*

Proof. This follows by combining Lemma 3.11 and Lemma 3.6. □

Lemma 3.13. $\mathrm{Int}(Y)$ *separates X into exactly two components. Similarly, each hypergraph Λ separates X into exactly two components.*

Proof. Indeed, $\mathrm{Int}(Y)$ separates X, because otherwise an application of van Kampen's theorem would contradict that X is simply-connected. The fact that there are exactly two components follows by observing that ∂Y has exactly two components. □

Remark 3.14. Let Λ be a hypergraph in X, and let H be the subgroup of G which stabilizes both components of $X - \Lambda$. Even when G acts properly discontinuously on X, it is possible for H to fail to be a codimension 1 subgroup. The problem is that H might act cocompactly on a component of $X - \Lambda$.

We note that this subsection has shown that each hypergraph in X is an example of what Dunwoody calls a *track* [Dun85].

3.3. **Convexity of hypercarriers and halfcarriers.**

Definition 3.15. A *geodesic* P in a complex Z is a combinatorial path in Z such that $|P| \leq |P'|$ for any path P' with the same endpoints as P. It is easy to see that if Z is connected, then any two points can be joined by at least one geodesic.

A subcomplex C of Z is *convex* if for every geodesic path P in Z, if the endpoints of P lie in C then P itself lies in C.

The following lemmas follow easily from the definitions:

Lemma 3.16. *The intersection of convex subcomplexes is convex.*

Lemma 3.17. *A convex subcomplex of a connected complex is connected.*

The following technical result will play a key role for the remainder of the paper.

Theorem 3.18. *If X is $B(6)$ [or $B(4)$-$T(4)$] then each hypercarrier Y in X is convex.*

Proof. Suppose there is a geodesic P whose endpoints lie in Y but $P \not\subset Y$. By Lemma 3.13, $\mathrm{Int}(Y)$ separates Y, so without loss of generality we may assume that the endpoints of P lie on the same component of ∂Y. Pick an immersed path S in Y with the same endpoints as P and let D be a minimal area disc diagram for PS^{-1}. Furthermore, suppose P, S and D are chosen such that the pair $\big(|P|, \mathsf{Area}(D)\big)$ is minimal in the lexicographical order.

By Corollary 2.8, either D contains an i-shell with $i \leq 3$ $[i \leq 2]$ whose outer path is a subpath of P, or D contains an i-shell R whose outer path Q is a subpath of S, or D is a ladder. Note that P' cannot be contained in Y, for then $P = Q$ but $P' = S$ is the concatenation of at most two [one] pieces by Corollary 3.10, so $|P'| < |P|$ and thus P is not a geodesic.

If there is an i-shell R in D whose outer path Q is a subpath of P, then we can push P across R to obtain a new path P'. Let D' be the diagram obtained from D by removing the interiors of R and Q. To see that $|P'| \leq |P|$ and hence P' is geodesic, observe that $|Q| \geq \frac{1}{2}|\partial R|$ because the remainder of ∂R is the concatenation of at most i pieces, whose total length is $\leq \frac{1}{2}|\partial R|$ by hypothesis. Since $\mathsf{Area}(D') < \mathsf{Area}(D)$, the new choice P', S, D' contradicts the minimality of P, S, D.

If there is an i-shell R in D whose outer path is a subpath of S, then by Corollary 3.10, we see that R must lift to Y. We now push S across R to obtain a new path S' in Y, and we again define D' to be the diagram obtained from D by removing the interiors of Q and R. Since $\mathsf{Area}(D') < \mathsf{Area}(D)$ the new choice P, S', D' contradicts the minimality of P, S, D.

Finally we consider the case where no i-shell in D has outer path in P or S, and so D is a ladder. Since P is not contained in Y, the path P passes through some point v in $X - Y$. Observe that the first and last cells of D are 2-cells and not 1-cells. For removing such a 1-cell results in a shorter geodesic counterexample and thus contradicts the minimality of P, S, D. In fact, D is nonsingular in the sense that it contains no cutpoint. For otherwise we let D_1 be a maximal nonsingular subladder of D whose boundary contains a subpath P_1 of P that passes through v, and let $S_1 = D_1 \cap S$, and then $|P_1| < |P|$ so P_1, S_1, D_1 contradicts the minimality of P, S, D.

Assume that D has at least two 2-cells. By choosing R to be either the first or last 2-cell of D, we can assume that v does not lie in the interior of $R \cap P$. By Lemma 3.8, the path $S_1 = R \cap S$ is the concatenation of two [one] or fewer pieces. But then we can produce a lower area counterexample by pushing P across R to obtain P', S, D' as illustrated in Figure 6. Specifically, if ∂R is the concatenation $S_1S_2S_3$, and P is the concatenation $S_3^{-1}Q$, then $P' = S_1S_2Q$, and D' is obtained from D by removing the open 2-cell R, and the open path S_2. Note that $S_2 \geq \frac{1}{2}|\partial R|$ by the $B(6)$ $[B(4)]$-$[T(4)]$ condition, since S_1 is the concatenation of at most two [one] pieces, and S_3 consists of a single piece. Therefore $|P'| \leq |P|$.

Finally, if D has a single 2-cell, then Lemma 3.8 states that S is the concatenation of at most two [one] pieces. Thus $|S| < |P|$ which contradicts that P is a geodesic. $\square$

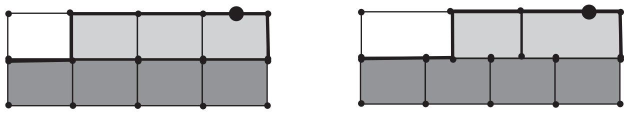

FIGURE 6. Illustrated above is $D \cup_S Y$ in the $B(4)$-$T(4)$ and $B(6)$ cases. The bold path is P, the bold point is v, the top ladder is D, and the bottom is Y. The shaded part of D is D', and the white 2-cell in D is R.

Definition 3.19. A *halfcarrier* is the union of Y and one of the two components of $X - \Lambda$, where Y is a hypercarrier associated to a hypergraph Λ. The term "halfcarrier" is short for "halfspace carrier". Halfspaces will be defined in Definition 5.2.

Corollary 3.20. *Every halfcarrier in X is convex.*

Proof. More generally, if C is convex in Z, and D is the union of C and some components of $Z - C$, then D is convex. Indeed, let P be the shortest geodesic counterexample. Then the endpoints of P must lie in C, for otherwise we could remove the first or last 1-cell and obtain a shorter counterexample. But then P itself must lie in C since C is convex. □

Corollary 3.21. *Let X be $B(4)$-$T(4)$ [or $B(6)$]. Then any nonempty intersection $\cap H_e$ of a set of halfcarriers of X is simply-connected.*

Proof. We first show that each halfcarrier is 2-closed [3-closed]. It will follow immediately that the intersection of halfcarriers is 2-closed [3-closed].

Let Q be the outer path of an i-shell R with $i \leq 2$ [or $i \leq 3$]. Suppose that Q is a path in a halfcarrier H associated with a hypercarrier Y, we will show that $R \subset H$.

If Q passes through a dual edge of Λ, then $R \subset Y$ because of the way we defined Y, so we may assume that Q is disjoint from Λ. Thus Q is a path in a component of $H - \Lambda$, and this component is either a component of $Y - \Lambda$, or a component of $X - \Lambda$ which is contained in H. In the latter case, R is obviously contained in H. In the former case, Q is a path in Y, and hence R is contained in Y because Y is i-closed by Lemma 3.10.

To prove simple connectivity we argue by contradiction as follows: Let P be an essential closed path in $\cap H_e$, and let $D \to X$ be a reduced disc diagram for P, and suppose P and D have been chosen so that $\big(\mathsf{Area}(D), |P|\big)$ is minimal in the lexicographical order among all possible counterexamples. Since D cannot be trivial, Theorem 2.7 states that either D contains a spur or D contains an i-shell.

If D contains a spur, then by removing this spur, we obtain a new disc diagram D' with a new boundary path P' such that $|P'| < |P|$, and P' is homotopic with P, and this contradicts the minimality of $\big(\mathsf{Area}(D), |P|\big)$.

Now suppose that D contains an i-shell $R \to X$ whose outer path Q is a subpath of P and ∂R is the concatenation QS. If $Q = P$, then S is the trivial path and $D = R$, and so since $R \subset \cap H_e$, we see that P is nullhomotopic in $\cap H_e$, and this contradicts that P is essential.

Otherwise, there is a homotopy in $\cap H_e$ obtained by pushing P across R to obtain a path P' which is the same as P except that Q is replaced by S. The path

$P' \to \cap H_e$ has a disc diagram $D' \to X$ that is obtained from D by removing the interior of R and the interior of Q. Since $\mathsf{Area}(D') < \mathsf{Area}(D)$ the minimality of $\big(\mathsf{Area}(D), |P|\big)$ is contradicted. $\square$

3.4. **Convex hulls of compact sets are compact.**

Definition 3.22 (Hyperclosure)**.** For a subset $W \subset X$, we define the *hyperclosure* of W to be $\text{Hyperclosure}(W) = \bigcap H_e$, where the intersection varies over the halfcarriers H_e corresponding to oriented hypergraphs Λ_e such that W is contained in the positive component of $X - \Lambda_e$. (Oriented hypergraphs are defined below in Definition 5.2.)

Theorem 3.23. *Let X be locally finite, and suppose there is an upper bound on the number of sides of a 2-cell in X. If W is compact then Hyperclosure(W) is compact.*

Proof. Since this is where all applications of interest lie, we will give the proof under the stronger hypothesis that there is a global bound on the number of sides of any 2-cell in X. At the end of the proof we will indicate how to relax this assumption to general local finiteness.

The idea of the proof is that if one sets out from some point (or indeed, some compact set) in any direction, then after a bounded distance (determined by the types of cells and the local finiteness) one is forced to cross a hypercarrier.

Choose m such that $2m > |\partial R|$ for any 2-cell R in X. Without loss of generality, we assume that W is connected. Let b denote the finite number of distinct hypercarriers in X that intersect W. Let $w \in W^0$. We will show that Hyperclosure(W) is contained in the ball B of radius $m(b+1)+1$ about w in X. To do this we will show that for each 0-cell $v \in X - B$, there is a hypercarrier Y such that W and v are contained in distinct components of $X - Y$.

Let $P = e_1 e_2 \dots e_{|P|}$ be a geodesic path from w to v which is expressed as the concatenation of edges. For an edge e_i we shall use the notation Y_i and H_i to denote the hypercarrier dual to e_i, and the associated halfcarrier containing w. Below we shall establish the claim that for $i < |P| - m$, there is an edge e_{n_i} with $i \leq n_i < i + m$ such that $e_{i+m} \not\subset Y_{n_i}$. It follows from Corollary 3.20 that $v \notin H_{n_i}$ for each i. Furthermore, observe that for $i + m < j$ we have $Y_{n_i} \neq Y_{n_j}$. Indeed otherwise, P would exit Y_{n_i}, and then later intersect $Y_{n_j} = Y_{n_i}$, and this would contradict the convexity of hypercarriers which was asserted in Theorem 3.18.

Since the Y_{n_i} are distinct for $i \in \{1, 1+m, 1+2m, \dots, 1+bm\}$, we see that at least one Y_{n_i} does not intersect W. Thus $W \subset H_{n_i}$ but $v \notin H_{n_i}$. Indeed, since W is connected and $W \cap Y_{n_i} = \varnothing$, we see that W lies in the component of $X - Y_{n_i}$ that contains $w \in W$. Thus $W \subset H_{n_i}$ as claimed.

We will now prove the claim made above. Pick an edge e_i of P and suppose that $e_{i+m} \subset Y_i$. Observe that e_{i+m} and e_i cannot be contained in the same 1-cell or 2-cell, for otherwise the definition of m shows that P fails to be a geodesic. We can therefore choose the smallest j with $i < j \leq i+m$ and with the property that $e_i e_{i+1} \dots e_j$ is not contained in the boundary of a single 2-cell of Y_i. Let R_1 denote a 2-cell whose boundary contains $e_i e_{i+1} \dots e_{j-1}$, let $R_2 \subset Y_i$ denote a 2-cell whose boundary contains e_j, and let r be the rung along which R_1 and R_2 intersect See the left diagram in Figure 7 which illustrates the situation for $i = 1$ and $j = 7$.

FIGURE 7.

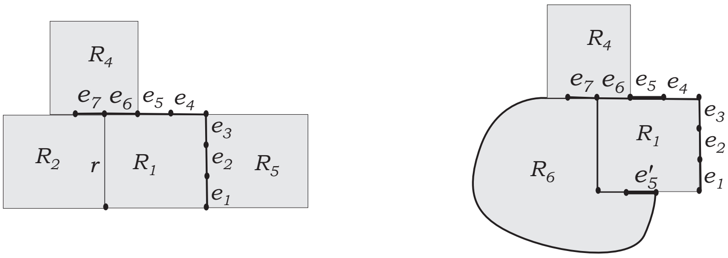

FIGURE 8.

In the $B(4)$-$T(4)$ case observe that $e_j \not\subset Y_{j-1}$. Indeed, if there were a 2-cell R_3 which contained both e_j and e_{j-1} then the triple R_1, R_2, R_3 would violate the $T(4)$ condition as illustrated on the right in Figure 7.

In the $B(6)$ case choose the least $k < j$ such that $e_{k+1} \ldots e_j$ is a subpath of ∂R_4 for some 2-cell R_4. Choose the greatest $\ell \geq 1$ such that $e_1 \ldots e_\ell$ is a subpath of some 2-cell $R_5 \neq R_1$. Observe that $k > \ell$, for otherwise $(e_1 \ldots e_k)(e_{k+1} \ldots e_{j-1})(r)$ would be the concatenation of three pieces in ∂R_1 whose length is greater than $\frac{1}{2}|\partial R_1|$ and this violates the $B(6)$ condition. See the left diagram in Figure 8 where $j = 7$, $k = 5$ and $\ell = 3$.

We claim that $e_j \not\subset Y_k$, and so since $i < k < j \leq i + m$, we can let $n_i = k$ to reach the desired conclusion because the convexity of Y_k and the fact that $e_j \not\subset Y_k$ implies that $e_i \not\subset Y_k$. If $e_j \subset Y_k$ then $e_j \subset \partial R_5$ for some 2-cell R_5 which contains a dual edge of Λ_k. But since $R_5 \cap R_1$ is nonempty, we see that R_5 must meet R_1 along a dual edge of Λ_k. Now, R_5 cannot meet R_1 along e_k since this would contradict the minimality of k. On the other hand, R_5 cannot meet R_1 along the antipodal edge e'_k of e_k in R_1, for then, since e_k is itself a piece by Convention 3.1, the $B(6)$ condition would be contradicted as illustrated on the right in Figure 8.

We now indicate how to prove Theorem 3.23 without placing a global bound on the number of sides of 2-cells. For each i, we let $m_i \in \mathbb{N}$ be such that $2m_i$ is greater than the length of the attaching map of any 2-cell lying in the ball of radius $2i$ in X centered at w. We then consider the hypercarriers Y_{n_i} where $i \in \{m_1, m_1 + m_2, m_1 + m_2 + m_3, \ldots\}$. Finally, we replace $m(b+1)+1$ by $\sum_{i=1}^{b+1} m_i + 1$. $\square$

3.5. **The construction can fail in $C(6)$ and $C(4)$-$T(4)$ cases.** The reader interested in understanding how the hypergraphs studied in Sections 3.3.1-3.3.3 can fail to lead to anything useful for general non-metric small-cancellation groups, is encouraged to consider the following presentation for a group G:

$$\langle\, a, b \mid a^m A, a^{m+1} A', b^m B, b^{m+1} B', a^k b^k C \,\rangle$$

By choosing A, A', B, B', and C to be long positive metric small-cancellation words, the presentation will satisfy $C(p)$-$T(4)$ for large p. In addition, suppose that $|A| = m-2 = |B|$, and $|A'| = m-1 = |B'|$, and $|C| = 2k-2$.

If we follow the construction of the hypergraphs in the universal cover X of the standard 2-complex of the presentation above, then we find that there is a single connected hypergraph Λ in X. The stabilizer of this hypergraph cannot be codimension 1, since it is the entire group. Similarly, any finite index subgroup of Stabilizer(Λ) gives a pair with zero ends.

To see this, observe that there is a subgraph of $G\backslash\Lambda$ which, when pushed into the 1-skeleton produces generators for the entire fundamental group. The subgraph contains five edges and four vertices and consists of two bigons joined by a fifth edge. The fifth edge maps to the path $(a^{k-1}b^{k-1})$. The generators are indicated by the following, where I have parenthesized subwords arising from each edge.

$$\left\langle\, (a^{m-2})(a^{-(m-1)}),\ (a^{k-1}b^{k-1})(b^{m-2})(b^{-(m-1)})(a^{k-1}b^{k-1})^{-1} \,\right\rangle.$$

4. The hypergraphs fill

For a cell c in X, let Star(c) denote the subcomplex of X which is the union of all closed cells containing c.

