



ELSEVIER

Contents lists available at ScienceDirect

Advances in Mathematics

www.elsevier.com/locate/aim



Tits Alternative for groups acting properly on 2-dimensional recurrent complexes



Damian Osajda ^{a,b,*}, Piotr Przytycki ^{c,1,2}, with an appendix by J. McCammond ^d, D. Osajda, and P. Przytycki

^a *Instytut Matematyczny, Uniwersytet Wrocławski, pl. Grunwaldzki 2/4, 50-384 Wrocław, Poland*

^b *Institute of Mathematics, Polish Academy of Sciences, Śniadeckich 8, 00-656 Warszawa, Poland*

^c *Department of Mathematics and Statistics, McGill University, Burnside Hall, 805 Sherbrooke Street West, Montreal, QC, H3A 0B9, Canada*

^d *Mathematics Department, UC Santa Barbara, Santa Barbara, CA 93106, United States of America*

ARTICLE INFO

Article history:

Received 17 March 2020

Received in revised form 14 July 2021

Accepted 21 July 2021

Available online xxxx

Communicated by Moon Duchin

Keywords:

Tits alternative

2-dimensional complex

CAT(0) complex

Artin group

Small cancellation

ABSTRACT

We prove the Tits Alternative for groups acting on 2-dimensional “recurrent” complexes with uniformly bounded cell stabilisers. This class of complexes includes, among others: 2-dimensional Euclidean buildings, 2-dimensional systolic complexes, $B(6)$ -small cancellation complexes, and standard Cayley complexes for Artin groups of extra-large type.

In the appendix written jointly with Jon McCammond we extend the result to a class of 2-dimensional Artin groups containing all large-type Artin groups.

© 2021 Elsevier Inc. All rights reserved.

* Corresponding author.

E-mail addresses: dosaj@math.uni.wroc.pl (D. Osajda), piotr.przytycki@mcgill.ca (P. Przytycki), jon.mccammond@math.ucsb.edu (J. McCammond).

¹ Partially supported by (Polish) Narodowe Centrum Nauki, UMO-2018/30/M/ST1/00668. This work was partially supported by the grant 346300 for IMPAN from the Simons Foundation and the matching 2015-2019 Polish MNiSW fund.

² Partially supported by NSERC grant RGPIN-2019-04458 and FRQNT grant 205186.

1. Introduction

Tits proved that every finitely generated linear group is either virtually solvable or contains a nonabelian free group [22]. In other words, each linear group $\mathbf{GL}_n(\mathbf{k})$ satisfies the *Tits Alternative*, saying that each of its finitely generated subgroups is virtually solvable or contains a nonabelian free group. It is believed that the Tits Alternative is common among ‘nonpositively curved’ groups. However, up to now it has been shown only for few particular classes of groups. Most notably, for: Gromov-hyperbolic groups [10], mapping class groups [12,16], $\text{Out}(F_n)$ [3,4], fundamental groups of closed 3-manifolds (by geometrisation, cf. [14]), fundamental groups of some nonpositively curved real-analytic 4-manifolds [27], CAT(0) cubical groups [21]. Whether CAT(0) groups satisfy the Tits Alternative remains an open question, even in the case of groups acting properly and cocompactly on 2-dimensional CAT(0) complexes.

In this article we prove the Tits Alternative for groups acting on triangle complexes that are ‘recurrent’. Here a *triangle complex* is a 2-dimensional simplicial complex X built of geodesic Euclidean triangles, see [5, I.7.2]. We postpone the general definition of ‘recurrent’ till Section 2, and here we discuss examples instead. All the group actions that we consider are by combinatorial isometries. An action of a group G is *without inversions* if each element of G stabilising a cell fixes it pointwise. The action is *almost free* if there is a bound on the order of cell stabilisers. Note that an almost free action on a triangle complex with finitely many isometry types of simplices is proper in the sense of [5, I.8.2].

Main Theorem. *Let X be a simply connected triangle complex that is recurrent w.r.t. a finitely generated group G acting almost freely and without inversions. Then G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

In particular, by Remark 2.2 the same conclusion will hold for any finitely generated subgroup of G . In other words, G satisfies the Tits Alternative.

For example, let X be a 2-dimensional Euclidean building or a 2-dimensional *systolic complex*, which is a CAT(0) triangle complex with all edges of length 1 and all triangles equilateral. We will show in Corollaries 2.6 and 2.7 that X has a subdivision recurrent with respect to any automorphism group of X . This implies the following for finitely generated subgroups of G .

Theorem A. *Let X be a 2-dimensional Euclidean building or a 2-dimensional systolic complex. Suppose that G acts almost freely on X (e.g. G acts on X properly and cocompactly). Then any subgroup of G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

Ballmann and Brin proved that if X is any 2-dimensional CAT(0) complex, and G acts on X properly and cocompactly, then G itself is virtually cyclic, or virtually \mathbb{Z}^2 , or

contains a nonabelian free group [2]. However, it was only very recently that we were able to perform with Norin a first step to understand the subgroups of G , by proving that each of them is finite or contains \mathbb{Z} [18].

Note that in Theorem A, as in many other applications of the Main Theorem, we will be able to remove the assumption that the group is finitely generated; see Section 5. However, we cannot remove the assumption on the uniform bound on the order of cell stabilisers, as the following example that we learnt from Pierre-Emmanuel Caprace shows. Namely, the wreath product $G = A_5 \wr \mathbb{Z}$, where A_5 denotes the alternating group on 5 elements, acts on a CAT(0) square complex with finite cell stabilisers [9, Prop 9.33]. However, G neither contains a nonabelian free subgroup, nor is virtually solvable.

Other classes of recurrent complexes arise from complexes with various combinatorial nonpositive-curvature-like features. This includes Cayley complexes for the standard presentations of Artin groups of extra-large type (see Subsection 4.2 for the definition).

Theorem B. *Let X be the Cayley complex for the standard presentation of an Artin group A_Γ of extra-large type. Suppose that G acts almost freely on X (e.g. $G = A_\Gamma$). Then any subgroup of G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

In Appendix A written jointly with Jon McCammond we extend Theorem B to a class of 2-dimensional Artin groups containing all large-type Artin groups. In the case where $G = A_\Gamma$ we extend Theorem B to all 2-dimensional A_Γ with W_Γ hyperbolic in [17]. We will give there an account on the current state of affairs concerning the Tits Alternative for other classes of Artin groups.

Another class of recurrent complexes arises from simply connected $B(6)$ -small cancellation complexes (see Subsection 4.1 for the definition and details).

Theorem C. *Let X be a simply connected $B(6)$ -small cancellation complex. Suppose that G acts almost freely on X . Then any finitely generated subgroup of G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

Let us note that Wise [26] associated to each simply connected $B(6)$ complex X a CAT(0) cube complex C . Furthermore, in [21] the Tits Alternative is shown for groups acting almost freely on finite dimensional CAT(0) cube complexes. However, the complex C associated to a simply connected $B(6)$ complex X might not be finitely dimensional — this happens e.g. when there is no bound on the size of the 2-cells in X . Therefore, the results from [21] do not imply Theorem C.

The method of proving the Tits Alternative presented in this paper raises the following natural questions.

Questions. Simplicial subdivisions of which 2-dimensional combinatorial complexes can be metrised as recurrent complexes? Can it be done for $C(6)$ -small cancellation complexes? What about Cayley complexes for standard presentations of 2-dimensional Artin groups?




Fig. 1. The dumbbell graph (left) and the geodesic segments it is mapped into (right).

Idea of proof of the Main Theorem. For simplicity we assume that the action of G is free and that X is systolic. Supposing additionally that X is countable, we exhaust the quotient $\bar{X} = X/G$ with compact subcomplexes $X_1 \subset X_2 \subset \dots$. The fundamental groups G_i of X_i have direct limit G .

