msp

Algebraic & Geometric Topology 16 (2016) 2045–2065

Resolutions of CAT(0) cube complexes and accessibility properties

BENJAMIN BEEKER
NIR LAZAROVICH

In 1985, Dunwoody defined resolutions for finitely presented group actions on simplicial trees, that is, an action of the group on a tree with smaller edge and vertex stabilizers. Moreover, he proved that the size of the resolution is bounded by a constant depending only on the group. Extending Dunwoody's definition of patterns, we construct resolutions for group actions on a general finite-dimensional CAT(0) cube complex. In dimension two, we bound the number of hyperplanes of this resolution. We apply this result for surfaces and 3–manifolds to bound collections of codimension-1 submanifolds.

20E08; 20F65

1 Introduction

An important aspect of group actions on trees is Dunwoody's theory of accessibility [6], and, in particular, finding bounds for "reasonable" actions on trees. The earliest result in this direction is Grushko's decomposition theorem [8], which implies in particular that there is a bound, depending only on the rank of the group G, on the number of edge-orbits in a G–tree with trivial edge stabilizers. An analogous result for 3–manifolds, known as the Kneser prime decomposition theorem [11] (see also Milnor [13]) implies that there is a bound, depending only on the compact 3–manifold M^3, on the number of embedded essential disjoint nonhomotopic spheres in M. In fact, Haken [9] proved that there is a bound on any collection of such two-sided subsurfaces (not necessarily spheres), under the assumption that the subsurfaces are incompressible and that the manifold is irreducible.

Grushko's result can be seen as a first result towards Dunwoody's accessibility theorem. As part of the proof, Dunwoody introduced two key tools: patterns and resolutions. He observed that any action of an almost finitely presented group G on a tree could be resolved to a G–tree obtained from a geometric pattern on the universal cover of the presentation complex of G. This resolution is simpler in certain aspects, eg the edge stabilizers are finitely generated and one can bound the number of parallelism classes of edges in the resolution. This result is known as Dunwoody's lemma [5, Lemma VI.4.4].

Published: 12 September 2016 DOI: 10.2140/agt.2016.16.2045

Sageev's seminal work on ends of group pairs [18] demonstrates how CAT(0) cube complexes can be used to generalize known results about group actions on trees. In this paper we aim to generalize Dunwoody's ideas to the realm of cube complexes.

In Section 3, we construct resolutions for cube complexes and prove the following.

Theorem 1.1 *Let G be a finitely presented group acting on a d–dimensional CAT(0) cube complex $\boldsymbol{X}$. There exists a d–dimensional CAT(0) cube complex and a G–equivariant map $F\colon \boldsymbol{X}' \to \boldsymbol{X}$ with the following properties:*

- *The hyperplane stabilizers in $\boldsymbol{X}'$ are finitely generated.*
- *Cube fixators and hyperplane stabilizers in $\boldsymbol{X}'$ are contained in those of $\boldsymbol{X}$. In particular, if the action $G \curvearrowright \boldsymbol{X}$ is free or proper then so is $G \curvearrowright \boldsymbol{X}'$.*

If, moreover, d is no more than 2, then the action of G on $\boldsymbol{X}'$ is cocompact.

In order to construct the resolution we use a d–dimensional analogue of Dunwoody's patterns, called d–patterns, which we define in Section 2. In Section 4 we restrict our attention mainly to square complexes, and obtain the following analogue of Dunwoody's lemma.

Theorem A *Let K be a 2–dimensional simplicial complex. Then there exists a constant C, depending only on K, such that any 2–pattern $\mathcal{P}$ on K has at most C parallelism classes of tracks.*

From the above, the following theorem is an immediate corollary.

Theorem 1.2 *Let G be a finitely presented group. There exists a constant C depending only on G such that for every G–action on a 2–dimensional CAT(0) cube complex $\boldsymbol{X}$, there exists a 2–dimensional CAT(0) cube complex $\boldsymbol{X}'$ with the properties of Theorem 1.1, and at most C parallelism classes of hyperplanes.* □

Remark 1.3 We note that the bound on the resolution was originally used to prove Dunwoody's accessibility theorem, but was also used by Bestvina and Feighn [2] to bound the number of edge-orbits in a reduced graph of group decompositions over small groups. We hope that our results will lead to analogous results for CAT(0) square complexes.

In Section 5, we turn to surfaces and 3–manifolds and prove the following 2–dimensional analogue of Haken's theorem.

Theorem 1.4 *Let M^n be an n–dimensional ($n = 2, 3$) compact manifold. There exists a constant C, depending only on M, such that if $\mathcal{S}$ is a collection of nonhomotopic π_1–injective co-dimension-1 two-sided embedded submanifolds such that the size of a pairwise intersecting collection of lifts to $\widetilde{M}$ is at most 2, then $|\mathcal{S}| \leq C$.*

Theorem A answers the following general question in the $d = 2$ case. However, to the best of our knowledge, it is open.

Question Is there a bound, depending only on the dimension d and the simplicial complex K, on the number of parallelism classes of tracks in d–patterns?

An affirmative answer would imply analogues of Theorem 1.2 and Theorem 1.4. Thus we also pose the following more specific question, which might be of interest on its own.

Question Is there a bound, depending only on d and the surface S, on the number of parallelism classes of curves in S such that the size of a pairwise intersecting collection of lifts to $\widetilde{S}$ is at most d?

A closely related question about collections of curves with bounded intersections on a surface was studies by Juvan, Malnič and Mohar [10]. They show that given a surface S of genus g and a number k, there is a bound on the number of curves on S that pairwise intersect at most k times. Lately this work has been further studied and some better bounds have been found by Malestein, Rivin and Theran [12], or more specifically in the case of single intersections, by Rivin [16], Przytycki [15] and Aougab and Gaster [1].

Acknowledgements The authors would like to thank Michah Sageev for his support and helpful comments. The second author is also thankful for fruitful discussions with Mladen Bestvina and Dani Wise.

The first author is supported in part by a Technion fellowship. The second author acknowledges the support received by the ETH Zürich Postdoctoral Fellowship Program and the Marie Curie Actions for People COFUND Program. While writing this paper, both authors were working at the Technion in Haifa, Israel.

2 Preliminaries

2A CAT(0) cube complexes

We begin with a short survey of definitions concerning CAT(0) cube complexes. For further details see, for example, Sageev [19].

A *cube complex* is a collection of euclidean cubes of various dimensions in which subcubes have been identified isometrically.

A simplicial complex is *flag* if every $(n+1)$–clique in its 1–skeleton spans a n–simplex. A cube complex is *nonpositively curved* if the link of every vertex is a flag simplicial complex. It is a *CAT(0) cube complex* if, moreover, it is simply connected.

A cube complex $\boldsymbol{X}$ can be equipped with two natural metrics, the euclidean and the L^1–metric. With respect to the former, $\boldsymbol{X}$ is nonpositively compact if and only if it is nonpositively compact à la Gromov; see Gromov [7] or Bridson and Haefliger [3]. The latter, on the other hand, is more natural to the combinatorial structure of CAT(0) cube complexes described below.

Given a cube $\boldsymbol{C}$ and an edge $\boldsymbol{e}$ of $\boldsymbol{C}$. The midcube of $\boldsymbol{c}$ associated to $\boldsymbol{e}$ is the convex hull of the midpoints of $\boldsymbol{e}$ and the edges parallel to $\boldsymbol{e}$. A *hyperplane* associated to $\boldsymbol{e}$ is the smallest subset containing the midpoint of $\boldsymbol{e}$ such that if it contains a midpoint of an edge it contains all the midcubes containing it. Every hyperplane $\hat{\mathfrak{h}}$ in a CAT(0) cube complex $\boldsymbol{X}$ separates $\boldsymbol{X}$ into exactly two components (see, for example, Niblo and Reeves [14]) called the *halfspaces* associated to $\hat{\mathfrak{h}}$. A hyperplane can thus also be abstractly viewed as a pair of complementary halfspaces. The *carrier* $N(\hat{\mathfrak{h}})$ of $\hat{\mathfrak{h}}$ is the union of the cubes intersecting $\hat{\mathfrak{h}}$. For a CAT(0) cube complex $\boldsymbol{X}$ we denote by $\widehat{\mathcal{H}} = \widehat{\mathcal{H}}(X)$ the set of all hyperplanes in $\boldsymbol{X}$, and by $\mathcal{H} = \mathcal{H}(X)$ the set of all halfspaces. For each halfspace $\mathfrak{h} \in \mathcal{H}$ we denote by $\mathfrak{h}^* \in \mathcal{H}$ its complementary halfspace, and by $\hat{\mathfrak{h}} \in \widehat{\mathcal{H}}$ its bounding hyperplane, which we also identify with the pair $\{\mathfrak{h}, \mathfrak{h}^*\}$.

A hyperplane in a CAT(0) cube complex *separates* two points if each one belongs to a different halfspace. Conversely two hyperplanes are *separated* by a point if there is no inclusion relation between the two halfspaces containing the point. If two hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ intersect, we write $\hat{\mathfrak{h}} \pitchfork \hat{\mathfrak{k}}$.

The *interval* between two vertices $\boldsymbol{x}$ and $\boldsymbol{y}$ of a CAT(0) cube complex is the maximal subcomplex $[\boldsymbol{x}\ \boldsymbol{y}]$ contained in every halfspace containing $\boldsymbol{x}$ and $\boldsymbol{y}$. Equivalently it can be seen as the union of all L^1–geodesics between $\boldsymbol{x}$ and $\boldsymbol{y}$.