Lemma 4.1. *Let v be a 0-cell, and for each closed 1-cell e, let Λ_e denote the hypergraph dual to e, and let H_e denote the halfcarrier corresponding to the orientation of Λ_e towards v. Then* $\mathrm{Star}(v) = \cap_{v\in e} H_e$.

Proof. Let $K = \cap_{v\in e} H_e$. It is clear that $\mathrm{Star}(v) \subset K$ since $\mathrm{Star}(v) \subset H_e$ for each e. Each H_e is convex by Corollary 3.20, and so K is convex by Lemma 3.16 and thus K is connected by Lemma 3.17. It therefore suffices to show that for any closed 1-cell f, if f is incident with a 0-cell of Star(v) but $f \not\subset \mathrm{Star}(v)$, then $f \not\subset K$.

Suppose first that f is incident with a 1-cell h that is incident with v. Since Λ_h obviously separates f and v, we see that f lies in H_h if and only if f lies in Y_h. But f lies in Y_h if and only if there is closed 2-cell g containing both f and h. But then $v \subset g$ so $f \subset g \subset \mathrm{Star}(v)$.

Now suppose that f is not incident with any 1-cell incident with v, but f intersects a closed 2-cell J in a 0-cell u. Without loss of generality, we assume that f is oriented away from u. Let P_1 be the longest counterclockwise subpath of ∂j ending at u, such that $P_1 f$ is a subpath of ∂r_1 for some closed 2-cell r_1. Similarly, let P_2 be the longest counterclockwise subpath of ∂j beginning at u, such that $f^{-1}P_2$ is a subpath of ∂r_2 for some closed 2-cell r_2. We refer the reader to Figure 9. Note that in the $B(4)$-$T(4)$ case, at least one of P_1 and P_2 is trivial. Indeed, otherwise the diagram consisting of r_1, r_2, and j glued together along the paths f, P_1, and P_2 would have an internal vertex of valence 3.

Observe that v is disjoint from P_1P_2, for if v is contained in P_i, then $v \in r_i$ and so f is contained in Star(v). But we have assumed that $f \not\subset \mathrm{Star}(v)$.

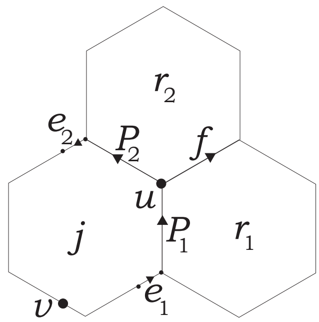

FIGURE 9. Separating f from $\mathrm{Star}(v)$.

Let e_1 and e_2 be edges such that $e_1P_1P_2e_2$ is a subpath of ∂j. Let b_1 and b_2 be the 1-cells in ∂j which are antipodal to e_1 and e_2. We will show that f fails to be in H_{e_i} for some i.

Since the argument is similar in the other case, we assume that $d(v, e_1) \leq d(v, e_2)$ where $d(a, b)$ denotes the distance between a and b along ∂j. Observe that f and v are separated by Λ_{e_1}, and so in order to prove that $f \notin H_{e_1}$, it is sufficient to show that f does not lie in Y_{e_1}.

Since P_1 is maximal, j is the only closed 2-cell in Y_{e_1} that contains e_1 and u. Since e_1 is a piece by Convention 3.1, $e_1P_1P_2$ is the concatenation of at most three [two] pieces, the $B(6)$ [$B(4)$-$T(4)$] condition guarantees that b_1 does not lie in P_1P_2, and b_1 therefore lies counterclockwise from P_2 at e_2 or beyond. Consequently, the maximality of P_2 implies that j is the only closed 2-cell in $Y_{e_1} = Y_{b_1}$ which contains b_1 and u. Thus j is the only closed 2-cell in Y_{e_1} which contains u. Since $f \not\subset j$, we conclude that $f \not\subset Y_{e_1}$. □

Definition 4.2 (fills)**.** The collection of hypergraphs *fills* X provided that for any distinct 0-cells u and v in X^0, there is a hypergraph Λ which separates u and v, in the sense that they lie in distinct components of $X - \Lambda$.

Lemma 4.3. *The collection of hypergraphs fills X.*

Proof. By Lemma 4.1, each 0-cell outside $\mathrm{Star}(v)$ is separated from v by some hypergraph. Therefore it suffices to show that w is separated from v by a hypergraph, where w is a 0-cell in $\mathrm{Star}(v)$ that is distinct from v.

Either w and v are connected by a single 1-cell h, or w lies on ∂j for some 2-cell j. In the latter case, let h denote the last 1-cell in a path from w to v that is at most half the length of ∂j.

In either case w and v are separated by the hypergraph Λ_h since there is a path from w to v that passes through a dual edge of Λ_h exactly once. □

For readers familiar with Sageev's thesis, we can now sketch a proof of a weak version of our main theorem in the word-hyperbolic case.

Theorem 4.4. *Let G be a word-hyperbolic group that acts properly discontinuously and cocompactly on a simply-connected $B(6)$ or $B(4)$-$T(4)$ complex X. Then G acts faithfully and cocompactly on a CAT(0) cube complex.*

Sketch. Let $\Lambda_1, \ldots, \Lambda_n$ be representatives of the finitely many orbits of hypergraphs in X. For each i, let H_i be the finite index subgroup of Stabilizer(Λ_i), which stabilizes both halfspaces of Λ_i.

Now, each codimension 1 subgroup H_i is quasiconvex by Theorem 3.18. Therefore, a multiple subgroup version of the results of [Sag97] shows that the action on Sageev's cube complex is cocompact. Finally, one can use Lemma 4.3, together with the hypothesis that G acts faithfully on X, and that $H_1, \ldots, H_n$ fill G to see that the action on the cube complex is faithful. $\square$

In the next few section we will give a complete proof of the cocompactness assertion in the previous theorem, by explicitly describing the maximal cubes. We will also add the conclusion that the action is properly discontinuous by analyzing the vertices in C. The $B(4)$-$T(4)$ and $B(6)$-$C(7)$ cases are treated in Theorem 11.1. The general $B(6)$ case is treated in Theorem 12.2, and the word-hyperbolic $B(6)$ case is obtained in Corollary 12.3.

5. The CAT(0) cube complex C

In this section we describe a variant of Sageev's construction [Sag95, Section 3]. The definitions that we give are catered to our situation which is simpler than the general situation that Sageev considered. The proofs that we give are simplified versions of his proofs. One explicit difference is that our cube complex derives from a 2-complex X, instead of from a group G and a set of codimension 1 subgroups. We begin by recalling the notion of a CAT(0) cube complex which was introduced in [Gro87].

Definition 5.1 (CAT(0) cube complex)**.** A 0-cube is a 0-cell, and a 1-cube consists of two 0-cells joined by a 1-cell. For $n \geq 2$, an *n-cube* is the combinatorial cell complex which is the n-fold cartesian product $I \times \cdots \times I$, where I is a 1-cube. For $n \geq 1$, each n-cube contains $2n$ distinct $(n-1)$-cubes, and these in turn may contain $(n-2)$-cubes and so on. A *cube complex* is a combinatorial cell complex which is the union of closed embedded cubes, such that the nonempty intersection of any two cubes is a cube.

A cube complex C satisfies the *CAT(0) condition* if it is simply-connected, and the following condition holds for each $n \geq 0$: For any three $(n+2)$-cubes c_1, c_2, c_3 such that $c_1 \cap c_2 \cap c_3$ is an n-cube, and such that $c_1 \cap c_2$, $c_2 \cap c_3$, and $c_3 \cap c_1$ are $(n+1)$-cubes, there exists an $(n+3)$ cube d which contains each c_i as a subcube. It is a theorem of Gromov that if a cube complex satisfies this combinatorial CAT(0) condition then it has a corresponding CAT(0) metric agreeing with the Euclidean metrics on its cubes.

Finally, we note that cube complexes contain convex codimension 1 subspaces called *hyperplanes* that only intersect cubes in *hypercubes* restricting one of their coordinates to $\frac{1}{2}$ (e.g. $I \times I \times \{\frac{1}{2}\} \times I$). The center of each n-dimensional cube is the intersection of n hyperplanes.

We shall now assume that X is a simply-connected 2-complex with a fixed set of embedded graphs each of which separates X into two components. We will refer to these as *hypergraphs*, and as in previous sections, we shall not assume that they are subcomplexes of X. The reader is welcome to think of hypergraphs as defined in Section 3, but the statements and proofs we give here work in this more general sense.

Definition 5.2 (Halfspaces and oriented hypergraphs)**.** As in Lemma 3.13, for each hypergraph Λ, there are exactly two connected components in $X - \Lambda$, and we shall refer to the closure of each of these components as a *halfspace*. If A is one halfspace associated to Λ, we use the notation A^c for the other halfspace, so $A^c = \Lambda \cup (X - A)$.

An *orientation* on a hypergraph Λ, is a choice of exactly one of its two associated halfspaces, and we shall say that Λ is *oriented towards* this halfspace. More generally, given a subset $S \subset X$, we shall say that Λ is *oriented towards* or *points towards* S if the halfspace that Λ is oriented towards, has nonempty intersection with S. We say it is *oriented away* if it is not oriented towards.

The complex C is a component of a cube complex $\hat{C}$ that we shall now construct. We will begin by defining the 1-skeleton of $\hat{C}$ and add the higher dimensional cubes later.

Let Σ denote the collection of halfspaces in X.

A *vertex* V of $\hat{C}$ is a subcollection of Σ satisfying two conditions:

1) For all $A \in \Sigma$, exactly one of A and A^c lies in V.

2) If $A \in V$ and $B \in \Sigma$, and $A \subset B$, then $B \in V$.

Equivalently, a vertex is a choice of orientation on each hypergraph in $\hat{X}$, subject to the condition that no two hypergraphs are oriented away from each other. (Note that if two hypergraphs intersect, they automatically point towards each other no matter how they are oriented.)

The *edges* of $\hat{C}$ are defined as follows: Given vertices V and W in $\hat{C}$, they are joined by an edge if and only if V and W disagree on the orientation of a single hypergraph. Equivalently, they are joined by an edge if and only if the symmetric difference of V and W consists of a single element of Σ.

An oriented hypergraph of a vertex V is *minimal* provided that a new vertex is obtained by reversing its orientation. Equivalently, it is minimal if its oriented halfspace is not contained in any other oriented halfspace of V. The edges of $\hat{C}$ at V are in one-to-one correspondence with the minimal hypergraphs of V.

Definition 5.3. We say distinct hypergraphs Λ_1 and Λ_2 *cross* if they intersect. Similarly, for use in later sections, we say that distinct hypercarriers Y_1 and Y_2 *cross* if $Y_1 \cap Y_2$ contains a 2-cell.

The cubes of $\hat{C}$ are added inductively by declaring that an n-dimensional cube is added whenever its $(n-1)$-skeleton appears. Consider a collection of n distinct pairwise crossing hypergraphs $\Lambda_1, \dots, \Lambda_n$. Furthermore, orient every other hypergraph Λ, so that Λ points towards Λ_i for each i. Having made these choices, the orientations of the Λ_i hypergraphs can be chosen freely to obtain 2^n vertices in $\hat{C}$ which lie on the 1-skeleton of a (missing) n-cube in $\hat{C}$. We attach an n-cube c to C along the previously attached $(n-1)$-cubes, each of which corresponds to fixing the orientation of one of the Λ_i, and letting the orientations of the remaining Λ_i vary. The set $\Lambda_1, \dots, \Lambda_n$ are the *independent hypergraphs* of c, and the remaining hypergraphs will be called *dependent* hypergraphs of c.

For any 0-cell $x \in X$, we let $V_x = \{A \in \Sigma : x \in A\}$ and observe that V_x satisfies the conditions listed above, so V_x is a vertex of $\hat{C}$. That is, we orient all hypercarriers in X so that they point towards x. Note that there are usually vertices in $\hat{C}$ that do not arise as V_x for some $x \in X^0$.

Lemma 5.4. *The vertices V_x are contained in a single component of $\hat{C}$.*

Proof. Since X is connected, it is sufficient to show that if x_1 and x_2 are joined by an edge e in X then V_{x_1} and V_{x_2} are joined by an edge in $\hat{C}$. Let Λ_e be the hypergraph dual to e, and observe that V_{x_1} and V_{x_2} only differ by the orientation chosen for Λ_e. Indeed, for any hypergraph Λ, if x_1 and x_2 lie in distinct halfspaces of Λ then e must be a dual edge of Λ. Therefore $\Lambda = \Lambda_e$. □

Lemma 5.4 allows us to define the complex C to be the component of $\hat{C}$ which contains all vertices V_x. The action of G on X determines the action on Σ which determines the action of G on $\hat{C}$. Since $gV_x = V_{gx}$ we see that G stabilizes C and hence there is an action of G on C.

Lemma 5.5. *Each component of $\hat{C}$ and thus C is simply-connected.*

We adopt some notation before beginning the proof: Given a vertex V and a sequence of hypergraphs $\Lambda_1, \Lambda_2, \ldots, \Lambda_k$, we use the notation $(V + \Lambda_1)$ to denote reversing the orientation of Λ_1. Similarly, $(V + \Lambda_1 + \Lambda_2)$ denotes reversing the orientation of Λ_1 followed by reversing the orientation of Λ_2. Note that if some Λ appears m times among $\Lambda_1, \ldots, \Lambda_k$, then the orientation of Λ in $(V + \Lambda_1 + \cdots + \Lambda_k)$ will remain unchanged if m is even and be reversed if m is odd. Furthermore, in general, $(V + \Lambda_1 + \cdots + \Lambda_k)$ may not be a vertex even though V is.

We use the notation $(V_0 : \Lambda_1, \Lambda_2, \ldots, \Lambda_n)$ for a path which starts at the vertex V_0, whose k-th vertex V_k equals $(V_0 + \Lambda_1 + \cdots + \Lambda_{k-1} + \Lambda_k)$, and whose k-th edge corresponds to changing the orientation Λ_k in V_{k-1} to obtain $V_k = (V_{k-1} + \Lambda_k)$.

Proof of Lemma 5.5. Consider a combinatorial path P in $\hat{C}$ which starts and ends at the vertex V_0. Suppose P is given by $P = (V_0 : \Lambda_1, \Lambda_2, \ldots, \Lambda_n)$. Since $V_0 = V_n$, we see that each hypergraph Λ appears an even number of times among $\{\Lambda_1, \ldots, \Lambda_n\}$. Suppose that $\Lambda_i = \Lambda_j$ for some $i < j$, and assume that i and j are chosen so that $(j - i)$ is minimal.

If $(j - i) = 1$, then P has a "backtrack" and can be homotoped to the shorter path corresponding to the sequence $(V_0 : \Lambda_1, \ldots, \Lambda_{i-1}, \Lambda_{j+1}, \ldots, \Lambda_n)$ whose sequence differs by the removal of Λ_i, Λ_j.

Otherwise, we will show below that Λ_{i+1} and Λ_i cross each other, and thus there is a 2-cube attached along the closed path $(V_i : \Lambda_i, \Lambda_{i+1}, \Lambda_i, \Lambda_{i+1})$. We may therefore homotope P through this 2-cube to obtain the path P' given by $(V_0 : \Lambda_1, \ldots, \Lambda_{i+1}, \Lambda_i, \ldots \Lambda_n)$ whose sequence of hypergraphs differs from the original sequence by interchanging the order of Λ_i and Λ_{i+1}. After repeating this interchange argument $j - i - 1$ times, we reach a path such that Λ_i and Λ_j appear consecutively, and may therefore be deleted as in the backtrack case above. After repeating this process finitely many times we reach the trivial path at V_0.

To see that Λ_i and Λ_{i+1} cross we argue as follows: Consider the vertices V_{i-1}, V_i, V_{i+1}, and V_{j+1}. Without loss of generality, assume that the orientations of Λ_i and Λ_{i+1} are $(+, +)$ at V_{i-1}. Then they are $(-, +)$ at V_i, and $(-, -)$ at V_{i+1}. The minimality of $(j-i)$ guarantees that both Λ_i and Λ_{i+1} are distinct from Λ_k for $(i+1) < k < j$. Therefore, they continue to have orientation $(-,-)$ all the way from V_{i+1} until V_{j-1}. Finally, since $\Lambda_j = \Lambda_i$, we see that their orientation changes to $(+,-)$ at V_j.

Since Λ_i and Λ_{i+1} can be oriented independently at various vertices, we see from the definition of a vertex, that they must cross each other. They therefore determine a square with a corner at V_i as claimed. □

Remark 5.6. The proof of Lemma 5.5 contains a proof that any length 4 simple cycle in $\hat{C}^1$ actually bounds a 2-cube corresponding to a pair of hypergraphs that cross, and can thus be oriented independently.