Collapsing we remove the free edges from X_i . We focus first on the case where some X_i is *thick* meaning that it has an edge e of degree ≥ 3 . Consider the space of all local geodesics in X_i that are concatenations of segments $\dots, a_{-1}a_0, a_0a_1, a_1a_2, \dots$ (where a_i denote points), as indicated in Fig. 1 on the right. We equip that space with a finite measure μ^* assigning to each ‘cylinder’ of geodesics passing through prescribed consecutive segments $a_0a_1, a_1a_2, \dots, a_{n-1}a_n$ the value $\prod_{i=1}^{n-1} \frac{1}{\deg a_i - 1}$, where $\deg a_i$ is the degree of the edge containing a_i . This is inspired by the work of Ballmann and Brin [2], who put a similar measure on a far larger space of geodesics. Using Poincaré recurrence à la [2] we can find a local isometric embedding f of the dumbbell graph $\bar{\Gamma}$ (see Fig. 1 on the left) into X_i with the following properties. Namely, f sends the vertices of $\bar{\Gamma}$ into e and the edges of $\bar{\Gamma}$ into concatenations of the segments in Fig. 1 on the right, terminating perpendicularly to e . In particular, $f(\bar{\Gamma})$ avoids the vertices of X_i and hence the stabiliser in G of the lift of $\bar{\Gamma}$ to X contains a nonabelian free group.

In the case where X_i is not thick, it is a union of a graph and a dimension 2 pseudomanifold. The components of the pseudomanifold are π_1 -injective in \bar{X} , since otherwise attaching compressing discs puts us back in the thick case. If one such component is a hyperbolic surface Σ , we find a nonabelian free group in $\pi_1(\Sigma) < G$. Otherwise, each G_i is a free product of some copies of \mathbb{Z}, \mathbb{Z}^2 and the Klein bottle group, which satisfies the Tits Alternative. One can arrange that the number of factors is bounded by a constant independent of i , and then use the Hopfian property to deduce that the sequence $G_1 \rightarrow G_2 \rightarrow \dots$ stabilises. This shows that G coincides with some G_i .

We believe that the overall method of our proof can be extended to treat all 2-dimensional CAT(0) complexes. In particular, the ‘no thick subcomplexes’ part (cf. Proposition 3.7) is valid in such a general setting. To treat thick subcomplexes, one needs to find a method of ‘closing’ geodesics without the use of the additional structure of the recurrent complex. It seems that finding free subgroups could work then also in higher dimensions, for analogues of thick subcomplexes. However, at the moment we do not know how to proceed in the ‘no thick subcomplexes’ case in such higher dimen-

sional setting, even for some restricted (combinatorial) classes of complexes, e.g. systolic complexes or Euclidean buildings.

Organisation. In Section 2 we define the main object of our interest, the recurrent complexes. We provide main examples and basic properties, and we show how to find nonabelian free subgroups given thick subcomplexes. In Section 3 we treat the case where there are no thick subcomplexes and we prove the Main Theorem. In Section 4 we provide applications to $B(6)$ -small cancellation complexes and Artin groups of extra-large type, proving Theorem C and the finitely generated case of Theorem B. We discuss the case of infinitely generated subgroups and we complete the proofs of Theorems A and B in Section 5. In Appendix A written jointly with Jon McCammond we present the aforementioned extension of Theorem B.

Acknowledgments

We thank Pierre-Emmanuel Caprace, Sam Shepherd, and the anonymous referees for valuable comments. This paper was written while D.O. was visiting McGill University. We would like to thank the Department of Mathematics and Statistics of McGill University for its hospitality during that stay.

2. Recurrent complexes

In this section we present a variant of the constructions introduced in [2].

Definition 2.1. Let X be a triangle complex, i.e. a 2-dimensional simplicial complex built of geodesic Euclidean triangles, with an action of a group G . Let $x \in X^1 - X^0$, let e be the edge containing x and let T be a triangle containing e . Then $\text{lk}_x T$ denotes the open half-circle of directions at x in T that are transverse to e . By $\deg x$ we denote the degree of e , i.e. the number of triangles containing x .

For $v \in \text{lk}_x T$, let $H(v)$ be the union of the directions $v' \in \text{lk}_x T'$ with $T' \neq T$ such that there is a geodesic through x in $T \cup T'$ with directions v and v' . Note that for each triangle T' containing x with $T' \neq T$ there exists a unique such v' . Thus $|H(v)| = \deg x - 1$. We have $v' \in H(v)$ if and only if $v \in H(v')$

Furthermore, for $v \in \text{lk}_x T$ suppose that the geodesic in T with the starting direction v terminates at a point $x' \in X^1 - X^0$. Then we denote its ending direction by $I(v) \in \text{lk}_{x'} T$. Note that $I(I(v)) = v$.

We say that X is *recurrent with respect to G* if there is a G -invariant subset A of the union of all $\text{lk}_x T$ such that all the following hold:

- (i) for each triangle T the set of $a \in A$ with a in some $\text{lk}_x T$ is finite,
- (ii) for each $a \in A$ we have $H(a) \subset A$,
- (iii) for each $a \in A$ we have that $I(a)$ is defined and belongs to A ,




Fig. 2. The set A for the isosceles right-angled triangle.

- (iv) for each edge e of degree ≥ 3 there exists $x \in e$ such that for some (hence any by (ii)) triangle T containing e the direction in $\text{lk}_x T$ perpendicular to e belongs to A ,
- (v) there is no finite sequence a_0, a_1, \dots, a_n , such that for all $0 \leq i < n$ we have $a_{i+1} \in H(I(a_i))$, and $a_0 = a_n$ or $a_0 = I(a_n)$.

Remark 2.2. If a triangle complex X is recurrent w.r.t. G and $X' \subseteq X$ is a G' -invariant subcomplex, for some $G' < G$, then X' is recurrent w.r.t. G' .

Remark 2.3. If X is CAT(0), then its local geodesics are global geodesics and hence embed, and consequently Definition 2.1(v) holds automatically for any A .

Example 2.4. Suppose that X admits a simplicial map ρ to a simplicial complex consisting of only one triangle with angles $\frac{\pi}{2}, \frac{\pi}{4}, \frac{\pi}{4}$ such that ρ restricted to each triangle of X is an isometry. Then X has A satisfying Definition 2.1(i)-(iv) w.r.t. any automorphism group of X . Indeed, it suffices to define $A \cap \text{lk}_x T$ with x in an edge e to be

- the vector perpendicular to e , for e the long edge and x dividing e in the ratio 1:3,
- the vectors at angles $\frac{\pi}{4}$ to e , for e the long edge and x the midpoint of e ,
- the vectors at angles $\frac{\pi}{4}$ and $\frac{\pi}{2}$ to e , for e the short edge and x the midpoint of e ,
- empty otherwise.

In other words, A is the union of the directions at the boundary of the two billiard trajectories in Fig. 2.

Example 2.5. Suppose that X admits a simplicial map ρ to a simplicial complex consisting of only one triangle with angles $\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{6}$ such that ρ restricted to each triangle of X is an isometry. Then X has A satisfying Definition 2.1(i)-(iv) w.r.t. any automorphism group of X . Indeed, it suffices to define A in each triangle as the union of the directions at the boundary of the two billiard trajectories in Fig. 3, where y is, say, the edge midpoint.

We have the following immediate consequence of Example 2.5 and Remark 2.3.

Corollary 2.6. Let X be a 2-dimensional systolic complex. Then the barycentric subdivision of X is recurrent with respect to any automorphism group.




Fig. 3. The set A for the triangle with angles $\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{6}$.

Corollary 2.7. *Let X be a 2-dimensional Euclidean building of type W with its usual geometric realisation, where each chamber is realised as a Euclidean triangle of angles $\frac{\pi}{m_{st}}$ in the cases $W = \tilde{A}_2, \tilde{C}_2, \tilde{G}_2$ or a square in the case $W = (\tilde{I}_2)^2 = D_\infty \times D_\infty$. Then X has a subdivision X^* that is recurrent with respect to any automorphism group G of X , and such that G acts on X^* without inversions.*

Proof. Since X is CAT(0), by Remark 2.3 we have that Definition 2.1(v) holds automatically. If $W = \tilde{G}_2$, it suffices to take $X^* = X$ and use Example 2.5. If $W = \tilde{A}_2$, we take X^* to be the barycentric subdivision of X and we use Example 2.5 as well. If $W = \tilde{C}_2$, let X^* be obtained from X by subdividing each triangle into two similar triangles along the altitude from the right angle. We then use Example 2.4. Finally, if $W = (\tilde{I}_2)^2$, let X^* be the barycentric subdivision of X and use Example 2.4. \square

Remark 2.8. In fact, if X is a triangle complex with finitely many isometry types of ‘simplices with specified directions in A' , recurrent w.r.t. G , then its barycentric subdivision X' is also recurrent w.r.t. G , and consequently in the Main Theorem one can remove the assumption that G acts without inversions.