Every interval of a d–dimensional CAT(0) cube complex admits an L^1–embedding into R^d; see Brodzki et al [4]. A hyperplane intersects the interval $[\boldsymbol{x}\ \boldsymbol{y}]$ if and only if it separates $\boldsymbol{x}$ and $\boldsymbol{y}$, and a cube belongs to the interval if every one of its hyperplanes separates them.

2B Pocsets to CAT(0) cube complex

We adopt Roller's viewpoint of Sageev's construction. Recall from Roller [17] that a *pocset* is a triple $(\mathcal{P}, \leq, *)$ of a poset $(\mathcal{P}, \leq)$ and an order-reversing involution $*\colon \mathcal{P} \to \mathcal{P}$ such that $\mathfrak{h} \neq \mathfrak{h}^*$ and $\mathfrak{h}$ and $\mathfrak{h}^*$ are incomparable for all $\mathfrak{h} \in \mathcal{P}$.

The set of halfspaces $\mathcal{H}$ of a CAT(0) cube complex has a natural pocset structure given by inclusion relation, and the complement operation $*$. Roller's construction starts with a locally finite pocset $(\mathcal{P}, \leq, *)$ of finite width (see Sageev [19] for definitions) and constructs a CAT(0) cube complex $\boldsymbol{X}(\mathcal{P})$ such that $(\mathcal{H}(X), \subseteq, *) = (\mathcal{P}, \leq, *)$. We briefly recall the construction; for more details see Roller [17] or Sageev [19].

An *ultrafilter* U on $\mathcal{P}$ is a subset verifying $\#(U \cap \{\mathfrak{k}, \mathfrak{k}^*\}) = 1$ for all $\mathfrak{k} \in \mathcal{P}$ such that for all $\mathfrak{h} \in U$, if $\mathfrak{h} \leq \mathfrak{k}$ then $\mathfrak{k} \in U$. If we denote $\widehat{\mathcal{P}} = \{\{\mathfrak{h}, \mathfrak{h}^*\} \mid \mathfrak{h} \in \mathcal{P}\}$, then U can be viewed as a choice function $U\colon \widehat{\mathcal{P}} \to \mathcal{P}$. Throughout the paper we will use both viewpoints.

An ultrafilter U satisfies the *descending chain condition (DCC)* if any descending chain $\mathfrak{k}_1 \supset \mathfrak{k}_2 \supset \cdots \supset \mathfrak{k}_n \supset \cdots$ of elements of U has finite length. The vertices of $\boldsymbol{X}(\mathcal{P})$ are the DCC ultrafilters of $\mathcal{P}$.

Two halfspaces are *compatible* if their intersection is not empty in the cube complex. A subset of $\mathcal{H}$ is an ultrafilter if and only if its halfspaces are pairwise compatible and it is maximal for this property.

2C Patterns

We adopt a somewhat similar definition for tracks as in Dunwoody [6], but we allow tracks to intersect, under some restrictions, to form d–patterns.

Definition 2.1 A *drawing* on a 2–dimensional simplicial complex K is a nonempty union of simple paths in the faces of K such that:

(1) On each face there is a finite number of paths.

(2) The two endpoints of each path are in the interior of distinct edges.

(3) The interior of a path is in the interior of a face.

(4) No two paths in a face have a common endpoint.

(5) If a point x on an edge e is an endpoint then in every face containing e, there exists a path with x as an endpoint.

A *pretrack* is a minimal drawing. A pretrack is *self-intersecting* if it contains two intersecting paths.

Denote by $\widetilde{K}$ the universal cover.

- A pretrack is a *track* if none of its pretrack lifts in $\widetilde{K}$ is self-intersecting.
- A *pattern* is a set of tracks whose union is a drawing.
- A *d–pattern* is a pattern where the size of any collection of lifts of its tracks in $\widetilde{K}$ that pairwise intersect is at most d.

We will sometimes view a pattern as the unions of its tracks in K.

2D The coarse and fine pocset structures associated to a pattern

Let $\widetilde{\mathcal{P}}$ be a pattern on a simply connected 2–simplex $\widetilde{K}$. For each track $\tilde{t}$ of $\widetilde{\mathcal{P}}$, the set $\widetilde{K} \setminus \tilde{t}$ has two connected components $\mathfrak{h}_{\tilde{t}}^{f}$ and $\mathfrak{h}_{\tilde{t}}^{f*}$; see Dunwoody [6]. We call these components the *fine halfspaces defined by* $\tilde{t}$, and the collection of all fine halfspaces is denoted by $\mathcal{H}^f = \mathcal{H}^f(\mathcal{P})$. This collection forms a locally finite pocset with respect to inclusion and complement operation $*$. If, moreover, $\widetilde{\mathcal{P}}$ is a d–pattern, then $\mathcal{H}^f$ has finite width. We denote by $\boldsymbol{X}^f = X(\mathcal{H}^f)$ the CAT(0) cube complex constructed from the pocset $\mathcal{H}^f$. Note that the dimension of $\boldsymbol{X}^f$ is at most d. With this definition we clearly have a bijective map sending $\tilde{t} \in \mathcal{P}$ to the hyperplane $\{\mathfrak{h}_{\tilde{t}}^{f}, \mathfrak{h}_{\tilde{t}}^{f*}\} \in \widehat{\mathcal{H}}^f = \widehat{\mathcal{H}}(\boldsymbol{X}^f)$.

We can also define the two *coarse halfspaces defined by* $\tilde{t}$ as the intersection $\mathfrak{h}_{\tilde{t}}^{c} = \widetilde{K}^0 \cap \mathfrak{h}_{\tilde{t}}^{f}$ and $\mathfrak{h}_{\tilde{t}}^{c*} = \widetilde{K}^0 \cap \mathfrak{h}_{\tilde{t}}^{f*}$ which are complementary in $\widetilde{K}^0$. The collection of all coarse halfspaces is denoted by $\mathcal{H}^c = X(\mathcal{H}^c)$. As above this set carries a locally finite pocset structure given by inclusion and complement. As before, if, moreover, $\widetilde{\mathcal{P}}$ is a d–pattern, then $\mathcal{H}^c$ has finite width, and we denote by $\boldsymbol{X}^c = X(\mathcal{H}^c)$ the CAT(0) cube complex constructed from the pocset $\mathcal{H}^c$. Note that the dimension of $\boldsymbol{X}^c$ is also at most d.

We call a connected component A of $K \setminus \mathcal{P}$ a *region* of the pattern. We define the *principal* ultrafilter corresponding to the region A to be the set $U_A = \{\mathfrak{k}^f \in \mathcal{H}^f \mid A \subseteq \mathfrak{k}^f\}$. Note that every principal ultrafilter verifies the DCC. Thus there is a map from the set of regions to $\boldsymbol{X}^{f0}$, and in particular a map from K^0 to $\boldsymbol{X}^{f0}$. The same way we define a map from K^0 to $\boldsymbol{X}^{c0}$.

Let $\phi_*\colon \mathcal{H}^f \to \mathcal{H}^c$ be the natural map $\mathfrak{h}_{\tilde{t}}^{f} \mapsto \mathfrak{h}_{\tilde{t}}^{c} = \widetilde{K}^0 \cap \mathfrak{h}_{\tilde{t}}^{f}$. The map ϕ_* respects the pocset structure and so defines $\widehat{\phi}\colon \widehat{\mathcal{H}}^f \to \widehat{\mathcal{H}}^c = \widehat{\mathcal{H}}(\boldsymbol{X}^c)$.

Definition 2.2 (parallelism) Two tracks of a pattern are *parallel* if they define the same coarse halfspaces, ie if they have the same image under the map $\widehat{\phi}$.

We have a natural map from the vertices (seen as ultrafilters) of $\boldsymbol{X}^c$ to those of $\boldsymbol{X}^f$. Indeed, the pullback of an ultrafilter by the map ϕ_* is also an ultrafilter. Thus, we can define the map $\Phi^{(0)}\colon X^{(0)}(\mathcal{H}^c) \to X^{(0)}(\mathcal{H}^f)$ by $\Phi^{(0)}(\boldsymbol{x}) = \phi_*^{-1}(\boldsymbol{x})$.

Proposition 2.3 *The map $\Phi^{(0)}$ can be extended to a canonical map $\Phi\colon \boldsymbol{X}^c \to \boldsymbol{X}^f$. Moreover, if a group G acts on K leaving the pattern $\mathcal{P}$ invariant, then G acts naturally on $\boldsymbol{X}^c$ and $\boldsymbol{X}^f$ and the map Φ is G–equivariant.*

Proof By construction, if two vertices $\boldsymbol{x}$ and $\boldsymbol{x}'$ of $\boldsymbol{X}^c$ are separated by the set of hyperplanes $\mathcal{S}$ then the set of hyperplane separating $\Phi^{(0)}(\boldsymbol{x})$ and $\Phi^{(0)}(\boldsymbol{x}')$ is $\widehat{\phi}^{-1}(\mathcal{S})$.