In the construction of $\hat{C}$, we implicity assumed that the only way the 1-skeleton of an n-cube appears in $\hat{C}$, is if there are n such independent hypergraphs. To see that this is true, suppose v is a vertex in a subgraph of $\hat{C}$ isomorphic with the 1-skeleton of an n-cube. Let $e_1, \ldots, e_n$ be the n edges of this subgraph incident with v. Then each e_i and e_j lie at the corner of a square, and thus, as explained above, their corresponding hypergraphs Λ_i and Λ_j cross. Thus the n-cube naturally corresponds with (choices of orientations of the) n pairwise crossing hypergraphs as claimed.

Lemma 5.7. *Each component of $\hat{C}$ and hence C is CAT(0).*

Proof. Suppose that c_1, c_2, and c_3 are $(n+2)$-cubes, that $c_1 \cap c_2$, $c_2 \cap c_3$ and $c_3 \cap c_1$ are $(n+1)$-cubes, and that $c_1 \cap c_2 \cap c_3$ is an n-cube.

Each c_i is associated with a set of independent hypergraphs. Consideration of the Venn diagram for these three sets shows that the total number of independent hypergraphs is $3(n+2) - 3(n+1) + n = (n+3)$. Furthermore, consideration of the Venn diagram shows that any two hypergraphs in the union of the three sets, actually lie together in one of the three sets, and thus cross. Indeed, this is obvious once we ignore all hypergraphs in the triple intersection, after which the Venn diagram corresponds to the three sides of a triangle. Finally, the desired $(n+3)$-cube has these $(n+3)$ independent hypergraphs, and all other hypergraphs are oriented towards the hypergraphs of the original $(n+2)$-cubes. □

6. Pairwise intersecting hypergraphs in the $B(4)$-$T(4)$ case

As defined in Section 5, the cubes in the CAT(0) cube complex C correspond to collections of pairwise crossing hypergraphs (see Definition 5.3). In this section, we begin to analyze the possible collections of pairwise intersecting hypergraphs. The main result in this section characterizes three pairwise crossing hypergraphs. We are able to apply this in this section to completely understand collections of pairwise crossing hypergraphs in the $B(4)$-$T(4)$ or $B(6)$-$C(7)$ cases. The analysis in the general $B(6)$ case is considerably more complicated, and is carried out in Section 8.

In this section we determine what the maximal cubes in C are. In the hyperbolic case we will show that when X is locally finite then so is C. In addition, in the $B(6)$-$C(7)$ or in the $B(4)$-$T(4)$ case we find that if the action on X is cocompact then the action on C is cocompact. However, in the $B(6)$ case, we find that the action on C is cofinite in a sense that we will describe in Section 8.

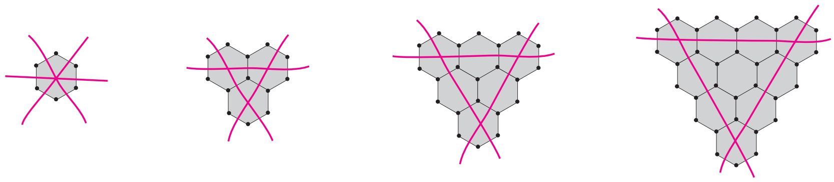

FIGURE 10. Hypergraphs meeting along tricombs.

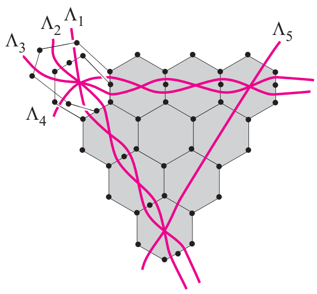

FIGURE 11. The above 2-complex contains eleven 2-cells. Each width 4 tricomb in the 2-complex is bordered by four hypergraphs.

Definition 6.1 (Tricombs, borders, & corners)**.** A *tricomb* Δ is a nonsingular disc diagram, such that every interior 2-cell meets exactly 6 other 2-cells, and all 2-cells that meet $\partial\Delta$ are 4-shells, with the exception of three 2-shells. These three exceptional 2-cells are the *corners* of Δ. We will also consider the case of a single 2-cell to be a degenerate tricomb whose three corners coincide. (The term 'tricomb' is meant to suggest 'triangular honeycomb'.)

We say a hypercarrier Y *borders* a tricomb Δ if Y passes through two corners of Δ. Note that if Y borders Δ then Y passes through an entire sequence of 2-cells along one side of Δ. The *width* of a tricomb is the number of 2-cells in one and hence any of its borders. We say hypercarriers Y_a, Y_b, and Y_c *meet along* Δ, provided that they each border Δ, and they pairwise cross at the three corners of Δ. We say that a hypergraph *borders* Δ if its associated hypercarrier borders it, and similarly, we say three hypergraphs *meet* along Δ if their associated hypercarriers meet along Δ. We refer the reader to Figure 10 for an illustration of the first four types of tricombs, as well as three hypergraphs corresponding to hypercarriers that meet along each tricomb. Finally, we note that when there are pieces of length ≥ 2, a tricomb may be bordered by more than three hypercarriers. For example, in Figure 11, each width 4 tricomb is bordered by four hypergraphs. Note that Λ_3 and Λ_4 do not cross at the upper left.

We briefly recall the Combinatorial Gauss-Bonnet Theorem. A proof can be found, for instance, in [MW02].

Theorem 6.2 (Combinatorial Gauss-Bonnet Theorem)**.** *Let X be a disc diagram with a real number called an* angle *assigned to each* corner *of each 2-cell of X. To be precise, the angles are assigned to the edges of the links of 0-cells of X, and in an obvious sense, each 2-cell and 0-cell has a set of corners with assigned angles.*

The curvature of a 0-cell $v \in X^0$ is defined to be

$$\kappa(v) = 2\pi - \pi\chi\big(\text{link}(v)\big) \; - \sum_{c\in Corners(v)} \angle c. \tag{1}$$

For a 2-cell f, the curvature $\kappa(f)$ is defined to be

$$\kappa(f) = \big(|\partial f| - 2\big)\pi \; - \sum_{c\in Corners(f)} \angle c. \tag{2}$$

The Combinatorial Gauss-Bonnet Theorem states that

$$2\pi \;=\; \sum_{v\in X^0} \kappa(v) \;+ \sum_{f\in 2\text{-}cells(X)} \kappa(f).$$

Lemma 6.3. *Suppose that Y_1, Y_2, and Y_3 pairwise cross.*
If X is $B(4)$-$T(4)$ then $Y_1 \cap Y_2 \cap Y_3$ contains a 2-cell.
If X is $B(6)$-$C(7)$ then $Y_1 \cap Y_2 \cap Y_3$ contains a 0-cell.

The proof of Lemma 6.3 will also prove the following:

Lemma 6.4. *Let X be $B(6)$. If the hypergraphs Λ_1, Λ_2, and Λ_3 pairwise cross then they meet along a tricomb in X.*

Note that the tricomb in Lemma 6.4 is not necessarily uniquely determined by Λ_1, Λ_2, and Λ_3. The ambiguity is at the corners, and for a width ≥ 3 tricomb, this ambiguity is addressed later in Lemma 8.2.

Proof. If $Y_1 \cap Y_2 \cap Y_3$ contains a 2-cell, then Lemma 6.3 obviously holds, and Lemma 6.4 holds with a degenerate tricomb consisting of a 2-cell. We therefore suppose that $Y_1 \cap Y_2 \cap Y_3$ does not contain a 2-cell. Choose a 2-cell U_{12} in $Y_1 \cap Y_2$, and similarly choose 2-cells U_{23} and U_{31}. Let L_1 be a ladder in Y_1 which starts in U_{31} and ends in U_{12}, and define L_2 and L_3 similarly. Furthermore, we suppose that the above choices have been made so that the the total number of 2-cells in L_1, L_2, and L_3 is minimal. For each i, let λ_i denote the geodesic in Λ_i corresponding to L_i.

There are three useful properties guaranteed by the minimality of the number of 2-cells. Firstly, each L_i has at least two 2-cells, indeed, if for instance $L_1 = U_{12} = U_{31}$ then this 2-cell is contained in $L_1 \cap L_2 \cap L_3$. Secondly, the 2-cells in the initial and terminal overlaps between the ladders consist of exactly U_{12}, U_{23} and U_{31}. Indeed, otherwise we can remove initial and terminal cells as in Figure 12.

Thirdly, the following holds for each fixed pair i, j: Let V_i and V_j denote the 2-cells in L_i and L_j which are attached to U_{ij}. Let $r_i = V_i \cap U_{ij}$ and $r_j = V_j \cap U_{ij}$. Then λ_i does not pass through r_j and similarly λ_j does not pass through r_i. Indeed, if λ_i passes through r_j then Λ_i passes through r_j and hence V_j, and so $V_j \subset Y_i$. We can therefore modify L_i by replacing U_{ij} by V_j, and modify L_j by simply removing U_{ij}. This would result in a new triple of ladders satisfying the properties above but with exactly one fewer 2-cell, which is impossible. (This third property is a strengthening of the second property above.)

Subdivide X so that each λ_i is contained in the 1-skeleton. Let u_{12}, u_{23}, and u_{31} be the new 0-cells at the centers of U_{12}, U_{23}, U_{31}. Thus λ_1 is a path from u_{21} to u_{13} and similarly for λ_2 and λ_3. Consider a reduced disc diagram D whose

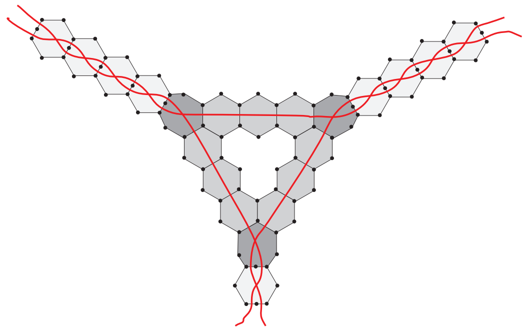

FIGURE 12. The three very dark 2-cells are U_{12}, U_{23}, and U_{31}. The ladders L_1, and L_2, and L_3 are the sequences of dark 2-cells joining them. The light 2-cells hanging off the ends of the triangle are the sequences of discarded 2-cells in the initial and terminal overlaps.

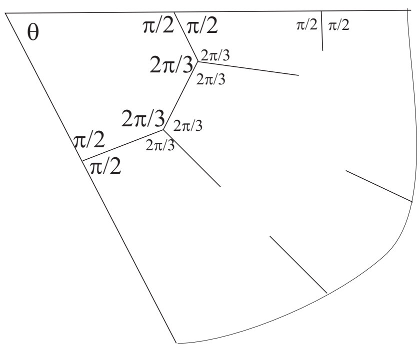

FIGURE 13. The angle assignments.

boundary path is the concatenation $\lambda_1\lambda_2\lambda_3$. We will ignore new 0-cells and 1-cells from the λ_i subdivisions except for when they occur on ∂D.

After ignoring all 0-cells of valence 2 (except for u_{12}, u_{23}, and u_{31}), we assign angles at corners of 2-cells of D in the following way: All interior corners receive an angle of $\frac{\pi}{2}$ in the $T(4)$ case and $\frac{2\pi}{3}$ in the $C(6)$ case. All corners containing a 1-cell of ∂D receive an angle of $\frac{\pi}{2}$.

In the $C(6)$ case, the corner of D at u_{12} receives the angle $\theta = \frac{(n-3)\pi}{3}$ so that $(n_{12}-2)\pi = \frac{\pi}{2} + \frac{\pi}{2} + (n_{12}-3)\frac{2\pi}{3} + \theta$ and hence the curvature of the n_{12}-sided 2-cell at this corner is nonpositive. Similarly, in the $B(4)$-$T(4)$ case, we assign an angle $\theta = \frac{(n-3)\pi}{2}$ so that $(n_{12}-2)\pi = \frac{\pi}{2} + \frac{\pi}{2} + (n_{12}-3)\frac{\pi}{2} + \theta$. The angles at the u_{23} and u_{31} corners are assigned similarly. We refer the reader to Figure 13 for the angle assignments.

Note that U_{12} is disjoint from λ_3 for otherwise $Y_1 \cap Y_2 \cap Y_3$ contains U_{12}. It follows that all other corners of the 2-cell of D containing u_{12} are either interior or are external and $\frac{\pi}{2}$. The analogous statements hold for u_{23} and u_{31}. Finally, we apply the third property that we deduced from the minimality above to see that each 2-cell containing a u_{ij} corner, has $n_{ij} \geq 4$ sides. Indeed, there is a side along part of each of λ_i and λ_j, and if there is only one other side, then either λ_i passes through r_j or λ_j passes through r_i.

It follows immediately from our angle assignments that internal 0-cells of D have nonpositive curvature. Similarly, external 0-cells not at u_{12}, u_{23} or u_{31} have nonpositive curvature since they contain at least two $\frac{\pi}{2}$ corners. We conclude that the sum of the curvatures at u_{12}, u_{23}, and u_{31} is exactly 2π.

This is clearly impossible in the $B(4)$-$T(4)$ case where since each $n_{ij} \geq 4$, these curvatures are all $\leq \frac{\pi}{2}$ and so our assumption that $Y_1 \cap Y_2 \cap Y_3$ contains no 2-cell is false.

In the $C(6)$ case, since each $n_{ij} \geq 4$, the curvatures at u_{12}, u_{23}, and u_{31} are all $\leq \frac{2\pi}{3}$, and therefore all equal to $\frac{2\pi}{3}$. There are therefore no 0-cells or 2-cells with negative curvature. Consequently every other 0-cell with valence > 2 has valence exactly three, and every interior 2-cell has exactly six sides, except for those that appear along the boundary which have exactly five, and the three 2-cells at the corners which have exactly four. Consequently, D can be obtained by cutting along the hypergraphs meeting along a tricomb as can be readily ascertained from the family of diagrams illustrated in Figure 10.

In the $C(7)$ case, any interior 2-cell would have at least 7 sides, so all but the first three tricombs are eliminated. The central 0-cell in the second and third cases is the desired 0-cell in $Y_1 \cap Y_2 \cap Y_3$. □

Definition 6.5. Let X be a $C(6)$ complex. Let u be the center of a 2-cell R. Let λ_1 and λ_2 be nontrivial geodesics in hypergraphs Λ_1 and Λ_2 that meet at u. Let P be a path in ∂R that is the concatenation of the fewest pieces such that the first edge in P is dual to λ_1 and the last edge is dual to λ_2. The *angle* between λ_1 and λ_2 at R is defined to be 0 if P consists of a single piece, $\frac{\pi}{3}$ if P consists of two pieces, $\frac{2\pi}{3}$ if P consists of three pieces, and π if P consists of four or more pieces.

We can now state the following corollary of the proof of Lemma 6.4.

Corollary 6.6. *Let X satisfy $B(6)$, and consider the concatenation $\lambda_1\lambda_2\lambda_3$ where each λ_i is a geodesic in a hypergraph Λ_i in X. If the angle between λ_1 and λ_2 at the 2-cell R that they meet in is $\geq \frac{2\pi}{3}$ then $\lambda_1\lambda_2\lambda_3$ is not closed unless λ_3 also passes through R.*

Definition 6.7 (Between)**.** It follows from Lemma 3.6 that the union L of a finite sequence of consecutive 2-cells $c_1, c_2, \ldots, c_r$ in a hypercarrier Y is a nonsingular ladder. Moreover, this ladder is nonsingular and has the property that $c_i \cap c_j = \emptyset$ for $i + 1 < j$. Each rung of L is an arc in a rung tree of Y.

For cells a and b in Y, there is a unique "shortest" ladder L in Y that contains a and b. It is the ladder *connecting* a and b.

A *rung* in Y is the intersection of a pair of neighboring 2-cells.

A 2-cell or rung of a ladder L lies *between* 0-cells u and v of L if u and v lies in the closures of distinct components of $L - r$.

Lemma 6.8. *Let u and v be 0-cells that lie in distinct components of ∂Y_1, and let L be the ladder between u and v. If u and v lie in Y_2 then L is contained in Y_2.*

Proof. Let P be a geodesic path in Y_1 from u to v. Since P is a geodesic, it is contained in some ladder L' in Y_1, and by possibly adding an additional 2-cell at

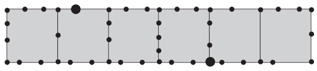

FIGURE 14. There are three rungs and two 2-cells between the two large 0-cells in the above ladder.

each side, we can assume that P does not pass through either dual edge on $\partial L'$. Since P is a geodesic, it cannot contain the complement of a rung in the boundary of some 2-cell.

Since u and v lie in distinct components of ∂Y_1 we see that P must pass through some rung r_1 of L'. Observe that P is also a path in Y_2 since Y_2 is convex by Theorem 3.18. Therefore, in particular, r_1 lies in Y_2.

By definition, the ladder L between u and v is contained in L'. Let $r_1, r_2, \ldots, r_m$ be a sequence of consecutive rungs in L' where r_m is a rung in L, and let $R_1, R_2, \ldots, R_{m-1}$ be the corresponding consecutive 2-cells. We will show that each R_i lies in Y_2. Indeed, consider the first 2-cell R_i that does not lie in Y_2, and note that r_i lies in Y_2. Observe that P must pass along an arc s_i which is the intersection of ∂R_i and $\partial L'$. Let Q be the concatenation $r_i s_i$, and observe that Q is a path in Y_2.