Indeed, first note that for ε sufficiently small, we can replace A by A' whose geodesic segments with starting direction $a' \in A'$ and ending direction $I(a')$ constitute the boundary of the ε -neighbourhood of analogous segments from $a \in A$ to $I(a)$. Secondly, except for finitely many of such ε , these segments from $a' \in A'$ to $I(a')$ do not pass through the vertices of X' , and hence they show the recurrence of X' w.r.t. G .

Recurrent complexes are designed to satisfy the following lemma.

Definition 2.9. A 2-dimensional simplicial complex is *essential* if every edge has degree at least 2, and none of connected components is a single vertex. An essential triangle complex is *thick* if it has an edge of degree at least 3.

Lemma 2.10. *Suppose that a triangle complex X has all edges of finite degree, is thick, and is recurrent with A/G finite. Then G contains a nonabelian free group.*

To prove Lemma 2.10 we will use the following method of [2].

Definition 2.11. Suppose that an essential triangle complex X has all edges of finite degree, and is recurrent with A/G finite. Consider the Markov chain with states A and the following transition function. Let $b \in A \cap \text{lk}_x T$. The transition probability $p(a, b)$ from $a \in A$ to b equals $\frac{1}{\deg x - 1}$ if $b \in H(I(a))$ and 0 otherwise.

Remark 2.12. A uniform measure μ on A is stationary for that Markov chain. Thus the space $A^{\mathbb{Z}}$ can be equipped with Markov measure μ^* invariant under the shift (see e.g. [24, Ex (8), page 21]). Since A/G is finite, the measure of the quotient $A^{\mathbb{Z}}/G$ by the diagonal action of G is finite, w.l.o.g. a probability measure. Note that the shift map descends to $A^{\mathbb{Z}}/G$ and is still measure preserving.

Lemma 2.13. *Suppose that an essential triangle complex X has all edges of finite degree, and is recurrent with A/G finite. Let $a \in A$ and $b \in H(a)$. Then there is a local geodesic $f: [0, l] \rightarrow X$ with the directions at $0, l$ mapping to b, ga for some $g \in G$.*

Proof. We have $p(I(a), b) \neq 0$. Thus by the Poincaré recurrence (see e.g. [24, Thm 1.4]) applied to $A^{\mathbb{Z}}/G$ we have a finite sequence $a_0 = I(a), a_1 = b, \dots, a_n = ga_0$, for some $g \in G$, such that for all $0 \leq i < n$ we have $a_{i+1} \in H(I(a_i))$. Define f as the concatenation of the geodesics from a_i to $I(a_i)$ for $1 \leq i < n$ and from $a_n = ga_0$ to $gI(a_0) = ga$. \square

Remark 2.14. An alternative, more combinatorial way of proving Lemma 2.13 was suggested to us by Sam Shepherd. For example, if G acts freely on X , then the finite set A/G is the set of states for appropriate Markov chain for which the uniform measure is stationary. Consequently, since we have a positive transition probability from the state $[I(a)]$ to the state $[b]$, we also have a positive probability of passing from $[b]$ to $[I(a)]$ after several steps.

Proof of Lemma 2.10. Let $\bar{\Gamma}$ (see Fig. 1 left) be the graph obtained from joining the basepoints of two closed paths \bar{C}, \bar{C}' by a path \bar{C}'' (their lengths will be determined later). Let Γ be the universal cover of $\bar{\Gamma}$ with the action of the deck transformation group F_2 . The main idea of the proof is to construct a homomorphism $\varphi: F_2 \rightarrow G$ and φ -equivariant local isometry $\Gamma \rightarrow X - X^0$ that is injective on the set of directions at the vertices of Γ .

Let C'' be a lift to Γ of \bar{C}'' with endpoints c, c' . Let C, C' be some lifts of the paths \bar{C}, \bar{C}' starting at c, c' . Let $h, h' \in F_2$ be the elements mapping c to the other endpoint of C , and c' to the other endpoint of C' , respectively. Observe that $C \cup C' \cup C''$ is a fundamental domain for the action of F_2 on Γ . Thus to define an equivariant map $\Gamma \rightarrow X - X^0$ it suffices to define a homomorphism $\varphi: F_2 \rightarrow G$ and a map $f: C \cup C' \cup C'' \rightarrow X - X^0$ with the property that $\varphi(h)$ maps $f(c)$ to the other endpoint of $f(C)$ and $\varphi(h')$ maps $f(c')$ to the other endpoint of $f(C')$.

Let e be an edge of X of degree ≥ 3 . Then for $i = 1, 2, 3$, there are distinct triangles T_i containing e . Since Y is recurrent, by Definition 2.1(iv) we have $x \in e$ such that for

any T containing e the direction in $\text{lk}_x T$ perpendicular to e belongs to A . Let v_i be that direction in $\text{lk}_x T_i$.

Apply Lemma 2.13 to $a = v_1, b = v_2$, to obtain a local geodesic $f: [0, l] \rightarrow X - X^0$ with ending directions v_2, gv_1 , for some $g \in G$. Identify C with $[0, l]$. Analogously, apply Lemma 2.13 to $a = v_1, b = v_3$ to obtain $f: C'' \rightarrow X - X^0$ with ending directions $v_3, g''v_1$, for some $g'' \in G$. Finally, apply Lemma 2.13 to $a = v_2, b = g''v_3$, to obtain $f: C' \rightarrow X - X^0$ with ending directions $g''v_3, g'g''v_2$, for some $g' \in G$. Define $\varphi: F_2 \rightarrow G$ by $\varphi(h) = g$ and $\varphi(h') = g'$. By the observation above, we can extend $f: C \cup C' \cup C'' \rightarrow X - X^0$ to a φ -equivariant map $\Gamma \rightarrow X - X^0$ for which we keep the same notation f . Note that for each vertex w of Γ the three directions at w are mapped under f to a G -translate of the triple $\{v_1, v_2, v_3\}$.

Let E be the set of directed edges of Γ . Consider the map $f_*: E \rightarrow A$ that maps each directed edge $wu \in E$ to the direction of $f(wu)$ at $f(w)$. We will prove that φ is injective by showing that f_* is injective. Suppose that there are two edges $wu, w'u' \in E$ with $f_*(wu) = f_*(w'u')$. Then we also have $f_*(uw) = f_*(u'w')$, so without loss of generality we can assume that the embedded edge-path γ in Γ from w to w' passes through u . For $i = 0, \dots, n$, let a_i be the images under f_* of consecutive edges of γ . In particular, $a_0 = f_*(wu)$ and in the case where u' lies on γ we have $a_n = f_*(u'w')$. If γ does not contain u' , then we add $a_{n+1} = f_*(w'u')$. Note that for $i = 0, 1, \dots$ we have $a_{i+1} \in H(I(a_i))$, and so $f_*(wu) = f_*(w'u')$ contradicts Definition 2.1(v). \square

3. Invariant cocompact subcomplexes

Definition 3.1. Let X be a simplicial complex with a simplicial action of a group G . We say that a subcomplex $Z \subseteq X$ is an *invariant cocompact subcomplex with respect to G* (shortly *G -c.s.*) if Z is G -invariant, and the quotient Z/G is compact. Note that a G -c.s. is not required to be connected.

A simplicial complex homeomorphic to the plane \mathbb{R}^2 (resp. to the 2-sphere S^2) is called a *simplicial plane* (resp. *simplicial 2-sphere*). A simplicial plane whose 1-skeleton is Gromov-hyperbolic (w.r.t. the metric where each edge has length 1) is called *hyperbolic*. By the classification of 2-dimensional orbifolds, if E is a non-hyperbolic simplicial plane with cocompact automorphism group H , then H is virtually \mathbb{Z}^2 . We call such E a *flat*.

