HYPERBOLIC ISOMETRIES OF SYSTOLIC COMPLEXES

Tomasz Prytuła

PhD Thesis
Department of Mathematical Sciences
University of Copenhagen

Tomasz Prytuła
Department of Mathematical Sciences
University of Copenhagen
Universitetsparken 5
DK-2100 København Ø
Denmark
t.prytula@math.ku.dk

PhD thesis submitted to the PhD School of Science, Faculty of Science,
University of Copenhagen, Denmark in February 2017.

Academic advisors:

Jesper Michael Møller	Damian Osajda
University of Copenhagen	University of Wrocław and
Denmark	Polish Academy of Sciences
	Poland

Assessment committee:

Ian J. Leary	Jacek Świątkowski	Nathalie Wahl (chair)
University of Southampton	University of Wrocław	University of Copenhagen
United Kingdom	Poland	Denmark

ISBN 978-87-7078-923-3

Abstract

The main topics of this thesis are the geometric features of systolic complexes arising from the actions of hyperbolic isometries. The thesis consists of an introduction followed by two articles.

Given a hyperbolic isometry h of a systolic complex X, our central theme is to study the minimal displacement set of h and its relation to the actions of h on X and on the systolic boundary ∂X. We describe the coarse-geometric structure of the minimal displacement set and establish some of its properties that can be seen as a form of quasi-convexity. We apply our results to the study of geometric and algebraic-topological features of systolic groups. In addition, we provide new examples of systolic groups.

In the first article we show that the minimal displacement set of a hyperbolic isometry of a systolic complex is quasi-isometric to the product of a tree and the real line. We use this theorem to construct a low-dimensional model for the classifying space $\underline{\underline{E}}G$ for a group G acting properly on a systolic complex, and to describe centralisers of hyperbolic isometries in systolic groups.

In the second article we are interested in the induced action of h on the systolic boundary, and particularly in the fixed points of this action. The main theorem gives a characterisation of the isometries acting trivially on the boundary in terms of their centralisers in systolic groups.

Resumé

Hovedemnerne i denne afhandling er de geometriske egenskaber ved systoliske komplekser, som opstår fra virkningerne af hyperbolske isometrier. Afhandlingen består af en introduktion efterfulgt af to artikler.

Givet en hyperbolsk isometri h af et systolisk kompleks X er vores hovedtema at studere h's minimale forskydningsmængde og dens relation til virkningerne af h på X og på den systoliske rand ∂X. Vi beskriver den minimale forskydningsmængdes grov-geometriske struktur og fastlægger nogle af dens egenskaber, som kan betragtes som en form for kvasi-konveksitet. Vi anvender vores resultater til at studere geometriske og algebraisk-topologiske egenskaber ved systoliske grupper. Desuden giver vi nye eksempler på systoliske grupper.

I den første artikel viser vi at den minimale forskydningsmængde for en hyperbolsk isometri af et systolisk kompleks er kvasi-isometrisk med produktet af et træ og den reelle akse. Vi bruger denne sætning til at konstruere en lavdimensional model af det klassificerende rum $\underline{\underline{E}}G$ for en gruppe G med proper virkning på et systolisk kompleks, og til at beskrive centralisatorerne af hyperbolske isometrier i systoliske grupper.

I den anden artikel interesserer vi os for den inducerede virkning af h på den systoliske rand og især for fixpunkterne under denne virkning. Hovedsætningen giver en karakterisering af de isometrier som virker trivielt på randen ud fra deres centralisatorer i systoliske grupper.

Acknowledgements

First of all I would like to thank my advisors, Jesper Michael Møller and Damian Osajda, for their help and support during the years of my studies. I am grateful to Jesper Michael for his constant optimism and encouragement, and to Damian for all the inspiring discussions and questions that guided me in the right direction, and for our fruitful collaboration.

I also thank Piotr Przytycki for suggesting the problem addressed in the second article of this thesis, and for the many insightful discussions we had during my stay at McGill University. My special thanks are due to Dieter Degrijse for generously sharing his knowledge, answering all of my questions, proofreading and commenting on this thesis. I am indebted to Kristian Moi and Irakli Patchkoria for educating me in algebraic topology and always finding time to explain things to me. I am grateful to Chris Davis and Kristian Olesen for their help in the final stages of my thesis work.

I thank fellow PhD students and the post-docs at the University of Copenhagen, and in particular my office mates Amalie, Isabelle, Manuel, Massimiliano, Matthias, Mauricio, Niek and Ryo for our friendship and for making my daily life in Copenhagen so enjoyable.

Finally, I would like to thank my parents Lidia and Kazimierz and my girlfriend Julie for always supporting me.

Tomasz Prytuła

Copenhagen, February 2017

Contents

Part I

Thesis overview

Introduction

This thesis is concerned with geometric aspects of systolic complexes connected to the actions of hyperbolic isometries. We are in particular interested in applications of these aspects to the study of systolic groups. The thesis consists of an introduction and the following two articles.

A. *Classifying spaces for families of subgroups for systolic groups.*

B. *Hyperbolic isometries and boundaries of systolic complexes.*

In the introduction we recall the theory of systolic complexes and highlight those of its aspects which are relevant to our study. We also discuss connections to other approaches to nonpositive curvature, particularly to the theory of CAT(0) spaces and groups, in order to put our results in a broader context. We conclude the introduction with a brief presentation of the two articles and a discussion of perspectives for future research.

Nonpositively curved groups

One of the main topics in geometric group theory is the study of nonpositively curved groups, understood in a broad sense. This includes classical notions of fundamental groups of nonpositively curved Riemannian manifolds, combinatorial theory of small cancellation groups, as well as more modern approaches, like CAT(0) or δ–hyperbolic groups and their generalisations. Due to an observation of Gromov, a special place among CAT(0) spaces is occupied by CAT(0) cubical complexes. Namely, for a simply connected cubical complex the standard piecewise Euclidean metric is CAT(0) if and only if a certain combinatorial condition is satisfied (flagness of vertex links). Many of the crucial theorems in the theory of CAT(0) cubical complexes can be proven without referring to the CAT(0) metric, and using the combinatorial structure only.

Furthermore, a combinatorial approach has been used to cubulate (i.e., equip with a geometric action on a CAT(0) cubical complex) numerous groups, as well as to establish various conjectures that remain open for arbitrary CAT(0) groups (e.g., Tits alternative, finiteness of asymptotic dimension).

The question arose whether there is a similar characterisation of nonpositive curvature for simplicial complexes. A partial, yet very inspiring answer to that question is the theory of systolic complexes and groups, sometimes referred to as *Simplicial Nonpositive Curvature* (SNPC).

Simplicial Nonpositive Curvature

A *systolic complex* is a simply connected simplicial complex whose vertex links satisfy a certain condition called 6*–largeness*. A *systolic group* is a group acting geometrically on a systolic complex. The condition of 6–largeness is in spirit very similar to the Gromov's link condition, i.e., it is local, combinatorial and easily checkable, and together with simple connectedness it implies contractibility. However, it does not imply the nonpositive curvature of the standard piecewise Euclidean metric.

Systolic complexes were first introduced by V. Chepoi under the name of *bridged complexes* [Che00], although their 1–skeleta, the *bridged graphs*, had appeared much earlier in graph theory [SC83, FJ87]. The definition of bridged graphs in a way captures the spirit of nonpositive curvature, as it can be stated in terms of the convexity of metric balls. Consequently, from the combinatorial viewpoint, systolic complexes are rather natural candidates for 'nonpositively curved' simplicial complexes.

The theory of systolic complexes and groups in its current shape was developed by Januszkiewicz-Świątkowski [JŚ06] and, independently, by Haglund [Hag03] and it has been intensively studied from a group-theoretical perspective ever since. The main advantage of the combinatorial definition is that it is relatively easy to verify whether a given simplicial complex is systolic, while on the other hand, it seems practically impossible to check whether its piecewise Euclidean metric is CAT(0). In particular, a combinatorial approach led to numerous constructions of interesting groups, e.g., hyperbolic Coxeter groups of large virtual cohomological dimension [JŚ03, Osa13a], groups having strong fixed point properties [ABJ+09], and spaces, e.g., negatively curved compact branched coverings of pseudomanifolds [JŚ06], finite developments of billiard tables [JŚ10].

It is worth noting that despite being similar to other classes of nonpositively curved spaces, systolic complexes display certain asymptotic properties which are rather unexpected and exotic (especially in higher dimensions). The most prominent of these properties is the Asymptotic Hereditary Asphericity discussed later in this introduction. In particular, these properties imply that in many aspects systolic groups behave very differently from the classical nonpositively curved groups like, for example, lattices in $\mathrm{Isom}(\mathbb{H}^n)$ (cf. [JŚ07]).

Over the last decade the theory of systolic complexes has been further developed, with a particular emphasis on its metric aspects. Many nonpositive curvature-like properties of systolic complexes and groups have been established and their exotic features have been formalised and put into a proper framework.

Relationship to δ–hyperbolic, small cancellation and CAT(0) groups

The theory of systolic complexes and groups is closely connected to other approaches to nonpositively curved groups. The most notable is the analogy between SNPC and the CAT(0) geometry which we discuss in more detail.