Now the $B(6)$ [$B(4)$-$T(4)$] condition applied in Y_1 shows that Q is not the concatenation of fewer than three [two] pieces. Therefore Lemma 3.8 implies that R_i cannot be missing in Y_2 along Q.

The degenerate case where u and v lie on distinct 2-cells and there is exactly one rung between them is handled as above. In case u and v lie on the same 2-cell of Y_1 and there is no rung between them, then if they are antipodal, the entire boundary and hence the 2-cell itself lies in Y_2, and otherwise, a rung on one side of this 2-cell is contained in Y_2. □

Theorem 6.9. *Let $\{Y_1, Y_2, \ldots\}$ be a set of hypercarriers in X which pairwise cross.*

If X is $B(6)$-$C(7)$ then $\cap_i Y_i$ contains a 0-cell.

If X is $B(4)$-$T(4)$ then $\cap_i Y_i$ contains a 2-cell.

The second part of Theorem 6.9 actually follows from the following theorem which will be proved together with Theorem 6.9:

Theorem 6.10. *Let $\{Y_1, Y_2, \ldots\}$ be a set of hypercarriers in the $B(4)$-$T(4)$ complex X such that $Y_i \cap Y_j \cap Y_k$ contains a 2-cell for all i, j, k. Then $\cap_i Y_i$ contains a 2-cell.*

Proof of Theorem 6.9 and Theorem 6.10. . We first consider the $B(6)$-$C(7)$ case. We will prove the theorem by induction. Suppose the theorem has been proven for any ℓ hypercarriers with $\ell \geq 3$, and note that the case $\ell = 3$ is stated in Lemma 6.3. Consider a set of hypercarriers $\{Y_i : i \in I\}$ where $I = \{1, \ldots, \ell + 1\}$. For a subset $S \subset I$, we will use the notation Y_S to denote $\cap_{i \in S} Y_i$, so that $Y_I \subset Y_{I-\{1\}} \subset Y_{I-\{1,2\}}$ and so on.

Since the statement is true for any ℓ hypercarriers, we see that $Y_{I-\{1\}}$, $Y_{I-\{2\}}$ and $Y_{I-\{3\}}$ are nonempty. For $1 \leq i \leq 3$ choose 0-cells $u_i \in Y_{I-\{i\}}$ and choose geodesic paths P_{ij} from u_i to u_j. By Theorem 3.18, P_{ij} is contained in $Y_{I-\{i,j\}}$ and in particular, in $Y_{\ell+1}$.

Each P_{ij} either lies entirely in one component of $\partial Y_{\ell+1}$, or by Lemma 6.8, every rung and 2-cell in $Y_{\ell+1}$ between u_i and u_j lies in $Y_{I-\{i,j\}}$. In the former case we can replace P_{ij} by an immersed path P'_{ij} in $Y_{I-\{i,j\}}$ that connects u_i and u_j and lies in $\partial Y_{\ell+1}$. It follows that if u_1, u_2, and u_3 all lie in the same component of $\partial Y_{\ell+1}$ then the triple intersection $P'_{12} \cap P'_{23} \cap P'_{31}$ in this tree component contains a 0-cell which necessarily lies in $Y_{I-\{1,2\}} \cap Y_{I-\{2,3\}} \cap Y_{I-\{3,1\}} = Y_I$.

We now consider the case where the u_i do not all lie in the same component of $\partial Y_{\ell+1}$. Without loss of generality, suppose that u_1 and u_2 lie in the 'top' component of $\partial Y_{\ell+1}$ and that u_3 lies in the 'bottom' component. The immersed path P'_{12} defined above lies entirely in the top component of $\partial Y_{\ell+1}$.

Let P'_{23} be the path in that component which corresponds to the sequence of rungs and 2-cells lying between u_2 and u_3. Note that this sequence of rungs and 2-cells forms a ladder L_{23}, and P'_{23} is the intersection of the upper boundary path of L_{23} and the top component of $\partial Y_{\ell+1}$. Define P'_{31} analogously.

Consider the sequences of rungs lying between u_3 and u_1 and between u_3 and u_2. If these sequences have no initial common subsequence then u_3 lies in the interior of the bottom of a 2-cell. Since the bottom as well as left and right rungs of this 2-cell lie in $Y_{I-\{3\}}$, the $B(6)$ condition and Lemma 3.8 imply that the entire 2-cell does. Therefore the top of this 2-cell lies in the intersection of $Y_{I-\{3\}}$ and P'_{12}. Otherwise, there is a last rung in this initial common subsequence. But then the top of this rung lies in the intersection of P'_{12} and $Y_{I-\{3\}}$.

The argument for Theorem 6.10 is similar to that given above. The inductive hypothesis provides us with 2-cells $R_i \in Y_{I-\{i\}}$ for $1 \leq i \leq 3$. For each i and j, we have a ladder L_{ij} in $Y_{I-\{i,j\}}$ which starts at R_i and ends at R_j.

Suppose that R_i lies in the ladder L_{jk} for some distinct i, j, and k. Then $R_i \subset \big(Y_{I-\{i\}} \cap L_{jk}\big) \subset Y_I$ and we are done.

Otherwise, let r be the dual edge in $Y_{\ell+1}$ along which the three ladders bifurcate. Note that r corresponds to the cut-vertex in $\Lambda_{\ell+1}$ that separates the three non-collinear vertices corresponding to R_1, R_2, and R_3. Let R'_1 denote the last 2-cell in the ladder from R_1 to r, and observe that $R'_1 \subset \big(L_{12} \cap L_{13}\big) \subset Y_{I-\{1\}}$. Define similarly $R'_2 \subset Y_{I-\{2\}}$ and $R'_3 \subset Y_{I-\{3\}}$.

Now, when we consider the way R'_1, R'_2 and R'_3 are folded together along r, we find that one of these 2-cells contains the intersection of the other two along its boundary. Indeed, r extends maximally "upwards" to a piece between two of these 2-cells, and r extends maximally "downwards" to a piece between two of these 2-cells. Since there are only three 2-cells, these maximal extensions both occur in the boundary of (at least) one of the 2-cells, say R'_1. In particular, the rung in L_{23} between R'_2 and R'_3 lies in $\partial R'_1$. But then every hypergraph passing through the rung between R'_2 and R'_3 must also pass through R'_1. Consequently, $R'_1 \subset Y_{I-\{2,3\}}$ and so since we already know that $R'_1 \subset Y_{I-\{1\}}$ we see that $R'_1 \subset Y_I$ and we are done.

Finally, the $B(4)$-$T(4)$ case of Theorem 6.9 follows from Theorem 6.10 because of Lemma 6.3. □

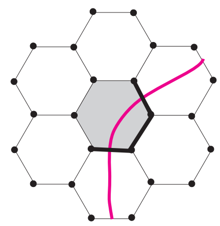

FIGURE 15. An internal hypercarrier in a $B(6)$ honeycomb cannot turn.

7. The flat planes in $B(6)$ complexes are isolated

Definition 7.1 (Flat Planes and Honeycombs)**.** A *flat plane* F is a $B(6)$ complex homeomorphic to the Euclidean plane such that each closed 2-cell intersects exactly six neighboring closed 2-cells. For a flat plane F there is an associated flat plane F' obtained by ignoring valence 2 0-cells of F, so F' is a regular hexagonal tiling of the Euclidean plane.

A reduced map of a nonsingular disc diagram $H \to X$ is a *honeycomb in* X if H is a subcomplex of a flat plane F. A hypercarrier in H is *internal*, if all its 2-cells are internal with the possible exception of its first and last closed 2-cells which contain edges on ∂H.

Given a honeycomb H embedded in the flat plane F, we let H' be the corresponding honeycomb embedded in F'. Note that H' is obtained from H by ignoring the 0-cells of H that have valence 2 in F.

Lemma 7.2. *Let H be a $B(6)$ honeycomb. All internal hypercarriers in H are standard in the sense that they correspond to internal hypercarriers in H'.*

Proof. Suppose the hypercarrier "turned" as in Figure 15, then the three shaded pieces in Figure 15 show that the $B(6)$ condition is violated. □

Lemma 7.3. *Let X satisfy $B(6)$, and let Δ_1 and Δ_2 be tricombs in X whose borders are hypercarriers. Suppose the hypercarrier Y borders both Δ_1 and Δ_2, and suppose that Δ_1 and Δ_2 lie on the same side of Y, and $\Delta_1 \cap Y = \Delta_2 \cap Y$. Then Δ_1 and Δ_2 are identical, except perhaps for the 2-cells at the corner opposite Y.*

Proof. The statement is trivial if $\text{width}(\Delta_i) \leq 2$. We will first prove the special case when $\text{width}(\Delta_i) = 3$ and we will then deduce the general case from this special case.

Consider the complex $\Delta_1 \cup_Y \Delta_2$ which is illustrated in Figure 16. Observe that the $B(6)$ condition implies that the ladders in Δ_1 and Δ_2 that lie along ∂Y are staggered as illustrated. Let R_1 and R_2 be the leftmost 2-cells in these ladders. By possibly interchanging Δ_1 and Δ_2 we may assume that R_2 does not meet ∂Y further to the left than R_1 meets ∂Y_1.

If R_1 and R_2 are not equivalent, then consideration of the three pieces in ∂R_1 that are illustrated in Figure 16, contradicts the $B(6)$ condition. The same argument holds for the 2-cells on the right, so the proof of the special case is complete.

The general case now follows by observing that the bottom of a large tricomb Δ is covered by smaller triangles whose sides have three 2-cells. Furthermore, the

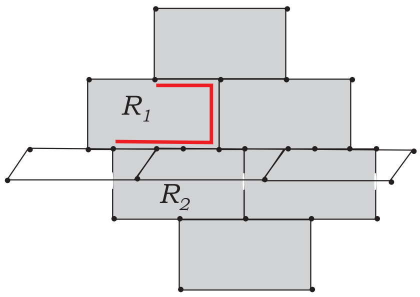

FIGURE 16. The complex $\Delta_1 \cup_Y \Delta_2$ is illustrated above. Note that Y consists of the horizontal strip of three 2-cells, and Δ_1 [Δ_2] consists of the union of this strip and the top [bottom] three 2-cells, Violation of $B(6)$ is indicated by the bold path of three pieces.

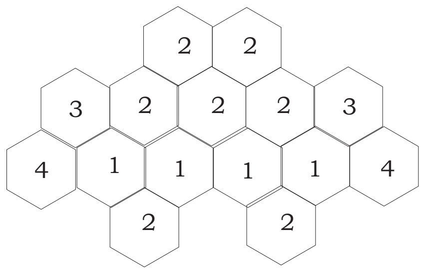

FIGURE 17. A chain of $n \geq 4$ 2-cells determines a chain of $n + 2$ 2-cells and so on. This sequence of chains determines a nested sequence of arbitrarily large regions in the flat plane F.

borders of these smaller triangles are genuine hypercarriers since they are either part of the border of the large triangle, or internal and thus hypercarriers by Lemma 7.2.

It follows that when we compare two large tricombs Δ_1 and Δ_2, the hypercarriers lying directly above Y are identical, and so Δ_1 and Δ_2 are identical all the way up until the top 2-cell which lies opposite Y. □

Lemma 7.4. *A flat plane F in the $B(6)$ complex X is uniquely determined by a chain of four 2-cells in F.*

Proof. Let Y_n be a chain of $n \geq 4$ 2-cells in a hypercarrier of E. By Lemma 7.3, this determines the cells labelled by '2' illustrated in Figure 17. Yet another application of Lemma 7.3, shows that the cells labelled by '3' are determined by the previous ones, and a further application shows that the cells labelled by '4' are determined.

Now observe that the 2-cells labelled by 1's and 4's form a chain of $n + 2$ 2-cells extending our original chain of n 2-cells. This process can be continued indefinitely, and so Lemma 7.3 shows that arbitrarily large nested regions in F are determined by the original chain of 4 cells. □

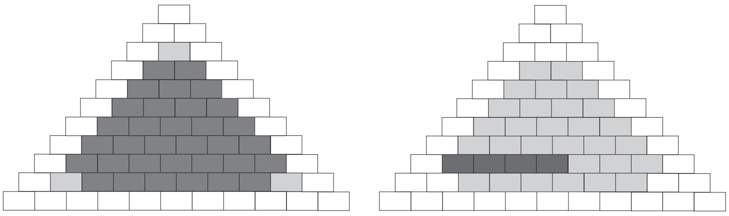

FIGURE 18. The truncated deputy, and a ladder of four 2-cells that uniquely determine it.

Theorem 7.5. *If G acts properly discontinuously and cocompactly on X then there are finitely many orbits of flat planes in X, and each of these corresponds to one of the finitely many conjugacy classes of maximal $\mathbb{Z} \times \mathbb{Z}$ subgroups in G.*

Proof. Since G acts properly discontinuously and cocompactly, there are only finitely many orbits of chains of four 2-cells, and therefore by Lemma 7.4, only finitely many orbits of flat planes. In particular, a flat plane E contains only finitely many 4-chains up to the action of G. Lemma 7.4 implies that any group element that takes one 4-chain in E to another 4-chain in E, must actually stabilize E. It follows that Stabilizer(E) acts cocompactly on E. Since the action of G on X is properly-discontinuous, so is the action of Stabilizer(E) on E. It follows that Stabilizer(E) is a finite index subgroup of the group of symmetries of the honeycomb tiling of the Euclidean plane, and therefore the translation subgroup of Stabilizer(E) is the desired maximal $\mathbb{Z} \times \mathbb{Z}$ subgroup.

We now sketch the converse: Given a $\mathbb{Z}\times\mathbb{Z}$ subgroup H of G generated by paths $\langle W_1, W_2\rangle$, there is a disc diagram $D \to H\backslash X$ whose boundary path is the commutator $W_1 W_2 W_1^{-1} W_2^{-1}$. We identifying opposite sides of D to obtain a singular surface whose fundamental group is isomorphic to $\mathbb{Z} \times \mathbb{Z}$, since it is generated by $\langle W_1, W_2\rangle$ and W_1 and W_2 commute because of D, and map surjectively to H in $\pi_1 X$. We remove all free faces and spheres from this singular surface, and successively removing cancellable pairs until we obtain a reduced map $T \to H\backslash X$ where T is homeomorphic to a torus. The Gauss-Bonnet theorem applies to show that T has exactly three hexagons meeting at a vertex, and so the universal cover $\widetilde{T}$ of T is the desired flat plane in X. □

Example 7.6. The standard 2-complex of the group G presented by

$$\langle a, b, t \mid t^{-1}a^3 t = b^3\rangle$$

satisfies $C(6)$ but does not satisfy $B(6)$. Furthermore, G has a free factor which is commensurable with $F_2 \times \mathbb{Z}$, so G has infinitely many maximal $\mathbb{Z} \times \mathbb{Z}$ subgroups.

Definition 7.7 (truncated deputy)**.** The *deputy* of a tricomb is the tricomb obtained by removing its borders. The *truncated deputy* is the deputy with its corner 2-cells removed. We refer the reader to the tricomb illustrated on the left in Figure 18. The shaded subdiagram there is the deputy of the entire tricomb, and the darkly shaded subdiagram is the truncated deputy.

Lemma 7.8. *The truncated deputy of a tricomb Δ in a $B(6)$ complex is determined uniquely by a sequence of four consecutive* 2*-cells in an internal hypercarrier in Δ.*

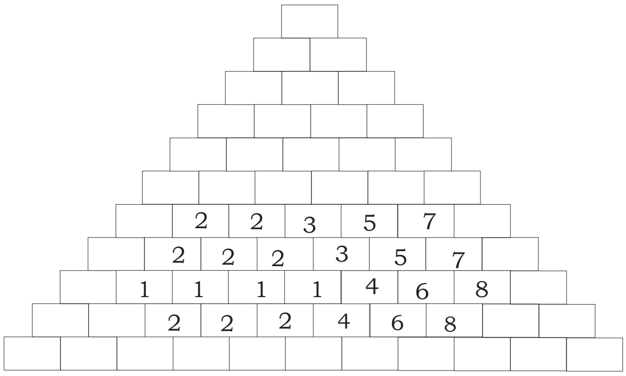

FIGURE 19. Four 2-cells which uniquely determine the truncated deputy.

Proof. In fact, this holds for *any* four consecutive 2-cells in an internal hypercarrier in the truncated deputy Δ^t, but we shall prove it for four consecutive 2-cells in a hypercarrier which is a border of the deputy of Δ. We refer the reader to the four heavily shaded 2-cells in the diagram on the right of Figure 18.

Consider the diagram in Figure 19. By Lemma 7.3, the four 2-cells labelled by 1's uniquely determine the 2-cells labeled by 2's. Similarly, the cells with label 1 and 2 determine the cells with label 3. The cells with label 1, 2, and 3 determine the cells with label 4. The cells with label 3 and 4 determine the cells with label 5, the cells with label 4 and 5 determine the cells with label 6, and so on.