Lemma 3.2. *Let X be a simply connected 2-dimensional simplicial complex with a finitely generated group G acting almost freely and without inversions. If each essential G -c.s. in X is a disjoint union of flats, then G is virtually a free product of some number of \mathbb{Z} and \mathbb{Z}^2 .*

Proof. Suppose first that X is countable. We may find an increasing sequence $X_1 \subset X_2 \subset \dots$ of connected G -c.s.'s exhausting X . (Start with a G -orbit of a vertex, connect it equivariantly by edge-paths, then at each step add equivariantly remaining cells.) The

action of G on X_i lifts to an action of a group G_i on the universal cover \tilde{X}_i of X_i . The corresponding maps $\tilde{X}_1 \rightarrow \tilde{X}_2 \rightarrow \dots$ and $\tilde{X}_i \rightarrow \tilde{X} = X$ induce homomorphisms $G_1 \rightarrow G_2 \rightarrow \dots$ and epimorphisms $G_i \twoheadrightarrow G$. Note that the vertex stabilisers of the action of G_i on \tilde{X}_i coincide with the vertex stabilisers of the action of G on X_i and thus have uniformly bounded order.

Since G is finitely generated, there is a finitely generated subgroup $H_1 < G_1$ such that $H_1 \rightarrow G$ is an epimorphism. For each $i > 1$, let H_i be the image of H_1 under the homomorphism $G_1 \rightarrow G_i$. We obtain an infinite sequence of epimorphisms

$$H_1 \twoheadrightarrow H_2 \twoheadrightarrow \dots \tag{1}$$

The epimorphism from the direct limit $\varinjlim H_i$ to G is in fact an isomorphism. Indeed, let $h_i \in \ker(H_i \rightarrow G)$ and let α be a path joining a basepoint $\tilde{x}_i \in \tilde{X}_i$ to $h_i\tilde{x}_i$. The projection of α to X_i is a closed path and it becomes contractible in some X_j , since X_j exhaust X , and X is simply connected. Consequently the image $h_j \in H_j$ of h_i fixes the image of \tilde{x}_i in \tilde{X}_j and thus $h_j \in \ker(H_j \rightarrow G)$ implies $h_j = 0$.

Since each essential G -c.s. in X is a disjoint union of flats and G acts without inversions, every X_i can be equivariantly collapsed (by removing triangles with free edges) to a space Y_i that is a union of a graph and a disjoint union of flats. The preimage $\tilde{Y}_i \subset \tilde{X}_i$ of Y_i under the covering map is thus a simply connected union of a graph and a disjoint union of flats, with a proper and cocompact action of G_i . Let Γ_i be the tree obtained from \tilde{Y}_i by quotienting each flat to a vertex. The quotient Γ_i/G_i is a finite graph of groups \mathcal{G}_i with $\pi_1\mathcal{G}_i = G_i$ and edge groups of uniformly bounded order. Its vertex groups G_v are also finite of uniformly bounded order, or have the following description for a vertex v obtained from quotienting a flat Z to v . Namely, let G'_v be the image of G_v in the isometry group of Z . We then have a short exact sequence $0 \rightarrow K \rightarrow G_v \rightarrow G'_v \rightarrow 0$, with K finite of uniformly bounded order. By the classification of 2-dimensional Euclidean orbifolds, there are only finitely many possible isomorphism types for G'_v . Consequently, there are only finitely many possible isomorphism types for G_v . Analogously, there are only finitely many possible isomorphism types for the subgroups of G_v .

If H_1 is generated by d elements, then so is each H_i for $i > 1$. Since each H_i is a subgroup of G_i , it is also the fundamental group of a finite graph of groups \mathcal{H}_i with edge groups of uniformly bounded order. It follows, by a result of Linnell [15, Thm 2], that there is a uniform bound on the number of edges in a minimal such graph with fundamental group H_i . (This is because the augmentation ideal in [15, Thm 2] is generated by at most d elements.) Edge and vertex groups of \mathcal{H}_i have only finitely many isomorphism types. Furthermore, there are finitely many possible injections from edge groups to vertex groups, up to conjugations in vertex groups. However, such conjugations do not change $\pi_1\mathcal{H}_i$. Hence there are only finitely many isomorphism types in $\{H_i\}_{i=1}^\infty$.

Let H be isomorphic to H_i for infinitely many i . Note that H is virtually a free product of some number of \mathbb{Z} and \mathbb{Z}^2 , and thus it is residually finite. Moreover, H is finitely generated, so it is Hopfian. It follows that if $H \cong H_i, H_{i+k}$, then $H_i \twoheadrightarrow H_{i+k}$ is

an isomorphism, and hence, for every $j = i, i + 1, \dots, i + k - 1$, the map $H_j \rightarrow H_{j+1}$ is an isomorphism. Therefore, the sequence (1) stabilises and $G = H$ is as required.

If X is not countable, we consider the collection X_λ of all connected G -c.s. containing a given connected G -c.s. X_1 , which form a directed set under inclusion. We define appropriate G_λ, H_λ as before, and we have again $G = \varinjlim H_\lambda$. There is still H such that for every X_λ there is $X_{\lambda'} \supset X_\lambda$ with $H_{\lambda'} \cong H$, and then replacing X_1 by such $X_{\lambda'}$ we obtain that all maps in our directed system are isomorphisms and consequently $G = H$. \square

Lemma 3.3. *Let $Z \subseteq X$ be a connected essential subcomplex in a connected simplicial complex X . If $\pi_1 Z \rightarrow \pi_1 X$ is not injective, then Z is contained in a thick subcomplex $Z' \subseteq X$ with $Z' - Z$ finite.*

Before we give the proof, we record the following consequences.

Corollary 3.4. *Let $Z \subseteq X$ be a connected essential G -c.s. in a connected simplicial complex X . If $\pi_1 Z \rightarrow \pi_1 X$ is not injective, then Z is contained in a thick G -c.s.*

Proof. Apply Lemma 3.3 to Z . Then GZ' is a thick G -c.s. \square

For a subcomplex $Z \subseteq X$ and a triangle T of Z let $\text{gal}_Z(T)$ denote the *gallery connected component* of Z containing T . That is, $\text{gal}_Z(T)$ is the subcomplex of Z consisting of T and all the triangles in Z that can be reached from T by passing from a triangle to a triangle adjacent along an edge.

Corollary 3.5. *Let X be a simply connected simplicial complex that does not contain simplicial 2-spheres. Let $Z \subseteq X$ be an essential G -c.s. that is not contained in a thick G -c.s. Then for each triangle T of Z , $\text{gal}_Z(T)$ is a simplicial plane.*

Proof. Since Z is essential and not thick, $\text{gal}_Z(T)$ is a 2-dimensional pseudomanifold. Thus $\text{gal}_Z(T)$ is homeomorphic with a connected surface with possibly some identifications on a discrete set of points. By Corollary 3.4, $\text{gal}_Z(T)$ is simply connected, so it is homeomorphic with S^2 or \mathbb{R}^2 . \square

We now pass to the proof of Lemma 3.3. A *disc diagram* D is a compact contractible simplicial complex with a fixed embedding in \mathbb{R}^2 . Its *boundary path* is the attaching map of the cell at ∞ . If X is a simplicial complex, a *disc diagram in X* is a nondegenerate simplicial map $\varphi: D \rightarrow X$, and its *boundary path* is the composition of the boundary path of D and φ . We say that φ is *reduced* if it maps triangles sharing an edge to two distinct triangles. The *area* of φ is the number of triangles of D .

Remark 3.6. Let γ be a closed edge-path in a simplicial complex X . If γ is contractible in X , then there is a reduced disc diagram in X with boundary path γ . For γ embedded, this is [13, Lem 1.6]. For γ not embedded, attach a triangulated annulus A to X along

γ to obtain X' . Applying the embedded case to the second boundary component γ' of A we obtain a reduced disc diagram $\varphi': D' \rightarrow X'$. Then $\varphi'^{-1}(X' - X)$ is the open combinatorial 1-ball around $\partial D'$. Consequently, φ' restricted to $D = D' - \varphi'^{-1}(X' - X)$ is a reduced disc diagram with boundary path γ .