Before doing so we need to recall some terminology. The condition of 6–largeness is a special case of the condition of k–largeness, which is defined for any $k \geqslant 6$. This leads to the notions of k–systolic complexes and k–systolic groups (we identify the terms 'systolic' and '6–systolic'). If $k \geqslant m$ then k–largeness implies m–largeness, and thus all k–systolic complexes and groups are in particular systolic.

δ–hyperbolic groups

A large subclass of systolic complexes, namely the 7–systolic complexes, are δ–hyperbolic. In fact, finding a simple combinatorial condition implying δ–hyperbolicity was one of the motivations for the condition of k–largeness (see [JŚ06]). In general a systolic complex does not have to be δ–hyperbolic, as it may contain flats. When one passes to systolic groups, the existence of flats is the only obstruction [Prz07] (the analogous characterisation holds for CAT(0) groups). On the other hand, most of the classical hyperbolic groups of dimension at least 3 are not systolic, thanks to the aforementioned Asymptotic Hereditary Asphericity property of systolic groups.

Let us mention that the existence of flats in systolic complexes is essential to our study. Namely, in the hyperbolic case various aspects of systolic complexes and groups reduce to the known properties and constructions for hyperbolic groups, e.g., the systolic boundary reduces to the Gromov boundary.

Small cancellation theory

In [JŚ06] the authors claim that SNPC may be seen as a higher-dimensional version of the small cancellation theory. This was made precise by D. Wise, who showed that any $C(k)$ small cancellation group is k–systolic [Wis03]. In article A we extend Wise's result to the case of *graphical small cancellation* groups. In dimension 2 the converse also holds, i.e., to a 2–dimensional k–systolic complex one associates a dual cell complex that satisfies the $C(k)$ graphical small cancellation condition.

These two approaches are also similar when it comes to the methods. The fundamental tools used to study systolic complexes are the minimal surfaces, which can be seen as simplicial analogues of the van Kampen diagrams used in the small cancellation theory.

CAT(0) spaces and groups

The theory of systolic complexes and groups is to a large extent parallel to the theory of CAT(0) spaces and groups. In dimension 2 it simply reduces to the CAT(0) geometry, i.e., a 2–dimensional simplicial complex is systolic if and only if its standard piecewise Euclidean metric is CAT(0). In higher dimensions these two theories differ, as one can easily construct the appropriate (counter) examples. The situation becomes much more difficult when one considers groups. There are plenty of CAT(0) groups which are not systolic, but there is no systolic group which is known not to be CAT(0). It is somewhat surprising how faithfully the condition of 6–largeness captures the key aspects of nonpositive curvature.

Consequently, over the years, for numerous properties of CAT(0) spaces, systolic analogues have been established. This includes, among others, the Cartan-Hadamard Theorem [Che00, JŚ06], the Fixed Point Theorem [Prz08, CO15], the Flat Torus Theorem [Els09a], the classification of isometries [Els09b] and the existence of a boundary which is an EZ–structure [OP09]. The results for systolic complexes and groups presented in this thesis are also to some degree analogous to the known results in CAT(0) geometry.

However, in some aspects systolic complexes behave better than CAT(0) spaces. This is mostly due to their combinatorial nature (which makes them actually more similar to the CAT(0) cubical complexes). For example, similarly to CAT(0) cubical groups, systolic groups are biautomatic [JŚ06, Świ06], while it is an open question for arbitrary CAT(0) groups.

Asymptotic Hereditary Asphericity

One of the most significant properties of systolic complexes, which distinguishes them from other classes of nonpositively curved spaces, is the so-called *Asymptotic Hereditary Asphericity* (AHA), introduced in [JŚ07]. It is a coarse property of metric spaces that describes asymptotic behaviour of spheres. In an AHA space, all spheres of dimension at least 2 are metrically uniformly 'slim'. This assertion can easily be shown to hold for spaces of (asymptotic) dimension at most 2, as in this case there is no room to map an n–sphere with $n \geqslant 2$ in an 'essential' way. What is surprising is that systolic complexes are AHA regardless of their dimension.

The AHA property implies many of the aforementioned exotic features of systolic complexes and groups, including the nonexistence of nonpositively curved manifold subgroups of dimension greater that 2, restrictions on AHA products, and a peculiar topology at infinity (see [JŚ07, Osa07]). One could say that AHA spaces in many ways behave like 2–dimensional objects. In particular, it follows that systolic groups do not contain subgroups isomorphic to $\mathbb{Z}^n$ for $n > 2$ and systolic complexes do not contain flats of dimension greater than 2. The further study of AHA spaces and groups was undertaken in [OŚ15]. Along with other results, it is shown there that AHA imposes certain restrictions on the topology of systolic boundaries.

In our work we do not explicitly invoke the AHA property; however, it is lurking behind many of our arguments. For example, in article A we use the S^kFRC property which is a direct consequence of AHA, see [JŚ07].

Hyperbolic isometries

A substantial part of the theory of systolic complexes concerns their metric aspects (the metric under consideration is the combinatorial metric on the 1–skeleton). An isometry (that is, a simplicial automorphism) h of a systolic complex X is called *hyperbolic* if it does not fix any simplex of X. The main object used to study the action of h on X is the *minimal displacement set* of h, which is a subcomplex of X spanned by all the vertices of X that are moved by h the minimal distance. In the case of CAT(0) spaces, the geometry of the minimal displacement set is well understood, and it determines many aspects of the action of h on X (see [BH99]). Particularly, if h is an element of a group G acting geometrically on X, then various algebraic properties of h in G are reflected in the geometric properties of the minimal displacement set. An important role in this correspondence is played by the centraliser of h in G whose action on X preserves the minimal displacement set of h.

The main theme of this thesis is to study the geometry of the minimal displacement set in the setting of systolic complexes, with a view to the particular applications to the classifying spaces for families of subgroups (article A) and to the study of actions of systolic groups on systolic boundaries (article B).

In the case where X is a CAT(0) space, the minimal displacement set of h is a convex (and hence isometrically embedded and CAT(0)) subspace of X and it splits as a product $Y \times \mathbb{R}$ where Y is a convex subspace, and such that h acts trivially on Y and as a translation on $\mathbb{R}$ (any h–invariant line $\{y\} \times \mathbb{R}$ is called an *axis* of h). The above result is the so-called Product Decomposition Theorem (see [BH99]) and it is a very useful tool that gives a control on the action of h on X. Among its applications there are classical theorems in the CAT(0) geometry like the Flat Torus Theorem or the Solvable Subgroup Theorem, as well as more modern results which we discuss later.

In the systolic setting the minimal displacement set does not have all of the desired metric properties. It is an isometrically embedded systolic subcomplex of X, however, it is not convex (in fact not even quasi-convex) [Els09b]. The lack of convexity is the main difficulty which we encounter in article B. Furthermore, the minimal displacement set does not have to consist of the axes of h, i.e., there are examples of hyperbolic isometries that have no invariant geodesic lines at all (however, for any such isometry one of its powers has an axis). Also, much less is known about the structure of the minimal displacement set. In article A we determine this structure up to quasi-isometry, thus establishing the systolic Product Decomposition Theorem.

Theorem (A, Theorem A)**.** *The minimal displacement set of a hyperbolic isometry of a uniformly locally finite systolic complex is quasi-isometric to the product $T \times \mathbb{R}$ of a tree and the line.*

Compared to the CAT(0) case our result may be seen as both weaker and stronger. The splitting in the CAT(0) case is isometric and it is obtained within the ambient space X. Ours is an abstract splitting, and only up to quasi-isometry. However, in the CAT(0) case the first factor can be an arbitrary CAT(0) space, whereas in our case it is a tree. This restriction is due to the AHA property of systolic complexes and it has deep consequences for certain algebraic properties of h, that we explore in article A. In article B, instead of its structure, we focus on the coarse-geometric aspects of the minimal displacement set that stem from the fact that it is not convex.

A. Classifying spaces for families of subgroups for systolic groups (joint with Damian Osajda)

A classifying space of a group G for a family of its subgroups $\mathcal{F}$ is a certain homotopical invariant of G which classifies actions of G with stabilisers in $\mathcal{F}$. If $\mathcal{F}$ consists of just the trivial subgroup, then the classifying space for $\mathcal{F}$ is the same as the free contractible G–space EG. Among nontrivial families, the two main cases of interest are the family of all finite subgroups of G and the family of all virtually cyclic subgroups of G, where the respective classifying spaces are denoted by $\underline{E}G$ and $\underline{\underline{E}}G$. The significance of $\underline{E}G$ and $\underline{\underline{E}}G$ comes from the fact that homology of these spaces appears in respectively Baum-Connes and Farrell-Jones isomorphism conjectures. Therefore, one would like to construct simple, finite-dimensional models for $\underline{E}G$ and $\underline{\underline{E}}G$, and in particular relate their dimension to other notions of dimension of G (virtual cohomological dimension, asymptotic dimension, etc.).