The remaining cells in the truncated deputy are determined by the ones below using Lemma 7.3. ☐

Lemma 7.9. *Let G act properly discontinuously and cocompactly on a $B(6)$ complex X. With at most finitely many exceptions, every tricomb in X has the property that its truncated deputy is contained in a periodic flat plane.*

Proof. Consider the determining chain L of four 2-cells in a truncated deputy Δ^t as illustrated on the left in Figure 18. Then either L is contained in a flat F of X, in which case $\Delta^t \subset F$ by Lemma 7.8. Or L is not contained in any flat, in which case Δ^t is not contained in any flat. We will complete the proof by showing that if L is not contained in any flat, then L appears in this way in only finitely many truncated tricombs.

Suppose L appears in this way in infinitely many tricombs $\Delta_1, \Delta_2, \ldots$ ordered by increasing width. Then Lemma 7.8 implies that their truncated deputies $\Delta_1^t \subset \Delta_2^t \subset \Delta_3^t$ are nested, and therefore there union is an infinite truncated deputy Δ_∞^t (which, by local finiteness, is actually contained in an infinite tricomb Δ_∞).

We think of Δ_∞^t as representing $\frac{1}{6}$ of a flat plane, and we will next show that it is actually contained in $\frac{1}{2}$ of a flat plane and finally in an entire periodic flat plane. Consider the hypercarrier Y in Δ_∞^t which contains L. By the pigeon hole principle, it contains two ladders of four two-cells, such that the leftmost one is translated to the rightmost one by some element $g \in G$. Applying Lemma 7.8 several times, we see that $g\Delta_\infty^t \subset \Delta_\infty^t$. Therefore Δ_∞^t is contained in half of a flat plane which is $\cup_{n\geq 1} g^{-n}\Delta_\infty^t$. We then apply the same argument with respect

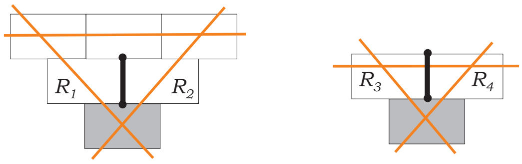

FIGURE 20.

to a hypercarrier that crosses Y to obtain a full flat plane stabilized by a $\mathbb{Z} \times \mathbb{Z}$ subgroup.

□

Theorem 7.10. *Suppose G acts properly discontinuously and cocompactly on a $B(6)$ complex X. The following are equivalent:*

(1) *G is not word-hyperbolic.*
(2) *X contains a flat plane.*
(3) *G contains a subgroup isomorphic to $\mathbb{Z} \times \mathbb{Z}$.*

Proof. If X contains a flat plane then G contains a $\mathbb{Z} \times \mathbb{Z}$ subgroup by Theorem 7.5. If G contains a $\mathbb{Z} \times \mathbb{Z}$ subgroup then it cannot be word-hyperbolic (see for instance, [Aea91]). If G is not word-hyperbolic, then X must contain arbitrarily large tricombs, and any sufficiently large tricomb is contained in a flat plane by Lemma 7.9. Indeed, the usual curvature argument for proving a linear isoperimetric function (see [Ger87, Appendix]) works if there is an upper bound on the radius of a flat region. Thus if there were a bound on the width of a tricomb then X and hence G would have a linear isoperimetric function. □

The special case of Theorem 7.10 where pieces are $\leq \frac{1}{6}$ was proven in [Wis96].

8. PAIRWISE INTERSECTING HYPERGRAPHS IN THE $B(6)$ CASE

In this section we characterize families of pairwise crossing hypergraphs in the $B(6)$ case. The two main results are Theorem 8.3 and Theorem 8.6 which conclude that the hypergraphs meet along a width $w \geq 3$ tricomb or meet along a vertex, according to whether there are three of these hypergraphs which meet along a width $w \geq 3$ or width 2 tricomb. These results will later enable us to classify all the cubes in C, and thus prove that the group acts properly discontinously and cocompactly/cofinitely in the word-hyperbolic/general case.

Lemma 8.1. *Let X be a $B(6)$ complex and let A and B be width 3 and width 2 tricombs along which hypercarriers meet. Suppose that, as illustrated in Figure 20, the shaded bottom corners and bold rungs in A and B are equal by a translation. Then the 2-cells labelled by R_1 and R_3 are equal, and similarly $R_2 = R_4$.*

Proof. The top path in ∂R_1 consist of two pieces, and R_1 and R_3 agree on the path along the identical rung and identical 2-cell. If $R_1 \neq R_3$, then this long path is a piece. Thus there is a path S consisting of the concatenation of three pieces, and since S passes through a pair of antipodal edges dual to the same hypercarrier, we see that $|S| > \frac{1}{2}|\partial R_1|$. Since this would violate the $B(6)$ condition, we conclude that $R_1 = R_3$. An analogous argument shows that $R_2 = R_4$. □

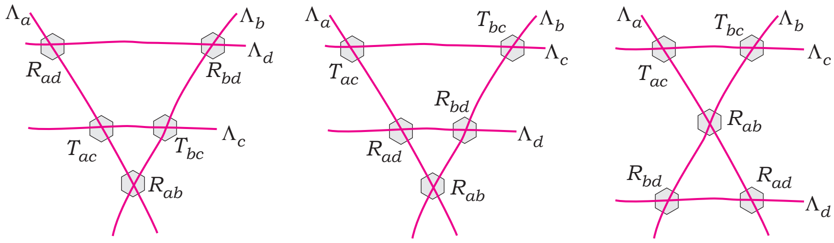

FIGURE 21. Scenarios where Λ_d cannot cross Λ_a, Λ_b, and Λ_c.

Lemma 8.2. *Let Y_a, Y_b, and Y_c be hypercarriers that meet along a nondegenerate tricomb $\Delta(abc)$ in the $B(6)$ complex X. Suppose the hypercarrier Y_d crosses each of these hypercarriers. Then there is a permutation $\{\bar{a}, \bar{b}, \bar{c}\}$ of $\{a, b, c\}$, and a tricomb $\Delta(\bar{a}\bar{b}d)$ such that $Y_{\bar{a}}$, $Y_{\bar{b}}$, and Y_d meet along $\Delta(\bar{a}\bar{b}d)$, and $Y_{\bar{c}}$ borders $\Delta(\bar{a}\bar{b}d)$. Furthermore, $\Delta(abc)$ equals $\Delta(\bar{a}\bar{b}d)$ except perhaps at the corners that Y_d passes through.*

Proof. Throughout the proof we will use the notation R_{ab}, R_{ad}, and R_{bd} for corners of $\Delta(abd)$, and likewise, we will use the notation T_{ab}, T_{bc}, and T_{ac} for corners of $\Delta(abc)$. We may assume without loss of generality, that Y_d does not contain any 2-cell which can replace the 2-cell T_{ab} at the ab-corner of $\Delta(abc)$.

We will begin by ruling out three possible cases illustrated in Figure 21. In the first two of these cases, the ab-corners of $\Delta(abc)$ and $\Delta(abd)$ are the same, so $R_{ab} = T_{ab}$. Furthermore, in the first case $\Delta(abc) \subsetneq \Delta(abd)$, and in the second case $\Delta(abd) \subsetneq \Delta(abc)$. In the third case, $\Delta(abd)$ and $\Delta(abc)$ are completely separated from each other by their ab-corners which may possibly be equal (as they are in Figure 21).

Observe that R_{ad} and R_{bd} cut Y_d into three subcomplexes called *parts*. One of these is bounded by R_{bd}, one is bounded by R_{ad}, and one is bounded by both R_{ad} and R_{bd}. This is clear if we think of the corresponding 1-cells in the hypergraph Λ_d which is a tree. There is also a degenerate possibility where $R_{ad} = R_{bd}$ in which case $\Delta(abd)$ is degenerate, and we shall ignore the part bounded by both 2-cells. We shall adopt similar terminology for the parts of Y_c. By examining the positions of 2-cells in the various parts, we will obtain the contradiction that Y_c and Y_d do not cross.

We first consider case 1. If L is a 2-cell in Y_c, then since $\Delta(abc)$ is nondegenerate, $T_{ad} \neq T_{bd}$, and so L either lies outside the part of Y_c bounded by T_{ac}, or L lies outside the part of Y_c bounded by T_{bc}. Since the latter possibility is analogous, we will show that the former possibility is impossible. Indeed, the geodesic path from L to T_{ac} in Λ_c, followed by the path from T_{ac} to R_{ad} in Λ_a, makes an angle of $\geq \frac{2\pi}{3}$ in T_{ac}. Therefore Corollary 6.6 implies that L cannot lie on Y_c, for otherwise $T_{ac} \subset Y_d$ and so Lemma 6.8 implies that the ladder in $\Delta(abd)$ between R_{ad} and T_{ac} is contained in Y_d which is impossible.

The argument for case 2 is identical except that the roles of Y_c and Y_d are exchanged. Note that $R_{ad} \neq R_{bd}$ in this case, because then they would both equal R_{ab}, and this would contradict our assumption that Y_d does not contain an ab-corner of $\Delta(abc)$.

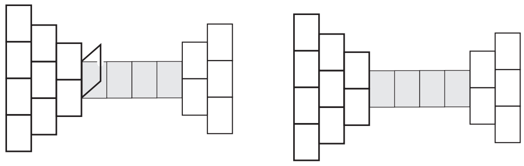

Figure 22.

We now consider case 3. If L is a 2-cell in the part of Y_c that is bounded by T_{ac}, then consideration of the geodesic path from L to T_{ac} in Λ_c, followed by the path from T_{ac} to R_{ad} in Λ_a, makes an angle of $\geq \frac{2\pi}{3}$ in T_{ac}. Therefore Corollary 6.6 implies that L cannot lie on Y_d. For then $T_{ac} \subset Y_d$ and so $R_{ab} \subset Y_d$ since it is in the ladder in Y_a between T_{ac} and R_{ad}. An analogous argument shows that if L is in the part of Y_c bounded by T_{bc}, then $L \not\subset Y_d$.

It follows that L must lie in the intersection of halfcarriers H_a and H_b which correspond to orientations of Λ_a and Λ_b so that they are oriented towards each other in the obvious way in $\Delta(abc)$.

Consider the immersed path P which starts in L, travels in Λ_c to the center of T_{ac}, travels in Λ_a until it reaches R_{ad}, and then travels in Λ_d until it terminates at the center of L. Observe that P passes through the 2-cell T_{ab} exactly once. Furthermore this 2-cell is a cut-cell in $H_a \cap H_b$, and so P is an essential closed path in $H_a \cap H_b$. This yields a contradiction since the intersection of halfcarriers is simply-connected by Corollary 3.21.

We now proceed to the main part of the argument. Choose the tricomb $\Delta(abd)$ to minimize the length of the ladder L in $Y_a \cap Y_b$ whose first 2-cell is T_{ab} and whose last 2-cell is the ab-corner R_{ab} of $\Delta(abd)$. As illustrated in Figure 22, there are essentially two ways for L to be connected to $\Delta(abc)$ (which is drawn boldly on the left in each case).

We first show that L consists of a single 2-cell $T_{ab} = R_{ab}$, and that $\Delta(abc)$ and $\Delta(abd)$ lie along the same side of this 2-cell. Indeed, otherwise, we would have one of the situations illustrated in Figure 23, where the case that $\Delta(abd)$ is nondegenerate is illustrated on the left, and the degenerate case is illustrated on the right, and both these cases correspond to case 3 of Figure 21 which was excluded above. In each case there are an arbitrary number (possibly zero) of white 2-cells lying in L between the two tricombs. In both cases, Y_d cannot contain the very dark 2-cell in L. Indeed, if $\Delta(abd)$ is degenerate, then this very dark 2-cell would be a choice of tricomb closer to $\Delta(abc)$ and this would contradict the minimality of the length of L. (So if $\Delta(abd)$ is degenerate then L consists of at least two 2-cells.) When $\Delta(abd)$ is nondegenerate, Y_d cannot contain the very dark 2-cell because the dark 2-cell is contained in $Y_a \cap Y_b$, and so it cannot be contained in Y_d for then $\Delta(abd)$ would have a degenerate tricomb which was closer to $\Delta(abc)$, and this again contradicts the minimality of L.

Now $\Delta(abc)$ and $\Delta(abd)$ lie on the same side of R_{ab} as illustrated on the left in Figure 22. Starting with R_{ab}, which is a length 1 row, we will regard these tricombs as consisting of parallel rows, $c_1, c_2, \ldots$ in $\Delta(abc)$ and $d_1, d_2, \ldots$ in $\Delta(abd)$ that are parallel to Y_c and Y_d. Note that $c_1 = T_{ab} = R_{ab} = d_1$. Let c_ℓ and d_ℓ denote the last rows in $\Delta(abc)$ and $\Delta(abd)$ which are equal. If there aren't any further

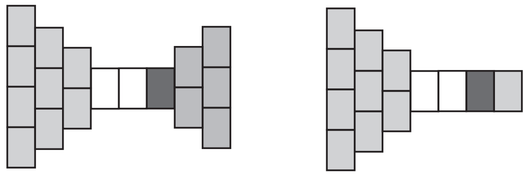

FIGURE 23.

rows, then $\Delta(abc) = \Delta(abd)$ and we are done. If exactly one of these tricombs doesn't have a next row, then we will obtain case 1 or case 2 in Figure 21 that we have excluded above.

Finally, we consider the case where they both have next rows $c_{\ell+1}$ and $d_{\ell+1}$, but $c_{\ell+1} \neq d_{\ell+1}$. If the hypercarriers Y'_c and Y'_d in X that correspond to $c_{\ell+1}$ and $d_{\ell+1}$ do not cross, then Y_c and Y_d are disjoint, since they are separated by either Y'_c or Y'_d (note that possibly $Y_c = Y'_c$ or $Y_d = Y'_d$). This contradicts our assumption that Y_c and Y_d cross. Observe that $c_{\ell+1}$ meets c_ℓ along the path between them in $\Delta(abc)$, and $d_{\ell+1}$ meets d_ℓ along the path between them in $\Delta(abd)$, and these paths are equal to each other since $c_\ell = d_\ell$. Now, since Y'_c and Y'_d cross, an application of Lemma 6.8 implies that $c_{\ell+1} = d_{\ell+1}$ except perhaps at the first and last 2-cells.

If $c_{\ell+1}$ and $d_{\ell+1}$ are the last rows of $\Delta(abc)$ and $\Delta(abd)$ then we are done. If one of them is not the last row, then by applying Lemma 8.1 at each end, we see that $c_{\ell+1} = d_{\ell+1}$ which contradicts our assumption that c_ℓ and d_ℓ are the last equal rows. □

Theorem 8.3. *Suppose X satisfies $B(6)$. Consider a set $\{Y_1, Y_2, Y_3, \dots\}$ of pairwise crossing hypercarriers. If Y_1, Y_2, and Y_3 meet along a tricomb of width $w \geq 3$ then there is a width w tricomb Δ such that Y_i borders Δ for each i.*

Proof. Suppose that Y_1, Y_2, and Y_3, meet along the tricomb $\Delta(123)$. Lemma 6.4 tells us that any three of these hypercarriers meet along a tricomb. For each $k \geq 4$, Lemma 8.2, provides us with a distinct pair $i_k < j_k \in \{1, 2, 3\}$ and a tricomb $\Delta(i_k j_k k)$ which Y_{i_k}, Y_{j_k}, and Y_k meet along, and such that $\Delta(i_k j_k k)$ is identical to $\Delta(123)$ except perhaps at the two corners that Y_k passes through. This partitions the hypercarriers Y_k into three equivalence classes, according to the values of i_k and j_k. The reader should think of this as whether Y_k can replace Y_1, Y_2, or Y_3.

We will now choose a tricomb Δ with the property that each Y_k borders Δ. The tricomb Δ is identical to $\Delta(123)$ and in fact to all $\Delta(i_k j_k k)$ except perhaps at its corners, so we will merely need to choose its corner 2-cells. We describe how to choose a $Y_1 \cap Y_2$ corner 2-cell for Δ, and note that the other corner 2-cells are chosen similarly.

Let R_1 denote the 2-cell in $Y_1 \cap \Delta(123)$ which intersects the 12-corner along a rung, and similarly, let R_2 denote the 2-cell in $Y_2 \cap \Delta(123)$ which intersects the 12-corner along a rung. Let p denote the 0-cell where R_1, R_2, and the 12-corner triply intersect.

Among all hypercarriers equivalent to Y_1, let $Y_{1'}$ denote the hypercarrier whose hypergraph passes through a dual 1-cell which is as far away from p as possible along ∂R_1 (travelling around R_1 in the direction of the path of intersection between R_1 and the 12-corner). Define $Y_{2'}$ analogously, so that its hypergraph passes through a 1-cell which is as far away from q as possible along ∂R_2.

