Algebraic & Geometric Topology 11 (2011) 1793–1820

Periodic flats in CAT(0) cube complexes

MICHAH SAGEEV
DANIEL T WISE

We show that the flat closing conjecture is true for groups acting properly and cocompactly on a CAT(0) cube complex when the action satisfies the cyclic facing triple property. For instance, this property holds for fundamental groups of 3–manifolds that act freely on CAT(0) cube complexes.

20E99, 20F65, 20F67

In memory of Bob Brooks

1 Introduction

One of the best known properties of a word-hyperbolic group, is that it cannot contain a subgroup isomorphic to $\mathbb{Z}\times\mathbb{Z}$. For a group G acting properly and cocompactly on a CAT(0) space X, it is known that G is word-hyperbolic if and only if X does not contain an isometrically embedded flat plane $\mathbb{E}^2$. The "flat torus theorem" asserts that if G contains a subgroup $H \cong \mathbb{Z}\times\mathbb{Z}$, then there is a flat plane stabilized by H. One is led to the following problem:

Problem 1.1 (Flat Closing) Let G act properly and cocompactly on a Hadamard space X. Suppose X contains a flat plane. Does X contain a "periodic" flat plane? Equivalently, does G contain a subgroup isomorphic to $\mathbb{Z}\times\mathbb{Z}$?

While it is widely believed that Problem 1.1 admits a negative solution in general, there is no known counterexample, nor even specific candidate counterexamples. In fact, it appears that in many geometrically interesting cases, Problem 1.1 actually admits a positive solution.

Groups acting on CAT(0) cube complexes are playing an increasingly prominent role in geometric group theory, and we are led to examine Problem 1.1 for CAT(0) cube complexes both because of the richness of examples, and because of their attractive simple nature as a test case. The characteristic feature of a CAT(0) cube complex are the "hyperplanes" which are lower-dimensional CAT(0) cube complexes that cut

Published: 12 June 2011 DOI: 10.2140/agt.2011.11.1793

it in half. For example, the hyperplanes in a tree are the centers of edges, and the hyperplanes in the usual CAT(0) cube structure on $\mathbb{E}^n$ are copies of $\mathbb{E}^{n-1}$ cutting orthogonally through cubes.

This paper hinges upon the following property:

Definition 1.2 A *facing triple* in a CAT(0) cube complex X, is a set of three disjoint hyperplanes H_1, H_2, H_3 such that no hyperplane separates the other two.

Let G act on the CAT(0) cube complex X. Then G has *cyclic facing triples* if for each facing triple H_1, H_2, H_3, the group $\bigcap_i \text{Stabilizer}(H_i)$ is either finite or virtually cyclic.

Our main result is:

Theorem 1.3 *Let X be a cocompact cube complex with cyclic facing triples. Then X contains a flat plane if and only if X contains a periodic flat plane.*

This result generalizes a theorem of Mosher [6] where the result was proven when X is a 3–dimensional manifold with a nonpositively curved cubing (in which case it follows that X has cyclic facing triples). The result also generalizes a result of the second author [9] where the theorem was proven in the case that X is 2–dimensional.

It appears unlikely that Problem 1.1 has an affirmative solution even in the limited category of CAT(0) cube complexes. Gromov proposed in [4] that the existence of aperiodic sets of Wang tiles suggests that there might even be a counterexamples to Problem 1.1 in the category of 2–dimensional CAT(0) cube complexes. Some efforts were made towards this by Kari and Papasoglu in [5] where examples were constructed whose periodic flat planes were more limited than the general flat planes.

There are some cases where Problem 1.1 can be strengthened to state that:

Problem 1.4 Let G act properly and cocompactly on the CAT(0) space X. Is every flat plane in X the limit of periodic flats?

This stronger statement does not hold in the category of CAT(0) 2–complexes, since in [10] the second author gave an example of a group acting on a 2–dimensional CAT(0) cube complex containing a flat plane that is not the limit of periodic flats. We are unaware of any further such examples. It was explicitly proven that flats are limits of periodic flats in groups acting properly and cocompactly on the product of trees in [9], and under the hypotheses considered by the second author in [11]. Such density of flats in periodic flats was further established for Euclidean buildings by Ballman and

Brin [1]. We believe that with a bit of further care, the method in this paper would show that under the cyclic facing triple hypothesis, every flat in X is actually the limit of periodic flats.

The results in this paper hold (and the proof is nearly identical) under the following more general definition of cyclic facing triples: $\bigcap_i \mathrm{Stabilizer}(H_i)$ is cyclic, whenever the facing triple H_1, H_2, H_3 satisfies $d(H_i, H_j) \geq C$ for some constant C.

Finally, because of its less significant group theoretical impact, we have not considered higher dimensional periodic flats, but we expect that a similar analysis to that done in this paper would yield analogous results under the hypothesis that facing triples are abelian.

1.1 Sketch of the argument

If some hyperplane Y in X contained a flat, then by induction on the dimension, Stabilizer(Y) would contain $\mathbb{Z} \times \mathbb{Z}$, so we may assume that each hyperplane is δ–hyperbolic. This has several important consequences that enable the arguments in the paper. For instance, we can assume that there is a uniform lower bound on the angle between any flat and hyperplane.

Letting F be a flat plane, we let Hull(F) denote the intersection of all halfspaces in X containing F. To facilitate further arguments, we show that Hull(F) lies in a finite neighborhood of F. Moreover, Hull(F) can be chopped into rectangular "blocks" by two infinite families of disjoint hyperplanes that intersect F in boundedly spaced "vertical" and "horizontal" lines. We can then view Hull(F) as the union of vertical "strips" consisting of infinite sequences of blocks bounded by consecutive vertical hyperplanes.

If there is a G–periodic strip of blocks in Hull(F), then it follows that each such strip of blocks is periodic. The cyclic facing triples condition then implies that all these strips have a uniform period. Consequently, Hull(F) is stabilized by a $\mathbb{Z}$ subgroup, and it is then easy to form periodic flats by finding distinct strips in Hull(F) that are in the same G–orbit.

We are left to show that the flat is singly periodic. A limiting argument shows that there must exist a hyperplane Y_o which contains a line ℓ_o parallel to a line ℓ in F. This hyperplane forms a facing triple with hyperplanes Y_1 and Y_2 that intersect F in lines ℓ_1 and ℓ_2 that are also parallel to ℓ. Applying the cyclic facing triple to Y_o, Y_1, Y_2 we obtain the desired periodic strip in Hull(F).

In a concluding section, we verify that if M is a 3–manifold and $\pi_1 M$ acts properly and cocompactly on a CAT(0) cube complex X, then the cyclic facing triple hypothesis is satisfied.

Acknowledgements The first author's research was supported by BSF and ISF. The second author's research was supported by FQRNT and NSERC.

2 Preliminaries: Hyperplanes in CAT(0) cube complexes

We recall basic terminology and facts about CAT(0) cube complexes. For more details, see the paper [7] by the first author.

A *CAT(0) cube complex* is a simply-connected combinatorial cell complex whose closed cells are Euclidean n–dimensional cubes $[0,1]^n$ of various dimensions such that:

(1) Any two cubes either have empty intersection or intersect in a single face of each.

(2) The link of each 0–cell is a *flag complex*, a simplicial complex such that any $(n+1)$ adjacent vertices belong to an n–simplex.

Since an n–cube is a product of n unit intervals, each n–cube comes equipped with n natural projection maps to the unit interval. A *hypercube* is the preimage of $\{\frac{1}{2}\}$ under one of these projections; each n–cube contains n hypercubes. A *hyperplane* in a CAT(0) cube complex X is a subspace intersecting each cube in a hypercube. Hyperplanes are said to *cross* if they intersect non-trivially; otherwise they are said to be *disjoint*.

Here are some basic facts about hyperplanes in CAT(0) cube complexes which we will use throughout our arguments.

- each hyperplane is embedded (that is, it intersects a given cube in a single hypercube)
- each hyperplane separates the complex into precisely two components, called *half-spaces*
- if $\{H_1, \dots, H_k\}$ is a collection of pairwise crossing hyperplanes, then $\bigcap_k H_k \neq \varnothing$
- each hyperplane is itself a CAT(0) cube complex

A triple of hyperplanes is said to be *facing* if they are disjoint from each other and the union of each pair of them is contained in a single halfspace of the third. Otherwise, the triple is said to be *nested*, which means that one of them separates the other two.

A CAT(0) cube complex has *cyclic facing triples* if for each facing triple H_1, H_2, H_3, the intersection of their stabilizers is virtually infinite cyclic or finite.

Finally, given a vertex v in $X^{(0)}$, we define the *dual block* containing v as follows. The hyperplanes of X provide a subdivision of X into another cube complex X', in which each n–cube of X is subdivided into 2^n subcubes. The dual block containing v is the union of the cubes of X' containing v.