The space $\underline{E}G$ is by now a classical object and for the majority of (interesting) groups, finite-dimensional models for $\underline{E}G$ have been constructed [Lüc05]. The space $\underline{\underline{E}}G$ gained recognition more recently, and it has been intensively studied over the last decade. The standard approach to models for $\underline{\underline{E}}G$ is the construction of Lück-Weiermann [LW12], which builds a model for $\underline{\underline{E}}G$ out of a model for $\underline{E}G$ and certain classifying spaces associated to infinite cyclic subgroups of G. This approach has been used in the vast majority of cases, including the case of CAT(0) groups [Lüc09]. For δ–hyperbolic groups a finite-dimensional model for $\underline{\underline{E}}G$ was constructed in [JPL06], where the authors used what later turned out to be a special case of the construction in [LW12].

In this article we construct a finite-dimensional model for the classifying space $\underline{\underline{E}}G$ where G is a group acting properly on a systolic complex. Given such an action, an infinite order element h of G is precisely a hyperbolic isometry of the systolic complex. The key point in the construction of [LW12] is to find a 'nice' model for $\underline{E}C_G(h)/\langle h\rangle$, where $C_G(h)$ is the centraliser of h in G. For this, one uses the action of $C_G(h)$ on the minimal displacement set of h. In the CAT(0) case, the Product Decomposition Theorem implies that $C_G(h)/\langle h\rangle$ acts properly on a CAT(0) space Y, and thus one can take Y as $\underline{E}C_G(h)/\langle h\rangle$ (see [Lüc09]). In our approach, the systolic Product Decomposition Theorem implies that $C_G(h)/\langle h\rangle$ is virtually free, which gives a 1–dimensional model for $\underline{E}C_G(h)/\langle h\rangle$ and in consequence a much nicer model for $\underline{\underline{E}}G$, when compared with a generic CAT(0) group.

Theorem (A, Theorem C)**.** *Let G be a group acting properly on a uniformly locally finite d–dimensional systolic complex. Then there exists a model for $\underline{\underline{E}}G$ of dimension*

$$\dim \underline{\underline{E}}G = \begin{cases} d+1 & \text{if } d \leqslant 3, \\ d & \text{if } d \geqslant 4. \end{cases}$$

We also study connections between small cancellation theory and SNPC. Following ideas of D. Wise we show that any group acting properly and/or cocompactly on a graphical small cancellation complex acts in the same way on a systolic complex. It follows that any such group has an at most 2–dimensional model for $\underline{E}G$. For groups with a 2–dimensional model for $\underline{E}G$ that satisfy some mild finiteness conditions, a 3–dimensional model for $\underline{\underline{E}}G$ was constructed in [Deg17]. However, there are groups acting properly on graphical small cancellation complexes that do not satisfy conditions needed in [Deg17], and thus our approach gives new models for $\underline{\underline{E}}G$. The main point of our theorem though, is that systolicity implies many properties of 'metric nature' that are difficult to establish using combinatorial methods of small cancellation theory.

B. Hyperbolic isometries and boundaries of systolic complexes

The boundary at infinity ∂X of a systolic complex X was constructed in [OP09]. The construction is ideologically very similar to the visual boundary of a CAT(0) space or to the Gromov boundary of a δ–hyperbolic space. In the two latter cases, the study of groups via their induced action on the boundary turned out to be a very fruitful approach. A general picture is that the topology of the boundary is reflected in the algebraic properties of a group, and often one obtains this correspondence by studying the dynamics of the action of individual (hyperbolic) elements. In particular this approach was used to connect the topology of the boundary with the existence of a splitting of a group as a direct product [Rua01] or as an amalgamated product [Bow98, PS09].

In the systolic setting the action of a group on the boundary has not yet been investigated. We undertake this task in our article, with the focus on the actions of hyperbolic isometries. In the CAT(0) or δ–hyperbolic case, the starting point of analysing the dynamics of a hyperbolic isometry h on the boundary is the observation that h has two canonical fixed points $h^{+\infty}$ and $h^{-\infty}$. These points are given by the endpoints of an axis of h in the first case and a *quasi-axis* of h in the second case.

We show the existence of these two fixed points in the systolic setting (which is not that straightforward, as in this case an axis of h does not immediately determine two points on the boundary). We also characterise hyperbolic isometries that act as the identity on the boundary, in a similar way as it is done for CAT(0) groups in [Rua01].

Theorem (B, Theorem 1)**.** *Let G be a group acting geometrically on a systolic complex X, and let $h \in G$ be a hyperbolic isometry. Then h acts trivially on the boundary ∂X if and only if the centraliser $C_G(h)$ has finite index in G.*

The approach in [Rua01] uses the action of $C_G(h)$ on the minimal displacement set of h. In our case the minimal displacement set is not convex and therefore it cannot be effectively used to study geodesic rays (that represent points of the boundary). We thus replace it with the *K–displacement set*, which consists of vertices that are moved by h the distance at most K, for some $K > 0$. Choosing the appropriate K allows us to circumvent the lack of convexity and proceed analogously as in [Rua01].

Perspectives for future research

Splittings of systolic groups over virtually cyclic subgroups

A theorem of Bowditch [Bow98] states that a one-ended hyperbolic group G splits over a virtually cyclic subgroup if and only if the Gromov boundary of G has a cut-pair. Papasoglu-Swenson [PS09] prove an analogous result for a group acting geometrically on a CAT(0) space X (in this case the Gromov boundary is replaced with the visual boundary of X). We would like to establish a systolic analogue of these theorems following the approach of [PS09]. This approach uses a variety of properties of CAT(0) boundaries which are not known in the systolic setting. In article B we establish some of these properties.

The key idea in [PS09] is to make use of a metric structure of the boundary and study the dynamics of the action of a hyperbolic isometry h on the *Tits boundary* of the CAT(0) space X (rather than the visual boundary). The main result says that this action is the so-called *π–convergence action.* To establish π–convergence the authors use, among others, the fact that an angle between two geodesic rays in a CAT(0) space, despite its global nature, can be defined locally.

In our case one has to first define a version of a Tits metric on the systolic boundary. This can be done in the same way as in the CAT(0) case, thanks to the contracting properties of geodesic rays used to define the systolic boundary (see [OP09]). However, since the systolic boundary has a 'coarse-geometric definition' it seems difficult to find a good local notion of an angle. This is the main (expected) difficulty in establishing a systolic version of π–convergence.

An interesting side question arises. A classical theorem in CAT(0) geometry states that the Tits boundary of a CAT(0) space is a CAT(1) space (see [BH99]). We believe that the same holds for the Tits boundary of systolic complexes. Moreover, we expect that Asymptotic Hereditary Asphericity imposes further restrictions on the Tits metric, namely that any geodesic triangle of circumference less than 2π is in fact a tripod. This would be a metric counterpart to a theorem of Osajda-Świątkowski [OŚ15] which says that the systolic boundary contains no 2–disks.

Minimal displacement set for generalisations of Simplicial Nonpositive Curvature

There are various generalisations of Simplicial Nonpositive Curvature, which include *weakly systolic complexes* [Osa13b] and *bucolic complexes* [BCC$^+$13]. One could define and study the minimal displacement set of hyperbolic isometries of these complexes and, e.g., try to establish a Product Decomposition Theorem.

In the setting of weakly systolic complexes, one does not have the entire minimal surface theory [Els09a] which we use in article A, but there are some partial results (e.g., the classification of minimal surfaces or the 'Projection Lemma'). Also, these complexes are not necessarily AHA, so in order to determine the coarse geometry of the factor Y in the product $Y \times \mathbb{R}$, one may need to use techniques different from ours. Nonetheless, it is still conceivable that Y is a tree. This would allow us to proceed as in article A and construct a finite-dimensional model for the classifying space $\underline{\underline{E}}G$ for weakly systolic groups.

Bibliography

[ABJ+09] Goulnara Arzhantseva, Martin R. Bridson, Tadeusz Januszkiewicz, Ian J. Leary, Ashot Minasyan, and Jacek Świątkowski, *Infinite groups with fixed point properties*, Geom. Topol. **13** (2009), no. 3, 1229–1263.

[Bow98] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.

[BCC+13] B. Brešar, J. Chalopin, V. Chepoi, T. Gologranc, and D. Osajda, *Bucolic complexes*, Adv. Math. **243** (2013), 127–167.

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999.

[Che00] Victor Chepoi, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179.

[CO15] Victor Chepoi and Damian Osajda, *Dismantlability of weakly systolic complexes and applications*, Trans. Amer. Math. Soc. **367** (2015), no. 2, 1247–1272.

[Deg17] Dieter Degrijse, *A cohomological characterization of locally virtually cyclic groups*, Adv. Math. **305** (2017), 935–952.

[Els09a] Tomasz Elsner, *Flats and the flat torus theorem in systolic spaces*, Geom. Topol. **13** (2009), no. 2, 661–698.

[Els09b] ———, *Isometries of systolic spaces*, Fund. Math. **204** (2009), no. 1, 39–55.

[FJ87] Martin Farber and Robert E. Jamison, *On local convexity in graphs*, Discrete Math. **66** (1987), no. 3, 231–247.