If two hyperplanes $\hat{\mathfrak{h}}_t^c$ and $\hat{\mathfrak{h}}_s^c$ cross, then $\hat{\mathfrak{h}}_t^f$ and $\hat{\mathfrak{h}}_s^f$ cross. Thus, given two opposite vertices in a cube $\boldsymbol{x}^c$ and $\boldsymbol{x}'^c$ separated by n pairwise intersecting hyperplanes $\{\hat{\mathfrak{h}}_1^c, \dots, \hat{\mathfrak{h}}_n^c\}$ in $\boldsymbol{X}^c$, the interval $[\Phi^{(0)}(\boldsymbol{x}^c)\, \Phi^{(0)}(\boldsymbol{x}'^c)]$ is isometric to the product cube complex

$$X(\widehat{\phi}^{-1}(\{\hat{\mathfrak{h}}_1^c, \dots, \hat{\mathfrak{h}}_n^c\})) = X(\widehat{\phi}^{-1}(\hat{\mathfrak{h}}_1^c)) \times \cdots \times X(\widehat{\phi}^{-1}(\hat{\mathfrak{h}}_n^c)).$$

Given a hyperplane $\hat{\mathfrak{h}}^c$ we define the map

$$\psi_{\hat{\mathfrak{h}}^c} \colon X(\{\hat{\mathfrak{h}}^c\}) \to X(\phi_*^{-1}(\{\hat{\mathfrak{h}}^c\}))$$

sending the vertices (seen as ultrafilters) $\{\mathfrak{h}^c\}$ and $\{\mathfrak{h}^{c*}\}$ to $\phi_*^{-1}(\mathfrak{h}^c)$ and $\phi_*^{-1}(\mathfrak{h}^{c*})$ and the edge $X(\{\hat{\mathfrak{h}}^c\})$ to the CAT(0) geodesic in between the endpoints of its image.

Given n intersecting hyperplanes $\hat{\mathcal{K}} = \{\hat{\mathfrak{h}}_1^c, \dots, \hat{\mathfrak{h}}_n^c\}$, we define the map $\psi_{\hat{\mathcal{K}}} = \psi_{\hat{\mathfrak{h}}_1^c} \times \cdots \times \psi_{\hat{\mathfrak{h}}_n^c}$ from the cube $X(\hat{\mathcal{K}})$ to the product

$$X\big(\phi_*^{-1}(\hat{\mathcal{K}})\big) = X\big(\widehat{\phi}^{-1}(\{\hat{\mathfrak{h}}_1^c\})\big) \times \cdots \times X\big(\widehat{\phi}^{-1}(\{\hat{\mathfrak{h}}_n^c\})\big).$$

We now define Φ from the cube $[\boldsymbol{x}^c\, \boldsymbol{x}'^c]$ to $[\Phi^{(0)}(\boldsymbol{x}^c)\, \Phi^{(0)}(\boldsymbol{x}'^c)]$ as the following composition:

$$[\boldsymbol{x}^c\, \boldsymbol{x}'^c] \xrightarrow{\sim} X\big(\widehat{\mathcal{H}}([\boldsymbol{x}^c\, \boldsymbol{x}'^c])\big) \xrightarrow{\psi_{\hat{\mathcal{K}}}} X\big(\phi_*^{-1}(\widehat{\mathcal{H}}([\boldsymbol{x}^c\, \boldsymbol{x}'^c]))\big) \xrightarrow{\sim} [\Phi^{(0)}(\boldsymbol{x}^c)\, \Phi^{(0)}(\boldsymbol{x}'^c)].$$

It is straightforward to verify that this extends the map $\Phi^{(0)}$. The G–equivariance follows from the canonicity of the map. □

3 Resolutions

Let G be a finitely presented group. Let K be a fixed finite triangle complex such that $G \simeq \pi_1(K, v_0)$ for some $v_0 \in K$. Let $\{\widetilde{v}_0, \dots, \widetilde{v}_l\}$, $\{\widetilde{e}_0, \dots, \widetilde{e}_m\}$ and $\{\widetilde{f}_1, \dots, \widetilde{f}_n\}$ be sets of representatives for the G–orbits of 0–, 1– and 2–cells in $\widetilde{K}$.

Now, let G act on a d–dimensional CAT(0) cube complex $\boldsymbol{X}^o$. For each $\widetilde{v}_i$ choose a vertex $\boldsymbol{x}_i^o$ in $\boldsymbol{X}^o$. Since G acts freely on $\widetilde{K}$, one can extend the map $\widetilde{v}_i \mapsto \boldsymbol{x}_i^o$ to a G–equivariant map $f \colon \widetilde{K}^{(0)} \to \boldsymbol{X}^{o(0)}$ by sending $g\widetilde{v}_i$ to $g\boldsymbol{x}_i^o$. We extend this map to a map on the 1–skeleton of $\widetilde{K}$ by sending each edge representative $\widetilde{e}_i$ linearly to a combinatorial geodesic connecting the images of its endpoints $f(i(\widetilde{e}_i))$ and $f(t(\widetilde{e}_i))$ in $\boldsymbol{X}^o$, and extend G–equivariantly to $\widetilde{K}^{(1)}$. Similarly, extend the map to a G–equivariant map $f \colon K \to \boldsymbol{X}^o$ by sending the 2–cells $\widetilde{f}_i$ to disks whose boundary is $f(\partial \widetilde{f}_i)$, which exist because $\boldsymbol{X}^o$ is simply connected. We may further assume that the image of f is in $\boldsymbol{X}^{o(2)}$, the 2–skeleton of $\boldsymbol{X}^o$, is transverse to the hyperplanes and

has minimal number of squares of $\boldsymbol{X}^o$. Such a map is called a *minimal disk*. For more details see Sageev [18].

Lemma 3.1 *Let $\widetilde{\mathcal{P}} = \bigcup_{\hat{\mathfrak{h}}^o \in \hat{\mathcal{H}}^o} f^{-1}(\hat{\mathfrak{h}}^o)$ be the pullback of the hyperplanes of $\boldsymbol{X}^o$ to $\widetilde{K}$. The set $\widetilde{\mathcal{P}}$ induces a d–pattern $\mathcal{P}$ on K.*

Proof Note that the pullback of each $\hat{\mathfrak{h}}^o \in \hat{\mathcal{H}}^o$ defines a 1–pattern on $\widetilde{K}$; see [5].

We are left to show that the size of a collection of pairwise crossing tracks is at most d. Let $\tilde{t}_1, \dots, \tilde{t}_k$ be distinct pairwise intersecting tracks in $\widetilde{\mathcal{P}}$. Each $\tilde{t}_i$ maps into the corresponding hyperplane $\hat{\mathfrak{h}}^o_i$. By the transversality to the hyperplanes, the hyperplanes $\hat{\mathfrak{h}}^o_1, \dots, \hat{\mathfrak{h}}^o_k$ are distinct intersecting hyperplanes. Thus $k \leq d$. □

The d–pattern defines the fine cube complex $\boldsymbol{X}^f$ (see Section 2) on which G acts. Note that the map f induces a map, which we denote by f_*, from $\mathcal{H}^f$, the set of halfspaces of $\boldsymbol{X}^f$, to $\mathcal{H}^o$. This map respects the complement operation, thus defines a map $\hat{f}_*\colon \hat{\mathcal{H}}^f \to \hat{\mathcal{H}}^o$ on hyperplanes. Note also that the image of $\hat{f}_*$ consists of all the hyperplanes that divide nontrivially the image of $f(K)$.

Proposition 3.2 *There exists a G–equivariant combinatorial map $F\colon \boldsymbol{X}^f \to \boldsymbol{X}^o$, which induces the map f_* on halfspaces.*

Proof Let us first define the map F on the vertices of $\boldsymbol{X}^f$. Let $\boldsymbol{x}^f$ be a vertex of $\boldsymbol{X}^f$, ie a DCC ultrafilter on the halfspaces, which we regard as the map $\boldsymbol{x}^f\colon \hat{\mathcal{H}}^f \to \mathcal{H}^f$ and which assigns to each hyperplane its halfspace that contains $\boldsymbol{x}^f$. We define $F(\boldsymbol{x}^f)$ as follows: for each $\hat{\mathfrak{h}}^o \in \hat{\mathcal{H}}^o$ either $\hat{\mathfrak{h}}^o$ belongs to the image of $\hat{f}_*$ or not. In the former case, let $\hat{\mathfrak{h}}^f \in \hat{\mathcal{H}}^f$ be a minimal hyperplane with respect to $\boldsymbol{x}^f$ among $(\hat{f}_*)^{-1}(\hat{\mathfrak{h}}^o)$, and define $F(\boldsymbol{x}^f)(\hat{\mathfrak{h}}^o) = f_*(\boldsymbol{x}^f(\hat{\mathfrak{h}}^f))$. In the latter case we choose $\boldsymbol{x}^o(\hat{\mathfrak{h}}^o)$ to be the halfspace $\mathfrak{h}^o$ which contains $f(K)$.

The map F does not depend on the choice of $\hat{\mathfrak{h}}^f$: for every $\hat{\mathfrak{h}}^o \in \hat{\mathcal{H}}^o$ and every $\hat{\mathfrak{h}}^f_1, \hat{\mathfrak{h}}^f_2 \in \hat{f}_*^{-1}(\hat{\mathfrak{h}}^o)$ minimal with respect to $\boldsymbol{x}^f$ among $\hat{f}_*^{-1}(\hat{\mathfrak{h}}^o)$, $f_*(\hat{\mathfrak{h}}^f_1) = f_*(\hat{\mathfrak{h}}^f_2)$. Otherwise a geodesic path that connects them in $\widetilde{K}$ is mapped by f to a path that passes from $\mathfrak{h}^o$ to $\mathfrak{h}^{o*}$ in $\boldsymbol{X}^o$ without crossing $\hat{\mathfrak{h}}^o$.