3 Flats and hyperplanes

3.1 The assumption that hyperplanes are hyperbolic

Let X be a CAT(0) cube complex and let G be a group which acts properly and cocompactly on X. By a *flat* in X we mean an isometric embedding of a 2–dimensional Euclidean plane into X. We will prove Theorem 1.3 by induction on the dimension of X. For 0–dimensional complexes, the theorem holds, so we focus on the inductive step. If any hyperplane H of X contains a flat, then viewing H as a CAT(0) cube complex in its own right, we see that the dimension of H is less than the dimension of X. Note that since G acts properly and cocompactly on X, and preserves the family of hyperplanes, it follows that for each H, stab(H) acts cocompactly on H. Since facing triples in H are simply the intersection with H of facing triples in X, it follows that H, together with the action of stab(H) on it, has cyclic facing triples. Thus, Theorem 1.3 holds for H, so that by induction, there exists a periodic flat in H. But then there exists a periodic flat in X and our theorem is proved. We will thus assume henceforth that H has no flats. It follows by the Flat Plane Theorem that each hyperplane is δ–hyperbolic; since there are finitely many orbits of hyperplanes, we may choose δ universally over all hyperplanes in X.

3.2 Intersections of flats and hyperplanes

A subset Y of a geodesic metric space Z is said to be *geodesically contained* if extensions of geodesics in Y are also contained in Y; more precisely, given a geodesic segment $I \subset Y$, if $J \subset X$ is a geodesic segment with $I \subset J$, then $J \subset Y$. A CAT(0) space Y is said to be *geodesically extendable* if every geodesic segment in Y can be extended to a bi-infinite geodesic in Y. Euclidean space, for example, is geodesically contained and geodesically extendable. However, a flat in an arbitrary CAT(0) space is geodesically extendable but need not be geodesically contained. (For example, imagine a space formed by gluing three half-planes glued along their boundary lines via isometries. In this case, each flat is not geodesically contained.)

Lemma 3.1 *Hyperplanes in a CAT(0) cube complex are convex and geodesically contained.*

Proof We will show that hyperplanes are geodesically contained by showing that they are “locally" geodesically contained. Let H be a hyperplane and that $I \subset H$ a geodesic segment. Suppose that J is a geodesic extension of I which is not contained in H. By possibly replacing I with a larger geodesic segment contained in J, we may assume that $J \cap H = I$. Let p be a boundary point of I. If there exists a neighborhood U of p in J so that U is contained in a cube of X, then U is contained in H because geodesic containment clearly holds for a hyperplane in a single cube. So suppose then that there is no such neighborhood of p. We then have two cubes σ_1 and σ_2 of X and two subintervals I_1 and I_2 of J such that $I_1 \cap I_2 = \{p\}$, $I_1 \subset \sigma_1$ and $I_2 \subset \sigma_2$ (see Figure 1).

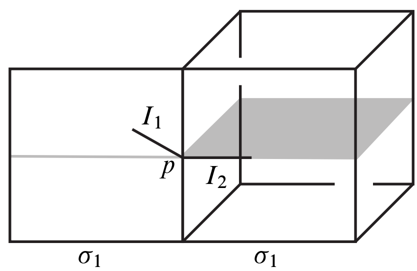

Figure 1: Hyperplanes are locally geodesically contained.

But now we see that $I_1 \cup I_2$ is not a local geodesic, a contradiction.

A similar argument shows that hyperplanes are locally convex and hence convex. $\square$

Remark If H is a hyperplane and $C(H)$ is its *carrier*, that is, the union of cubes meeting H, then $C(H) \cong H \times I$ and there exists a natural projection $q : C(H) \to I$ with $H = q^{-1}(1/2)$. The above arguments show that $q^{-1}(t)$ is convex and geodesically contained, for any $t \in (0, 1)$.

We now show that intersections of hyperplanes and flats are what we expect.

Proposition 3.2 *Let F be a flat in X and H a hyperplane in X. Then $F \cap H$ is either empty or a line.*

Proof By Lemma 3.1, hyperplanes are convex and geodesically contained; flats are convex. It follows that $F \cap H$ is a convex and geodesically contained subset of F. Hence $F \cap H$ is either empty, a point, a line or all of F. Since H does not contain flats, $F \cap H \neq F$. We thus need to rule out the possibility that $F \cap H$ is a point.

Recall that $C(H)$, the carrier of H, has a product structure $C(H) \cong H \times I$, with a projection map $q : C(H) \to I$ for which $H = q^{-1}(1/2)$. So now suppose that $F \cap H$ is a single point, $F \cap H = \{p\}$. Note that H separates X, so that if $F \setminus \{p\}$ met both components of $X \setminus H$, then $\{p\} = H \cap F$ would separate F, a contradiction. Thus, we have that $F \setminus \{p\}$ meets only one of the components of $X \setminus H$. Now suppose that

$l \subset F$ is a line in F containing p. Then l meets a single component U of $X \setminus H$. But then it follows that there exists $t \in [0,1]$, $t \neq 1/2$, so that l meets $q^{-1}(t)$ in at least two points a and b, with p between a and b. But $q^{-1}(t)$ is convex, a contradiction to the remark following Lemma 3.1. □

The next basic fact that we will need is that there is a lower bound on the angle between the hyperplanes and F. First, we define the angle between a flat and a hyperplane (see Figure 2). Suppose that H is a hyperplane which intersects F. Let $l = H \cap F$ and choose $x \in l$. Let $\mathbf{n}$ denote one of the two normal vectors to H at x and let $\mathbf{n}'$ denote the normal vector to l in F at x, which lies on the same halfspace as $\mathbf{n}$. Then we define the *angle between H and F* to be $\angle(H,F) = \pi/2 - \angle(\mathbf{n},\mathbf{n}')$. ($F$ lies in H if and only if $\angle(H,F) = 0$.)

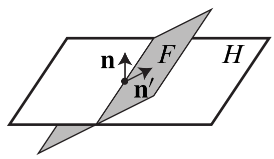

Figure 2: The angle between a flat and a hyperplane.

Note that we may similarly define the angle between a line and a hyperplane. If l is a line meeting H, we let $x = H \cap l$ and replace $\mathbf{n}'$ by the tangent vector to l in the previous definition.

Lemma 3.3 *There exists a lower bound on the angle between a hyperplane and a flat.*

Proof Suppose there exists a sequence of hyperplanes H_i meeting flats F_i, with $\theta_i = \angle(H_i, F_i)$ and $\theta_i \to 0$. Then since there are finitely many orbits of hyperplanes, we may assume, after translation and passing to a subsequence, that we have a single hyperplane H and a sequence of flats F_i such that $F_i \cap H \neq \varnothing$ and the angle θ_i between F_i and H approaches 0. Since the stabilizer of H acts on H cocompactly, we may further assume that there exists a closed ball B such that $F_i \cap H \cap B \neq \varnothing$ for every i. Each F_i can be viewed as an isometric embedding $f_i : \mathbf{E}^2 \to X$, with $f_i(0) \in B$. This is a sequence of maps and we may now pass to a subsequence of $\{f_i\}$ which converges uniformly on compact sets. The limiting map is an isometric embedding of a Euclidean plane in H, a contradiction. □

The above argument can be adapted slightly to prove the following technical lemma. Given two geodesic segments (possibly lines or rays) I_1 and I_2 meeting at a single point, we define $\angle(I_1, I_2)$ to be the minimal angle which they subtend.

Lemma 3.4 *Given a number $\theta > 0$, there exists $\nu = \nu(\theta) > 0$, such that the following holds. Let F be a flat and H a hyperplane in X and suppose that F and H meet in a line l. Let p be a point on l. If l' is a line in F containing p such that $\angle(l, l') > \theta$, then $\angle(l', I) > \nu$ for all geodesic intervals I in H containing p.*

3.3 Parallelism classes of lines in a flat

Consider the intersection of F with the collection of hyperplanes in X. This intersection is a collection $\mathcal{L}$ of lines in F. As noted earlier, a basic fact about CAT(0) cube complexes is that if a collection of hyperplanes pairwise intersect, then they intersect. Since the number of hyperplanes that can intersect is bounded by the dimension of the complex, this puts a bound on the collection of hyperplanes that can pairwise intersect. Since X is finite dimensional, it follows that $\mathcal{L}$ contains a finite number of parallelism classes. Note also that each complementary region of $\mathcal{L}$ is mapped into a dual block of X. Since dual blocks have bounded diameter, each complementary region of $\mathcal{L}$ in F is bounded. A parallelism class $\mathcal{L}_i$ of lines in $\mathcal{L}$ is *boundedly spaced* if there exists $k > 0$ such that the k–neighborhood of $\cup_{L \in \mathcal{L}_i} L$ contains F.