Proof of Lemma 3.3. Let $\varphi: D \rightarrow X$ be a minimal area reduced disc diagram with boundary path in Z representing a nontrivial element of $\pi_1(Z)$. Note that D cannot have area 0. Let $Z' = Z \cup \varphi(D)$. Observe that Z' is essential, since Z is essential and φ is reduced. Furthermore, let T be a triangle of D adjacent to a boundary edge e . By the minimality of area, $\varphi(e)$ has degree ≥ 3 in Z' . \square

Proposition 3.7. *Let X be a simply connected 2-dimensional simplicial complex that contains no simplicial 2-spheres with a finitely generated group G acting almost freely and without inversions. If X contains no thick G -c.s., then G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

Proof. Consider possible essential G -c.s. $Z \subseteq X$. By Corollary 3.5, for each triangle T of Z we have that $Y = \text{Gal}_Z(T)$ is a simplicial plane. If any such Y is not a flat, then it is a hyperbolic simplicial plane. Consequently, the stabiliser $\text{Stab}_G(Y)$, which acts properly and cocompactly on Y , contains a nonabelian free group (by e.g. [8, Thm 8.37]). If Y and Y' are two such intersecting flats, then by Corollary 3.4, the connected component W of Z containing $Y \cup Y'$ is an infinite valence tree of flats and thus $\text{Stab}_G(W)$ contains a nonabelian free group. It remains to consider the case where each Z is a disjoint union of flats. Then by Lemma 3.2 we have that G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group. \square

Proof of the Main Theorem. If X contains a simplicial 2-sphere Σ , then there is no triangle T_0 in Σ with an $a_0 \in A$ in some $\text{lk}_{x_0}T_0$. Indeed, otherwise using Definition 2.1(ii) and (iii) we could construct in A an infinite sequence a_0, a_1, \dots such that for all $i \geq 0$ we have $a_{i+1} \in H(I(a_i))$, with a_i in some $\text{lk}_{x_i}T_i$ and $T_i \subset \Sigma$. This would contradict Definition 2.1(i) or (v). Consequently, by Definition 2.1(iv) each edge of Σ has degree 2 in X . We can thus remove all triangles and edges of each Σ and replace them by a cone over the vertex set Σ^0 . After this operation X is still simply connected. We can thus assume that X does not contain simplicial 2-spheres.

If X contains a thick G -c.s. Z , then by Remark 2.2 the triangle complex Z is recurrent with respect to G . Since G acts cocompactly on Z , we have that A/G is finite. Moreover, since G acts properly on Z , all edges of Z have finite degree in Z . Thus by Lemma 2.10 we have that G contains F_2 . If X does not contain a thick G -c.s., then the Main Theorem follows from Proposition 3.7. \square

Remark 3.8. The assumption in Proposition 3.7 that X contains no simplicial 2-spheres could be removed at the cost of allowing, instead of flats, trees of 2-spheres and trees of




Fig. 4. The set A for the barycentric subdivision of the regular octagon.

a flat and 2-spheres in the statement of Lemma 3.2. This would complicate slightly the proof of Lemma 3.2, so we decided to keep this assumption.

4. More applications

Example 4.1. Let X be a 2-dimensional combinatorial complex with all edges of length 1 and all 2-cells regular Euclidean $2n$ -gons, where n might vary. Suppose that all cells embed in X . Then the barycentric subdivision X' of X is a triangle complex satisfying Definition 2.1(i)-(iv) w.r.t. any automorphism group of X . Indeed, we define A in the triangles forming a given polygon as the directions coming from the union of line segments perpendicular to opposite edge pairs and dividing them in the ratio 1 : 3, see Fig. 4.

Here we study two classes of complexes, where we have also Definition 2.1(v), and consequently X' is recurrent w.r.t. any automorphism group of X .

4.1. $B(6)$ complexes

The following notion was introduced by Wise [26]. Let X be a 2-dimensional combinatorial complex. X satisfies the $B(6)$ (*small cancellation*) condition if for each 2-cell R , and for each path $S \rightarrow \partial R$ which is the concatenation of at most 3 pieces, we have $|S| \leq |\partial R|/2$, where $|\cdot|$ denotes the number of edges in a path. (See [26] for defini-

tions of paths, pieces, and further details.) In particular, $B(6)$ complexes satisfy the combinatorial $C(6)$ small cancellation condition [26, Prop 2.11].

Example 4.2. Suppose that X is simply connected and satisfies $B(6)$. By subdividing each edge into two, we can assume that for every 2-cell R of X the length $|\partial R|$ is even. A *hypergraph* in X is then a connected component Λ of a graph whose vertices correspond to edges of X and whose edges correspond to pairs of antipodal edges in R , with the obvious map $\Lambda \rightarrow X$ [26, Def 3.2 and Rem 3.4]. Equip the barycentric subdivision X' of X with the metric and A of Example 4.1. Note that all the 2-cells of X embed by [26, Cor 2.9].

Observe that any sequence a_0, a_1, \dots, a_n of elements of A , with $a_{i+1} \in H(I(a_i))$ for $0 \leq i < n$ extends to such a sequence with $a_0 \in \text{lk}_x T, a_n \in \text{lk}_{x'} T'$ with x, x' in the edges of X (and dividing them in the ratio 1 : 3). Joining consecutive a_i by geodesics we obtain a local geodesic segment $\gamma \rightarrow X$, which factors (up to a distance $\frac{1}{4}$ translation) through a hypergraph $\Lambda \rightarrow X$. Since $\Lambda \rightarrow X$ is an embedding [26, Cor 3.12] we have Definition 2.1(v), and thus X' is recurrent with respect to any automorphism group of X . Consequently the Main Theorem applies to X' , implying Theorem C.

Our arguments do not extend directly to $C(6)$ -small cancellation complexes because, as shown in [26, §3.5], in that case a hypergraph might not embed. It is an open question (see e.g. [26, Prob 1.4]) whether one can define any reasonable ‘walls’ in that case.

4.2. Artin groups of extra-large type

Let Γ be a finite simple graph with each of its edges labelled by an integer ≥ 2 . Let $V\Gamma$ be the vertex set of Γ . The *Artin group* A_Γ is given by the following presentation, where $p_m(a, b)$ denotes the word $\underbrace{aba \cdots a}_m$:

$$\langle V\Gamma \mid p_m(a, b) = p_m(b, a) \text{ for each edge } ab \text{ labelled by } m \rangle.$$

We call the presentation above the *standard presentation* for A_Γ . The Artin group A_Γ is of *extra-large type* if all the labels satisfy $m \geq 4$. The *Coxeter group* W_Γ is obtained from A_Γ by adding the relations $a^2 = 1$ for all $a \in V\Gamma$.

Let X be the Cayley complex of A_Γ for the standard presentation. It consists of cells that are $2m$ -gons with m the labels of Γ . A *hypergraph* $\Lambda \rightarrow X$ is defined as in Example 4.2. Note that $\Lambda \rightarrow X$ is an embedding, since it projects to a hypergraph in the Cayley complex of W_Γ , which is embedded. Similarly, $X - \Lambda$ has two connected components.

Proposition 4.3. *Suppose that A_Γ is of extra-large type. Then each hypergraph Λ in X is a tree.*

By Example 4.1 and Proposition 4.3, the Main Theorem applies to X' implying Theorem B for finitely generated subgroups of G .

The remaining part of the section is devoted to the proof of Proposition 4.3.