Consider the triangle $\Delta(1'2'3)$, and let C_{12} denote its 12-corner. We choose C_{23} and C_{13} in a manner analogous to the way we chose C_{12}. Let Δ be the tricomb obtained from $\Delta(123)$ by replacing its corners by C_{12}, C_{23}, and C_{31}. To see that each of our hypercarriers borders Δ, observe that if Y is equivalent to Y_1, then its hypergraph Λ must pass through an edge in R_1 which is at least as close to p as the edge that $Y_{1'}$ passes through. Consequently, this edge must lie on C_{12}. The analogous argument works in the other five cases. $\square$

Theorem 8.3 characterizes the collections of pairwise crossing hypergraphs such that some three of these hypergraphs meet along a width ≥ 3 tricomb. Theorem 6.10 characterizes the situation where each three of these hypergraphs triply intersect in a 2-cell. We note that a hypergraph dual to an isolated 1-cell (which does not lie on the boundary of any 2-cell) cannot cross any other hypergraph. The only remaining case to consider is a collection of hypergraphs $\{\Lambda_1, \Lambda_2, \Lambda_3, \dots\}$ such that three of these hypergraphs meet along a width 2 tricomb. In this case we arrive at a conclusion somewhat more complicated than the conclusion of Theorem 8.3. The reason for the discrepancy between the width 2 and width ≥ 3 cases is that in the width 2 case, the tricomb $\Delta(\bar{a}\bar{b}d)$ provided by Lemma 8.2 has the property that $Y_d \cap \Delta(\bar{a}\bar{b}d)$ equals the union of the two corners that Y_d passes through, whereas there are additional 2-cells between these two corners in the width ≥ 3 case.

Definition 8.4 (Meet along a vertex)**.** We say a collection of hypergraphs $\{\Lambda_1, \Lambda_2, \dots\}$ *meet along a vertex* $v \in X^0$ provided that they pairwise cross, and that for each i, Λ_i passes through an edge in a piece P_i of X^1 whose initial vertex is v. A nice example to consider is a 2-complex consisting of six hexagons glued together along a vertex v, so that $\text{link}(v)$ is a tetrahedron.

In general it is possible for $\{\Lambda_1, \Lambda_2, \dots\}$ to meet along more than one vertex in X^0. However, this cannot happen when all pieces have length 1.

Note that if $\Lambda_1, \Lambda_2, \dots$ meet along v and $Y_1, Y_2, \dots$ are the corresponding hypercarriers then $v \in \cap_i Y_i$. However, the converse may not hold since the Y_i may not pairwise cross.

Definition 8.5 ($L(v)$ and hypergraphs meeting along a vertex)**.** Let $v \in X^0$, let $T(v) \subset X^1$ be the subgraph which is the union of all paths starting at v that consist of a single piece. Note that if X satisfies $C(3)$ then $T(v)$ is a tree. We now define a graph $L(v)$ as follows: The vertices of $L(v)$ are in one-to-one correspondence with the centers of 1-cells of $T(v)$, and for each pair of vertices v_1 and v_2 corresponding to 1-cells c_1 and c_2, there is an edge in $L(v)$ joining v_1 and v_2 if and only if c_1 and c_2 lie on the boundary of a 2-cell $R \in X$. Note that $L(v) \cong \text{link}(v)$ if all pieces in X have length 1.

Any complete subgraph of $L(v)$ obviously corresponds to a collection of pairwise intersecting hypergraphs meeting along v. We will show in Theorem 8.6 that the reverse correspondence holds as well.

Let $A \subset L(v)$ be a complete subgraph, and let $\bar{A}$ denote the subtree of $T(v)$ corresponding to A. It will be convenient to distinguish between two cases: If $\bar{A}$ is homeomorphic to a real interval then the corresponding collection of intersecting hypergraphs actually pass through a single 2-cell containing v. The more interesting case is when $\bar{A}$ contains a vertex of valence ≥ 3, and letting Λ_1, Λ_2, and Λ_3 be hypergraphs dual to three distinct edges at this valence ≥ 3 vertex, we see that Λ_1, Λ_2, and Λ_3 meet along a width 2 tricomb. Consequently, Lemma 8.2 implies that any three hypergraphs in the collection meet along a tricomb of width ≤ 2.

Theorem 8.6. *Let $\Lambda_1, \Lambda_2, \Lambda_3, \ldots$ be a collection of pairwise crossing hypergraphs, and suppose that Λ_1, Λ_2, and Λ_3 meet along a width 2 tricomb $\Delta(123)$, and let v be the central vertex of $\Delta(123)$. Then $\Lambda_1, \Lambda_2, \Lambda_3, \ldots$ meet along the vertex v.*

Proof. Specifically, we will show that there is a complete graph $A \subset L(v)$ such that $\Lambda_1, \Lambda_2, \Lambda_3, \ldots$ correspond to the vertices of A.

Let $T'(v)$ denote the subgraph of X^1 that is the union of all paths starting at v which are the concatenation of at most two pieces. That $T'(v)$ is a tree follows from the $C(5)$ condition.

Let Λ_d denote one of the hypergraphs in the collection. We first show that Λ_d passes through an edge e_d in $T'(v)$. Indeed, the central vertex v appears in the interior of the tricomb $\Delta(\bar{1}\bar{2}d)$ provided by Lemma 8.2, and therefore Λ_d passes through a 1-cell e_d in the interior of the tricomb $\Delta(\bar{1}\bar{2}d)$ where e_d lies on a path starting at v that is the concatenation of at most two pieces.

We will now show that each e_d is contained in $T(v)$. Using the fact that $T'(v)$ is a tree that branches at v, choose $i, j \in \{1, 2, 3\}$ such that the geodesic from v to e_d does not pass through the interior of the geodesics from v to the edges e_i and e_j. Now consider the tricomb $\Delta(ijd)$, and observe that the central vertex of $\Delta(ijd)$ must be v, for otherwise the above choice is contradicted. Since v is the central vertex of $\Delta(ijd)$ we see that there is a path from v to e_d consisting of a single piece.

Finally, let e_c and e_d be a pair of edges in $T(v)$ corresponding to Λ_c and Λ_d. We will show that e_c and e_d lie on the boundary of the same 2-cell, and therefore the corresponding vertices are joined by an edge in the graph Γ. Choose $i \in \{1, 2, 3\}$ such that the geodesics from v to e_c and e_d do not pass through the interior of the geodesic from v to e_i. Then the hypergraphs Λ_i, Λ_c, and Λ_d meet along a width 2 tricomb $\Delta(icd)$ and so e_c and e_d lie on the boundary of the 2-cell at the cd-corner. □

9. The maximal cubes in C

Definition 9.1 (Basic cubes)**.** Let $\hat{C}$ be the cube complex arising from a $B(4)$-$T(4)$ or $B(6)$ complex X. We now describe the *basic cubes* which are the cubes in $\hat{C}$ corresponding to collections of hypergraphs associated with either an isolated 1-cell, a 2-cell, a complete graph in the link of a vertex, or a tricomb.

The basic cube corresponding to an isolated 1-cell is a 1-cube whose independent hypergraph is the hypergraph dual to the isolated 1-cell. All other hypergraphs are oriented towards the isolated 1-cell.

The basic cube corresponding to a 2-cell, has the hypergraphs passing through this 2-cell as its independent set, and has all other hypergraphs oriented towards the 2-cell.

The remaining two types of basic cubes arise only in the $B(6)$ case.

The basic cube corresponding to a collection of pairwise crossing hypergraphs meeting along a vertex v, has these as its independent hypergraphs and has all other hypergraphs oriented towards v.

The basic cube corresponding to a width ≥ 3 tricomb, has its border hypergraphs as independent, and has every other hypergraph oriented towards those border hypergraphs that it doesn't intersect.

We shall refer to the four types of basic cubes as *edge-cubes*, *cell-cubes*, *link-cubes*, and *tricomb-cubes*. We say a tricomb-cube has *width* w if its associated tricomb has width w. A link-cube has *width* 2 if three of its independent hypergraphs intersect in a width 2 tricomb, or equivalently, if it contains a width 2 tricomb-cube.

Lemma 9.2. *Every edge-cube and cell-cube is maximal. Every [width* 2 *link-cube] tricomb-cube of width* $w \geq 3$ *is contained in a maximal cube which is also a [width* 2 *link-cube] width* w *tricomb-cube.*

Proof. The maximality of edge-cubes follows immediately from the definition of $\hat{C}$. The maximality of cell-cubes follows from Theorem 6.9, together with the observation that two distinct 2-cells cannot have the same boundary 1-cells unless they have the same attaching map (in which case they are identical anyway by Convention 3.1). Similarly, the tricomb-cube statement follows from Theorem 8.3. Finally, the link-cube statement follows from Theorem 8.6 by choosing a maximal complete subgraph of $L(v)$. □

Lemma 9.3. *Every basic cube lies in* C*.*

Proof. In the $B(4)$-$T(4)$ case this is obvious, since one of the 0-cubes on an edge-cube or cell-cube corresponding to a 1-cell or 2-cell incident with the 0-cell x, is V_x.

We now consider the additional $B(6)$ cases. A width 2 link-cube centered at a vertex x actually lies near V_x. Indeed, beginning with a 0-cube on this link-cube, we can arrive at V_x by reversing the orientations on its independent hypergraphs so that they all point towards x. Therefore V_x is actually one of the 0-cubes on this link-cube.

Finally, for a tricomb-cube, we choose one side of the tricomb, and repeatedly change orientations on borders, until after passing through cubes corresponding to successively smaller tricombs, we eventually reach a tricomb consisting of a single 2-cell, in which case the above argument holds. □

Finally, the following shows that there are no other cubes besides those contained in basic cubes, so the complex C may be constructed by gluing together basic cubes.

Lemma 9.4. *Every cube in* C *is a subcube of a basic cube.*

Proof. We first treat the $B(4)$-$T(4)$ case. Consider an n-cube c with independent set of hypergraphs $\Lambda_1, \ldots, \Lambda_n$. Let f be a 2-cell that they pass through. Because

c lies in C, all but finitely many hypergraphs are oriented towards f. Let S denote the collection of oriented hypergraphs which are oriented away from f. Any hypergraph Λ of c in S, must cross each of $\Lambda_1, \ldots, \Lambda_n$. Indeed, if it was disjoint from Λ_i then it would point away from Λ_i since it points away from f. But then Λ_i could not be independent.

Let $\Lambda'_1, \ldots, \Lambda'_m$ be a set of oriented hypergraphs in S, which is maximal subject to the condition that they pairwise cross. Let f' denote a 2-cell such that $\Lambda_1, \ldots, \Lambda_n, \Lambda'_1, \ldots, \Lambda'_m$ all pass through f'. Then c is a subcube of the cube corresponding to f'. Indeed, every dependent hypergraph of c is either disjoint from some Λ_i in which case it is oriented towards it and hence towards f, or else, it crosses each Λ_i in which case by maximality, it is eihter equal to or disjoint from some Λ'_i. If it equals Λ'_i then it is oriented towards f' since it passes through f', and otherwise it is oriented towards Λ'_i and thus towards f'.

We now treat the $B(6)$ case. Suppose that c is a cube whose independent hypergraphs $\Lambda_1, \ldots, \Lambda_n$ meet along a width ≥ 3 tricomb Δ. Let S denote the finitely many dependent hypergraphs of c that are oriented away from some vertex in Δ. As in the $B(4)$-$T(4)$ case, any such hypergraph must cross each of $\Lambda_1, \ldots, \Lambda_n$. Now we choose $\Lambda'_1, \ldots, \Lambda'_m \subset S$ to be maximal subject to pairwise crossing. They therefore pairwise cross as above. We can therefore use the tricomb corresponding to $\Lambda_1, \ldots, \Lambda_n, \Lambda'_1, \ldots, \Lambda'_m$, and c is a subcube of the cube corresponding to this tricomb. An analogous argument holds to show that if the independent hypergraphs of c meet along a vertex then c is a subcube of a link-cube. □

10. Local finiteness and finite stabilizers

In this section we work under the hypothesis that the simply-connected complex X is either $B(6)$ or $B(4)$-$T(4)$.

Definition 10.1. For a vertex $v \in C$, we define $K_v \subset X$ to be the intersection of halfcarriers in X corresponding to the oriented hypergraphs of v. So

$$K_v = \cap_{H_e \in v} H_e.$$

Lemma 10.2. *Let $v \in C$. The 2-cells in $K_v \subset X$ are in one-to-one correspondence with the cell-cubes containing v. Each edge-cube containing v gives rise to a unique isolated 1-cell of X in K_v.*

Proof. Let c be a maximal n-cube in C that contains v. Suppose that c is not a (nondegenerate) tricomb-cube, and let f be the 2-cell or isolated 1-cell in X corresponding to c. Let $\Lambda_1, \ldots, \Lambda_n$ be the independent hypergraphs of c, and note that each Λ_i is minimal in v. Consider an arbitrary hypergraph Λ_e. The minimality of the independent hypergraphs of c ensures that either Λ_e crosses each Λ_i, or Λ_e is disjoint from some Λ_i but is oriented towards it. In the first case, the maximality of c implies that Λ_e equals Λ_i for some i, and so $f \subset H_e$ where H_e is the hypercarrier corresponding to the oriented hypergraph Λ_e. In the second case, $f \subset H_i \subset H_e$ where H_i denotes the oriented hypergraph corresponding to Λ_i.

Now let f be a 2-cell in K_v. For each oriented hypergraph Λ_e in v, if Λ_e passes through the center of f then Λ_e is minimal in v. Indeed each other hypergraph Λ_d is oriented towards $f \subset K_v$ and thus towards Λ_e. Observe that v is incident with

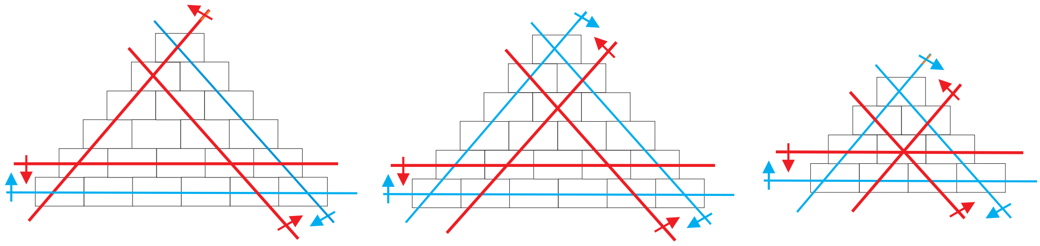

FIGURE 24. In each of the three cases above, the tricomb Δ contains a tricomb Δ' whose hypergraphs are oriented outwards relative to Δ.

the cell-cube c corresponding to f, since its independent hypergraphs are minimal in v.

Finally, we note that distinct edge cubes in C have distinct hypergraphs, which therefore correspond to distinct edges in X. □

I do not know whether isolated edges in K_v are in one-to-one correspondence with edge-cubes containing v, in the $B(4)$-$T(4)$ case. However, this fails in the $B(6)$ case, where isolated edges can arise in K_v, where v is a certain vertex on a width 2 tricomb-cube. Nevertheless, while these edges are isolated in K_v they are not isolated in X.

Lemma 10.3. *Let X be a $B(6)$ complex. Let $v \in C$ be a vertex which lies on a maximal cube c, and suppose that c is a link-cube corresponding to a collection of crossing hypergraphs meeting along x, or let c be a width 3 tricomb-cube corresponding to a tricomb Δ whose central vertex is x. Then $x \in K_v$.*

Proof. We adopt the notation employed in the proof of Lemma 10.2. Let Λ_e be a directed hypergraph in v. If Λ_e intersects each of $\Lambda_1, \ldots, \Lambda_n$, then the maximality of c implies that $\Lambda_e = \Lambda_i$ for some i, and so $x \in Y_e \subset H_e$. Otherwise, Λ_e is nested with some Λ_i, in which case $x \in Y_i \subset H_e$. □

Lemma 10.4. *Let X be $B(6)$. Let $v \in C$ be a vertex that lies on a maximal cube c that is a tricomb-cube of width $w \geq 4$ corresponding to a tricomb Δ. There exists a tricomb $\Delta' \subset \Delta$, such that every edge e at v corresponds to a hypergraph Λ_e that borders a tricomb $\Delta(ejk)$ of width $\leq w+1$ such that $\Delta' \subseteq \Delta(ejk)$.*

In particular, if X is locally finite, then C is locally finite at v.

Proof. For $1 \leq i \leq 3$, consider the i-th side of Δ. Either some hypergraph bordering the i-th side is directed outwards in which case we let Λ_i denote this hypergraph. Or every hypergraph bordering the i-th side is directed inwards towards the remainder of Δ. In this case every hypergraph passing through Δ exactly one layer inwards from the i-th side is directed outwards. We then let Λ_i denote one such hypergraph. Let $\Delta' \subset \Delta$ be the tricomb along which Λ_1, Λ_2, and Λ_3 meet. Several scenarios are illustrated in Figure 24. In each case the tricomb Δ' is obtained from Δ by substituting the bold outwards pointing hypergraphs for some (or all) of the hypergraphs meeting along Δ. The last case indicates the degenerate situation arising when Δ' is of width 1. In all cases, the orientations of the hypergraphs are indicated by arrows passing through them transversely.