Lemma 3.5 *Every parallelism class is boundedly spaced.*

To prove this we use the following basic lemma about CAT(0) cube complexes.

Lemma 3.6 *Let $x \in X$ and H be a hyperplane in X. Suppose that $\alpha = [x, y]$ is a shortest geodesic from x to H. Then every hyperplane crossed by α is disjoint from H.*

Proof Let H' be a hyperplane crossed by α and suppose that $H \cap H' \neq \varnothing$ (see Figure 3). Let $z = \alpha \cap H'$. Consider a shortest geodesic segment $[z, w]$ from z to $H \cap H'$. Now the normal vector to H at w lies in H' and hence agrees with $[z, w]$.

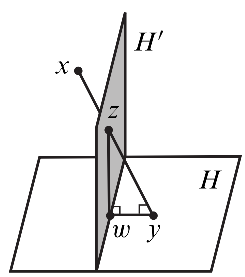

Figure 3: A geodesic from a point to a hyperplane.

Therefore the $\triangle(z, w, y)$ is a triangle with two right angles, a contradiction. $\square$

Proof of Lemma 3.5 Refer to Figure 4. Consider an equivalence class of lines $\mathcal{L}_i$ in $\mathcal{L}$, and let l be a line in $\mathcal{L}_i$. Let H be a hyperplane such that $H \cap F = l$. Let D denote the maximal diameter of a dual block in X. By Lemma 3.3, there is a lower bound on the angle between F and H. Thus there exists $k = k(D)$ such that if $x \in F$ and $d(x, l) > k$, then $d(x, H) > D$.

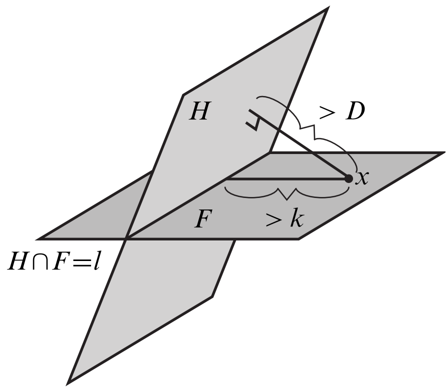

Figure 4: If a point in a flat is far from the intersection of a hyperplane and the flat, then the point is far from the hyperplane.

Now given $x \in F$ such that $d(x, l) > k$, we claim that there exists another line $l' \in \mathcal{L}_i$ such that $d(x, l') < d(x, l)$ and $d(l, l') > 1$. This proves the claim, for then every point of F is within k of some line in $\mathcal{L}_i$. Since $d(x, H) > D$, the geodesic α between x and H is crossed by some hyperplane H'. By Lemma 3.6, H' is disjoint from H. Moreover, since H' separates x from H, it follows that H' intersects F non-trivially. But since H and H' are disjoint, $H' \cap F$ is a line parallel to l and is thus in $\mathcal{L}_i$. Moreover, since disjoint hyperplanes are at least distance 1 apart, we have that $d(l, l') > 1$, as required. □

4 Utilizing hyperbolicity

In this section we collect the lemmas that show that all sorts of things do not intersect due to hyperbolicity. We suppose as usual that F is a flat in X and that all of the hyperplanes of X are hyperbolic. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two parallelism classes of lines in F as above. For convenience, we call $\mathcal{L}_1$ *vertical* and $\mathcal{L}_2$ *horizontal*. We also refer to the hyperplanes that produce the vertical lines as *vertical hyperplanes* and those that produce the horizontal lines as *horizontal hyperplanes*.

We start with the following lemma, which provides a lower bound on the angle between lines in $\mathcal{L}_1$ and lines in $\mathcal{L}_2$.

Lemma 4.1 *The parallelism classes $\mathcal{L}_1$ and $\mathcal{L}_2$ can be chosen so that the angle between lines in $\mathcal{L}_1$ and lines in $\mathcal{L}_2$ is at least $\theta_0 = \arcsin(1/\sqrt{d})$, where $d = \dim(X)$.*

For this we we need the following basic linear algebra fact.

Lemma 4.2 *Consider $\mathbb{R}^d$ with the standard basis $\mathcal{E} = \{e_1, \dots, e_d\}$. If R is a ray emanating from the origin, then there exists a codimension 1 hyperplane H spanned by some collection of $d-1$ vectors in $\mathcal{E}$, such the angle between R and H is at least θ_0.*

Proof In fact, one shows that the angle with one of the hyperplanes is at least $\arcsin(1/\sqrt{d})$.

Consider the unit vector v in the direction of R. Then in the standard basis $v = (v_1, \dots, v_d)$. Since $\Sigma v_i^2 = 1$, there exists some i such that $v_i \geq 1/\sqrt{d}$. So if η denotes the angle between v and e_i, then $\eta \leq \arccos(1/\sqrt{d})$ and so the angle between v and the plane spanned by all the remaining elements of $\mathcal{E}$ is at least $\arcsin(1/\sqrt{d})$. □

Proof of Lemma 4.1 Let $\mathcal{L}_1$ be some parallelism class and let l_1 be a line in $\mathcal{L}_1$. Let C be some d–cube such that l_1 intersects C in an interval I. Let l denote a line containing the barycenter of C and parallel to I We identify C with the standard unit cube in $\mathbb{R}^d$. By Lemma 4.2, there exists some hyperplane H such that $\angle(H, l) \geq \theta_0 = \arcsin(1/\sqrt{d})$. Consider now the carrier $C(H)$ of H. Since the metric on $C(H)$ is the product metric of H and I, we have that $\angle(H, l_1) = \angle(H, l) \geq \theta_0$. Thus the angle between l_1 and any line in H containing $H \cap l_1$ is at least θ_0 in particular, the angle between l_1 and $H \cap F$ is at least θ_0, as required. □

From this point on we assume that the angle between the vertical and horizontal lines in F is always at least θ_0.

Lemma 4.3 (Bounded Prisms) *There exists a number $K > 0$ (depending only on X), such that the following holds. Let H and H' be two hyperplanes that meet F in parallel lines l and l'. If $H \cap H' \neq \varnothing$, then there exists a line $l'' \subset H \cap H'$ parallel to l and l' such that $d(l, l'') < K$ and $d(l', l'') < K$.*

Remark. It follows from the lemma that if l and l' are distance at least $2K$ apart then H and H' are disjoint.

Proof Without loss of generality, we may assume that l and l' are vertical. Suppose that $H \cap H' \neq \varnothing$. Consider a horizontal line m in F corresponding to an intersection of F with a horizontal hyperplane J (see Figure 5). Let $p = m \cap l$ and $q = m \cap l'$.

Let $\theta \geq \theta_0$ denote the angle between l and m, and let $\nu = \nu(\theta) > 0$ be the required angle appearing in Lemma 3.4. Then the angle between m and any geodesic in H containing p is greater than ν and the angle between m and any geodesic segment in H' containing q is greater than ν. Let $[p, r]$ denote the geodesic segment obtained by dropping a perpendicular from p to $H \cap H'$.

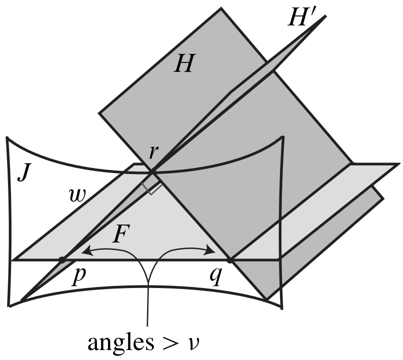

Figure 5: A prism formed by a flat and two hyperplanes.

Now since all the angles in the triangle $\triangle(p, q, r)$ are bounded below by $min\{\nu, \pi/2\} > 0$, and $\triangle(p, q, r)$ is contained in J, which is δ–hyperbolic, we have a bound $K = K(\nu, \delta)$ on the lengths of the edges of $\triangle(p, q, r)$.

Now let us enumerate the lines in $\mathcal{L}_2$, $\{\ldots m_{-1}, m_0, m_1, \ldots\}$. By the above argument, if $H \cap H' \neq \varnothing$, for each line m_i, we get a triple of points p_i, q_i, r_i of bounded diameter K, with $p_i = m_i \cap l$, $q_i = m_i \cap l'$ and $r_i \in H \cap H'$. Now we consider the sequence of geodesic segment $\alpha_i = [r_{-i}, r_i]$. This sequence of geodesic segments lies in a K–neighborhood of l and l'. Thus $\{\alpha_i\}$ limits on a geodesic in $H \cap H'$ in the K–neighborhood of l and l', as required. □

We will also need the following basic lemma which controls how close a flat can get to a hyperplane when it is disjoint from it. It is very similar to Lemma 3.3.