The map $F(\boldsymbol{x}^f)\colon \hat{\mathcal{H}}^o \to \mathcal{H}^o$ is an ultrafilter: let $\hat{\mathfrak{h}}^o, \hat{\mathfrak{k}}^o \in \hat{\mathcal{H}}^o$ which do not cross, let $\hat{\mathfrak{h}}^f, \hat{\mathfrak{k}}^f$ be the hyperplanes in the definition of $F(\boldsymbol{x}^f)$, and let $\mathfrak{h}^f, \mathfrak{k}^f$ be their orientation in $\boldsymbol{x}^f$. Clearly the orientation of the halfspaces $\mathfrak{h}^f, \mathfrak{k}^f \subset \widetilde{K}$ is such that they have a nontrivial intersection (otherwise, $\boldsymbol{x}^f$ is not an ultrafilter). If p is a point in this intersection then both $F(\boldsymbol{x}^f)(\hat{\mathfrak{h}}^o), F(\boldsymbol{x}^f)(\hat{\mathfrak{k}}^o)$ contain $f(p)$, showing that they form a compatible pair.

The map $F\colon \boldsymbol{X}^{f\,(0)} \to \boldsymbol{X}^{o(0)}$ extends to $F\colon \boldsymbol{X}^f \to \boldsymbol{X}^o$: if $\boldsymbol{x}_1^f, \dots, \boldsymbol{x}_{2^k}^f$ are the vertices of a k–dimensional cube in $\boldsymbol{X}^f$, ie their ultrafilters differ on exactly k hyperplanes $\hat{\mathfrak{h}}_1^f, \dots, \hat{\mathfrak{h}}_k^f$, then their images will differ exactly on the collection of distinct pairwise transverse hyperplanes $\hat{f}_*(\hat{\mathfrak{h}}_1^f), \dots, \hat{f}_*(\hat{\mathfrak{h}}_k^f)$. □

The pair $(\boldsymbol{X}^f, F)$ is called a *fine resolution* of $\boldsymbol{X}^o$. In Section 2, we also constructed the coarse CAT(0) cube complex $\boldsymbol{X}^c$ and a G–equivariant map $\Phi\colon \boldsymbol{X}^c \to \boldsymbol{X}^f$ between the two complexes. The pair $(\boldsymbol{X}^c, F\circ\Phi)$ is called a *coarse resolution* of $\boldsymbol{X}^o$. Note that both resolution depend on the choice of K and on the equivariant map $f\colon K \to \boldsymbol{X}^o$.

Proposition 3.3 *The fine resolution $(\boldsymbol{X}^f, F)$ has the following properties:*

- *The hyperplane stabilizers in $\boldsymbol{X}^f$ are finitely generated.*
- *Cube and hyperplane stabilizers in $\boldsymbol{X}^f$ are contained in those of $\boldsymbol{X}^o$. In particular, if the action $G \curvearrowright \boldsymbol{X}^o$ is proper or free then so is $G \curvearrowright \boldsymbol{X}^f$.*

Proof The stabilizer of a track $\tilde{t} \in \widetilde{\mathcal{P}}$ is the image of $\pi_1(t)$ in $\pi_1(K) \simeq G$ under the inclusion map. The track t is a finite graph in K, and thus finitely generated.

The map $F\colon \boldsymbol{X}^f \to \boldsymbol{X}^o$ is G–equivariant and combinatorial, thus for all cube $\boldsymbol{C}^f \in \boldsymbol{X}^f$ we have $\mathrm{Stab}_G(\boldsymbol{C}^f) < \mathrm{Stab}_G(F(\boldsymbol{C}^f))$, and similarly for hyperplanes. □

Since the map $\Phi\colon \boldsymbol{X}^c \to \boldsymbol{X}^f$ is not combinatorial, the properties of the coarse resolution are slightly weaker.

Proposition 3.4 *The coarse resolution $(\boldsymbol{X}^c, F\circ\Phi)$ has the following properties:*

(1) *The hyperplane stabilizers in $\boldsymbol{X}^c$ are finitely generated.*

(2) *If the hyperplanes in $\boldsymbol{X}^o/G$ are embedded, ie for all $g \in G$ and $\hat{\mathfrak{h}}^o \in \hat{\mathcal{H}}^o$ either $g\hat{\mathfrak{h}}^o = \hat{\mathfrak{h}}^o$ or $g\hat{\mathfrak{h}}^o \cap \hat{\mathfrak{h}}^o = \varnothing$, then the oriented hyperplane stabilizers in $\boldsymbol{X}^c$, ie stabilizers of the hyperplane that do not exchange the halfspace are contained in those of $\boldsymbol{X}^o$.*

(3) *Cube fixators in $\boldsymbol{X}^c$ are contained in fixators of cubes of the same dimension in $\boldsymbol{X}^o$.*

(4) *Cube stabilizers in $\boldsymbol{X}^c$ act elliptically on $\boldsymbol{X}^o$.*

(5) *If the action $G \curvearrowright \boldsymbol{X}^o$ is proper or free then so is $G \curvearrowright \boldsymbol{X}^c$.*

Proof Note that the implications (3) $\Longrightarrow$ (4) $\Longrightarrow$ (5) are trivial.

Recall from Section 2 that the preimage $\widehat{\phi}^{-1}(\hat{\mathfrak{h}}_t^c)$ is finite, and thus the stabilizer of $\hat{\mathfrak{h}}_t^c$ is generated by those of the tracks $\widehat{\phi}^{-1}(\hat{\mathfrak{h}}_t^c)$ and finitely many elements permuting this collection. This completes the proof of (1).

If the hyperplanes in $\boldsymbol{X}^o/G$ are embedded then so are the hyperplanes in $\boldsymbol{X}^c/G$ and in $\boldsymbol{X}^f/G$, otherwise there is an element $g \in G$ and a track $\tilde{t} \in \widetilde{\mathcal{P}}$ such that $\tilde{t}$ and $g\tilde{t}$ cross, which by the definition of f would imply that $\hat{f}_*(\hat{\mathfrak{h}}_{\tilde{t}}^f)$ and $g\,\hat{f}_*(\hat{\mathfrak{h}}_{\tilde{t}}^f)$ cross.

Also note that the collection of $\phi_*^{-1}(\mathfrak{h}^c)$ is a poset on which the stabilizer of $\mathfrak{h}^c$ acts. Moreover, two hyperplanes in $\phi_*^{-1}(\mathfrak{h}^c)$ cross if and only if they are incomparable. Since no hyperplane can be sent to a crossing hyperplane, the stabilizer of $\mathfrak{h}^c$ fixes this poset. Thus the stabilizer of $\mathfrak{h}^c$ is included in the stabilizer of each of $\mathfrak{h}^f \in \phi_*^{-1}(\mathfrak{h}^c)$. This completes the proof of (2).

By the construction of the map $\Phi\colon \boldsymbol{X}^c \to \boldsymbol{X}^f$ we see that each cube $\boldsymbol{C}^c$ of $\boldsymbol{X}^c$ is affinely embedded into a product region of higher dimension than that of $\boldsymbol{C}^c$, thus a generic point in $\boldsymbol{C}^c$ is sent to a generic point of a cube of higher dimension. This completes the proof of (3). □

Though one might expect that the resolution of a cocompact G–action would be cocompact, this is not always the case as the following example shows. Let $G = \mathbb{Z}^2 = \langle e_1 = (1,0), e_2 = (0,1)\rangle$, with the presentation complex K obtained by gluing two triangles along an edge to form a square and then identifying opposite edges to form a torus. The group G acts on $\boldsymbol{X}^o = \mathbb{R}^2 \times [0,1]$ (with the standard cubulation by unit cubes) by $e_i(x,t) = (e_i + x, 1-t)$ for $i = 1,2$, ie it acts by translations on the first factor and by inversions on the second.

The pattern obtained on $\widetilde{K} = \mathbb{R}^2$ consists of 3 infinite sets of tracks of different parallelism classes of lines (see Figure 1). Therefore, the associated CAT(0) cube complexes $\boldsymbol{X}^c$ (in this case $\boldsymbol{X}^f = \boldsymbol{X}^c$) is the standard cubulation of $\mathbb{R}^3$, on which $G = \mathbb{Z}^2$ does not act cocompactly.

However, such an example cannot occur in dimension 2 (or smaller). In fact an even stronger statement holds in this case.

Proposition 3.5 *Let K be a compact triangle complex, and let $G = \pi_1(K)$. If $G \curvearrowright \boldsymbol{X}^o$, a 2–dimensional CAT(0) cube complex, then $G \curvearrowright \boldsymbol{X}^f$ cocompactly.*

Proof By assumption G acts cocompactly on $\widetilde{K}$, and the pattern on each triangle of K is finite. Thus, G acts cocompactly on the set of regions, ie the set of connected

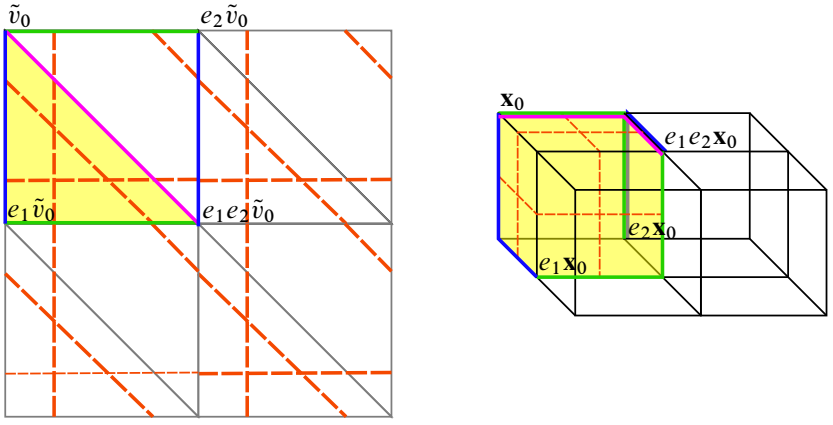

Figure 1: An example for a noncocompact resolution of a cocompact action

components of $K \setminus \mathcal{P}$. Hence, it is enough to show that every vertex $\boldsymbol{x}^f \in \boldsymbol{X}^f$ is a principal ultrafilter, ie corresponds to a region in $\widetilde{K}$.