[Hag03] F. Haglund, *Complexes simpliciaux hyperboliques de grande dimension*, Prepublication Orsay **71** (2003), preprint, available at `http://www.math.u-psud.fr/~haglund/cpl_hyp_gde_dim.pdf`.

[JŚ03] Tadeusz Januszkiewicz and Jacek Świątkowski, *Hyperbolic Coxeter groups of large dimension*, Comment. Math. Helv. **78** (2003), no. 3, 555–583.

[JŚ06] ———, *Simplicial nonpositive curvature*, Publ. Math. Inst. Hautes Études Sci. **104** (2006), 1–85.

[JŚ07] ———, *Filling invariants of systolic complexes and groups*, Geom. Topol. **11** (2007), 727–758.

[JŚ10] ———, *Non-positively curved developments of billiards*, J. Topol. **3** (2010), no. 1, 63–80.

[JPL06] Daniel Juan-Pineda and Ian J. Leary, *On classifying spaces for the family of virtually cyclic subgroups*, Recent developments in algebraic topology, Contemp. Math., vol. 407, Amer. Math. Soc., Providence, RI, 2006, pp. 135–145.

[Lüc05] Wolfgang Lück, *Survey on classifying spaces for families of subgroups*, Infinite groups: geometric, combinatorial and dynamical aspects, Progr. Math., vol. 248, Birkhäuser, Basel, 2005, pp. 269–322.

[Lüc09] ———, *On the classifying space of the family of virtually cyclic subgroups for CAT(0)-groups*, Münster J. of Math. **2** (2009), 201–214.

[LW12] Wolfgang Lück and Michael Weiermann, *On the classifying space of the family of virtually cyclic subgroups*, Pure Appl. Math. Q. **8** (2012), no. 2, 497–555.

[Osa07] Damian Osajda, *Connectedness at infinity of systolic complexes and groups*, Groups Geom. Dyn. **1** (2007), no. 2, 183–203.

[Osa13a] ———, *A construction of hyperbolic Coxeter groups*, Comment. Math. Helv. **88** (2013), no. 2, 353–367.

[Osa13b] ———, *A combinatorial non-positive curvature I: weak systolicity* (2013), preprint, available at `http://www.math.uni.wroc.pl/~dosaj/trav/cnpc-web.pdf`.

[OP09] Damian Osajda and Piotr Przytycki, *Boundaries of systolic groups*, Geom. Topol. **13** (2009), no. 5, 2807–2880.

[OŚ15] Damian Osajda and Jacek Świątkowski, *On asymptotically hereditarily aspherical groups*, Proc. Lond. Math. Soc. (3) **111** (2015), no. 1, 93–126.

[PS09] Panos Papasoglu and Eric Swenson, *Boundaries and JSJ decompositions of CAT(0)-groups*, Geom. Funct. Anal. **19** (2009), no. 2, 559–590.

[Prz07] Piotr Przytycki, *Systolic groups acting on complexes with no flats are word-hyperbolic*, Fund. Math. **193** (2007), no. 3, 277–283.

[Prz08] ———, *The fixed point theorem for simplicial nonpositive curvature*, Math. Proc. Cambridge Philos. Soc. **144** (2008), no. 3, 683–695.

[Rua01] Kim E. Ruane, *Dynamics of the action of a* CAT(0) *group on the boundary*, Geom. Dedicata **84** (2001), no. 1-3, 81–99.

[SC83] V. P. Soltan and V. D. Chepoĭ, *Conditions for invariance of set diameters under d-convexification in a graph*, Kibernetika (Kiev) **6** (1983), 14–18 (Russian, with English summary); English transl., Cybernetics **19** (1983), no. 6, 750–756 (1984).

[Świ06] Jacek Świątkowski, *Regular path systems and (bi)automatic groups*, Geom. Dedicata **118** (2006), 23–48.

[Wis03] Daniel T. Wise, *Sixtolic complexes and their fundamental groups* (2003), unpublished manuscript.

Part II

Articles

CLASSIFYING SPACES FOR FAMILIES OF SUBGROUPS FOR SYSTOLIC GROUPS

DAMIAN OSAJDA AND TOMASZ PRYTUŁA

Abstract. We determine the large scale geometry of the minimal displacement set of a hyperbolic isometry of a systolic complex. As a consequence, we describe the centraliser of such an isometry in a systolic group. Using these results, we construct a low-dimensional classifying space for the family of virtually cyclic subgroups of a group acting properly on a systolic complex. Its dimension coincides with the topological dimension of the complex if the latter is at least four. We show that graphical small cancellation complexes are classifying spaces for proper actions and that the groups acting on them properly admit three-dimensional classifying spaces with virtually cyclic stabilisers. This is achieved by constructing a systolic complex equivariantly homotopy equivalent to a graphical small cancellation complex. On the way we develop a systematic approach to graphical small cancellation complexes. Finally, we construct low-dimensional models for the family of virtually abelian subgroups for systolic, graphical small cancellation, and some CAT(0) groups.

Contents

Date: February 26, 2017.
2010 *Mathematics Subject Classification.* 20F65, 55R35 (Primary), 20F67, 20F06 (Secondary).
Key words and phrases. Classifying space, systolic complex, small cancellation.

1. Introduction

Let G be a group and let $\mathcal{F}$ be a family of subgroups of G, that is, a collection of subgroups which is closed under taking subgroups and conjugation by elements of G. A *classifying space for the family $\mathcal{F}$ is a G–CW–complex $E_{\mathcal{F}}G$* with stabilisers in $\mathcal{F}$, such that for any subgroup $F \in \mathcal{F}$ the fixed point set $(E_{\mathcal{F}}G)^F$ is contractible. When the family $\mathcal{F}$ consists of just a trivial subgroup, the classifying space $E_{\mathcal{F}}G$ is the universal free G–space EG, and if $\mathcal{F}$ consists of all finite subgroups of G then $E_{\mathcal{F}}G$ is the so-called *classifying space for proper actions*, commonly denoted by $\underline{E}G$. Recently, much attention has been attracted by the classifying space $\underline{\underline{E}}G$ for the family of all virtually cyclic subgroups. One reason for studying $\underline{\underline{E}}G$ is its appearance in the formulation of the Farrell-Jones conjecture concerning algebraic K- and L-theory (see e.g. [Lüc05]).

It can be shown that the classifying space $E_{\mathcal{F}}G$ always exists and that it is unique up to a G–homotopy equivalence [Lüc05]. A concern is to provide specific models that are as "simple" as possible. One, widely used, measure of such simplicity is the (topological) dimension. For example, having a low-dimensional model for $\underline{\underline{E}}G$ enables one to better understand its homology that appears in the left-hand side of the assembly map in the formulation of the Farrell-Jones conjecture. Low-dimensional models for $\underline{\underline{E}}G$ were constructed for many classes of groups including hyperbolic groups [JPL06], groups acting properly on CAT(0) spaces [Lüc09], and many two-dimensional groups [Deg17]. In all of these constructions the minimal dimension of $\underline{\underline{E}}G$ is related to the minimal dimension of $\underline{E}G$. However, the discrepancy between these two may be arbitrarily large [DP14]. Finally, it is an open question whether finite-dimensional models for $\underline{\underline{E}}G$ exist for all groups admitting a finite dimensional model for $\underline{E}G$.

The main purpose of the current article is to construct low-dimensional models for $\underline{\underline{E}}G$ in the case of groups acting properly on systolic complexes.

A simply connected simplicial complex is *systolic* if it is a flag complex, and if every embedded cycle of length 4 or 5 has a diagonal. This condition may be treated as an upper curvature bound, and therefore systolic complexes are also called "complexes of simplicial non-positive curvature". They were first introduced by Chepoi [Che00] under the name *bridged complexes*. Their 1–skeleta, *bridged graphs*, were studied earlier in the frame of Metric Graph Theory [SC83, FJ87]. Januszkiewicz-Świątkowski [JŚ06] and Haglund [Hag03] rediscovered, independently, systolic complexes and initiated the exploration of groups acting on them. Since then the theory of automorphisms groups of systolic complexes has been a powerful tool for providing examples of groups with often unexpected properties, see e.g. [JŚ07, OŚ15].

An exotic large-scale geometric feature of systolic complexes is their "asymptotic asphericity" – asymptotically, they do not contain essential spheres. Such a behaviour is typical for complexes of asymptotic dimension one or two, but systolic complexes exist in arbitrarily high dimensions. The asphericity property is the crucial phenomenon used in our approach in the current article. First, we use it to determine the large-scale geometry of the minimal displacement set of a hyperbolic isometry of a systolic complex.

Recall that an isometry (i.e. a simplicial automorphism) h of a systolic complex X is called *hyperbolic* if it does not fix any simplex of X. For such an isometry one defines the *minimal displacement set* $\mathrm{Min}(h)$ to be the subcomplex of X spanned by the vertices which are moved by h the minimal combinatorial distance.