Lemma 4.4 *There exists a number C such that if $F \cap H = \varnothing$, then $d(F, H) > C$.*

Proof Suppose not. Then, as in the proof of Lemma 3.3 there exists a sequence of flats F_i and a hyperplane H such that $F_i \cap H = \varnothing$ and $d(F_i, H) \to 0$. Moreover, by the compactness of the G action we may assume that there exists a ball B such that a nearest point of F_i to H lies in B. Let α_i denote a geodesic from F_i to H. We now note two things. First, as in Lemma 3.3, we get that the F_i's converge uniformly to a limiting flat F. Next, if $\angle(H, F) > 0$, then we would get that F crosses H

transversely, and hence so would F_i for sufficiently large i, it follows that $F \subset H$, a contradiction. $\square$

The above lemma tells us that flats cannot get too close to hyperplanes. Here is another version of this fact that we will use as well.

Lemma 4.5 *Given a number $R > 0$ there exists $C = C(R) > 0$ such that if F is a flat, H is a hyperplane and D is a disk of radius r in F with $D \subset N_R(H)$, then $r < C$.*

Proof Suppose we have an R such that $N_R(H)$ contains arbitrarily large flat disks: disks $\{B_n\} \subset F$ with B_n of radius n. In each disk B_n, choose a regular geodesic triangle $T_n = \triangle(a_n, b_n, c_n)$ of side length n. For each edge of each such triangle, we can find points $a'_n, b'_n, c'_n \in H$ distance R from a_n, b_n and c_n respectively. Since X is CAT(0), the edges of the geodesic triangle $\triangle(a'_n, b'_n, c'_n)$ are within R of their respective edges in $\triangle(a_n, b_n, c_n)$. But now for sufficiently large n, this produces non-δ–thin triangles in H, a contradiction. $\square$

Remark 4.6 Note that the above lemma does not require the existence of the entire flat F. The hyperbolicity of H provides a bound on the size of a Euclidean disk that can be embedded in $N_R(H)$.

Here is a specific corollary which we will make use of.

Corollary 4.7 *Suppose that $R > 0$ is given. Let F be a flat and H be a hyperplane disjoint from F. Suppose that H contains a line l' parallel to a line l in F, and so that $d(l, l') < R$. Suppose that m is a line in F transverse to l. Then there exists a bound (depending on R and the angle between l and m) on the length of a segment of m which can lie in $N_R(H)$.*

Proof Suppose that $\alpha = [x, y]$ is a segment of m lying in $N_R(H)$. Then the hull Hull(α, l) of α and l is contained in $N_R(H)$. Since the distance function is convex. But then if α is long, Hull(α, l) contains large disks in $N_R(H)$. Thus we get a bound on the length of α. $\square$

5 Convex Hulls

Recall that a *halfspace* in X is the closure of the complement of a hyperplane of X.

Definition 5.1 (Hull) For a subset $S \subset X$ we define $\mathrm{Hull}(S)$ to be the intersection of halfspaces of X containing S. If no halfspace of X contains S then we define $\mathrm{Hull}(S)$ to be X.

Theorem 5.2 *Let $F \subset X$ be a flat plane. Then there exists $K > 0$, such that* $\mathrm{Hull}(F)$ *is contained in a K–neighborhood of F.*

We will need the following technical lemma for hyperbolic CAT(0) cube complexes.

Lemma 5.3 *Suppose that X is a δ–hyperbolic CAT(0) cube complex, with uniformly bounded geometry in the sense that there is a uniformly bounded number of cells in a ball of a given radius. Let l be a line in X and $R > 0$ a given number. Then there exists a number $n > 0$, such that if p is a point in X such that $d(p,l) > n$, then there exists a hyperplane H separating p and l such that $H \cap N_R(l) = \varnothing$.*

Proof First, recall that the hyperplanes subdivide X into bounded complementary regions called dual blocks (for brevity, we will call these just *blocks*). Let D be the maximal diameter of a block. It follows that if β is a geodesic segment of length at least Dn, it meets at least n hyperplanes. Let M be a number larger than the maximal number of hyperplanes which meet a ball of radius 2δ in X. This number is bounded since X has bounded geometry.

Now suppose that $d(p,l) > DM + R + 2\delta$. Let α be the shortest geodesic from p to l and let $q = \alpha \cap l$. If some hyperplane crossed by α does not meet $N_R(l)$, we have found our desired hyperplane. So let us assume that every hyperplane crossed by α meets $N_R(l)$.

Suppose that H is a hyperplane crossing α, so that there exists some point r_H with $H \cap N_R(l) = r_H$. Let $p_H = H \cap \alpha$. Let t_H be the point along l closest to r_H. Now we have a geodesic rectangle $\Box_H(p_H, q, r_H, t_H)$.

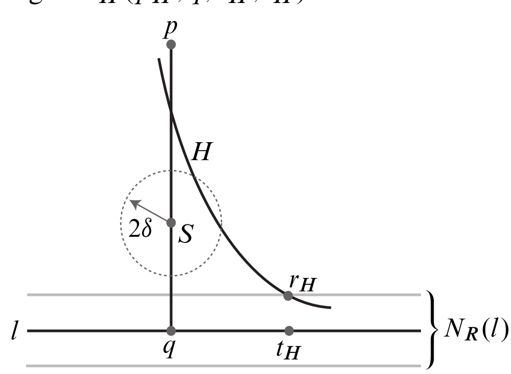

Now geodesic rectangles are 2δ thin; each point on a side is within 2δ of one of the other sides. Since α is a geodesic from p to l, it follows that the piece of the geodesic $[p_H, q]$ which is within δ of $[q, t_H]$ has length no more than δ. Furthermore the length of the edge $[r_H, t_H]$ is no longer than R. Thus, if we let s denote a point along $[p, q]$ such that $d(s, q) = 2\delta + R$, we have that the geodesic segment $[p_H, s]$ lies within 2δ of $[p_H, r_H]$ and hence within 2δ of H. Now by our assumption about $d(p, l)$, we have that the length of $[p, s]$ is at least DM. Thus $[p, s]$ crosses at least M hyperplanes. But each of these meets the ball of radius 2δ about s, a contradiction to our choice of M. □

Proof of Theorem 5.2 As in the previous lemma, let D denote the maximal diameter of a complementary region to the hyperplanes in X. Note that this same D serves as such a bound for each hyperplane in X (when each hyperplane is viewed as a CAT(0) cube complex). Let M be the bound in the previous lemma on the number of hyperplanes which meet a ball of radius 2δ. By Lemma 4.3, there is a bound on the size of prisms. More precisely, there exists some $K > 0$, such that if H and H' are hyperplanes meeting F in parallel lines m_1 and m_2, and $H \cap H' \neq \varnothing$, then there exists a line l in $H \cap H'$, parallel to the m_i's and such that $d(l, m_i) < K$.

We consider a point x such that $d(x, F) > D(M + 1) + K + 2\delta$. We wish to find a hyperplane not meeting F separating x from F. Consider the geodesic α in X from x to F. If α is contained in a hyperplane, let H denote that hyperplane. If α is not contained in a hyperplane, do the following. Consider the first hyperplane H crossed by α and let $y = \alpha \cap H$. If this hyperplane does not meet F, we are done. If H does meet F, then replace the part of α from y to F in X by the geodesic path from y to F in H. We now consider the hyperplane H as a cube complex in its own right. We will use that it is δ–hyperbolic. Let $l = H \cap F$. Note that the geodesic from y to l has length at least $DM + K + 2\delta$. Thus, by Lemma 5.3, we have that there exists some hyperplane J in H crossed by α, such that H does not meet $N_K(l)$.

Now in X, J is the intersection of H with some hyperplane H'. Since H' does not intersect l, if it intersects F, it does so in a line l' parallel to l. But now the hyperplanes H and H' would form a prism. But this would mean, by Lemma 4.3 (Bounded Prisms), that there is a line l' in $J = H \cap H'$ parallel to l such that $d(l, l') < K$, a contradiction to the fact that J does not meet $N_K(l)$. □

Besides knowing that hulls of flats lie within a bounded neighborhood of a flat, we will also need to know that the hyperplanes that meet the flat cut up the hull into bounded pieces. More precisely, let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two parallelism classes of lines in F which correspond to the intersection of F with two classes of hyperplanes $\mathcal{H}_1$ and

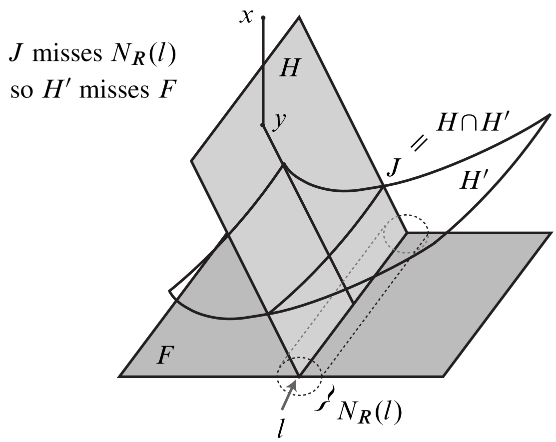

$\mathcal{H}_2$. We will often refer to one of these classes as vertical and the other as horizontal. By applying Lemma 4.3, we may cull $\mathcal{H}_1$ and $\mathcal{H}_2$ so that the each parallelism class is still boundedly spaced, but the hyperplanes in $\mathcal{H}_1$ are disjoint and the hyperplanes in $\mathcal{H}_2$ are disjoint. Now the hyperplanes in $\mathcal{H}_1$ and $\mathcal{H}_2$ subdivide Hull(F) into regions which we call *F–blocks*.