Lemma 4.4. *Let Γ be a single edge ab with label $m \geq 3$, and let $\Lambda \subset X$ be a hypergraph. Suppose that we have an edge-path γ in X with only the first and the last edge corresponding to vertices of Λ . Then γ is not labelled by a word of the form a^kba^l, b^kab^l or a^kb^l with $k, l \in \mathbb{Z} - \{0\}$.*

Proof. Let $w \in W_\Gamma$ denote the element represented by the longest word $p_m(a, b)$. Consider an edge of Λ . Its vertices correspond to opposite edges e, e' in a 2-cell D of X . Thus e and e' are labelled by the same (resp. distinct) letters for m even (resp. odd) and oriented towards the same side of Λ . A path in ∂D starting with e and ending with $(e')^{-1}$ is labelled by a word projecting to wa or wb in W_Γ . If we consider another edge of Λ with vertices corresponding to e', e'' , a path starting with e and ending with $(e'')^{-1}$ is labelled by a word projecting to $1 \in W_\Gamma$. Consequently, a concatenation of such paths that is starting with (an edge labelled by) a and ending with a^{-1} is labelled by a word projecting to 1 or possibly to wa for m even, which is distinct from the possible projections aba, ab, ba, b of a^kba^l . Similarly, if such a concatenation of paths is starting with a and ending with b^{-1} , then m is odd, and the path is labelled by a word projecting to wb , which is distinct from the possible projections $ab, a, b, 1$ of a^kb^l . \square

Lemma 4.5. *Let Γ be a single edge ab with label $m \geq 3$. Let u be a cyclically reduced word representing $1 \in A_\Gamma$. Then, possibly after a cyclic permutation of u , there are two subwords w_1, w_2 of u of forms $p_m(a, b), p_m(a, b)^{-1}, p_m(b, a)$, or $p_m(b, a)^{-1}$, none of whose letters lie in the same syllable of u , nor any of the cyclic permutations of u .*

In the proof we will use the small cancellation techniques of [1]. Here we recall only the notions that are less standard and we refer the reader to [1] for more details. An \mathcal{R} -diagram M is a van Kampen diagram, with the boundary path of each 2-cell (called a *region*) labelled by a word in a set of relators \mathcal{R} . If we ignore the labelling, M is just called a *diagram*. A *spike* of M is a boundary vertex of valence 1. The *interior degree* $i(D)$ of a region D is the number of interior edges of ∂D (after forgetting vertices of valence 2). D is a *simple boundary region* if its *outside boundary* $\partial D \cap \partial M$ is nonempty, and $M - \overline{D}$ is connected. A *singleton strip* is a simple boundary region with $i(D) \leq 1$. A *compound strip* is a subdiagram R of M consisting of regions D_1, \dots, D_n , with $n \geq 2$ with $D_{k-1} \cap D_k$ a single interior edge of R (after forgetting vertices of valence 2), satisfying $i(D_1) = i(D_n) = 2, i(D_k) = 3$ for $1 < k < n$ and $M - R$ connected.

We will use the following [1, Lem 2], where case (ii) needs to be added to account for a minor error in the third paragraph of their proof where the singleton strip D might be glued to a region contained in two distinct strips of M' .

Lemma 4.6. *Let M be a simply connected diagram with no spikes and more than one region. If M satisfies $C(4)$ and $T(4)$, then*

- (i) M contains two singleton strips, or
- (ii) M contains one singleton strip and two compound strips, or
- (iii) M contains four compound strips.

Proof of Lemma 4.5. Let M be a minimal $\hat{\mathcal{R}}_{ab}$ -diagram for u , where $\hat{\mathcal{R}}_{ab}$ is the symmetrised set obtained from the relator $p_m(a, b)(p_m(b, a))^{-1}$. Each region of M has two *separating vertices* separating the paths in its boundary labelled $p_m(a, b)$ and $(p_m(b, a))^{-1}$. We prove inductively on the number of regions of M a slight strengthening of Lemma 4.5 where we add the requirement that each w_i labels a path in the outside boundary of a region of M (which joins its separating vertices).

If M has a single region, then (after possibly a cyclic permutation and/or an inversion) we have $u = p_m(a, b)(p_m(b, a))^{-1}$ and we are done. Otherwise we can apply Lemma 4.6, since by [1, Lem 3], M satisfies $C(4)$ and $T(4)$.

In cases (ii) and (iii) we claim that there are regions D_1, D_2 in the strips of M with both of their separating vertices on their outside boundaries (we will call such regions *exposed*), and such that the outside boundaries of D_1, D_2 are separated in ∂M by outside boundaries of other simple boundary regions of interior degree at most 2.

To justify the claim, by [1, Lem 5] each strip S has an exposed region D . Consequently, in case (ii), if we denote by D_s the singleton strip and by S, S' its consecutive (clockwise and counterclockwise) compound strips, there are exposed regions D in S and D' in S' . We can take $D_1 = D_s$ and D_2 to be D or D' unless D and D' are the regions consecutive to D_s among the interior degree 2 regions of M lying in S and S' . But then we can take $D_1 = D, D_2 = D'$ instead. Finally, in case (iii), if we call any exposed region D_s , we have two strips S, S' disjoint from D_s and the same procedure as in case (ii) yields the required D_1, D_2 . This justifies the claim.

Choose w_1, w_2 labelling the paths in the outside boundaries of D_1, D_2 between their separating vertices. Note that each simple boundary region of interior degree at most 2 has outside boundary of length at least 2, and hence witnesses a syllable change in u . Thus by the claim the letters of w_1, w_2 do not lie in the same syllable.

It remains to consider case (i). Let D be a singleton strip and let M' be the diagram obtained from M by removing D and possibly a spike or a sequence of spikes, so that M' has cyclically reduced boundary word u' . By induction hypothesis, u' has appropriate subwords w'_1, w'_2 in the outside boundaries of single regions. We can choose $w_1 = w'_1, w_2 = w'_2$ for u , unless the intersection β of D or the final spike with M' contains an interior vertex of the path α labelled by one of the w'_i , say w'_1 . There is an endpoint x of β that is not an endpoint of α . Let $w_2 = w'_2$. For the choice of w_1 , suppose first that β is not a single vertex. Then x is a separating vertex of D . Thus we can take w_1 to be the word labelling the length m subpath of $\partial D - \beta$ starting at x . It remains to consider $\beta = x$. In

that case we can take w_1 to be the word labelling any path (there might be two) in the outside boundary of D joining its separating vertices. \square

We are finally ready for the following.

Proof of Proposition 4.3. We first focus on the case where Γ is a single edge. (For future reference in the proof of Proposition A.1, note that the following argument works also for $m = 2, 3$.) Suppose that there is a cycle γ in Λ with edges corresponding to 2-cells $D_0, D_1, \dots, D_k = D_0$ and vertices corresponding to edges $e_0, e_1, \dots, e_k = e_0$ of X with e_{i-1}, e_i opposite in D_i for $i = 1, \dots, k$. Let u_i be the labels on the length $m - 1$ paths joining in ∂D_i either the initial vertices or the terminal vertices of the directed edges e_{i-1}, e_i . Then $u = u_1 u_2 \cdots u_k$ represents the trivial element in A_Γ and hence there is an $\hat{\mathcal{R}}_{ab}$ -reduced diagram M for u . Suppose that M has minimal possible number n of regions among all cycles γ of Λ .

Attach to M along its boundary all the D_i , and glue D_i to D_{i+1} along e_i , to form a diagram M' . Note that M' is still reduced, since if D_i would cancel with a region D of M , we would have (using the observation that Λ does not self-intersect) a cycle γ' of Λ inside M , contradicting the minimality of n . Furthermore, D_i cannot cancel with D_j , since this would also contradict the minimality of n . By Lemma 4.6 and [1, Lem 5], there is D_i with its outside boundary of length at least m , which contradicts the definition of a hypergraph. This finishes the case where Γ is a single edge.