Let $\mathcal{H}^f_{\boldsymbol{x}^f}$ be the collection of minimal halfspaces in $\boldsymbol{x}^f$. It is enough to show that the intersection

$$A_{\boldsymbol{x}^f} = \bigcap_{\mathfrak{h}^f \in \mathcal{H}^f_{\boldsymbol{x}^f}} \mathfrak{h}^f \subseteq \widetilde{K}$$

is nonempty since the intersection is a region in $\widetilde{K}$ corresponding to $\boldsymbol{x}^f$. Fix $\mathfrak{h}^f_{\tilde{t}} \in \mathcal{H}^f_{\boldsymbol{x}^f}$. There are two cases to consider:

(1) The track $\tilde{t}$ does not intersect any other track. In this case, every other $\mathfrak{h}^f_{\tilde{v}} \in \mathcal{H}^f_{\boldsymbol{x}^f}$ contains $\tilde{t}$, for otherwise $\mathfrak{h}^f_{\tilde{t}}$ is not minimal. Thus, any point in $\mathfrak{h}^f_{\tilde{t}} \subseteq \widetilde{K}$ close enough to $\tilde{t}$ will be in the intersection above.

(2) There exists $\mathfrak{h}^f_{\tilde{u}} \in \mathcal{H}^f_{\boldsymbol{x}^f}$ such that $\tilde{t}$ and $\tilde{u}$ intersect. In this case, for all $\mathfrak{h}^f_{\tilde{v}} \in \mathcal{H}^f_{\boldsymbol{x}^f}$ the track $\tilde{v}$ cannot intersect both $\tilde{t}$ and $\tilde{u}$. Hence, the corresponding fine halfspace $\mathfrak{h}^f_{\tilde{v}}$ contains either $\tilde{t}$ or $\tilde{u}$. Thus, any point in $\mathfrak{h}^f_{\tilde{t}} \cap \mathfrak{h}^f_{\tilde{u}} \subseteq \widetilde{K}$ close enough to $\tilde{t} \cap \tilde{u}$ will be in the intersection above. □

Recall from Dicks and Dunwoody [5] that a group is *almost finitely presented* if it acts freely cocompactly on a simplicial complex $\overline{K}$ with $H^1(\overline{K}, \mathbb{Z}/2\mathbb{Z}) = 0$. All of the above works when replacing $\widetilde{K}$ with $\overline{K}$. Hence, by Proposition 3.3 and Proposition 3.5 we get the following corollary.

Corollary 3.6 *Any almost finitely presented group that acts properly (resp. freely) on a 2–dimensional CAT(0) cube complex acts properly (resp. freely) and cocompactly on a 2–dimensional CAT(0) cube complex, and in particular it is finitely presented.* □

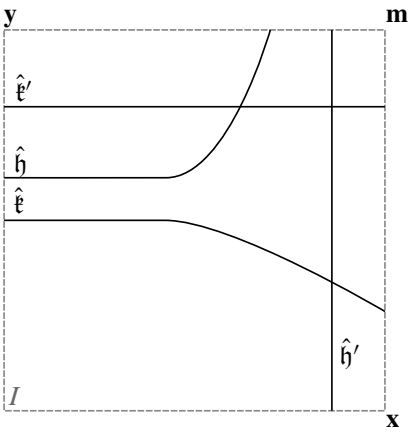

Figure 2: Parallelism obstructing pair (POP)

4 Bounding the number of tracks in a pattern

This section focuses on proving Theorem A which can be formulated as follows.

Theorem A′ *Let K be a 2–dimensional simplicial complex and $\mathcal{P}$ be a 2–pattern on K with no parallel tracks. Then there exists an integer D, depending only on K, such that the number of tracks in $\mathcal{P}$ is bounded by D.*

We begin by defining a weak notion of parallelism for adjacent hyperplanes in an interval. First, recall that a pair of noncrossing hyperplanes are *adjacent* if their carrier contains a common vertex.

Definition 4.1 Given a pair of noncrossing hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$, a *parallelism obstructing pair (POP)* is a pair of crossing hyperplanes $(\hat{\mathfrak{h}}', \hat{\mathfrak{k}}')$ such that $\hat{\mathfrak{h}}' \pitchfork \hat{\mathfrak{k}}$ but $\hat{\mathfrak{h}}' \not\pitchfork \hat{\mathfrak{h}}$, and $\hat{\mathfrak{k}}' \pitchfork \hat{\mathfrak{h}}$ but $\hat{\mathfrak{k}}' \not\pitchfork \hat{\mathfrak{k}}$; see Figure 2. It is a POP in an interval I if the four hyperplanes intersect I.

Two noncrossing hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ in an interval I are *adj-P* in I, if they are adjacent and do not have a POP in I.

In a CAT(0) cube complex we write $\hat{\mathfrak{h}} <_{\boldsymbol{x}} \hat{\mathfrak{k}}$ when $\hat{\mathfrak{h}}$ separates $\boldsymbol{x}$ from $\hat{\mathfrak{k}}$. Here, $\boldsymbol{x}$ can be a vertex or a hyperplane. If $\boldsymbol{x}$ is a vertex it is equivalent to say that the ultrafilter corresponding to $\boldsymbol{x}$ satisfies $\boldsymbol{x}(\hat{\mathfrak{h}}) < \boldsymbol{x}(\hat{\mathfrak{k}})$

Lemma 4.2 *For every d, there exists $C(d)$ such that for any interval I of a CAT(0) cube complex of dimension d and for every vertex $\boldsymbol{m}$ of I, the set of (nonordered) pairs of adj-P hyperplanes separated by $\boldsymbol{m}$ has cardinality at most $C(d)$.*

Proof We start by proving that if a pair of adj-P hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ in I is separated by $\boldsymbol{m}$, then at least one of the two hyperplanes is adjacent to $\boldsymbol{m}$. Assume not, then there exists $\hat{\mathfrak{h}}' <_{\boldsymbol{m}} \hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}' <_{\boldsymbol{m}} \hat{\mathfrak{k}}$. Since the four hyperplanes belong to the interval I and there are no facing triples in an interval, we get $\hat{\mathfrak{k}} \pitchfork \hat{\mathfrak{h}}'$, $\hat{\mathfrak{h}}' \pitchfork \hat{\mathfrak{k}}'$, and $\hat{\mathfrak{k}}' \pitchfork \hat{\mathfrak{h}}$. Contradicting that $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$ are adj-P.

Without loss of generality let $\hat{\mathfrak{h}}$ be the hyperplane adjacent to $\boldsymbol{m}$. Since I is an interval of dimension d, thus can be embedded in $\mathbb{R}^d$, there are at most $2d$ hyperplanes adjacent to $\boldsymbol{m}$ in I.

Now assume $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ and $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}}')$ are adj-P and separated by $\boldsymbol{m}$. The hyperplanes $\hat{\mathfrak{k}}$ and $\hat{\mathfrak{k}}'$ cannot be facing since otherwise $\hat{\mathfrak{h}}$, $\hat{\mathfrak{k}}$ and $\hat{\mathfrak{k}}'$ would be a facing triple of I. The hyperplanes $\hat{\mathfrak{k}}$ for which $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ is an adj-P pair separated by $\boldsymbol{m}$, pairwise cross. Hence, there are at most d of them. Thus we can set $C(d) = 2d^2$. □

Lemma 4.3 *For any pair of intervals $I_1 = [\boldsymbol{x}\ \boldsymbol{y}_1]$ and $I_2 = [\boldsymbol{x}\ \boldsymbol{y}_2]$ in a CAT(0) cube complex $\boldsymbol{X}$ of dimension 2, there are at most four pairs of hyperplanes intersecting I_1 and I_2, adj-P in I_1 but not in I_2.*

Proof Denote by $\boldsymbol{m}$ the median point of $\boldsymbol{x}$, $\boldsymbol{y}_1$ and $\boldsymbol{y}_2$. Note that every hyperplane intersecting I_1 and I_2 separates $\boldsymbol{m}$ from $\boldsymbol{x}$ and that no hyperplane separates $\boldsymbol{m}$ from $\boldsymbol{x}$ and $\boldsymbol{y}_2$. Without loss of generality we assume throughout that a pair of adj-P hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ is such that $\hat{\mathfrak{h}} <_{\boldsymbol{m}} \hat{\mathfrak{k}}$ (or equivalently $\hat{\mathfrak{k}} <_{\boldsymbol{x}} \hat{\mathfrak{h}}$).