Proposition 5.4 *F–blocks are uniformly bounded.*

Proof Recall by Lemma 3.5, that each parallelism class of lines in F is boundedly spaced, in particular this holds for $\mathcal{L}_1$ and $\mathcal{L}_2$. Furthermore, this remains true after culling as above. Thus the components of $F - (\mathcal{L}_1 \cup \mathcal{L}_2)$ are a uniformly bounded collection of rhombi. Let R be a region of Hull$(F) \setminus \mathcal{H}_1 \cup \mathcal{H}_2$ and let D be the corresponding rhombus. We will now show that R lies in some uniformly bounded neighborhood of one such rhombus.

Let $x \in R$. We know by Theorem 5.2 that $d(x, F) < C$ for some constant C, depending on X. Now drop a perpendicular α from x to F. Since the length of α is bounded by C, it follows that α crosses a bounded number $N = N(C)$ of hyperplanes.

Let $y \in F$ be the other endpoint of α. Since α crosses at most N hyperplanes, and x is not separated from D by any hyperplanes in $\mathcal{H}_1 \cup \mathcal{H}_2$, it follows that y and z are separated by at most N lines in $\mathcal{L}_1 \cup \mathcal{L}_2$. The rhombi are uniformly bounded, so there exists $K > 0$, such that all the rhombi are of diameter less than K. It follows that $d(y, D) < KN$. Thus $d(x, D) < KN + C$, as required. □

6 Dippers

Given a parallelism class $\mathcal{L}$ of lines on F, an *n–dipper in X relative to $\mathcal{L}$* is a hyperplane H which satisfies

- H does not intersect F.
- H contains a line l parallel to $\mathcal{L}$
- $d(l, F) \leq n$ (Here distance between sets is the mindistance between pairs of the points in the sets.)

A *vertical (horizontal) n–dipper* is an n–dipper relative to the vertical (horizontal) direction. A *dipper* is an n–dipper for some n. The following lemma will allow us to restrict our later arguments to dippers that are a uniformly bounded distance from the flat F.

Lemma 6.1 *There exists a number n_0, such that if H is a vertical (horizontal) n–dipper, then there exists a vertical (horizontal) m–dipper H', such that H' separates H from F and such that $m \leq n_0$.*

Proof Let H be a vertical n–dipper. We wish to see that if n is sufficiently large, then we can find a vertical dipper H' separating H from F. Let $l \in H$ and $l' \in F$ be parallel lines with l' vertical in F, and $d(l, l') \leq n$. We assume that l and l' realize the minimal distance between lines in H and vertical lines in F. Let S denote the strip bounded by l and l'. By our choice of l and l', S is perpendicular to F, so that the angle between any segment in S perpendicular to l' and any line in F meeting l' is $\pi/2$.

Choose a point $p \in l'$ that lies in the intersection of l' and a horizontal hyperplane J. Let α denote the segment in S containing p and perpendicular to l'. We first claim that there is a lower bound on the angle $\angle(H, \alpha)$. This follows from Remark 4.6, for if $\angle(H, \alpha)$ were small, we would obtain a fat Euclidean triangle (that is, containing a large Euclidean disk) in S that is close to H. Such a triangle can be constructed by taking the segment α together with a large segment of l as two of the sides of the triangle. It follows that there is a constant $C > 0$ depending on X, such that if $length(\alpha) > n$, then $d(p, H) > Cn$.

Now note that if $d(p, H) > Cn$, and D is the maximal diameter of a dual block, then p is separated from H by at least $N = [Cn/D]$ hyperplanes. Denote these hyperplanes by $H_1, \dots, H_N$. Since each H_i separate p from H, each H_i must meet S. Moreover, since each H_i is disjoint from H, it follows that $H_i \cap S$ is a line parallel to l.

We now need to find such an H_i that is disjoint from F. So suppose that all the H_i's intersect F. Then for each i, H_i meets F in some line vertical line l'' parallel to l'. We aim to bound $d(l', l'')$. Consider the horizontal hyperplane J and let $m = J \cap F$. By Lemma 4.1, $\angle(l, m) \geq \theta_0$. We now have three points: $p = J \cap l'$, $q = J \cap l''$

and $r = J \cap l$. These bound a triangle $\triangle = \triangle(p, q, r) \subset J$. Now since $\angle(F, H_i)$ is bounded below, so too is the angle at q in $\triangle$. Moreover, the angle of $\triangle$ at p is $\pi/2$. Thus, by hyperbolicity of J, the segment $[p, q]$ is bounded by a constant depending only on X. Thus, there is a bound C' depending only on X, such that $d(l', l'') < C'$. By bounded geometry, we have a bound $N = N(C)$ on the number of hyperplanes meeting F in vertical lines at distance less than C away from l'. This gives us an upper bound on N. Since N is bounded above, so is n, as required. □

The following proposition is reminiscent of Lemma 4.3.

Proposition 6.2 *Let F be a flat. Let $\mathcal{H}$ denote the collection of hyperplanes which meet F in a particular parallelism class $\mathcal{L}$. Suppose that H is a dipper relative to $\mathcal{L}$. Then there exists a bound, depending only on δ, on the number of hyperplanes of $\mathcal{H}$ which intersect H.*

Proof This will be an application of Corollary 4.7. Label the elements of $\mathcal{H}$ in order $\{\dots, H_{-1}, H_0, H_1 \dots\}$. We regard the elements of $\mathcal{H}$ as vertical hyperplanes. Consider a horizontal hyperplane J. Now since J crosses all the vertical lines, it crosses the line l in H which is parallel to the vertical direction of F. Thus J crosses H. Now we suppose that we have enumerated the vertical hyperplanes so that H crosses the hyperplanes $\{H_{-k}, \dots H_k\}$. We will show that this will mean that F lies close to H along a subdisk of F whose radius depends on k. This will then bound k. We consider the intersection of this pattern with the hyperplane J. That is, let $H' = H \cap J$, $H'_i = H_i \cap J$ and let $x = F \cap H_k \cap J$ and let $y = F \cap H_{-k} \cap J$. We drop a perpendicular in $J \cap H_k$ from x to H' and let z denote the foot of the perpendicular. Similarly, we let w denote the foot of the perpendicular from y in $J \cap H_{-k}$ to H'. We then have a rectangle in J, $R = R(x, y, w, z)$. Note that we

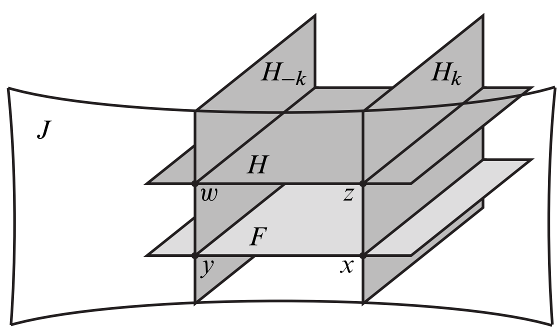

have right angles at the vertices w and z. At the vertices x and y we have angles greater than $\nu = \nu(\theta)$, where θ is the angle between the horizontal and vertical lines and ν is the angle provided in Lemma 3.4. Now applying hyperbolicity, we have (as in Lemma 4.3) that the lower bound on the angles of the rectangle provides a bound on the length of the subsegments of $[x,y]$ which are within δ of $[x,z]$ or $[y,w]$. It follows that as k gets larger, we obtain larger segments of $[x,y]$ which are within δ of H. But then Corollary 4.7 bounds this segment as well, so that we obtain a bound on k. □

Suppose that H is an n–dipper relative to $\mathcal{L}$. Then by the above lemma, if we choose a line $l_0 \in \mathcal{L}$ closest to H, we may choose lines l_1, l_2 sufficiently far away from l_0 on either side of l_0, so that the corresponding hyperplanes H_1 and H_2 are disjoint and disjoint from H. Since there are paths from l_0 to each of the H, H_1 and H_2, they form a facing triple. Moreover, note these hyperplanes contain lines, l, l_1 and l_2, which are a bounded distance from one another (that is, bounded only by the hyperbolicity constant δ).