Theorem A (Theorem 3.4 and Corollary 4.7)**.** *The minimal displacement set of a hyperbolic isometry of a uniformly locally finite systolic complex is quasi-isometric to the product $T \times \mathbb{R}$ of a tree and the line.*

This theorem can be viewed as a systolic analogue of the so-called Product Decomposition Theorem for CAT(0) spaces [BH99, Theorem II.6.8]. Unlike the CAT(0) case where the splitting is isometric and it is realised within the ambient space, we provide only an abstract quasi-isometric splitting. This is mainly due to the lack of a good notion of products in the category of simplicial complexes. However, our theorem may be seen as more restrictive, since in the CAT(0) case instead of a tree one can have an arbitrary CAT(0) space.

This restriction is used to determine the structure of certain groups acting on the minimal displacement set. If a group G acts properly on a systolic complex, then one can easily see that the centraliser of a hyperbolic isometry acts properly on the minimal displacement set. If the action of G is additionally cocompact, i.e. G is a *systolic group*, Theorem A allows us to describe the structure of such centraliser. This establishes a conjecture by D. Wise [Wis03, Conjecture 11.6].

Theorem B (Corollary 5.8)**.** *The centraliser of an infinite order element in a systolic group is commensurable with $F_n \times \mathbb{Z}$, where F_n is the free group on n generators for some $n \geqslant 0$.*

Theorem B extends also some results from [JŚ07, OŚ15, Osa15] concerning normal subgroups of systolic groups. Theorem A is the key result in our approach to constructing low-dimensional models for $\underline{\underline{E}}G$ for groups acting properly on systolic complexes. We follow a "pushout method" to construct the desired complex using low-dimensional models for $\underline{E}G$. In [CO15] it is shown that if a group acts properly on a d–dimensional systolic complex X then the barycentric subdivision of X is a model for $\underline{E}G$. We then follow the strategy of W. Lück [Lüc09] used for constructing models for $\underline{\underline{E}}G$ for CAT(0) groups. The key point there is, roughly, to determine the structure of the quotient $N_G(h)/\langle h\rangle$ of the normaliser of a hyperbolic isometry h. Using similar arguments as in the proof of Theorem B, we show that this quotient is locally virtually free. This is a strong restriction on $N_G(h)/\langle h\rangle$ which leads to better dimension bounds, when compared with the CAT(0) case.

Theorem C (Theorem 5.1)**.** *Let G be a group acting properly on a uniformly locally finite d–dimensional systolic complex. Then there exists a model for $\underline{\underline{E}}G$ of dimension*

$$\dim \underline{\underline{E}}G = \begin{cases} d+1 & \text{if } d \leqslant 3, \\ d & \text{if } d \geqslant 4. \end{cases}$$

In Section 8 we provide several classes of examples to which our construction applies.

As a follow-up, we consider the family $\mathcal{VAB}$ of all virtually abelian subgroups. To the best of our knowledge there have been no known general constructions of low-dimensional classifying spaces for this family, except for cases reducing to studying the family of virtually cyclic groups (as in the case of hyperbolic groups). In the realm of systolic groups we are able to provide such constructions in the full generality.

Theorem D (Theorem 5.9)**.** *Let G be a group acting properly and cocompactly on a d–dimensional systolic complex. Then there exists a model for $E_{\mathcal{VAB}}G$ of dimension $\max\{4, d\}$.*

The most important tool used in the latter construction is the systolic Flat Torus Theorem [Els09a]. As an immediate consequence of the methods developed for proving Theorem D, we obtain the following.

Theorem E (Corollary 5.13)**.** *Let G be a group acting properly and cocompactly by isometries on a* CAT(0) *space X of topological dimension $d > 0$. Furthermore, assume that for $n > 2$ there is no isometric embedding $\mathbb{E}^n \to X$ where $\mathbb{E}^n$ is the Euclidean space. Then there exists a model for $E_{\mathcal{VAB}}G$ of dimension* $\max\{4, d+1\}$.

In particular, this result applies to lattices in rank–2 symmetric spaces thus answering a special case of a question by J.-F. Lafont [Cha08, Problem 46.7].

Classical examples of groups acting properly on systolic complexes are small cancellation groups [Wis03]. Note that small cancellation groups are not always hyperbolic and only for some of them a CAT(0) structure is provided. There is a natural construction by D. Wise of a systolic complex associated to a small cancellation complex. Therefore, Theorem C and Theorem D apply in the small cancellation setting.

In the current article we explore the more general and more powerful theory of *graphical small cancellation*, attributed to Gromov [Gro03]. Furthermore, instead of studying small cancellation presentations, we consider a slightly more general situation of graphical small cancellation complexes and their automorphism groups. Our approach is analogous to the one by McCammond-Wise [MW02] in the classical small cancellation theory. We initiate the systematic study of graphical small cancellation complexes, in particular in Section 6 we prove their basic geometric properties. The theory of groups acting properly on graphical small cancellation complexes provides a powerful tool for constructing groups with prescribed features. Examples include finitely generated groups containing expanders, and non-exact groups with the Haagerup property, both constructed in [Osa14].

Towards our main application, which is constructing low-dimensional models for classifying spaces, we define the *Wise complex* of a graphical small cancellation complex. It is the nerve of a particular cover of the graphical complex. We show that the Wise complex of a simply connected $C(p)$ graphical complex is p–systolic (Theorem 7.12), and that the two complexes are equivariantly homotopy equivalent in the presence of a group action (Theorem 7.11). The latter result is achieved by the use of an equivariant version of the Borsuk Nerve Theorem, which we prove on the way (Theorem 7.3). As a corollary we obtain the following.

Theorem F (Corollary 7.14)**.** *Let G be a group acting properly and cocompactly on a simply connected $C(p)$ graphical complex for $p \geqslant 6$. Then G acts properly and cocompactly on a p–systolic complex, i.e. G is a p–systolic group.*

The result above allows one to conclude many strong features of groups acting geometrically on graphical small cancellation complexes. Among them is biautomaticity, proved for classical small cancellation groups in [GS91b], and for systolic groups in [JŚ06].

Using the above techniques we are able to construct low-dimensional models for classifying spaces for groups acting properly on graphical small cancellation complexes.

Theorem G (Theorem 7.15)**.** *Let a group G act properly on a simply connected uniformly locally finite $C(6)$ graphical complex X. Then:*

(1) the complex X is a 2–dimensional model for $\underline{E}G$,
(2) there exists a 3–dimensional model for $\underline{\underline{E}}G$,
(3) there exists a 4–dimensional model for $E_{\mathcal{VAB}}G$, provided the action is additionally cocompact.

Organisation. The article consists of two main parts. The first part (Sections 2–5) is concerned mostly with geometry of systolic complexes. In Section 2 we give a background on systolic complexes and on classifying spaces for families of subgroups. We also recall a general method of constructing classifying spaces for the family of virtually cyclic subgroups developed in [LW12]. In Section 3 we show that the minimal displacement set of a hyperbolic isometry of a systolic complex splits up to quasi-isometry as a product of a real line and a certain simplicial graph. Section 4 is devoted to proving that this graph is quasi-isometric to a simplicial tree. The proof relies on the aforementioned asymptotic asphericity properties of systolic complexes. Finally in Section 5, using the method described in Section 2 we construct models for $\underline{\underline{E}}G$ and $E_{\mathcal{VAB}}G$ for groups acting properly on systolic complexes.

In the second part (Sections 6–7) we study graphical small cancellation theory. In Section 6 we initiate systematic studies of small cancellation complexes and show their basic geometric properties. In Section 7 we prove that the dual complex of a graphical small cancellation complex is systolic. Then we use this fact to construct models for $\underline{E}G$, $\underline{\underline{E}}G$ and $E_{\mathcal{VAB}}G$ for groups acting properly on graphical small cancellation complexes.

We conclude with Section 8 where we provide various examples of groups to which our theory applies.

Acknowledgements. We would like to thank Dieter Degrijse for many helpful discussions, in particular pointing us towards the proof of Theorem 5.9, and for a careful proofreading of the manuscript.

D.O. was supported by (Polish) Narodowe Centrum Nauki, grant no. UMO-2015/-18/M/ST1/00050. T.P. was supported by the Danish National Research Foundation through the Centre for Symmetry and Deformation (DNRF92).

2. Preliminaries

2.1. **Classifying spaces with finite or virtually cyclic stabilisers.** The main goal of this section is, given a group G, to describe a method of constructing a model for a classifying space $\underline{\underline{E}}G$ out of a model for $\underline{E}G$. The presented method is due to W. Lück and M. Weiermann [LW12]. After giving the necessary definitions we describe the steps of the construction, some of which we adjust to our purposes.

A collection of subgroups $\mathcal{F}$ of a group G is called a *family* if it is closed under taking subgroups and conjugation by elements of G. Three examples which will be of interest to us are the family $\mathcal{FIN}$ of all finite subgroups, the family $\mathcal{VCY}$ of all virtually cyclic subgroups, and the family $\mathcal{VAB}$ of all virtually abelian subgroups. Let us define the main object of our study.

Definition 2.1. Given a group G and a family of its subgroups $\mathcal{F}$, a *model for the classifying space* $E_{\mathcal{F}}G$ is a G–CW–complex X such that for any subgroup $H \subset G$ the fixed point set X^H is contractible if $H \in \mathcal{F}$, and empty otherwise.