Note that two crossing hyperplanes that intersect a common interval, cross inside the interval. Thus if a pair of hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ is adj-P in I_1 but not in I_2, then at least one of the hyperplanes of the POP $(\hat{\mathfrak{h}}', \hat{\mathfrak{k}}')$ in I_2 does not intersect I_1. Since $\hat{\mathfrak{h}}' <_{\boldsymbol{x}} \hat{\mathfrak{h}}$, the hyperplane $\hat{\mathfrak{h}}'$ has to intersect I_1. Therefore the hyperplane $\hat{\mathfrak{k}}'$ does not intersect I_1 and hence separates $\boldsymbol{y_2}$ from $\boldsymbol{x}$ and $\boldsymbol{m}$ (see Figure 3).

We claim that for a pair of hyperplanes $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}})$ which are adj-P in I_1 but not in I_2,

(1) there are at most two possibilities for the hyperplane $\hat{\mathfrak{h}}$, and

(2) for each $\hat{\mathfrak{h}}$ there are at most two possibilities for the hyperplane $\hat{\mathfrak{k}}$.

This will prove that there are at most four pairs of hyperplanes which are adj-P in I_1 but not in I_2.

We begin by proving the second claim. We observe that if $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}}_1)$ and $(\hat{\mathfrak{h}}, \hat{\mathfrak{k}}_2)$ are pairs of adjacent hyperplanes with $\hat{\mathfrak{k}}_1 \neq \hat{\mathfrak{k}}_2$, then $\hat{\mathfrak{k}}_1 \pitchfork \hat{\mathfrak{k}}_2$ since they are both adjacent to $\hat{\mathfrak{h}}$. The second claim follows since the CAT(0) cube complex is 2–dimensional.

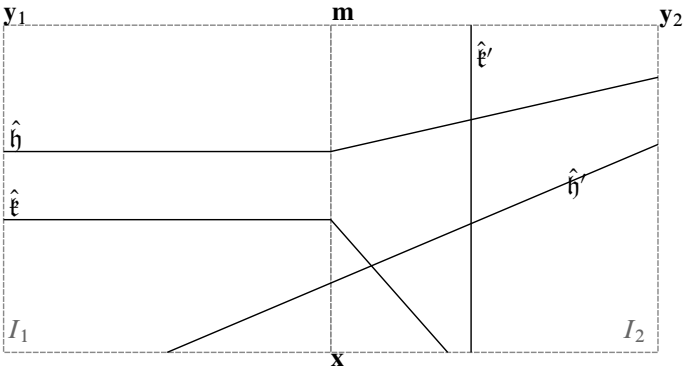

Figure 3: The hyperplane $\hat{\mathfrak{k}}'$ has to separate $\boldsymbol{m}$ and $\boldsymbol{y}_2$

To prove the first claim, let $\hat{\mathfrak{h}}_1 \neq \hat{\mathfrak{h}}_2$, and let $(\hat{\mathfrak{h}}_1, \hat{\mathfrak{k}}_1)$ and $(\hat{\mathfrak{h}}_2, \hat{\mathfrak{k}}_2)$ be two pairs of hyperplanes which are adj-P in I_1 but not in I_2. As before, the claim would follow from the dimension of the complex once we prove that $\hat{\mathfrak{h}}_1 \pitchfork \hat{\mathfrak{h}}_2$. We do so by contradiction.

Without loss of generality we may assume $\hat{\mathfrak{h}}_1 <_{\boldsymbol{m}} \hat{\mathfrak{h}}_2$.

The pairs $(\hat{\mathfrak{h}}_1, \hat{\mathfrak{k}}_1)$ and $(\hat{\mathfrak{h}}_2, \hat{\mathfrak{k}}_2)$ are adjacent and noncrossing in I_1 and thus also in I_2. Therefore, there exist POPs $(\hat{\mathfrak{h}}'_1, \hat{\mathfrak{k}}'_1)$ and $(\hat{\mathfrak{h}}'_2, \hat{\mathfrak{k}}'_2)$ for $(\hat{\mathfrak{h}}_1, \hat{\mathfrak{k}}_1)$ and $(\hat{\mathfrak{h}}_2, \hat{\mathfrak{k}}_2)$, respectively. Moreover they can be chosen such that $\hat{\mathfrak{k}}'_1$ and $\hat{\mathfrak{k}}'_2$ are minimal with respect to $\boldsymbol{x}$ among all other POPs.

From $\hat{\mathfrak{h}}_1 <_{\boldsymbol{m}} \hat{\mathfrak{h}}_2$ it follows that the interval $[\boldsymbol{m}\ \boldsymbol{y}_2]$ and the hyperplane $\hat{\mathfrak{h}}_2$ are separated by $\hat{\mathfrak{h}}_1$. The hyperplane $\hat{\mathfrak{k}}'_2$ intersects both $[\boldsymbol{m}\ \boldsymbol{y}_2]$ and $\hat{\mathfrak{h}}_2$ hence it crosses $\hat{\mathfrak{h}}_1$. Thus, since $\boldsymbol{X}$ is 2–dimensional, $\hat{\mathfrak{k}}'_1$ and $\hat{\mathfrak{k}}'_2$ cannot cross. We have two cases:

(1) If $\hat{\mathfrak{k}}'_2 <_{\boldsymbol{x}} \hat{\mathfrak{k}}'_1$, then the interval $[\boldsymbol{m}\,\boldsymbol{x}]$ and the hyperplane $\hat{\mathfrak{k}}'_1$ are separated by $\hat{\mathfrak{k}}'_2$. Thus $\hat{\mathfrak{h}}'_1$ and $\hat{\mathfrak{k}}'_2$ cross. Again, because the CAT(0) cube complex is 2–dimensional, the hyperplanes $\hat{\mathfrak{k}}'_2$ and $\hat{\mathfrak{k}}_1$ cannot cross. Thus the pair $(\hat{\mathfrak{h}}'_1, \hat{\mathfrak{k}}'_2)$ is a POP for $(\hat{\mathfrak{h}}_1, \hat{\mathfrak{k}}_1)$, contradicting the minimality of $\hat{\mathfrak{k}}'_1$.

(2) If $\hat{\mathfrak{k}}'_1 \leq_{\boldsymbol{x}} \hat{\mathfrak{k}}'_2$, first notice that $\hat{\mathfrak{h}}_2 \not>_{\boldsymbol{x}} \hat{\mathfrak{k}}_1$ since otherwise $\hat{\mathfrak{h}}_2$ would separate the adjacent hyperplanes $\hat{\mathfrak{k}}_1$ and $\hat{\mathfrak{h}}_1$. Therefore, either $\hat{\mathfrak{h}}_2 \leq_{\boldsymbol{x}} \hat{\mathfrak{k}}_1$ or $\hat{\mathfrak{h}}_2 \pitchfork \hat{\mathfrak{k}}_1$. We proceed by showing that neither of these is possible.

If $\hat{\mathfrak{h}}_2 \leq_{\boldsymbol{x}} \hat{\mathfrak{k}}_1$, then we have $\hat{\mathfrak{h}}_2 \leq_{\boldsymbol{x}} \hat{\mathfrak{k}}_1 <_{\boldsymbol{x}} \hat{\mathfrak{k}}'_1 \leq_{\boldsymbol{x}} \hat{\mathfrak{k}}'_2$, which contradicts $\hat{\mathfrak{h}}_2 \pitchfork \hat{\mathfrak{k}}'_2$.

Suppose $\hat{\mathfrak{h}}_2 \pitchfork \hat{\mathfrak{k}}_1$. From $\hat{\mathfrak{h}}'_2 <_{\boldsymbol{x}} \hat{\mathfrak{h}}_2 \pitchfork \hat{\mathfrak{k}}_1$ we have $\hat{\mathfrak{k}}_1 \not\leq_{\boldsymbol{x}} \hat{\mathfrak{h}}'_2$, and from $\hat{\mathfrak{k}}_1 <_{\boldsymbol{x}} \hat{\mathfrak{k}}'_1 \leq \hat{\mathfrak{k}}'_2 \pitchfork \hat{\mathfrak{h}}'_2$ we have $\hat{\mathfrak{h}}'_2 \not\leq_{\boldsymbol{x}} \hat{\mathfrak{k}}_1$. Thus $\hat{\mathfrak{h}}'_2 \pitchfork \hat{\mathfrak{k}}_1$. Hence, the pair $(\hat{\mathfrak{k}}_1, \hat{\mathfrak{h}}'_2)$ forms a POP for $(\hat{\mathfrak{h}}_2, \hat{\mathfrak{k}}_2)$ in I_1 (see Figure 4), contradicting our assumptions. □

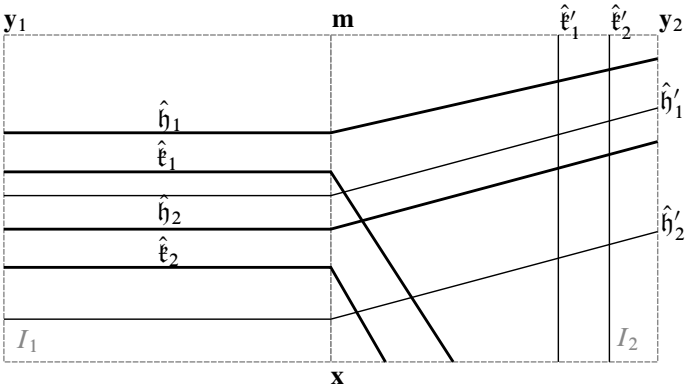

Figure 4: Configuration in the case $\hat{\mathfrak{k}}'_1 \leq_{\boldsymbol{m}} \hat{\mathfrak{k}}'_2$

Proof of Theorem A′ Let $\widetilde{K}$ be the universal cover of K and $\widetilde{\mathcal{P}}$ the pattern on $\widetilde{K}$ associated to $\mathcal{P}$. Let $\boldsymbol{X}^c$ be the coarse CAT(0) cube complex associated to the pattern $\widetilde{\mathcal{P}}$. Since $\mathcal{P}$ is a 2–pattern, the CAT(0) cube complex $\boldsymbol{X}^c$ is a square complex.