Now we show that the existence of dippers yields periodicity. Here our cyclic facing triple condition comes into play.

Lemma 6.3 *Suppose that $\{H_1, H_2, H_3\}$ is a disjoint facing triple, with each H_i containing a line l_i, such that for each i, j, l_i is parallel to l_j and $d(l_i, l_j) < n$. Then $\cap \operatorname{stab}(H_i)$ is an infinite virtually cyclic group. Moreover, there exists $C(n)$, such that $\cap \operatorname{stab}(H_i)$ contains an infinite order element whose translation length bounded above by C.*

Proof Since the action of G on X is proper and cocompact, there are finitely many conjugacy classes of finite subgroups. Thus there exists a number $J = J(G)$ such that the order of every finite subgroup is bounded by J.

By assumption, we have that $N_n(l_1)$ contains each of the lines l_1, l_2, l_3. Choose a sequence of points $\{p_i\}$ along l_1 such that $d(p_i, p_{i+1}) = 1$. As before, we let D denote the maximal diameter of a dual block. We then have that every point in X is within D of some vertex and hence some edge of X. So for each p_i, we may find edges e_i^k (for $k = 1, 2, 3$) transverse to H_k and such that $d(e_i^k, p_i) < n + D$. Thus we obtain an infinite sequence of distinct triples of edges $\{e_i^1, e_i^2, e_i^3\}$ such that

- e_i^k is transverse to the hyperplane H_k for $k = 1, 2, 3$,
- for each i, $d(e_i^k, e_i^j) < 2n + 2D$.

Now by cocompactness, up to the action of the group there exists a bound on the number of such triples. So there is a single orbit containing more then J triples. Thus two of these triples differ by a non-torsion element. This element $g \in \cap \operatorname{stab}(H_i)$, as required.

We now claim that the translation length can be chosen to be bounded by some uniform constant $C = C(n)$. There are finitely many facing triples intersecting any given n–ball. Thus, since there are only finitely many orbits of n–balls, there are only finitely many orbits of cyclic facing triples satisfying the hypothesis of the lemma. For each orbit, there exists some translation and hence a translation length. Choose some bound for all of them. □

Remark 6.4 Note that there is at most one parallelism class of lines lying in a bounded neighborhood of all three hyperplanes in a facing triple of hyperplanes. In particular, the axis of an infinite order element stabilizing all three hyperplanes lies in this parallelism class.

We will employ the above lemma in Section 7 to obtain periodicity along strips in the hull of flats.

7 The main argument

7.1 The setup

We first describe our setup a bit more precisely. As before, we have two parallelism classes of lines $\mathcal{L}_1$ and $\mathcal{L}_2$ in F corresponding to collections of hyperplanes $\mathcal{H}_1$ and $\mathcal{H}_2$. By applying Lemma 4.3 (Bounded Prisms), we may remove some lines in $\mathcal{L}_1$ and $\mathcal{L}_2$ and their corresponding hyperplanes in $\mathcal{H}_1$ and $\mathcal{H}_2$ so that no two hyperplanes in $\mathcal{H}_1$ intersect and no two hyperplanes in $\mathcal{H}_2$ intersect. As before, we will call the lines in $\mathcal{L}_1$ vertical and the ones in $\mathcal{L}_2$ horizontal. We use this same language for the elements of $\mathcal{H}_1$ and $\mathcal{H}_2$. Recall that these families cut up $\operatorname{Hull}(F)$ into regions called *F–blocks* with compact closure. Let $\mathcal{B}$ denote the collection of F–blocks. Since F–blocks have uniformly bounded size and G acts cocompactly on X, there are finitely many orbits of F–blocks.

Let us develop a bit more terminology which will give us a slightly more restrictive notion of blocks being in the same orbit. We imagine that our flat F is laid out in front of us so that there is a notion of right, left, top and bottom. Suppose B is an F–block bounded by hyperplanes $H_1, H_1' \in \mathcal{H}_1$ with H_1 to the left of H_1' and $H_2, H_2' \in \mathcal{H}_2$, with H_2 below H_2'. We say that B is *bounded on the left* by H_1, on the *right* by H_1', on the *bottom* by H_2 and on the *top* by H_2'. The F–block B together with these bounding hyperplanes is called an *oriented* block. We say that two F–blocks B and

B' are in the same orbit if there exists a group element $g \in G$ such that $B = gB'$, so that B' is bounded on the left by gH_1, on the right by gH_1', on the bottom by gH_2 and the top by gH_2'. There are only finitely many orbits of oriented F–blocks, so we view the oriented F–blocks as being labeled by a finite labeling set. Henceforth, when we say that two F–blocks are in the same orbit or have the same label, we will mean that they are in the same orbit as oriented blocks.

If H_1 and H_2 are disjoint vertical (horizontal) hyperplanes, then the union of all F–blocks lying between by H_1 and H_2 is called a *vertical (horizontal) strip* of $\mathrm{Hull}(F)$. The distance between the hyperplanes H_1 and H_2 is called the *width* of this strip. If there exists an infinite order element stabilizing a strip we say that the strip is *periodic*. As with blocks, we define an *oriented* strip as a strip together with the hyperplanes H_1 and H_2. So if S is a vertical strip, then it is *bounded on the left* by H_1 and *bounded on the right* by H_2. If g stabilizes S as well as H_1 and H_2 we say that g stabilizes the oriented strip. The *period of S relative to g* is the number of blocks in $S/\langle g\rangle$. The *period of S* is the minimal number of elements in such a quotient for any g stabilizing S.

We now see how dippers give rise to periodicity in the strips of the hull. We state the following for vertical dippers and note that it is equally valid for horizontal ones.

Lemma 7.1 *Suppose that $g \in G$ has an axis l parallel to the vertical direction of F. Let S be a vertical strip in* $\mathrm{Hull}(F)$*. Then there exists $n > 0$, depending only on the distance from l to S and the width of S, such that $\langle g^n\rangle$ stabilizes S.*

Proof Let S be a vertical strip of $\mathrm{Hull}(F)$ between two vertical hyperplanes J_1 and J_2. This strip is a complementary region of J_1, J_2 and a collection of hyperplane $\widehat{\mathcal{H}}$ which meet the boundary of $\mathrm{Hull}(F)$. Let $\mathcal{H}$ denote the union of the orbits of elements of $\widehat{\mathcal{H}}$ under the action of $\langle g\rangle$. This collection of hyperplanes $\mathcal{H}$ consists of two subsets:

$$\mathcal{H}_1 = \{H \in \mathcal{H} : H \text{ contains a line parallel to } l\}$$

and $\mathcal{H}_2 = \mathcal{H} \setminus \mathcal{H}_1$.

Now for each $H \in \mathcal{H}_1$, let l_H denote a line in H which is parallel to l. Since each H lies in the orbit of a hyperplane adjacent to S, there is an upper bound R on the distance from l_H to l, which depends only on the distance from l to S. By local finiteness, there is a bounded number of hyperplanes which contain a line lying in the R–neighborhood of l. Thus $\mathcal{H}_1$ is finite. We can thus choose n so that $\langle g^n\rangle$ stabilizes each element of $\mathcal{H}_1$ as well as J_1 and J_2. We now show the following claim

Claim *For each $H \in \mathcal{H}$, $g^n H \cap F \neq \varnothing$ if and only if $H \cap F \neq \varnothing$.*

To see this claim, first note that it is trivial for the elements of $\mathcal{H}_1$ since they are stabilized by g^n. So now consider an element $H \in H_2$. If $g^n H$ meets F, it must then do so in a line m not parallel to l. Thus, $g^n H$ intersects a vertical line. Since l is parallel to a vertical line in F, it follows that $g^n H$ meets l. But l is invariant under g and so H meets l and hence any vertical line in F, a contradiction. This proves the claim.

Now consider a vertical line l', equidistant from the lines $J_1 \cap F$ and $J_2 \cap F$. Since $\mathcal{H}_1$ is finite, there exists a lower bound on the distance between l' and hyperplanes in $\mathcal{H}_1$. The hyperplanes in $\mathcal{H}_2$ do not meet F, so that by Lemma 4.4, there is a lower bound on the distance between l' and hyperplanes in $\mathcal{H}_2$. Thus there is a lower bound C on the distance between l' and hyperplanes in $\mathcal{H}$. Since l' is parallel to l, which is stabilized by g, we choose some power m of n so that $d(l', g^m(l')) < C$. We then have that no hyperplanes of $\mathcal{H}$ separate l' and $g^m(l')$.