In order to simplify the notation, throughout the article let $\underline{E}G$ denote $E_{\mathcal{FIN}}G$ and let $\underline{\underline{E}}G$ denote $E_{\mathcal{VCY}}G$. This is a standard, commonly used notation.

A model for $E_{\mathcal{F}}G$ exists for any group and any family; moreover, any two models for $E_{\mathcal{F}}G$ are G–homotopy equivalent. For the proofs of these facts see [Lüc05]. However, general constructions always produce infinite dimensional models. We will now describe the aforementioned method of constructing a finite dimensional model for $\underline{\underline{E}}G$ out of a model for $\underline{E}G$ and appropriate models associated to infinite virtually cyclic subgroups of G. Before doing so, we need one more piece of notation. If $H \subset G$ is a subgroup and $\mathcal{F}$ is a family of subgroups of G, let $\mathcal{F} \cap H$ denote the family of all subgroups of H which belong to the family $\mathcal{F}$. More generally, if $\psi \colon H \to G$ is a homomorphism, let $\psi^*\mathcal{F}$ denote the smallest family of subgroups of H that contains $\psi^{-1}(F)$ for all $F \in \mathcal{F}$.

Consider the collection $\mathcal{VCY} \setminus \mathcal{FIN}$ of infinite virtually cyclic subgroups of G. It is not a family since it does not contain the trivial subgroup. Define an equivalence relation on $\mathcal{VCY} \setminus \mathcal{FIN}$ by

$$H_1 \sim H_2 \iff |H_1 \cap H_2| = \infty.$$

Let $[H]$ denote the equivalence class of H, and let $[\mathcal{VCY} \setminus \mathcal{FIN}]$ denote the set of equivalence classes. The group G acts on $[\mathcal{VCY} \setminus \mathcal{FIN}]$ by conjugation, and for a class $[H] \in \mathcal{VCY} \setminus \mathcal{FIN}$ define the subgroup $N_G[H] \subseteq G$ to be the stabiliser of $[H]$ under this action, i.e.

$$N_G[H] = \{g \in G \mid |g^{-1}Hg \cap H| = \infty\}.$$

The subgroup $N_G[H]$ is called the *commensurator of* H, since its elements conjugate H to the subgroup commensurable with H. For $[H] \in [\mathcal{VCY} \setminus \mathcal{FIN}]$ define the family $\mathcal{G}[H]$ of subgroups of $N_G[H]$ as follows

$$\mathcal{G}[H] = \{K \subset G \mid K \in [\mathcal{VCY} \setminus \mathcal{FIN}], [K] = [H]\} \cup \{K \in \mathcal{FIN} \cap N_G[H]\}.$$

In other words $\mathcal{G}[H]$ consists of all infinite virtually cyclic subgroups of G which have infinite intersection with H and of all finite subgroups of $N_G[H]$. In this setting we have the following.

Theorem 2.2. [LW12, Theorem 2.3] *Let I be a complete set of representatives $[H]$ of G–orbits of $[\mathcal{VCY} \setminus \mathcal{FIN}]$ under the action of G by conjugation. Choose arbitrary models for $\underline{E}N_G[H]$ and $E_{\mathcal{G}[H]}N_G[H]$ and an arbitrary model for $\underline{E}G$. Let the space X be defined as a cellular G–pushout*

$$\begin{array}{ccc} \coprod_{[H]\in I} G \times_{N_G[H]} \underline{E}N_G[H] & \xrightarrow{\ i\ } & \underline{E}G \\ \Big\downarrow \coprod_{[H]\in I} \mathrm{id}_G \times_{N_G[H]} f_{[H]} & & \Big\downarrow \\ \coprod_{[H]\in I} G \times_{N_G[H]} E_{\mathcal{G}[H]}N_G[H] & \longrightarrow & X \end{array}$$

such that $f_{[H]}$ is a cellular $N_G[H]$–map for every $[H] \in I$, and i is an inclusion of G–CW–complexes. Then X is a model for $\underline{\underline{E}}G$.

Existence of such pushout follows from universal properties of appropriate classifying spaces, and the fact that if the map i fails to be injective, one can replace it with an inclusion into the mapping cylinder. For details see [LW12, Remark 2.5]. This observation leads to the following corollary.

Corollary 2.3. [LW12, Remark 2.5] *Assume that there exists an n–dimensional model for $\underline{E}G$, and for every $[H] \in [\mathcal{VCY} \setminus \mathcal{FIN}]$ there exists an n–dimensional model for $E_{\mathcal{G}[H]}N_G[H]$, and an $(n-1)$–dimensional model for $\underline{E}N_G[H]$. Then there exists an n–dimensional model for $\underline{\underline{E}}G$.*

In what follows we need our groups to be finitely generated. The commensurator $N_G[H]$ in general does not have to be finitely generated. The following proposition allows us to reduce the problem of finding various models for $E_{\mathcal{F}}N_G[H]$ to the study of its finitely generated subgroups.

Proposition 2.4. *If for every finitely generated subgroup $K \subset N_G[H]$ there exists a model for $E_{\mathcal{G}[H]\cap K}K$ with $\dim E_{\mathcal{G}[H]\cap K}K \leqslant n$, then there exists an $(n+1)$–dimensional model for $E_{\mathcal{G}[H]}N_G[H]$. The same holds for models for $\underline{E}N_G[H]$.*

Proof. The proof is a straightforward application of Theorem 4.3 in [LW12], which treats colimits of groups. The group $N_G[H]$ can be written as a colimit $N_G[H] = \mathrm{colim}_{i\in I}K_i$ of a directed system $\{K_i\}_{i\in I}$ of all of its finitely generated subgroups (since $N_G[H]$ is countable, this system is countable as well). Since the structure maps are injective and since every subgroup $F \in \mathcal{G}[H]$ is finitely generated, it is contained in the image of some $\psi_i \colon K_i \hookrightarrow G$. Again by the injectivity of ψ_i, we have $\psi_i^*\mathcal{G}[H] = \mathcal{G}[H] \cap K_i$. The claim follows from Theorem 4.3 in [LW12]. □

The following condition will allow us to find infinite cyclic subgroups which are normal in K.

Definition 2.5. [Lüc09, Condition 4.1] A group G satisfies condition (C) if for every $g, h \in G$ with $|h| = \infty$ and any $k, l \in \mathbb{Z}$ we have

$$gh^kg^{-1} = h^l \implies |k| = |l|.$$

Lemma 2.6. *Let $K \subset N_G[H]$ be a finitely generated subgroup that contains some representative of $[H]$ and assume that the group G satisfies condition* (C)*. Choose an element $h \in H$ such that $[\langle h\rangle] = [H]$ (any element of infinite order has this property). Then there exists $k \geqslant 1$, such that $\langle h^k\rangle$ is normal in K.*

Proof. Let $s_1, \ldots, s_m$ be generators of K. Since $K \subset N_G[H]$, for any s_i we have $s_i^{-1}h^{k_i}s_i = h^{l_i}$ for some $k_i, l_i \in \mathbb{Z} \setminus \{0\}$. Then the condition (C) implies that $l_i = \pm k_i$ for all $i \in \{1, \ldots, m\}$. Thus k defined as $\prod_{i=1}^n k_i$ has the desired property. □

In order to treat short exact sequences of groups we need the following result.

Proposition 2.7. [DP14, Corollary 2.3] *Consider the short exact sequence of groups*

$$0 \longrightarrow N \longrightarrow G \xrightarrow{\pi} Q \longrightarrow 0.$$

Let $\mathcal{F}$ be a family of subgroups of G and $\mathcal{H}$ be a family of subgroups of Q such that $\pi(\mathcal{F}) \subseteq \mathcal{H}$. Suppose that there exists a integer $k \geqslant 0$, such that for every subgroup $H \in \mathcal{H}$ there exists a k–dimensional model for $E_{\mathcal{F}\cap\pi^{-1}(H)}\pi^{-1}(H)$. Then given a model for $E_{\mathcal{H}}Q$, there exists a model for $E_{\mathcal{F}}G$ of dimension $k + \dim E_{\mathcal{H}}Q$.

2.2. **Systolic complexes.** In this section we establish the notation and define systolic complexes and groups. We do not discuss general properties of systolic complexes, the interested reader is referred to [JŚ06]. We also give basic definitions regarding metric on simplicial complexes, including notation which is slightly different from the one usually used.

Let X be a simplicial complex. We assume that X is finite dimensional and uniformly locally finite. For a subset of vertices $S \subseteq X^{(0)}$ define the subcomplex *spanned by* S to be the maximal subcomplex of X having S as its vertex set. We denote this subcomplex by $\operatorname{span} S$. We say that X is *flag* if every set of pairwise adjacent vertices spans a simplex of X. For a simplex $\sigma \in X$, define the *link* of σ as the subcomplex of X that consists of all simplices of X which do not intersect σ, but together with σ span a simplex of X. A *cycle* in X is a subcomplex $\gamma \subset X$ homeomorphic to the 1–sphere. The length $|\gamma|$ of the cycle γ is the number of its edges. A *diagonal* of a cycle is an edge connecting two of its nonconsecutive vertices.