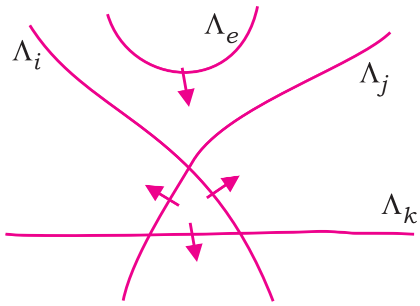

FIGURE 25. If Λ_e doesn't cross Λ_i then it must cross both Λ_j and Λ_k.

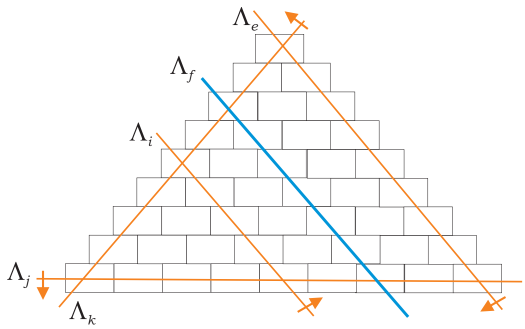

FIGURE 26.

Let Λ_e be a directed hypergraph that is minimal in v. Suppose that Λ_e does not cross Λ_i where $\{i, j, k\}$ is a permutation of $\{1, 2, 3\}$. We will show that Λ_e, Λ_j, and Λ_k meet along a tricomb $\Delta(ejk)$ which contains Δ'.

We first show that Λ_e must cross Λ_j for otherwise as suggested by Figure 25, Λ_e could not cross Λ_k. Indeed, the minimality of Λ_e guarantees that Λ_i and Λ_j are oriented towards Λ_e. If Δ' is nondegenerate then Λ_i and Λ_j are oriented towards disjoint parts of Λ_k, and so Λ_e cannot cross Λ_k. In the case where Δ' is degenerate, the only point of Λ_k which Λ_i and Λ_j are both oriented towards is the center of Δ', and Λ_e cannot pass through this point since we have assumed it does not cross Λ_i. Since Λ_k and Λ_e do not cross, we see that $H_k \subsetneq H_e$ which contradicts the minimality of Λ_e. The same argument shows that Λ_e must also cross Λ_k.

Lemma 6.4 shows that Λ_e, Λ_j and Λ_k meet along a tricomb $\Delta(ejk)$, and we may assume that $\Delta(ejk)$ and Δ' have the same jk-corner. Since both $Y_e \cap Y_j$ and $Y_e \cap Y_k$ are contained in H_i, but the jk-corner of $\Delta(ejk)$ is not contained in H_i, we see that the ladders in $\Delta(ejk)$ from the jk-corner to the ek-corner and ej-corners pass through Y_i. Therefore $\Delta' \subset \Delta(ejk)$. Note that when Δ' is degenerate, the position of Λ_i, Λ_j, and Λ_k in Δ ensures that $Y_i \cap Y_j \cap Y_k$ consists of a single 2-cell.

If the difference between the widths of $\Delta(ejk)$ and Δ' is > 1, then there is a hypergraph Λ_f lying between Λ_e and Λ_i and crossing neither of them, and which meets with Λ_j and Λ_k along a tricomb contained in $\Delta(ejk)$, (see Figure 26). If Λ_f is oriented towards Λ_e, then the minimality of Λ_i is contradicted, and if Λ_f is oriented towards Λ_i, then the minimality of Λ_e is contradicted.

It follows that the difference between the widths of $\Delta(ejk)$ and Δ' is either zero or one. In the former case, Λ_e passes through Λ_i itself, and in the latter case Λ_e passes parallel to Λ_i but through a 2-cell just outside Δ'. The local finiteness at v follows because the number of possible hypercarriers which travel just outside Δ' or actually border Δ' is bounded by the number of 2-cells incident with fixed 0-cells x_i, x_j, x_k contained in the three sides of $\partial\Delta'$. □

Theorem 10.5. *If X is locally finite then C is locally finite.*

Proof. By the definition of C, for each vertex v in C, there is $x \in X^0$ such that v and V_x are connected by a path in $\hat{C}$, and so v and V_x differ on the orientations of finitely many hypergraphs. Let $\Lambda_1, \ldots, \Lambda_n$ denote these hypergraphs, let $Y_1, \ldots, Y_n$ denote their hypercarriers and let $H_1, \ldots, H_n$ denote the corresponding halfcarriers in v. For $1 \leq i \leq n$, choose $x_i \in H_i - Y_i$. Let W be a compact connected set containing $\{x_1, \ldots, x_n\} \cup \{x\}$.

We claim that $K_v \subset \text{Hyperclosure}(W)$. For any hypergraph Λ which is disjoint from W, the orientation of Λ towards W is the same as the orientation of Λ in v. Indeed, $\Lambda \neq \Lambda_i$ since $\Lambda_i \cap W \neq \emptyset$. Otherwise, the orientation of Λ in v, is the same as the orientation of Λ in V_x, which is the orientation of Λ towards x, which is the same as the orientation of Λ towards $W \ni x$. Thus the collection of halfcarriers of which $\text{Hyperclosure}(W)$ is the intersection, is a proper subcollection of the collection of halfcarriers whose intersection defines K_v.

Observe that K_v is compact since $K_v \subset \text{Hyperclosure}(W)$ and $\text{Hyperclosure}(W)$ is compact by Theorem 3.23. Finally, by Lemma 10.2, the cell-cubes at v are in one-to-one correspondence with the 2-cells in K_v, and the edge-cubes at v give rise to distinct isolated 1-cells in K_v. In the $B(4)$-$T(4)$ case, this completes the proof because there are finitely many cells in K_v, and so C is locally finite at v as claimed.

For the $B(6)$-$C(7)$ case, Lemma 10.3 asserts that for any width 3 maximal tricomb-cube corresponding to a tricomb Δ, the central 0-cell of Δ lies in K_v. Likewise, for any maximal link-cube its corresponding central 0-cell lies in K_v.

The local finiteness of X shows that finitely many width 3 tricombs have the same central 0-cell, and finitely many width 2 link-cubes are centered at the same 0-cell. Therefore, in addition to the finitely many cell-cubes and edge-cubes at v which we determined above, there are finitely many link-cubes and width 3 tricomb-cubes at v.

Finally, for the general $B(6)$ case, we already showed in Lemma 10.4 that C is locally finite if v lies on a cube corresponding to a tricomb of width ≥ 4 (in which case K_v is empty). Therefore C is locally finite at v as claimed. □

Lemma 10.6. *The homomorphism $Aut(X) \to Aut(C)$ is injective.*

Proof of Lemma 10.6. Let $g \in Aut(X)$, and suppose that $gx \neq x$ for some $x \in X^0$. Then by Lemma 4.3, there is a hypergraph Λ which separates x and gx. Therefore $gV_x = V_{gx} \neq V_x$. Consequently g maps to a nontrivial element of $Aut(C)$.

Suppose that g fixes each 0-cell of X. By Convention 3.1, X has no 2-cells with the same boundary, and so if g is nontrivial there is an edge e such that e and $f = ge$ have the same endpoints. But then ef^{-1} is a closed immersed path in X, and hence bounds a disc diagram D. An application of Theorem 2.7 shows that

D consists of a single 2-cell, or a pair of 2-cells meeting along a single edge. But then the $B(4)$ condition shows that either e or f is a free face of a bigon, which contradicts Convention 3.1. □

Corollary 10.7. *Let X be locally finite. If the action of G on X has finite cell stabilizers, then the action of G on C has finite cell stabilizers. If the action of G on X has no global fixed point, then the action of G on C has no global fixed point.*

Proof. If some cell c in C has infinite stabilizer, then since $\mathrm{Aut}(c)$ is finite, Stabilizer(v) is infinite for some vertex v on c. By definition of C, there is a path from v to V_x for some $x \in X^0$. By Theorem 10.5, C is locally finite, and so Stabilizer(v) and Stabilizer(V_x) are commensurable. But Stabilizer(V_x) equals Stabilizer(x) which is finite.

Let $\Lambda_1, \Lambda_2, \ldots, \Lambda_n$ be the independent hypergraphs of c. In the $B(4)$-$T(4)$ case, $\cap_1^n Y_i$ is a nonempty convex subcomplex of Y_1. So G acts on this intersection, which is essentially a tree, and so $\bar{G}$ and thus G has a global fixed point. In the $B(6)$ case, $Y_1, Y_2, \ldots, Y_n$ bound a collection of tricombs which have the same central point, which is thus stabilized by G. □

11. G acts properly-discontinuously and cocompactly on C in $B(4)$-$T(4)$ case

In this section we observe that the ingredients thus proven can be assembled to prove our main result in the $B(4)$-$T(4)$ or even the $B(6)$-$C(7)$ case. We note that while the $C(7)$ condition implies word-hyperbolicity, the $B(4)$-$T(4)$ does not.

Theorem 11.1. *If G acts properly discontinuously and cocompactly on a $B(6)$-$C(7)$ or $B(4)$-$T(4)$ complex X, then G acts properly discontinuously and cocompactly on the CAT(0) cube complex C.*

Proof. Since G acts cocompactly on X, there are finitely many orbits of cell-cubes and edge-cubes in C, and for the $B(6)$-$C(7)$ case there are finitely many width 3 tricombs and finitely many width 2 link-cubes. Since any cube is a subcube of such a cube, by Lemma 9.4, we see that G acts cocompactly on C. Since X is locally finite, G acts with finite stabilizers by Corollary 10.7. □

12. G acts cofinitely on C in $B(6)$ case

Definition 12.1 (cofinite and cusps)**.** Let A be a group that acts properly discontinuously and cocompactly on the Euclidean plane $\mathbb{E}^2$. We say Q is an *A-cusp* if A acts on Q and there is an equivariant quasi-isometry $Q \to \mathbb{E}^2 \times \mathbb{R}$ where A acts by isometries on $\mathbb{E}^2 \times \mathbb{R}$ and $g(e, 0) = (ge, 0)$ for all $g \in A$ and $(e, 0) \in \mathbb{E}^2 \times \mathbb{R}$. Note that $g(e, r)$ may be equal to either (ge, r) or $(ge, -r)$.

We say the action of G on C is *cofinite* if for $1 \leq i \leq n$ there are subgroups $A_i \subset G$ and subspaces $Q_i \subset C$ such that Q_i is an A_i-cusp; and there is a compact set K such that C is the union of translates of K and the Q_i by elements of G; and the intersection $g_i Q_i \cap g_j Q_j$ of any two translates of cusps is contained in the union of translates of K.

The word 'cusp' is normally used to suggest something compact$\times[0, \infty)$ but compact$\times(-\infty, \infty)$ is natural in our context.

Theorem 12.2. *If G acts properly discontinuously and cocompactly on a $B(6)$ complex X, then G acts properly discontinuously and cofinitely on the locally finite CAT(0) cube complex C. Furthermore, the conjugacy classes of maximal virtually $\mathbb{Z} \times \mathbb{Z}$ subgroups of G are represented by the cusp subgroups.*

The proof of Theorem 12.2 will be deferred until after our analysis of the cusps arising from flats. However, we shall note the following important corollary:

Corollary 12.3. *If X is $B(6)$ and G is word-hyperbolic then G acts properly discontinuously and cocompactly on the CAT(0) cube complex C.*

Proof. Theorem 12.2 gives a cofinite action of G on C. When G is word-hyperbolic, X cannot contain a $\mathbb{Z} \times \mathbb{Z}$ subgroup. Therefore, the G action on C has no cusps, and is thus cocompact.

Note that we could argue more directly as follows: If G is word-hyperbolic then as in Theorem 7.10, there is a bound on the width of tricombs. Consequently there are finitely many orbits of tricombs in X, and hence finitely many orbits of cubes in C. □

Example 12.4. We list some examples of $B(6)$ presentations containing various flat planes in the universal covers of their standard 2-complexes. The flat planes are illustrated from left to right in Figure 27.

1. $\langle a, b, c \mid abca^{-1}b^{-1}c^{-1}\rangle$ contains a flat with stabilizer $\mathbb{Z} \times \mathbb{Z}$.
2. $\langle a, b, c \mid aabccb^{-1}\rangle$ contains a flat with stabilizer isomorphic to the fundamental group of the klein bottle.
3. $\langle a, b, c \mid cbabca^{-1}\rangle$ again contains a flat which is the universal cover of an immersed klein bottle.
4. $\langle a, b, c, d, e, f \mid (abc)^2, (ade)^2, (bfd^{-1})^2, (fec^{-1})^2\rangle$ contains a flat whose stabilizer is the fundamental group of the pillow case, that is, the sphere with four cone points with angle π.

A simple way to construct more complex examples is to form the based union of several copies of the above complexes together with a bouquet of circles and form their based union, and then add some 2-cells so that the $B(6)$ condition is still satisfied. For instance, the universal cover X of the standard 2-complex of the presentation below has two flat planes stabilized by klein bottle groups. According to the proof of Theorem 12.2, the quotient of the cube complex $C(X)$ by the group G of the presentation, is quasi-isometric to the union of $(-\infty, \infty)$ and $[0, \infty)$ identified along $\{0\}$. This is because the homomorphism $A \to \{\pm 1\}$ discussed prior to Lemma 12.7, is trivial for the stabilizer of one flat, but nontrivial for the other.

$$\langle a, b, c, x, y, z \mid cbabca^{-1}, xxyzzy^{-1}, axbycz, (ay)^3\rangle \tag{3}$$

Lemma 12.5. *Let F be a flat plane in the $B(6)$ complex X, and let $A = Stabilizer(F)$. Then there is an A-equivariant isometric embedding $C(F) \hookrightarrow C(X)$.*

Proof. We now describe the natural embedding $C(F) \hookrightarrow C(X)$. First observe that each hypergraph Λ^+ of X is either disjoint from F, or intersects F in a hypergraph

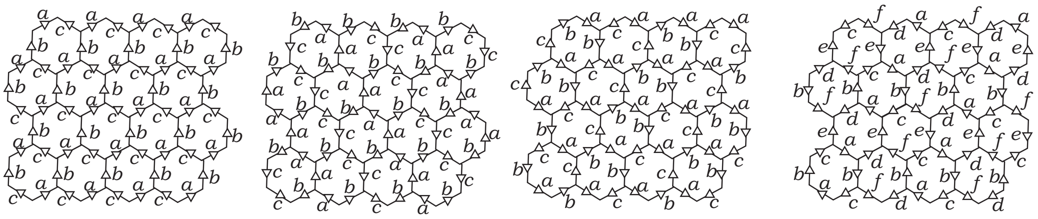

FIGURE 27. Some flat planes in $B(6)$-complexes.

Λ. Any vertex in $C(F)$ is a choice of orientations on the hypergraphs of F, and this determines a vertex in $C(X)$ by orienting each hypergraph Λ^+ consistently with Λ if Λ^+ intersects F, and orienting Λ^+ towards F if Λ^+ is disjoint from F. Likewise, cubes in $C(F)$ are sets of independent pairwise crossing hypergraphs $\Lambda_1, \dots, \Lambda_n$ which map to cubes in $C(X)$ corresponding to $\Lambda_1^+, \dots, \Lambda_n^+$ in an analogous manner.

The image of $C(F)$ is a convex subspace, since it is the intersection of 'real' halfspaces in $C(X)$. Specifically, let e be an edge in $C(X)$ with exactly one vertex on $C(F)$. Let P_e be the hyperplane in $C(X)$ dual to e. Then P_e is disjoint from $C(F)$ for otherwise there would be a geodesic triangle with three $\frac{\pi}{2}$ angles. For $\alpha \in (0,1) = \text{Int}(e)$, let $H_{e,\alpha}$ be the halfspace in $C(X)$ containing $C(F)$ and bounded by the hyperplane parallel to P_e that passes through $\alpha \in e$ (so P_e itself corresponds to $\alpha = \frac{1}{2}$). Then $C(F) = \cap_{e,\alpha} H_{e,\alpha}$ where e and α vary as above. Thus $C(F)$ is convex since it is the intersection of convex subspaces of a CAT(0) space.

Alternatively, we can prove that the map is a local and hence global isometry by observing that two edges of $C(F)$ which meet along the corner of a square in $C(X)$ actually meet along a square in $C(F)$. Indeed, if two hypergraphs Λ_1 and Λ_2 in F do not cross, then their extensions in X do not cross either. This follows from Corollary 6.6 by considering a geodesic triangle connecting centers of 2-cells, which has legs in Λ_1 and Λ_2, and a base in a hypergraph Λ_3 in F that crosses both of these hypergraphs, and that makes a $\frac{2\pi}{3}$ angle on one side. □

Lemma 12.6. *Let A act properly discontinuously and cocompactly on a $B(6)$ flat plane F, and let E be the corresponding $B(6)$ flat plane with all valence 2 0-cells removed. The homeomorphism $F \to E$ induces an A-equivariant quasi-isometry $C(F) \to C(E)$.*

Proof. The hypergraphs in F are partitioned into equivalence classes according to their associated hypercarriers. There is thus a map from these hypergraphs to the corresponding hypergraphs in E.