Now for each hyperplane $H \in \mathcal{H} \cup \{J_1, J_2\}$, let H^+ denote the half space containing l'. From the above, it follows that $g^m(H^+) = (g^m(H))^+$. Since

$$S = J_1^+ \cap J_2^+ \cap \bigcap_{H \in \mathcal{H}} H^+,$$

it follows that $g^m(S) = S$, as required. □

This gives the following corollary

Corollary 7.2 *Suppose that H is a vertical n–dipper. Let l be a line in F, such that $d(l, F) \leq n$. Then there exist a vertical strip S in $\mathrm{Hull}(F)$ between vertical hyperplanes H_1 and H_2 with the following properties:*

- *S contains l*
- *the width of S is bounded by a constant depending only on n*
- *$\{H, H_1, H_2\}$ form a facing triple*
- *there exists a number $m(n) > 0$ such that S has period less than m.*

Proof By Proposition 6.2, there exists a bound (depending only on δ) on the number of vertical hyperplanes crossed by H. Thus, we may find vertical hyperplanes H_1 and H_2 intersecting F in lines on either side of l, so that H_1 and H_2 are a bounded distance apart and which are disjoint from H. By Lemma 4.3 , we may further choose H_1 and H_2 to be disjoint. The cyclic facing triple condition applied to H, H_1, H_2 gives us that there is cyclic element which stabilizes all three hyperplanes. Moreover, we have a bound on the distances between these hyperplanes, so that we obtain a bound on the translation length of their common stabilizing element. By Remark 6.4, this infinite order element has an axis which must be parallel to a vertical line in F. We now apply Lemma 7.1. □

7.2 One periodic strip suffices.

The aim of this section is to prove the following.

Theorem 7.3 *Suppose that* $\mathrm{Hull}(F)$ *contains a periodic strip. Then* G *contains a* $\mathbb{Z}\times\mathbb{Z}$ *subgroup.*

Without loss of generality, we may assume that the strip is vertical. We prove several lemmas, which we will need in the course of the proof of this theorem. First, we note that by Lemma 7.1, if one vertical strip is periodic then they all are. Thus we may assume that all vertical strips are periodic. Now we see that to get a $\mathbb{Z}\times\mathbb{Z}$ in G, all we need are two equivalent vertical strips. (Two strips are *equivalent* if they are in the same G–orbit as oriented strips.)

Lemma 7.4 *Suppose that* $\mathrm{Hull}(F)$ *contains two equivalent vertical strips. Then* G *contains a* $\mathbb{Z}\times\mathbb{Z}$ *subgroup.*

Proof Suppose that S and S' are equivalent vertical strips. Let g denote an infinite order element in the stabilizer of S. Since S is periodic, for any strip wider than S which includes it, there exists a power of g which stabilizers it. It follows that there exists a power $h=g^n$ which stabilizes both S and S'. Now let $k\in G$ be an element such that $kS=S'$.

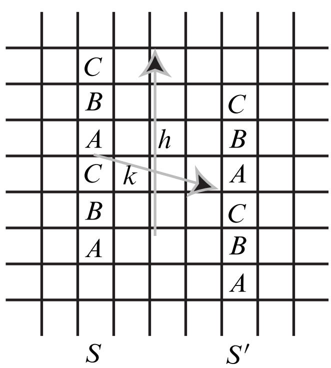

Note that h and $k^{-1}hk$ both stabilize S, so that h and $k^{-1}hk$ are commensurable. Thus there exists $m,n\in Z$ such that $h^m=k^{-1}h^nk$. Since X is a CAT(0) group, we have that $m=\pm n$. In either case, we obtain that h^m and k^2 commute.

Finally, we need to see that h^m and k^2 generate a $\mathbb{Z} \times \mathbb{Z}$. To see this, note first that k carries the oriented vertical strip S to the oriented strip S'. Let H be the hyperplane bounding S so that S and S' are both contained in closure of the same halfspace defined by H. Thus k carries this halfspace into itself. This implies that k is an infinite order element. It further implies that the axis of k is transverse to H, which means that k and h^m have non-parallel axes. Thus k and h^m generate a $\mathbb{Z} \times \mathbb{Z}$ subgroup, as required. $\square$

Now consider an oriented vertical strip S_0 stabilized by some infinite order element g. Let l denote the line running through the middle of $S_0 \cap F$; it divides F into two half planes F_R, the right half-plane, and F_L, the left half-plane. We will need to examine what g does to F_R or F_L. To this end, we have the following lemma.

Lemma 7.5 *Suppose that $g(F_R)$ is contained in some tubular neighborhood of F_R. Then there are two equivalent vertical strips.*

Proof Let Hull(F_R) denote the convex hull of F_R (that is, the intersection of all the halfspaces containing F_R). Since g has an axis parallel to l, for any $\epsilon > 0$, there exists n such that $g^n(F_R) \subset N_\epsilon(F_R)$. Now as in Lemma 4.4, we can not have F_R too close to a hyperplane. Thus, we may choose n sufficiently large so that no hyperplane separates F_R and $g(F_R)$. It follows that F_R and $g^n(F_R)$ have the same convex hull, which means that $g^n \in \mathrm{stab}(\mathrm{Hull}(F_R))$. We thus have that g^n stabilizes all the vertical strips. Since there are finitely many orbits of strips of a given period, we have that two of them are in the same orbit and hence equivalent. $\square$

Proof of Theorem 7.3 We assume that the periodic strip in question is vertical, so that all the vertical strips are periodic. The goal will be to produce two equivalent vertical strips and then apply Lemma 7.4.

We let S_0 be a vertical strip bounded on the left by H_L and on the right by H_R, and we consider a vertical halfspace F_R as in Lemma 7.5. By Lemma 7.5 we may assume that $g(F_R)$ is not contained in any tubular neighborhood of F_R. Note also that $g(F_R)$ cannot lie in a tubular neighborhood of F_L since it preserves the oriented strip S_0 and hence H_L and H_R. Thus $g(F_R)$ does not lie in any tubular neighborhood of F.

It follows that there exists some line $l_1 \in g(F_R)$ such that:

- l_1 is the intersection of some hyperplane H_1 with $g(F_R)$,
- l_1 is parallel to l (as all lines in F_R are parallel to l),
- l_1 is separated from l by H_R,

- l_1 is not contained in the M–neighborhood of F, where $M = M(\delta)$ is the upper bound on the distance between two parallel geodesics in a CAT(0), δ–hyperbolic space,
- l_1 is contained in the N–neighborhood of F, where $N = M + k$, where k is the maximal distance between two lines in a parallelism class for F.

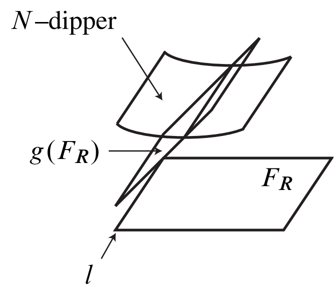

From the above, it follows that H_1 is an N–dipper relative to F. Moreover, there exists a line $l_2 \in F$ to the right of $H_R \cap F$ which satisfies the hypothesis of Corollary 7.2. Thus by Corollary 7.2, there exists a vertical hyperplane H_2 to the right of H_R which bounds a strip S_1 of period less than $m(N)$. We then repeat this argument with S_1 and produce a sequence of vertical strips all of bounded period. Thus two of them are equivalent and so by Lemma 7.4, we have a $\mathbb{Z} \times \mathbb{Z}$ subgroup as required.

We now repeat this argument with the strip S_1. □

7.3 Producing a single periodic strip

The aim of this section will be to produce a periodic strip. By the previous section, this will imply that there exists a $\mathbb{Z} \times \mathbb{Z}$ subgroup.

Theorem 7.6 *If X contains a flat plane, then X contains a flat plane whose hull contains a periodic strip.*

Proof In order to produce a periodic strip, we need to show that there exists a dipper. In order to do so we will need a more refined notion of a dipper. Given a flat F and a parallelism class $\mathcal{L}$ in F, a (k,n)*–dipper* relative to $\mathcal{L}$ is a hyperplane which contains a geodesic segment of length k, which in turn is contained in an n–neighborhood of a geodesic segment of F parallel to $\mathcal{L}$.

The strategy of our proof will be to examine the following two cases:

(1) There exist arbitrarily long vertical (k,n)–dippers. That is, there exists some fixed n, a sequence $k_i \to \infty$, and a sequence of (k_i,n)–dippers. In this case, we use a limiting argument to produce a vertical dipper.

(2) There exists an some $n > 0$ and an upper bound M, such that any (k,n)–dipper has $k < M$. Here we choose a horizontal strip of width much larger than k_0. We then seek horizontal periodicity.

So suppose that there exists a number $M > 0$ so that there exist vertical (k, M)–dippers with arbitrarily large k. This means that for each $n > n_0$, there exists a vertical strip S_n of length n, and a hyperplane H_n, which satisfies

(1) $H_n \cap F = \varnothing$.

(2) F contains a vertical geodesic segment l_n of length n, and H_n contains a geodesic segment l'_n of length n, such that $d_{Haus}(l_n, l'_n) < M$.