Definition 2.8. [JŚ06, Definition 1.1] Let $k \geqslant 6$ be a positive integer. A simplicial complex X is k*–large* if it is flag and every cycle γ of length $4 \leqslant |\gamma| < k$ has a diagonal.

We say that X is k*–systolic* if it is connected, simply connected and the link of every simplex of X is k–large.

One can show that k–systolic complexes are in fact both k–large and flag. In the case when $k = 6$, which is the most interesting to us, we abbreviate 6*–systolic* to *systolic*. Note that if $m > k$ then m–systolicity implies k–systolicity.

Now we introduce the convention used throughout this article regarding the metric on simplicial complexes. Some of our definitions are slightly different from the usual ones, however they seem to be more convenient here in order to avoid technical complications.

Convention 2.9. (Metric on simplicial complexes). Let X be a simplicial complex. Unless otherwise stated, when we refer to the *metric* on X, we mean its 0–skeleton $X^{(0)}$, where the distance between two vertices is the minimal number of edges of an edge-path joining these two vertices. Notice that for flag complexes, the 0–skeleton together with the above metric entirely determines the complex. By an *isometry* we mean a simplicial map $f \colon Y \to X$, which restricted to 0–skeleta is an isometry with respect to the metric defined above. In particular, any simplicial isomorphism is an isometry. A *geodesic* in a simplicial complex is defined as a sequence of vertices $(v_i)_{i \in I}$ such that $d(v_i, v_j) = |i - j|$ for all $i, j \in I$, where $I \subseteq \mathbb{N}$. Note that we allow $I = \mathbb{N}$, i.e. a geodesic can be infinite in both directions. A *graph* is a 1–dimensional simplicial complex. In particular, graphs do not contain loops and multiple edges.

Remark 2.10. For a graph the usually considered metric is the geodesic metric where every edge is assigned length 1. If X is a simplicial complex, then the restriction of the geodesic metric on the graph $X^{(1)}$ to $X^{(0)}$ is precisely the metric we defined above.

Let v_0 be a vertex of X. Define the *combinatorial ball* of radius $r \in \mathbb{N}$, centred at v_0 as a subcomplex

$$B_n(v_0, X) = \operatorname{span}\{v \in X \mid d(v_0, v) \leqslant n\},$$

and a *combinatorial sphere* as

$$S_n(v_0, X) = \operatorname{span}\{v \in X \mid d(v_0, v) = n\}.$$

We finish this section with basic definitions regarding group actions on simplicial complexes. Unless stated otherwise, all groups are assumed to be discrete and all actions are assumed to be *simplicial*, i.e. groups act by simplicial automorphisms. We say that the action of G on a simplicial complex X is *proper* if for every vertex $v \in X$ and every integer $n \geqslant 0$ the set

$$\{g \in G \mid gB_n(v, X) \cap B_n(v, X) \neq \emptyset\}$$

is finite. When X is uniformly locally finite this definition is equivalent to all vertex stabilisers being finite. We say that the action is *cocompact* if there is a compact subset $K \subset X$ that intersects every orbit of the action.

A group is called *systolic* if it acts properly and cocompactly on a systolic complex. However, most of the time we are concerned with proper actions that are not necessarily cocompact.

2.3. **Quasi-isometry.** Let (X, d_X) and (Y, d_Y) be metric spaces. A (not necessarily continuous) map $f\colon X \to Y$ is a *coarse embedding* if there exist real non-decreasing functions ρ_1, ρ_2 with $\lim_{t\to+\infty}\rho_1(t) = +\infty$, such that the inequality

$$\rho_1(d_X(x_1, x_2)) \leqslant d_Y(f(x_1), f(x_2)) \leqslant \rho_2(d_X(x_1, x_2))$$

holds for all $x_1, x_2 \in X$. If both functions ρ_1, ρ_2 are affine, we call f a *quasi-isometric embedding*. Given two maps $f, g\colon X \to Y$ we say that f and g are *close* if there exists $R \geqslant 0$ such that the inequality $d_Y(f(x), g(x)) \leqslant R$ holds for all $x \in X$. We say that the coarse embedding $f\colon X \to Y$ is a *coarse equivalence*, if there exists a coarse embedding $g\colon Y \to X$ such that the composite $g \circ f$ is close to the identity map on X and $g \circ f$ is close to the identity map on Y. Analogously, a quasi-isometric embedding $f\colon X \to Y$ is called a *quasi-isometry*, if there exists a quasi-isometric embedding $g\colon Y \to X$ such that the appropriate composites are close to identity maps.

The following criterion will be very useful to us: a coarse embedding (quasi-isometric embedding) $f\colon X \to Y$ is a coarse equivalence (quasi-isometry) if and only if it is *quasi-onto*, i.e. there exists $R \geqslant 0$ such that for any $y \in Y$ there exists $x \in X$ with $d_Y(y, f(x)) \leqslant R$.

3. Quasi-product structure of the minimal displacement set

In this section we describe the structure of the minimal displacement set associated to a hyperbolic isometry of X. We prove that this subcomplex of X is quasi-isometric to the product of the so-called *graph of axes* and the real line. This may be viewed as a coarse version of the Product Decomposition Theorem for CAT(0) spaces (see [BH99]). Our arguments rely on the work of T. Elsner in [Els09a] and [Els09b].

Let h be an isometry of a simplicial complex X. Define the *displacement function* $d_h \colon X^{(0)} \to \mathbb{N}$ by the formula $d_h(x) = d_X(x, h(x))$. The *translation length* $L(h)$ is defined as

$$L(h) = \min_{x \in X^{(0)}} d_h(x).$$

If h does not fix any simplex of X, then h is called *hyperbolic*. In such case one has $L(h) > 0$. For a hyperbolic isometry h, define the *minimal displacement set* $\mathrm{Min}(h)$ as the subcomplex of X spanned by the set of vertices where d_h attains its minimum. Clearly $\mathrm{Min}(h)$ is invariant under the action of h. If X is a systolic complex, we have the following.

Lemma 3.1. [Els09b, Propositions 3.3 and 3.4] *Let h be a hyperbolic isometry of a systolic complex X. Then the subcomplex $\mathrm{Min}(h)$ is a systolic complex, isometrically embedded into X.*

An h–invariant geodesic in X is called an *axis* of h. We say that $\mathrm{Min}(h)$ is the *union of axes*, if for every vertex $x \in \mathrm{Min}(h)$ there is an h–invariant geodesic passing through x, i.e. $\mathrm{Min}(h)$ can be written as follows

$$\mathrm{Min}(h) = \mathrm{span}\{\bigcup \gamma \mid \gamma \text{ is an } h\text{–invariant geodesic}\}. \tag{3.1}$$

In this case, the isometry h acts on $\mathrm{Min}(h)$ as a translation along the axes by the number $L(h)$.

Proposition 3.2. *Let h be a hyperbolic isometry of a systolic complex X. Then the following hold:*

(i) [Els09b, Proposition 3.1] *For any $n \in \mathbb{Z} \setminus \{0\}$, the isometry h^n is hyperbolic.*
(ii) [Els09b, Theorem 3.5] *There exists an $n \geqslant 1$ such that there is an h^n–invariant geodesic in X.*
(iii) [Els09b, Remark, p. 48] *If there exists an h–invariant geodesic then for any vertex $x \in \mathrm{Min}(h)$ there is an h–invariant geodesic passing through x, i.e. the isometry h satisfies* (3.1)*.*
(iv) If h satisfies (3.1) *then so does h^n for any $n \in \mathbb{Z} \setminus \{0\}$.*

Proof. Observe that (iv) follows from (iii) and the fact that an h–invariant geodesic is h^n–invariant. □

For two subcomplexes $X_1, X_2 \subset X$, the distance $d_{min}(X_1, X_2)$ is defined to be

$$d_{min}(X_1, X_2) = \min\{d_X(x_1, x_2) \mid x_1 \in X_1, x_2 \in X_2\}.$$

Note that in general d_{min} is not even a pseudometric. We are ready now to define the graph of axes.

Definition 3.3. (Graph of axes). For a hyperbolic isometry h satisfying (3.1), define the simplicial graph $Y(h)$ as follows

$$Y(h)^{(0)} = \{\gamma \mid \gamma \text{ is an } h\text{–invariant geodesic in } \mathrm{Min}(h)\},$$
$$Y(h)^{(1)} = \{\{\gamma_1, \gamma_2\} \mid d_{min}(\gamma_1, \gamma_2) \leqslant 1\}.$$

Let $d_{Y(h)}$ denote the associated metric on $Y(h)^{(0)}$ (see Convention 2.9).

The main goal of this section is to prove the following theorem.

Theorem 3.4. *Let h be a hyperbolic isometry of a uniformly locally finite systolic complex X, such that the translation length $L(h) > 3$, and the subcomplex* Min(h) *is the union of axes. Then there is a quasi-isometry*

$$c\colon (Y(h) \times \mathbb{Z}, d_h) \to (\mathrm{Min}(h), d_X), \tag{3.2}$$

where the metric d_h is defined as

$$d_h((\gamma_1, t_1), (\gamma_2, t_2)) = d_{Y(h)}(\gamma_1, \gamma_2) + |t_1 - t_2|,$$

and d_X is the metric induced from X.