For a vertex $\tilde{x}$ in $\widetilde{K}$ call $\bar{\boldsymbol{x}}^c$ the corresponding vertex in $\boldsymbol{X}^c$. Similarly the hyperplane corresponding to a track $\tilde{t}$ in $\widetilde{\mathcal{P}}$ is called $\hat{\mathfrak{h}}^c_{\tilde{t}}$. A triangle in $\boldsymbol{X}^c$ is a triplets of vertices $(\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c, \bar{\boldsymbol{z}}^c)$ coming from a triangle $(\tilde{x}, \tilde{y}, \tilde{z})$ of $\widetilde{K}$.

Two tracks $\tilde{t}$ and $\tilde{t}'$ of $\widetilde{\mathcal{P}}$ are adj-P if they cross an edge $[\tilde{x}, \tilde{y}]$ such that $\hat{\mathfrak{h}}^c_{\tilde{t}}$ and $\hat{\mathfrak{h}}^c_{\tilde{t}'}$ are adj-P in the interval defined by $\bar{\boldsymbol{x}}^c$ and $\bar{\boldsymbol{y}}^c$.

Note that if a pair of hyperplanes $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ in $\boldsymbol{X}^c$ intersect an interval $[\bar{\boldsymbol{v}}^c, \bar{\boldsymbol{w}}^c]$ in which they are adj-P but not parallel, then:

(1) either there exists some triangle $(\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c, \bar{\boldsymbol{z}}^c)$ such that $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is adj-P in $[\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c]$ but is separated by the midpoint of $(\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c, \bar{\boldsymbol{z}}^c)$, or

(2) there exists some triangle $(\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c, \bar{\boldsymbol{z}}^c)$ such that $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is adj-P in $[\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{y}}^c]$, intersects $[\bar{\boldsymbol{x}}^c, \bar{\boldsymbol{z}}^c]$ but is not adj-P in it.

Indeed, if they are not parallel, there exists a triangle in which they are separated. Take a sequence of triangles $T_1, \dots, T_n$ such that the vertices $\bar{\boldsymbol{v}}^c$ and $\bar{\boldsymbol{w}}^c$ are vertices of T_1, two consecutive triangles T_i and T_{i+1} share an edge $(\bar{\boldsymbol{v}}^c_i, \bar{\boldsymbol{w}}^c_i)$ crossed by $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$, and $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ are separated by a vertex of T_n. If it exists take the smallest i such that $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is not adj-P in the interval defined by $\bar{\boldsymbol{v}}^c_i$ and $\bar{\boldsymbol{w}}^c_i$. Then the triangle T_i fits the second criterion. If such an i does not exist, then the pair $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is adj-P in $(\bar{\boldsymbol{w}}^c_n, \bar{\boldsymbol{v}}^c_n)$, but is separated by some vertex of T_n thus by the midpoint of T_n.

We show that for every interval $[\bar{\boldsymbol{v}}^c\ \bar{\boldsymbol{w}}^c]$ and every hyperplane $\hat{\mathfrak{h}}^c$ of $[\bar{\boldsymbol{v}}^c\ \bar{\boldsymbol{w}}^c]$, either $\hat{\mathfrak{h}}^c$ is adjacent to $\bar{\boldsymbol{v}}^c$ or there exists another hyperplane $\hat{\mathfrak{k}}^c$ adj-P to $\hat{\mathfrak{h}}^c$ in $[\bar{\boldsymbol{v}}^c\ \bar{\boldsymbol{w}}^c]$.

If $\hat{\mathfrak{h}}^c$ is not adjacent to $\bar{v}^c$ then there exists an adjacent hyperplane $\hat{\mathfrak{k}}_1^c <_v \hat{\mathfrak{h}}^c$. If $\hat{\mathfrak{k}}^c$ is not adj-P to $\hat{\mathfrak{h}}^c$ then there exists a POP $(\hat{\mathfrak{k}}', \hat{\mathfrak{h}}')$. We can assume that $\hat{\mathfrak{h}}'$ is minimal with respect to $\hat{\mathfrak{h}}^c$. If $\hat{\mathfrak{h}}'$ is not adj-P to $\hat{\mathfrak{h}}^c$, there exists another POP $(\hat{\mathfrak{h}}_1^c, \hat{\mathfrak{h}}_1')$. But as $\hat{\mathfrak{k}}^c$ is minimal with respect to $\hat{\mathfrak{h}}^c$, we have $\hat{\mathfrak{h}}_1' \pitchfork \hat{\mathfrak{k}}^c$. Hence, we obtain five distinct hyperplanes $\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}', \hat{\mathfrak{h}}', \hat{\mathfrak{k}}^c, \hat{\mathfrak{h}}_1'$, which cyclically cross, which is forbidden in a interval of a 2–dimensional CAT(0) cube complex.

Notice that in a 2–dimensional interval $[\bar{v}^c\ \bar{w}^c]$, there are at most two hyperplanes adjacent to $\bar{v}^c$.

Since there are no parallel tracks in $\mathcal{P}$, the above discussion shows that a hyperplane $\hat{\mathfrak{h}}^c$ in $\boldsymbol{X}^c$ belongs to one of the following categories:

(1) The hyperplane $\hat{\mathfrak{h}}^c$ is adjacent to a (fixed) extremity of an edge.

(2) There exist some hyperplane $\hat{\mathfrak{k}}^c$ and some triangle $(\bar{x}^c, \bar{y}^c, \bar{z}^c)$ such that $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is adj-P in $[\bar{x}^c, \bar{y}^c]$ but is separated by the midpoint of $(\bar{x}^c, \bar{y}^c, \bar{z}^c)$.

(3) There exist some hyperplane $\hat{\mathfrak{k}}^c$ and some triangle $(\bar{x}^c, \bar{y}^c, \bar{z}^c)$ such that $(\hat{\mathfrak{h}}^c, \hat{\mathfrak{k}}^c)$ is adj-P in $[\bar{x}^c, \bar{y}^c]$, intersects $[\bar{x}^c, \bar{z}^c]$ but is not adj-P in it.

By Lemma 4.2 each midpoint of triangle separates at most 3×8 pairs of adj-P hyperplanes. By Lemma 4.3, each triangle contains at most 24 pairs of hyperplanes adj-P with respect to one edge but not with respect with another. For each edge there are at most two hyperplanes adjacent to a (fixed) extremity.

Denote by E and T be the numbers of edges and triangles, respectively, of K. There are at most $2\times(8\times3+24)\times T+2\times E = 96T+2E$ orbits of hyperplanes in $\boldsymbol{X}^c$ under the action of $\pi_1(K)$. Therefore, the pattern $\mathcal{P}$ contains at most $96T+2E$ tracks. □

5 Bounding submanifolds in surfaces and 3–manifolds

5A Patterns on surfaces

Let S be a compact surface, with a triangulation $K \cong S$. Let $\mathcal{P}$ be a collection of nonhomotopic essential, two-sided, properly immersed curves and arcs such that the size of a pairwise intersecting collection of their lifts to $\widetilde{S}$ is at most d. We would like to homotope the pattern $\mathcal{P}$ such that $\mathcal{P}$ will be a d–pattern in K.

By homotoping each curve to the corresponding geodesic curve in S, one assures that the lift to the universal cover of each curve is an embedded line, while still having that the size of a pairwise intersecting collection of their lifts to $\widetilde{S}$ is at most d.

Assume $\mathcal{P}$ is not a d–pattern; this can only occur if some arc or curve $t \in \mathcal{P}$ has a self-returning segment, ie a segment $\gamma \subseteq t$ properly embedded in some 2–simplex $f \subset K$ such that $\partial\gamma = \gamma \cap e$ for some edge $e \subset f$. Note that in this case $e \not\subseteq \partial M$ otherwise t is nonessential. Let $\gamma \subseteq t \cap f$ be a self-returning segment with innermost endpoints amongst all curves and arcs in $\mathcal{P}$, ie a self-returning segment whose endpoints do not bound two endpoints of another self-returning segment. Homotope t by pushing γ to the 2–simplex on the other side of e. Note that since we chose an innermost γ, this homotopy does not create any new intersections in $\widetilde{S}$. Hence, the collection still satisfies all the assumptions, and the total number of intersections of the curves in S, and arcs of $\mathcal{P}$ with the one-skeleton of K decreased. Thus, after finitely many such moves the resulting collection is a d–pattern.

Now, by the main theorem we have the following.

Theorem 5.1 *Let M be a compact surface, with a triangulation $K \cong M$. There exists a constant C such that if $\mathcal{S}$ is a collection of nonhomotopic essential two-sided properly immersed curves and arcs such that the size of a pairwise intersecting collection of lifts to $\widetilde{M}$ is at most 2, then $|\mathcal{S}| \leq C$.*

Proof By the above, one can replace the original collection of surfaces with homotopic curves that form a pattern on K. By the main theorem, there exists C such that the number of parallelism classes of tracks in $\mathcal{P}$ is bounded by C. Now it remains to note that if two curves (or arcs) are parallel then they are homotopic. And indeed, if they are parallel then (up to homotopy) we may assume that they bound an I–bundle region. But since they are two-sided, it must imply that there exists a homotopy between them. □

We note that this theorem could also be proven by a simpler argument using the Euler characteristic.