By translation we may assume that $S_n \subset S_{n+1}$. This comes at the cost of changing of F, so that S_n is a strip in a convex hull of a flat F_n. Now F_n limits on a flat F which has an M–dipper. So by Corollary 7.2, one of the vertical strips is periodic and we are done.

Thus, suppose that there exists a bound K, such that all the vertical (k, M)–dippers are of length $k < K$. We choose a horizontal strip S of width w much larger than K (we will say later how much larger). Now S is of finite width, so there exist two equivalent horizontal segments S_1 and S_2 along S. We consider the group element g which carries S_1 to S_2. Now if g carried a horizontal line to one parallel to a horizontal line, then we would have horizontal periodicity, by Lemma 7.1. So we can assume that if l is a horizontal line, $g(l)$ is not parallel to a horizontal line. Now if $g(l)$ is parallel to another line in F, then by choosing the the strips S_1 and S_2 sufficiently far apart, we would obtain two parallelism classes in $F \cap \mathcal{H}$ with arbitrarily small angles between them, contradicting the fact that there are finitely many parallelism classes of lines of intersection of $F \cap \mathcal{H}$.

Thus we may assume that for any horizontal line l, $g(l)$ escapes every neighborhood of F.

We need some names for some objects now. Let $F' = gF$. Let H_R denote the vertical hyperplane bounding S_1 on the right. So gH_R bounds gS_1 on the right. The image under g of the vertical direction in F is called the vertical direction in F'. Let H_1 and H_2 denote the two horizontal hyperplanes which bound the strip S. For $i = 1, 2$, let $l_i = H_i \cap F$ and $l'_i = H_i \cap F'$.

We now consider the hyperplane H_1, which, as we recall, is hyperbolic. Now since l'_i is escaping F, and there exists a lower bound on the angle between l_i and the vertical hyperplanes in F'. There will be one such vertical hyperplane H in F', such that the line $l' = H \cap F'$ meets H_1 in a point distance larger than M from F , which ensures that l_1 and H do not intersect. We know, without loss of generality, that this happens in H_1 before it happens in H_2. Now we claim that H does not meet F. For suppose $H \cap F = m$. Since H does not meet l_1, it follows that m is a horizontal line in F. But now F meets H along m and is within M of H along the line segment between $F' \cap H$ between H_1 and H_2. If the strip S is wide enough, this contradicts Lemma 4.5. So now by choosing w large enough, we obtain that H is a (k, M) for $k > K$, a contradiction. □

8 Application to 3–manifolds

Given a group G and a collection of codimension-1 subgroups, the construction given by the first author in [7], produces an action of G on a CAT(0) cube complex C whose hyperplanes have stabilizers commensurable with the codimension-1 subgroups. One of the most geometric examples of this construction arises from a manifold and a collection of immersed codimension-1 submanifolds that lift to 2–sided embeddings in the universal cover. In this section we examine the cube complex C which arises when the manifold is 3–dimensional.

Let us first examine the model situation that would arise most naturally from the construction in [7]. We emphasize that though M is 3–dimensional, the cube complex C would usually have dimension $\gg 3$.

Theorem 8.1 *Let M be a 3–manifold, and suppose that $G = \pi_1 M$ acts properly on a CAT(0) cube complex C, with a G–equivariant map $\phi\colon \widetilde{M} \to C$ with the property that for each hyperplane $H \subset C$, the preimage $\phi^{-1}(H)$ is a nonempty simply-connected surface. Then G acts on C with cyclic facing triples.*

Proof Consider a facing triple H_1, H_2, and H_3 in C. Let C' be the subspace of C bounded by this triple. Consider the equivariant map $\phi\colon \widetilde{M} \to C$. Consider the preimage $\widetilde{M}' = \phi^{-1}(C')$ which is bounded by surfaces $\widetilde{S}_i = \phi^{-1}(H_i)$. Let $K = \cap_i \mathrm{Stabilizer}(H_i)$, and let $M' = K \backslash \widetilde{M}'$. For each i let $S_i = K \backslash \widetilde{S}_i$.

Let N' be the double of M' along S_1, S_2, and S_3. Then $\pi_1 N' = K \times F_2$ where F_2 is a rank 2 free group, and K is a surface group so we will now exclude the second and third of the following possibilities:

(1) K is $\mathbb{Z}$ or $\mathbb{Z}_2$ or 1,

(2) K contains $\mathbb{Z}^2$,

(3) K contains a copy of F_2.

In the second case, $F_2 \times \mathbb{Z}^2$ would be a subgroup of $\pi_1 N'$ which is impossible, since then N' would have an infinite degree cover $\widehat{N}'$ with fundamental group $\mathbb{Z}^3$, which leads to a contradiction. Indeed, we can assume without loss of generality that N' is irreducible since $\pi_1 N'$ has no free factor, and then $\mathsf{H}_3(\widehat{N}') = 0$ which is impossible.

In the third case, note that 3–manifold fundamental groups are *coherent* (see Scott [8]) which means that every finitely generated subgroup is finitely presented. However, $\pi_1 N'$ contains $F_2 \times F_2$ which is impossible since $F_2 \times F_2$ is not coherent. Indeed, it is well-known and readily verified that the kernel of the homomorphism $F_2 \times F_2 \cong \langle a_1, a_2 \rangle \times \langle a_3, a_4 \rangle \to \mathbb{Z}$ induced by $a_i \mapsto 1$ is finitely generated but not finitely presented. It seems the earliest reference to this subgroup pathology is in the paper by Baumslag, Boone and Neumann [2] but it has been reproven numerous times especially and most recently in the context of Bestvina–Brady Morse theory (see the paper by Bestvina and Brady [3]). □

Remark In fact, one can use the construction in [7] to show that whenever G acts properly and cocompactly on a CAT(0) cube complex, then it acts properly and cocompactly on a possibly different CAT(0) cube complex satisfying the conditions of Theorem 8.1 are satisfied. As in the proof of the Theorem, one considers the equivariant map from $\widetilde{M} \to C$ as above and to produce a collection of surfaces in $\widetilde{M}$. One then builds the cube complex associated associated to this collection of surfaces. This new CAT(0) cube complex comes equipped with an action of $\pi_1(M)$ and satisfies the conditions of the Theorem 8.1.

One can use this to prove the following statement, were "sufficiently many" means a finite collection of surface subgroups so that the construction of [7] leads to a proper action on a CAT(0) cube complex.

Corollary 8.2 *An atoroidal compact* 3*–manifold with sufficiently many surface subgroups has word-hyperbolic fundamental group.*

We end with a question.

Question 8.3 Let G act faithfully and cocompactly on a CAT(0) cube complex C with cyclic facing triples. Suppose the stabilizer of each hyperplane of C is a quasiconvex subgroup of G in some appropriate sense. Then G is word-hyperbolic if and only if G does not contain a subgroup isomorphic to $\langle a, b \mid t^{-1}(a^n)t = a^m \rangle$ where $nm \neq 0$.

References

[1] **W Ballmann**, **M Brin**, *Orbihedra of nonpositive curvature*, Inst. Hautes Études Sci. Publ. Math. (1995) 169–209 (1996) MR1383216

[2] **G Baumslag**, **W W Boone**, **B H Neumann**, *Some unsolvable problems about elements and subgroups of groups*, Math. Scand. 7 (1959) 191–201 MR0163948

[3] **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470 MR1465330

[4] **M Gromov**, *Hyperbolic groups*, from: "Essays in group theory", Math. Sci. Res. Inst. Publ. 8, Springer, New York (1987) 75–263 MR919829

[5] **J Kari**, **P Papasoglu**, *Deterministic aperiodic tile sets*, Geom. Funct. Anal. 9 (1999) 353–369 MR1692474

[6] **L Mosher**, *Geometry of cubulated 3–manifolds*, Topology 34 (1995) 789–814 MR1362788

[7] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) 71 (1995) 585–617 MR1347406

[8] **G P Scott**, *Finitely generated 3–manifold groups are finitely presented*, J. London Math. Soc. (2) 6 (1973) 437–440 MR0380763

[9] **D T Wise**, *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*, PhD thesis, Princeton University (1996) MR2694733

[10] **D T Wise**, *A flat plane that is not the limit of periodic flat planes*, Algebr. Geom. Topol. 3 (2003) 147–154 MR1997317

[11] **D T Wise**, *Approximating flats by periodic flats in CAT(0) square complexes*, Canad. J. Math. 57 (2005) 416–448 MR2124924

Department of Mathematics, Technion
Haifa 32000, Israel

Mathematics and Statistics, McGill University
Montreal, Quebec, Canada H3A 2K6

sageevm@techunix.technion.ac.il, wise@math.mcgill.ca

Received: 24 January 2010