Now we consider arbitrary Γ . By [23] for each edge ab of Γ the group A_{ab} embeds in A_Γ . Thus we have in X copies of the Cayley complex X_{ab} of A_{ab} , which we call *blocks*, corresponding to the cosets of A_{ab} in A_Γ . Given a cycle γ in Λ , define e_i, D_i as before. Let $B_0, B_1, B_2, \dots, B_m = B_0$ be the consecutive blocks visited by (D_i) . Note that $m > 0$ by the case of Γ a single edge. For $j = 1, \dots, m$, let $D_{i_j} \subset B_{j-1}, D_{i_{j+1}} \subset B_j$ be the cells where we transition from B_{j-1} to B_j . Let L_j be the line that is the connected component of $B_{j-1} \cap B_j$ containing e_{i_j} (in fact $L_j = B_{j-1} \cap B_j$ but we do not need that). Consider closed immersed edge-paths $\delta = \delta_1 \delta_2 \cdots \delta_m$ in X such that each δ_j is a path in B_j from L_j to L_{j+1} (where $L_{m+1} = L_1$). Note that each δ_j is nontrivial since otherwise a path in $L_j \cup L_{j+1}$ labelled by $a^k b^l$ would contradict Lemma 4.4. The word $v = v_1 v_2 \cdots v_m$, with v_j the label of δ_j , represents the trivial element in A_Γ and hence there is an \mathcal{R} -reduced diagram M for v , where \mathcal{R} is the symmetrised set obtained from the standard presentation of A_Γ . Choose δ so that M has the minimal possible number of regions.

By [1, Lem 8] M satisfies $C(6)$, and so if M has more than one region, it has a simple boundary region D with interior degree at most 3. Suppose w.l.o.g. that ∂D is labelled by a word in a and b . Since the words labelling the intersections of ∂D with its adjacent regions cannot exceed one syllable, by Lemma 4.5 there is an occurrence of $p_{m-1}(a, b)$ or $p_{m-1}(b, a)$ in the outside boundary of D . Since $m \geq 4$, this shows that v has a syllable consisting of a single letter a or b , say b . Suppose that this b occurs in v_j . By the minimality of M , we have that B_j is not a copy of X_{ab} and so δ_j consists of a

single edge. This contradicts Lemma 4.4 applied to B_j . If M is a single region, the proof is analogous. \square

5. Infinitely generated groups

In this section we explain when we can extend the Main Theorem to infinitely generated G .

Let G be a group acting on a CAT(0) triangle complex X with finitely many isometry types of simplices. By [5, I.7.19], we have that X is complete, and so by [5, II.2.8] every finite subgroup of G fixes a cell of X . Consequently, if G acts almost freely, then there is a bound on the order of finite subgroups of G . Since G acts properly in the sense of [5, I.8.2], the following lemma is an immediate consequence of [5, II.7.5 and II.7.7(2)].

Lemma 5.1. *Let G be a group acting almost freely on a CAT(0) triangle complex with finitely many isometry types of simplices. Then every sequence $G_1 < G_2 < \dots$ of virtually abelian subgroups of G stabilises, that is, there is n such that for all $i \geq n$ the inclusion $G_i < G_{i+1}$ is an isomorphism.*

In view of Corollaries 2.6 and 2.7, the following completes the proof of Theorem A.

Corollary 5.2. *If X is a CAT(0) triangle complex with finitely many isometry types of simplices, then the Main Theorem holds also for infinitely generated G .*

Proof. Consider the family of finitely generated subgroups G_λ of G . If any G_λ contains F_2 , then we are done. Otherwise, by the Main Theorem every G_λ is virtually \mathbb{Z}^2 or virtually cyclic. It remains to observe that for some λ we have $G_\lambda = G$. Indeed, otherwise we can inductively define a sequence $G_{\lambda_1} \leq G_{\lambda_2} \leq \dots$ contradicting Lemma 5.1. \square

In [25] Wise presents a procedure of constructing a systolic complex associated to every simply connected $B(6)$ complex (in fact, to every simply connected $C(6)$ complex). This construction is described in details in [19] and we follow the notation from there. Without loss of generality we may assume that X is the union of its 2-cells: otherwise we attach equivariantly a 2-cell to every edge of degree 0. Then, the *Wise complex* $W(X)$ of X is defined as the nerve of the covering of X by closed 2-cells. The Wise complex of a simply connected $C(6)$ complex is systolic, see [25, Thm 6.7] and [19, Thm 7.10].

Lemma 5.3 ([20, Lem 2.2]). *Let G be a group acting properly on a uniformly locally finite systolic complex. Suppose that there is a bound on the order of finite subgroups of G . Then every sequence $G_1 < G_2 < \dots$ of virtually abelian subgroups of G stabilises.*

Similarly as in the CAT(0) setting, finite subgroups of isometries of systolic complexes fix points [6]. Therefore, a group acting almost freely on a systolic complex has a bound on the order of finite subgroups. Consequently, Lemma 5.3 implies the following:

Corollary 5.4. *If X is a simply connected $B(6)$ complex with uniformly locally finite $W(X)$, then Theorem C holds also for infinitely generated subgroups of G .*

We believe that Theorem C holds for infinitely generated subgroups of G without the assumption of the uniform local finiteness of $W(X)$. The reason for that assumption in Lemma 5.3 is that it is deduced from the systolic Flat Torus Theorem proved at the moment only for uniformly locally finite systolic complexes [7, Thm 6.1].

Finally, we complete the proof of Theorem B. By [11, Thm 5.6], a group acting properly on the Cayley complex for the standard presentation of an Artin group of extra-large type acts properly on a uniformly locally finite systolic complex. Therefore, we can also apply Lemma 5.3 to deduce the case of infinitely generated subgroups in Theorem B from the finitely generated case covered in Section 4.2.

Appendix A. When hypergraphs are trees for 2-dimensional Artin groups, by Jon McCammond, Damian Osajda, and Piotr Przytycki

In this appendix, we generalise Proposition 4.3 to the following. Here X is the Cayley complex for the standard presentation of an Artin group A_Γ .

Proposition A.1. *Suppose that A_Γ is a 2-dimensional Artin group. Then each hypergraph Λ in X is a tree if and only if Γ has no triangle with an edge labelled by 2.*

By Example 4.1 and Proposition A.1, the Main Theorem applies to X' implying the following.

Theorem A.2. *Let A_Γ be a 2-dimensional Artin group such that Γ has no triangle with an edge labelled by 2. Suppose that G acts almost freely on X . Then any finitely generated subgroup of G is virtually cyclic, or virtually \mathbb{Z}^2 , or contains a nonabelian free group.*

Note that if additionally Γ has no square with at least three edges labelled by 2, then by [11, Thm 5.6] G acts properly on a uniformly locally finite systolic complex and therefore we can apply Lemma 5.3 to show that Theorem A.2 holds also for G infinitely generated.

We first justify the ‘only if’ part of Proposition A.1.

Example A.3. Suppose that Γ has a triangle abc with the edge ab labelled by 2. Consider the following \mathcal{R} -reduced diagram M consisting of 12 regions. The 4 central regions have boundaries labelled by $aba^{-1}b^{-1}$. The 4 top and bottom regions have boundaries labelled by a and c , and the 4 left and right regions have boundaries labelled by b and c (see Fig. 5). M contains a cycle of Λ , which is thus not a tree.

For the ‘if’ part of Proposition A.1, we need the following.




Fig. 5. Example A.3. A cycle in Λ is marked by the dashed line.

Lemma A.4. *Let Σ be the Cayley complex of the Coxeter group W_Γ such that Γ has no triangle with an edge labelled by 2. Then for any 2-cells σ, τ of Σ sharing an edge, and a hypergraph Λ_σ intersecting σ , there is a hypergraph Λ_τ intersecting τ that is disjoint from or equal to Λ_σ .*

Proof. Denote $e = \sigma \cap \tau$. Note that we can assume that Λ_σ is disjoint from e , since otherwise we can take $\Lambda_\tau = \Lambda_\sigma$.

Suppose first that τ is a square. Then take Λ_τ to be the hyperplane intersecting τ but not e . Let f be an edge of σ sharing a vertex v with e . If Λ_σ intersects f , then let α be the path obtained by concatenating at v a half-edge of f and a half-edge in τ ending at Λ_τ . Note that α is a geodesic in the Moussong metric on Σ , by the assumption that Γ has no triangle with an edge labelled by 2. Moreover, the Alexandrov angle between the endpoints of α and Λ_σ and Λ_τ are $\frac{\pi}{2}$. Consequently Λ_σ and Λ_τ are disjoint. Analogously, for the hyperplane Λ'_σ intersecting the other edge of σ sharing a vertex with e , we have that Λ'_σ and Λ_τ are disjoint. Consequently, Λ_τ is contained in the same component of $\Sigma - \Lambda_\sigma \cup \Lambda'_\sigma$ as e . Then Λ_τ is also disjoint from all other hyperplanes intersecting σ but not e .