5B Patterns on 3–manifolds

Recall the following definition from Dicks and Dunwoody [5].

Definition 5.2 Let M be a 3–manifold, and let K be a triangulation of M. An embedded surface $S \subset M$ is *patterned* if $S \cap K^{(2)}$ is a 1–pattern in $K^{(2)}$ (the two skeleton of K), and if for every 3–simplex $\sigma \subset K$, $\sigma \cap S$ is a collection of disjoint embedded disks.

Such a patterned surface S is determined by $S \cap K^{(2)}$ [5].

For our discussion let M be a compact irreducible boundary irreducible 3–manifold. Let $\mathcal{S}$ be a collection of nonhomotopic incompressible ∂–incompressible two-sided properly embedded surfaces (in general position) in M such that the size of a pairwise intersecting collection of their lifts to $\widetilde{M}$ is at most d.

We would like, as in the previous section, to homotope the surfaces such that each surface is patterned and $\mathcal{S} \cap K^{(2)}$ is a d–pattern. Since the surfaces are embedded and satisfy that the size of a pairwise intersecting collection of their lifts to $\widetilde{M}$ is at most d, it is enough to prove that we can homotope them (while preserving the above properties) to surfaces such that the intersection of each one with $K^{(2)}$ is a pattern.

In the procedure defined in [5], the authors describe three types of moves, A, B and C, transforming the embedding $f\colon S \to M$ to an embedding $f'\colon S' \to M$, possibly changing the surface.

Under our assumptions, we note that if D is an embedded disk such that $D \cap S = \partial D$, then by incompressibility there exists a disk $D' \subset S$ such that $D' \cap D = \partial D' = \partial D$, thus, by irreducibility $D \cup D'$ bounds a 3–ball and one can homotope S to S' where $S' = (S \setminus D) \cup D'$.

Similarly, if D is an embedded disk whose boundary consists of two arcs α, β such that $\alpha = D \cap S$, $\beta = D \cap \partial M$, then by ∂–incompressibility there exists a disk $D' \subset S$ such that $D' \cap D = \alpha$ and $\partial(D \cup D') \subset \partial M$. Call $\tilde{D}$ the disk $D \cup D'$. Now by boundary irreducibility there exists a disk $D'' \subset \partial M$ such that $D'' \cap \tilde{D} = \partial D'' = \partial \tilde{D}$ and $D'' \cup \tilde{D}''$ bounds a 3–ball and one can homotope S to S' where $S' = (S \setminus D) \cup D'$.

Thus it follows that some of the moves defined in [5] are not relevant in our case, and the every other can be made by a homotopy. In Figures 5–10, we give a schematic description of some of the moves, adjusted to our case. For further details, refer to [5].

By choosing innermost curves or arcs in the procedure, we guarantee that the surfaces obtained will remain embedded and will not have more intersections. This procedure terminates [5]. Thus we obtain the following.

Lemma 5.3 *Let M and $\mathcal{S}$ be as above. Then one can choose representatives of $\mathcal{S}$ such each $S \in \mathcal{S}$ is patterned and $\mathcal{S} \cap K^{(2)}$ is a d–pattern.*

Now, by a similar proof to that of Theorem 5.1 using the main theorem and the previous lemma we get the following.

Theorem 5.4 *Let M be a compact irreducible boundary-irreducible 3–manifold. There exists a number C, depending only on M such that if $\mathcal{S}$ is a collection of non-homotopic, π_1–injective, two-sided, properly embedded surfaces (in general position) in M, such that the size of a pairwise intersecting collection of lifts to $\widetilde{M}$ is at most 2, then $|\mathcal{S}| \leq C$.*

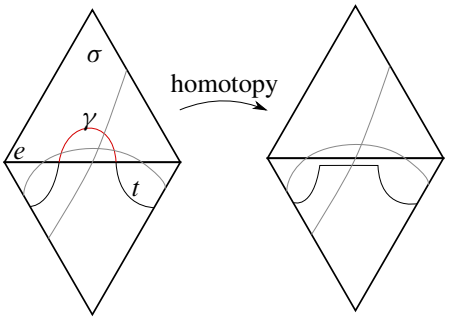

Figure 5: Fixing a self-returning curve in a surface

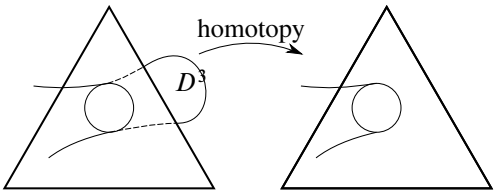

Figure 6: Fixing a closed simple curve in a face

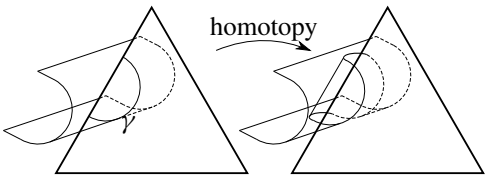

Figure 7: Fixing a self-returning curve γ, which returns to a nonboundary edge

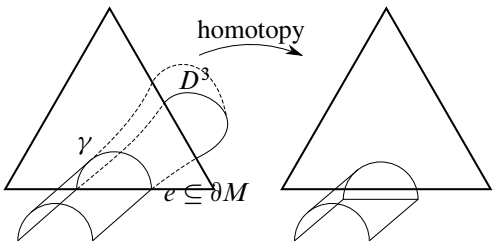

Figure 8: Fixing a self-returning curve γ, which is not in the boundary but returns to a boundary edge $e \subset \partial M$

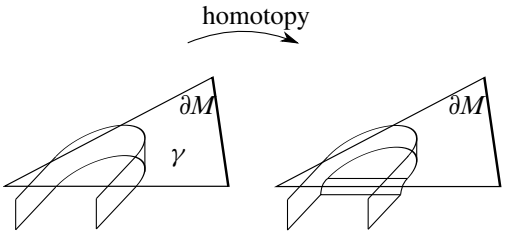

Figure 9: Fixing a self-returning curve γ, which is contained in the boundary of M

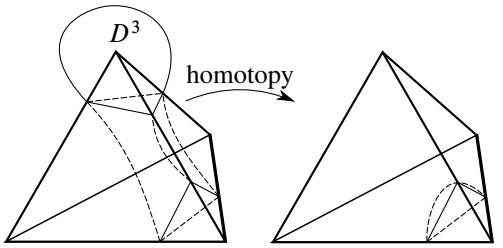

Figure 10: Fixing a non-disc component of $\mathcal{S} \setminus M^{(2)}$

References

[1] **T Aougab**, **J Gaster**, *Curves intersecting exactly once and their dual cube complexes*, preprint (2015) arXiv:1502.00311

[2] **M Bestvina**, **M Feighn**, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. 103 (1991) 449–469 MR1091614

[3] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer, Berlin (1999) MR1744486

[4] **J Brodzki**, **S J Campbell**, **E Guentner**, **G A Niblo**, **N J Wright**, *Property A and* CAT(0) *cube complexes*, J. Funct. Anal. 256 (2009) 1408–1431 MR2490224

[5] **W Dicks**, **M J Dunwoody**, *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics 17, Cambridge Univ. Press (1989) MR1001965

[6] **M J Dunwoody**, *The accessibility of finitely presented groups*, Invent. Math. 81 (1985) 449–457 MR807066

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

[8] **I Gruschko**, *On the bases of a free product of groups*, Rec. Math. [Mat. Sbornik] N.S. 8 (50) (1940) 169–182 MR0003412 In Russian

[9] **W Haken**, *Ein Verfahren zur Aufspaltung einer* 3*–Mannigfaltigkeit in irreduzible* 3*–Mannigfaltigkeiten*, Math. Z. 76 (1961) 427–467 MR0141108

[10] **M Juvan**, **A Malnič**, **B Mohar**, *Systems of curves on surfaces*, J. Combin. Theory Ser. B 68 (1996) 7–22 MR1405702

[11] **H Kneser**, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jahresbericht der Deutschen Mathematiker-Vereinigung 38 (1929) 248–259

[12] **J Malestein**, **I Rivin**, **L Theran**, *Topological designs*, Geom. Dedicata 168 (2014) 221–233 MR3158040

[13] **J Milnor**, *A unique decomposition theorem for* 3*–manifolds*, Amer. J. Math. 84 (1962) 1–7 MR0142125

[14] **G A Niblo**, **L D Reeves**, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology 37 (1998) 621–633 MR1604899

[15] **P Przytycki**, *Arcs intersecting at most once*, Geom. Funct. Anal. 25 (2015) 658–670 MR3334237

[16] **I Rivin**, *Geodesics with one self-intersection, and other stories*, Adv. Math. 231 (2012) 2391–2412 MR2970452

[17] **M Roller**, *Poc sets, median algebras and group actions: an extended study of Dunwoody's construction and Sageev's theorem*, Habilitationsschrift, University of Southampton (1998) `arXiv:1607.07747`

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

[19] **M Sageev**, CAT(0) *cube complexes and groups*, from: "Geometric group theory", (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc. (2014) 7–54 MR3329724

Department of Mathematics, Hebrew University
9190401 Jerusalem, Israel

Department of Mathematics, ETH Zürich
Rämistrasse 101, CH-8092 Zürich, Switzerland

`beeker@tx.technion.ac.il, nir.lazarovich@math.ethz.ch`

Received: 26 February 2015 Revised: 11 June 2015