We choose a fixed 'positive' orientation for each of the three parallelism classes of hypergraphs in E. Now we define a map $C(F)^0 \to C(E)^0$ by mapping each vertex f in $C(F)$ to a vertex e in $C(E)$ as follows: Consider a hypergraph Λ in E, and let $\Lambda_1, \dots, \Lambda_n$ be its preimages in F. If the vertex f has $\Lambda_1, \dots, \Lambda_n$ all oriented in the same direction, then the vertex e has Λ oriented in this direction as well. Otherwise, e has Λ oriented in the positive direction.

The map has a right inverse $C(E)^0 \to C(F)^0$ which sends each vertex e of $C(E)$ to the vertex e' of $C(F)$ whose hypergraphs are oriented the way their images are oriented in e.

To see that this is a quasi-isometry, first note that $C(F)^0 \to C(E)^0$ is obviously surjective. The map $C(F)^0 \to C(E)^0$ is distance non-increasing since vertices in $C(F)$ joined by an edge, are either mapped to the same vertex in $C(E)$ or are mapped to vertices joined by an edge. The map $C(E)^0 \to C(F)^0$ expands distance by a factor of at most $M/2$ where M is a bound on the number of sides of any 2-cell of X. Indeed, the edge between two vertices in $C(E)^0$ corresponds to changing the orientation on exactly one hypergraph. There is a path between their images in $C(F)^0$ which corresponds to successively changing orientations on the preimages of this hypergraph, and there are at most $M/2$ of these. □

We now describe a homomorphism which will arise in Lemma 12.7. Let E denote the 2-complex consisting of the regular hexagonal tiling of the Euclidean plane. Observe that there are two equivalence classes, 'left' and 'right', of nondegenerate tricombs in E, where two tricombs are equivalent, if their respective corners point in the same directions. For a group A acting on E, there is thus a homomorphism $A \to \{+1, -1\}$ denoted by $a \mapsto \bar{a}$ where $\bar{a} = \pm 1$ according to whether a fixes or interchanges these equivalence classes. For the flat planes illustrated in Figure 27, the homomorphism $A \to \{+1, -1\}$ is trivial in the first and third cases, and is nontrivial in the second and fourth cases.

Lemma 12.7. *Let E be the 2-complex consisting of the regular tiling of the Euclidean plane by hexagons, and let A be a group acting on E. Then $C(E)$ is isomorphic to the universal cover of the 3-torus $T^3 = S^1 \times S^1 \times S^1$. Furthermore, $C(E)$ is A-equivariantly isometric to $E \times \mathbb{R}$ where for $a \in A$ the action is defined by $a(e,t) = \big(ae, \bar{a}t\big)$.*

Proof. The universal cover $\widetilde{T}^3$ can be identified with $\mathbb{R}^3$, and its 3-cubes are the translates of the unit cube I^3 in $\mathbb{R}^3$ by elements of $\mathbb{Z}^3$. Accordingly, the hypercarriers in $\widetilde{T}^3$ correspond to translates of the planes $x = \frac{1}{2}$, $y = \frac{1}{2}$ and $z = \frac{1}{2}$.

Consider the plane D in $\mathbb{R}^3$ given by the equation $x + y + z = 0$. The set of lines of intersection of the hyperplanes with D is identical with the set of hypergraphs in the plane E. Furthermore, hypergraphs in E intersect if and only if the corresponding hyperplanes in $\widetilde{T}^3$ intersect.

Moreover, there is a one-to-one correspondence between choices of orientations on hypergraphs in D that yield vertices in $C(D)$, and choices of orientations on hyperplanes in $\widetilde{T}^3$ so that all hyperplanes are oriented towards some vertex in $\widetilde{T}^3$. This induces the isomorphism $C(E) = C(D) \to \widetilde{T}^3$.

Finally, the A-action on E determines an A-action on D. This extends uniquely to an action on $D \times \mathbb{R}$ that preserves the cell structure of $\widetilde{T}^3$. Note that a vertex lies on the left or right of D according to whether the corresponding tricomb in D is left or right. This explains why the corrective homomorphism $a \to \bar{a}$ is employed.

□

Proof of Theorem 12.2. Since X is locally finite, G acts with finite stabilizers by Corollary 10.7. By Theorem 7.5, there are finitely many orbits of flat planes, and let $F_1, \ldots, F_n$ denote representatives of these orbits. For each i, let $A_i = \text{Stabilizer}(F_i)$. By Lemma 7.9, there exists a number W such that for any tricomb

Δ of width $\geq W$, its truncated deputy lies in a flat plane. For use below, we shall also assume that $W \geq 10$.

Since X is locally finite, and since G acts cocompactly on X, there are finitely many orbits of isolated edges, 2-cells, complete subgraphs in $L(v)$ and tricombs of each width. Let K denote a compact connected subcomplex of C which contains representatives of each orbit of edge-cube, cell-cube, link-cube, and width $\leq W$ tricomb-cube in C. For each i let Q_i be the component of $C(F_i) \cup (C - GK)$ containing $C(F_i)$, where $GK = \cup_{g \in G}(gK)$ and $C(F_i)$ is the subcomplex described in Lemma 12.5.

Let v be a vertex in $gQ_i - GK$ and let Δ be the tricomb corresponding to the minimal width tricomb-cube containing v. We will now show that the truncated deputy Δ^t of Δ lies in gF_i. Note that by definition of Q_i, the vertices in $gQ_i \cap GK$ already lie in $gC(F_i) = C(gF_i)$, and so they are obviously contained in maximal cubes corresponding to (possibly degenerate) tricombs in F_i. Let v_1 be a vertex in gQ_i and let P be a shortest path in gQ_i which starts at v_1 and ends on a vertex of $C(gF_i)$. Note that only the last vertex of P can lie on GK, for since gQ_i is a component of $C(gF_i) \cup (C - GK)$, the moment P enters GK it must lie in $C(gF_i)$ and so terminates. Let $v_1, \dots, v_q$ be the vertices of P, and for $1 \leq \ell \leq q$, let Δ_ℓ be a tricomb corresponding to the smallest width tricomb-cube containing v_ℓ. Note that for each ℓ, the width of Δ_ℓ is $\geq W$. For each ℓ the truncated deputy Δ^t_ℓ of the tricomb Δ_ℓ lies in a flat plane. Furthermore these flat planes are identical since consecutive truncated deputies share a chain of four 2-cells, and these determine the flat plane uniquely by Lemma 7.4. Finally, since Δ^t_n lies in gF_i we see that Δ^t_1 also lies in gF_i.

We now show that $(gQ_i \cap hQ_j) \subset GK$ unless $gF_i = hF_j$. Indeed, let v be a vertex in $gQ_i \cap hQ_j$ which does not lie in GK. Let Δ be a minimum width tricomb whose associated tricomb-cube contains v. Then as shown above, the truncated deputy Δ^t of Δ is contained in gF_i and in hF_j. But by Lemma 7.4, Δ^t is contained in a unique flat plane and so $gF_i = hF_j$.

We now show that Q_i is contained in a $3M$ neighborhood of $C(F_i)$, where M is an upper bound on the number of sides in any 2-cell of X. As shown above, a vertex v in Q_i either already lies in $C(F_i)$ or lies on a tricomb-cube corresponding to a tricomb Δ of width $> W$ whose truncated deputy Δ^t is contained in F_i. We will show that there is a path P of length $\leq 3M$ in X from v to $C(F_i)$. The path P begins with a sequence of edges that correspond to successively changing orientations of hypergraphs on the three borders of Δ so that they all point inwards. We arrive at a vertex on the tricomb-cube corresponding to the deputy Δ' of Δ. Repeating this procedure, we arrive at a vertex which lies on the tricomb-cube Δ'' corresponding to the deputy of Δ'. Note that Δ'' is nonempty since its width is $> W - 6 \geq 4$, and $\Delta'' \subset \Delta^t \subset F_i$ so the final vertex lies on a cube in $C(F_i)$. Since each border of a tricomb contains at most $M/2$ hypergraphs, and since we have dealt with at most 6 borders, the path P has length $\leq 3M$, and we are done.

By Lemma 12.5, $C(F_i) \subset C$ is an A_i-equivariant isometric embedding for each i. Since Q_i lies in a bounded neighborhood of $C(F_i)$, the map $C(F_i) \to Q_i$ is quasi-surjective and hence a quasi-isometry. Moreover, it is A_i-equivariant because A_i stabilizes Q_i since it is the component of $C(F_i) \cup (X - GK)$ containing $C(F_i)$.

By Lemma 12.6, there is an A_i-equivariant quasi-isometry $C(F_i) \to C(E)$. Finally, by Lemma 12.7, there is an A_i-equivariant isometry $C(E) \to E \times \mathbb{R}$ such that the A_i action on $E \times \mathbb{R}$ stabilizes $E \times \{0\}$. It follows that Q_i is an A_i-cusp.

Finally, as in the main statement of Theorem 7.5, the conjugacy classes of maximal virtually $\mathbb{Z} \times \mathbb{Z}$ subgroups are in correspondence with the finitely many orbits of flat planes represented by $F_1, \ldots, F_n$, and as in the converse of Theorem 7.5, the stabilizers of these flat planes represent these distinct conjugacy classes. □

Remark 12.8. The action of G on C restricts to a cocompact action on a convex subspace C' of C, so G actually acts properly discontinuously and cocompactly on a CAT(0) space. One obtains C' by choosing a sufficiently large distance d from GK, choosing a flat plane E_i in GK stabilized by A_i, and truncating the cusp Q_i by diagonally cutting the cubes a distance d from E_i along hyperplanes parallel to E_i.

We are then able to give the following simple consequence of Theorem 12.2 which partially answers a question in [Gro93, p. 161] about amalgamated free products of $C(6)$ groups. An analogous statement holds for $B(4)$-$T(4)$ groups.

Corollary 12.9. *The free product of finitely presented $B(6)$ groups amalgamating a cyclic subgroup is CAT(0) and hence has a quadratic isoperimetric function.*

14. Codimension 1 subgroups and the absence of property (T)

Theorem 14.1. *Let G act properly discontinuously on a simply-connected $B(6)$ or $B(4)$-$T(4)$ complex X, and suppose the action on X^1 is cocompact. Then G contains a codimension 1 subgroup.*

Proof. There are finitely many orbits of hypergraphs in X and hence finitely many orbits of hyperplanes in C, so one would hope to produce a codimension 1 subgroup from one of the hyperplanes of C directly from the results of [Sag95, NR98b, Ger97]. However, the results of Section 10 were not sufficiently strong to prove that the action of G on C has no global fixed point, without assuming X is locally finite (see Corollary 10.7).

Instead we shall provide the codimension 1 subgroup directly from Section 3.

Without loss of generality, we may assume that there are no free faces. For otherwise, X equivariantly deformation retracts to a subcomplex satisfying the same small-cancellation conditions, but with fewer orbits of free faces. (Note that we may have to subdivide the retracted 2-cell.) After repeating this finitely many times we arrive at a complex without free faces, but satisfying the same small-cancellation conditions.

Let $Y \subset X$ be a hypercarrier, and let $H \subset G$ be the subgroup of Stabilizer(Y) which preserves the two halfcarriers. The 2-complex $H\backslash X$ is quasi-isometric with the coset graph $H\backslash X$. So we will show that $H\backslash X$ has at least two ends.

Pick a 2-cell R in Y, and let e_1 and e_2 be edges on opposite sides of ∂Y such that neither e_1 nor e_2 lie on any other 2-cell of Y. Let Y_1 and Y_2 be the hypercarriers dual to e_1 and e_2. Since Y_1 has no free faces, we can choose λ_1 to be an infinite ray in Λ_1 which starts at the center of R, and passes through e_1. We define λ_2 similarly, so λ_1 and λ_2 extend into opposite halfcarriers of Y.

We claim that the component of $H\backslash X - \Lambda$ which contains the image of λ_1 is infinite, and in fact λ_1 eventually leaves any compact subset of this component. Indeed, if λ_1 remained in some compact subset, then λ_1 would have to pass through the center of some 2-cell twice in the same direction. That is, λ_1 would have an initial subpath $PaQa$, where a is the image of an edge of Λ_1 in $H\backslash X - \Lambda$. Let U be the longest path so that $P = P'U$ and $Q = Q'U$ (note that U might be trivial.) Then $\lambda_1' = PaQ'P'^{-1}$ is a path in Λ_1 which starts at the center of R and ends at a translate hR where $h \in H$. Thus $hR \subset (Y \cap Y_1)$. But then by Lemma 6.8, all the 2-cells between R and hR lie in $Y \cap Y_1$, and so $Y \cap Y_1$ contains 2-cells on both sides of e_1. This contradicts our choice of e_1. Since the analogous statement holds for λ_2, we see that $H\backslash X$ has at least two ends and we are done. □

Property (T) for an infinite finitely generated small-cancellation group follows from Theorem 14.1 together with the result of [NR98b]. However, we can deduce a statement with a more general flavor from the finitely presented case.

Theorem 14.2. *Let G be a finitely generated group that acts freely on a simply-connected $B(6)$ or $B(4)$-$T(4)$ complex X. Then G does not have property (T) unless it is finite cyclic.*

Note that we drop Convention 3.1 for the beginning of the proof of this theorem, so the theorem covers the case where G is a small-cancellation group (possibly with torsion) acting freely on the universal cover of its standard 2-complex.

Proof. The quotient $G\backslash X$ is a 2-complex that satisfies the same small-cancellation condition that X does.

Suppose that G had property (T), then by [Sha00], there is an epimorphism $J \to G$ where J is a finitely presented group satisfying property (T).

By [Ger96], $J \cong \pi_1 W$ where W is a compact 2-complex satisfying the same small-cancellation property as $G\backslash X$.

But if J is infinite then it cannot have Property (T) because Theorem 11.1 or Theorem 12.2 shows that J acts on a finite dimensional CAT(0) 2-complex without a fixed point, and hence the result of [NR98b] shows that J does not have property (T). □

Acknowledgment: This research graduated from codimension 1 subgroups to actions on cube complexes while the author was visiting at the Technion during December 2000. I am grateful to the late Bob Brooks for supporting this visit, and to Michah Sageev for helpful explanations about cube complexes. Finally, I am grateful to the referee for many helpful comments and corrections.

References

[Aea91] 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.

[Bro94] Kenneth S. Brown. *Cohomology of groups*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[BŚ97] W. Ballmann and J. Świątkowski. On L^2-cohomology and property (T) for automorphism groups of polyhedral cell complexes. *Geom. Funct. Anal.*, 7(4):615–645, 1997.

[Dun85] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.

[Ger87] S. M. Gersten. Reducible diagrams and equations over groups. In *Essays in group theory*, pages 15–73. Springer, New York-Berlin, 1987.

[Ger96] S. M. Gersten. Subgroups of word hyperbolic groups in dimension 2. *J. London Math. Soc. (2)*, 54(2):261–283, 1996.

[Ger97] 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.

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

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[Gro03] M. Gromov. Random walk in random groups. *Geom. Funct. Anal.*, 13(1):73–146, 2003.

[GS90] S. M. Gersten and H. B. Short. Small cancellation theory and automatic groups. *Invent. Math.*, 102(2):305–334, 1990.

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

[MW02] Jonathan P. McCammond and Daniel T. Wise. Fans and ladders in small cancellation theory. *Proc. London Math. Soc. (3)*, 84(3):599–644, 2002.

[NR98a] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[NR98b] Graham A. Niblo and Martin A. Roller. Groups acting on cubes and Kazhdan's property (T). *Proc. Amer. Math. Soc.*, 126(3):693–699, 1998.

[Pri83] Stephen J. Pride. Some finitely presented groups of cohomological dimension two with property (FA). *J. Pure Appl. Algebra*, 29(2):167–168, 1983.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.

[Sco78] Peter Scott. Ends of pairs of groups. *J. Pure Appl. Algebra*, 11(1-3):179–198, 1977/78.

[Sha00] Yehuda Shalom. Rigidity of commensurators and irreducible lattices. *Invent. Math.*, 141(1):1–54, 2000.

[Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.

[Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[vK33] E. R. van Kampen. On some lemmas in the theory of groups. *Amer. J. Math.*, 55:268–273, 1933.

[Wis] Daniel T. Wise. The residual finiteness of qausi-positive one-relator groups. *Jour. London Math. Soc. (3)*. to appear.

[Wis96] Daniel T. Wise. *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*. PhD thesis, Princeton University, 1996.

[Wis01] Daniel T. Wise. The residual finiteness of positive one-relator groups. *Comment. Math. Helv.*, 76(2):314–338, 2001.

[Żuk96] Andrzej Żuk. La propriété (T) de Kazhdan pour les groupes agissant sur les polyèdres. *C. R. Acad. Sci. Paris Sér. I Math.*, 323(5):453–458, 1996.

Dept. of Math., McGill University, Montreal, Quebec, Canada H3A 2K6
E-mail address: wise@math.mcgill.ca