If σ is a square, then the same argument shows that we can take Λ_τ to be any hyperplane intersecting τ but not e .

It remains to consider the case where neither of σ, τ are squares. We modify the Moussong metric on Σ in the following way. Every square remains a Euclidean square of side length 1. Every 2-cell that is not a square is subdivided into triangles by (Moussong) geodesic segments joining vertices with the centre and we turn each triangle into the

Euclidean equilateral triangle of side length 1. Note that because Γ has no triangle with an edge labelled by 2, this metric is still CAT(0). As before, let f be an edge of σ sharing a vertex v with e . Let g be the edge of τ sharing a vertex with e distinct from v , and let α be the path obtained by concatenating at v a half-edge of f and the geodesic segment from v to the centre of τ . If Λ_σ intersects f , then take Λ_τ to be the hypergraph intersecting g . Again α is a geodesic meeting Λ_σ and (since we modified the metric) Λ_τ at Alexandrov angle $\frac{\pi}{2}$. Consequently Λ_σ and Λ_τ are disjoint.

Finally, consider a hyperplane Λ'_σ intersecting σ but none of its edges sharing vertices with e . Let w be the centre of σ , and let λ_σ be the component of $\Lambda_\sigma - w$ intersecting f . Let λ'_σ be the component of $\Lambda'_\sigma - w$ intersecting $\partial\sigma$ earlier if we traverse it starting from e and ending with f . Since we modified the metric, $\Lambda'' = \lambda_\sigma \cup w \cup \lambda'_\sigma$ is convex. Using the path α we obtain that Λ'' is disjoint from Λ_τ . Since Λ'_σ is contained in the closure of the component of $\Sigma - \Lambda''$ that does not contain e , we have that Λ'_σ and Λ_τ are disjoint. \square

Proof of Proposition A.1. The ‘only if’ part follows from Example A.3. For the ‘if’ part, given a cycle γ in Λ , define D_i, B_j, δ_j and M as in the proof of Proposition 4.3. Note that δ_j are still nontrivial, since Lemma 4.4 obviously holds for $m = 2$ and the word $a^k b^l$.

We claim that M has no 2-cells. Indeed, otherwise let τ_X be a 2-cell of X in the image of M containing an edge e of some δ_j , and let σ_X be a 2-cell of B_j containing e . Let σ, τ be the projections of σ_X, τ_X to the Cayley complex Σ of W_Γ . Let Λ_σ be the projection to Σ of Λ , which intersects σ since D_{i_j+1} lies in the same block as σ_X , and hence projects also to σ . By Lemma A.4 there is a hypergraph Λ_τ intersecting τ that is disjoint from or equal to Λ_σ . Let M' be a diagram obtained from M by attaching along δ_j diagrams in B_j and such that γ traverses consecutively its boundary 2-cells. The component of the preimage of Λ_τ in M' intersecting τ_X forms a cycle contradicting the minimality of M . This justifies the claim.

Thus M is a tree. Each leaf of that tree is a trivial δ_j , which is a contradiction. \square

References

- [1] K.I. Appel, P.E. Schupp, Artin groups and infinite Coxeter groups, *Invent. Math.* 72 (2) (1983) 201–220, MR700768.
- [2] W. Ballmann, M. Brin, Orbihedra of nonpositive curvature, *Publ. Math. Inst. Hautes Études Sci.* 82 (1995) 169–209, 1996, MR1383216.
- [3] M. Bestvina, M. Feighn, M. Handel, The Tits alternative for $\text{Out}(\text{Fn})$. I. Dynamics of exponentially-growing automorphisms, *Ann. Math. (2)* 151 (2) (2000) 517–623, MR1765705.
- [4] M. Bestvina, M. Feighn, M. Handel, The Tits alternative for $\text{Out}(\text{Fn})$. II. A Kolchin type theorem, *Ann. Math. (2)* 161 (1) (2005) 1–59, MR2150382.
- [5] M.R. Bridson, A. Haefliger, Metric spaces of non-positive curvature, *Grundlehren der Mathematischen Wissenschaften (Fundamental Principles of Mathematical Sciences)*, vol. 319, Springer-Verlag, Berlin, 1999, MR1744486.
- [6] V. Chepoi, D. Osajda, Dismantlability of weakly systolic complexes and applications, *Trans. Am. Math. Soc.* 367 (2) (2015) 1247–1272, MR3280043.
- [7] T. Elsner, Flats and the flat torus theorem in systolic spaces, *Geom. Topol.* 13 (2) (2009) 661–698, MR2469526.

- [8] É. Ghys, P. de la Harpe, Sur les groupes hyperboliques d'après Mikhael Gromov, in: *Progress in Mathematics*, vol. 83, Birkhäuser Boston, Inc., Boston, MA, 1990, Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988, MR1086648.
- [9] A. Genevois, Cubical-like geometry of quasi-median graphs and applications to geometric group theory, available at arXiv:1712.01618, 2017.
- [10] M. Gromov, Hyperbolic groups, in: *Essays in group theory*, 1987, pp. 75–263, MR919829.
- [11] J. Huang, D. Osajda, Large-type Artin groups are systolic, *Proc. Lond. Math. Soc.* (3) 120 (1) (2020) 95–123, MR3999678.
- [12] N.V. Ivanov, Algebraic properties of the Teichmüller modular group, *Dokl. Akad. Nauk SSSR* 275 (4) (1984) 786–789, MR745513.
- [13] T. Januszkiewicz, J. Świątkowski, Simplicial nonpositive curvature, *Publ. Math. Inst. Hautes Études Sci.* 104 (2006) 1–85, MR2264834.
- [14] D.H. Kochloukova, P.A. Zalesskii, Tits alternative for 3-manifold groups, *Arch. Math.* (Basel) 88 (4) (2007) 364–367, MR2311844.
- [15] P.A. Linnell, On accessibility of groups, *J. Pure Appl. Algebra* 30 (1) (1983) 39–46, MR716233.
- [16] J. McCarthy, A “Tits-alternative” for subgroups of surface mapping class groups, *Trans. Am. Math. Soc.* 291 (2) (1985) 583–612, MR800253.
- [17] A. Martin, P. Przytycki, Acylindrical actions for two-dimensional Artin groups of hyperbolic type, *Int. Math. Res. Not.* (2021), <https://doi.org/10.1093/imrn/rnab068>, available at arXiv:1906.03154, in press.
- [18] S. Norin, D. Osajda, P. Przytycki, Torsion groups do not act on 2-dimensional CAT(0) complexes, *Duke Math J.*, available at arXiv:1902.02457, 2021, in press.
- [19] D. Osajda, T. Prytuła, Classifying spaces for families of subgroups for systolic groups, *Groups Geom. Dyn.* 12 (3) (2018) 1005–1060, MR3845715.
- [20] T. Prytuła, Solvable subgroup theorem for simplicial nonpositive curvature, *Int. J. Algebra Comput.* 28 (4) (2018) 605–611, MR3819825.
- [21] M. Sageev, D.T. Wise, The Tits alternative for CAT(0) cubical complexes, *Bull. Lond. Math. Soc.* 37 (5) (2005) 706–710, MR2164832.
- [22] J. Tits, Free subgroups in linear groups, *J. Algebra* 20 (1972) 250–270, MR0286898.
- [23] H. van der Lek, The homotopy type of complex hyperplane complements, PhD thesis, Katholieke Universiteit te Nijmegen, 1983.
- [24] P. Walters, An introduction to ergodic theory, *Graduate Texts in Mathematics*, vol. 79, Springer-Verlag, New York-Berlin, 1982, MR648108.
- [25] D.T. Wise, Sixtolic complexes and their fundamental groups, 2003, unpublished manuscript.
- [26] D.T. Wise, Cubulating small cancellation groups, *Geom. Funct. Anal.* 14 (1) (2004) 150–214, MR2053602.
- [27] X. Xie, Tits alternative for closed real analytic 4-manifolds of nonpositive curvature, *Topol. Appl.* 136 (1–3) (2004) 87–121, MR2023412.