In the remainder of this section let h be a hyperbolic isometry of X such that Min(h) is the union of axes. By Lemma 3.1 it is enough to prove Theorem 3.4 in case where Min$(h) = X$. In order to define the map c we parameterise geodesics in $Y(h)$, i.e. to each $\gamma \in Y(h)$ we assign an origin $\gamma(0)$ and a direction. After this is done, the geodesic γ can be viewed as an isometry $\gamma\colon \mathbb{Z} \to X$, and the map c can be defined as $c(\gamma, t) = \gamma(0 + t)$. Before we describe the procedure of parameterising γ, we need to establish the following metric estimate between $d_{Y(h)}$ and d_X.

Lemma 3.5. *Let γ_1 and γ_2 be h–invariant geodesics. Then:*

(i) For any vertices $x_1 \in \gamma_1$ and $x_2 \in \gamma_2$ we have $d_{Y(h)}(\gamma_1, \gamma_2) \leqslant d_X(x_1, x_2) + 1$.
(ii) For any vertex $x_1 \in \gamma_1$ there exists a vertex $x_2 \in \gamma_2$ with $d_X(x_1, x_2) \leqslant (L(h) + 1) d_{Y(h)}(\gamma_1, \gamma_2)$.

Proof. (i) We proceed by induction on $d_X(x_1, x_2)$. If $d_X(x_1, x_2) = 0$ then γ_1 and γ_2 intersect, hence $d_{Y(h)}(\gamma_1, \gamma_2) \leqslant 1$. Assume the claim is true for $d_X(x_i, x_j) \leqslant n$, and let $d_X(x_1, x_2) = n+1$. Let x_n be the vertex on a geodesic in X between x_1 and x_2, with $d_X(x_1, x_n) = n$. Choose a geodesic $\gamma_n \in Y(h)$ passing through x_n (such a geodesic exists since h satisfies (3.1)). By inductive hypothesis we have $d_{Y(h)}(\gamma_1, \gamma_n) \leqslant n+1$ and clearly $d_{Y(h)}(\gamma_n, \gamma_2) \leqslant 1$, hence the claim follows from the triangle inequality.

(ii) It suffices to prove the claim in the case where $d_{Y(h)}(\gamma_1, \gamma_2) = 1$. Let $x_1 \in \gamma_1$ be any vertex. The vertex x_2 is chosen as follows. Let $x_1' \in \gamma_1$ be the vertex realising the distance between γ_1 and γ_2 (i.e. it is either the vertex of intersection, or the vertex on the edge joining these two geodesics). Choose x_1' to be the closest vertex to x_1 with this property. Due to h–invariance of γ_1 and γ_2, the distance $d_X(x_1, x_1')$ is not greater than $L(h)$ (even $\frac{L(h)}{2}$ in fact). If γ_1 and γ_2 intersect, define x_2 to be x_1'. If not then γ_1 and γ_2 are connected by an edge, one of whose endpoints is x_1'. Define the vertex $x_2 \in \gamma_2$ to be the other endpoint of that edge. □

For an h–invariant geodesic $\gamma \subset X$ define the linear order $\prec$ on the set of vertices of γ, by setting $x \prec h(x)$ for some (and hence all) $x \in \gamma$. Fix a geodesic γ_0 and identify the set of its vertices with $\mathbb{Z}$, such that the order $\prec$ agrees with the natural order on $\mathbb{Z}$.

Consider the family of combinatorial balls $\{B_n(-n, X)\}_{n\in\mathbb{N}}$, where $-n \in \gamma_0$. Notice that we have $B_n(-n, X) \subseteq B_{n+1}(-(n+1), X)$, i.e. the family $\{B_n(-n, X)\}_{n\in\mathbb{N}}$ is ascending. The following lemma is crucial in order to define the origin $\gamma(0)$ of γ.

Lemma 3.6. *Let γ be an h–invariant geodesic. Then there exists a vertex $v \in \gamma$ such that for any vertex w contained in the intersection*

$$(\bigcup_{n\geqslant 0} B_n(-n, X)) \cap \gamma$$

we have $w \prec v$.

Proof. Observe first that for any $n \geqslant 0$ we have

$$\sup_{\prec}\{B_n(-n, X) \cap \gamma_0\} = 0 \in \gamma_0,$$

therefore taking 0 as v proves the lemma for $\gamma = \gamma_0$. For an arbitrary γ, we proceed by contradiction. Assume conversely, that the supremum $\sup_{\prec}\{B_n(-n, X) \cap \gamma\}$ is not attained at any vertex of γ. Therefore every vertex of γ belongs to $B_n(-n, X)$ for some $n > 0$.

Let $x \in \gamma$ be a vertex which is at distance at most $K = (L(h)+1)d_{Y(h)}(\gamma_0, \gamma)$ from $0 \in \gamma_0$. Lemma 3.5.(ii) assures that such a vertex exists. For any $m > 0$ consider vertices $h^m(x)$ and $h^m(0)$. By our assumption there exists $n > 0$ such that $h^m(x)$ is contained in $B_n(-n, X)$ (see Figure 1). Therefore by the triangle inequality we get

$$d_X(-n, h^m(0)) \leqslant n + K.$$

On the other hand we have

$$d_X(-n, h^m(0)) = n + L(h^m),$$

since γ_0 is a geodesic. For $L(h^m) > K$ this gives a contradiction. □

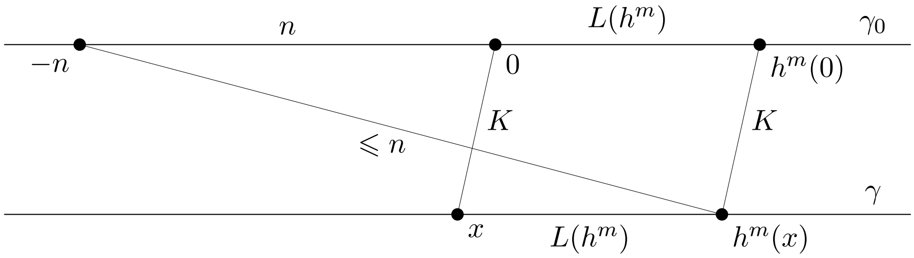

FIGURE 1. The vertex $h^m(x)$ cannot belong to $B_n(-n, X)$ and be arbitrarily far from x at the same time.

Definition 3.7. Let $\gamma \subset X$ be an h–invariant geodesic. Define the vertex $\gamma(0)$ as

$$\gamma(0) = \max_{\prec}\Big(\bigcup_{n\geqslant 0} B_n(-n, X) \cap \gamma\Big). \tag{3.3}$$

The set of vertices of $(\bigcup_{n\geqslant 0} B_n(-n, X) \cap \gamma)$ is bounded from above by Lemma 3.6 hence the maximal vertex exists. Observe that for γ_0 we have $\gamma_0(0) = 0$.

Having geodesics parameterised, we need the following two technical lemmas that describe certain metric inequalities, which are needed to prove Theorem 3.4.

Lemma 3.8. *Let γ_0, γ be as in Definition 3.7 and assume that $L(h) > 3$ and $d_{Y(h)}(\gamma_0, \gamma) \geqslant 2L(h) + 4$. Then there exists an $n_0 \in \mathbb{N}$, such that for all $n \geqslant n_0$ and for all $t \geqslant 0$ we have*

$$n + t - 1 \leqslant d_X(-n, \gamma(t)) \leqslant n + t + 1. \tag{3.4}$$

Proof. The idea is to reduce the problem to the study of $\mathbb{E}^2_\Delta$, the equilaterally triangulated Euclidean plane. To do so, we construct an h–equivariant simplicial map $f\colon P \to X$, which satisfies the following properties:

(1) the complex P is an h–invariant triangulation of a strip $\mathbb{R} \times I$, and P can be isometrically embedded into $\mathbb{E}^2_\Delta$,

(2) the boundary ∂P is mapped by f onto the disjoint union $\gamma_0 \sqcup \gamma$ such that the restriction of f to each boundary component is an isometry,
(3) for every pair of vertices $u, v \in P$ there is an inequality

$$d_P(u,v) - 1 \leqslant d_X(f(u), f(v)) \leqslant d_P(u,v).$$

The construction of such a map is given in the proofs of Theorems 2.6 and 3.5 (case 1) in [Els09b]. It requires that $L(h) > 3$ and $d_{Y(h)}(\gamma_0, \gamma) \geqslant 2L(h) + 4$, and this is the only place where we need these assumptions.

Fix an embedding of P into $\mathbb{E}^2_\Delta$. Since the restriction of f to each boundary component is an isometry, let us keep the same notation for the preimages under f of γ_0 and γ. For a vertex $v \in \partial P$ let $\angle(v)$ denote the number of triangles of P which contain v. We have the following two cases to consider:

(i) for every vertex v of γ we have $\angle(v) = 3$,
(ii) there exists a vertex v of γ with $\angle(v) \neq 3$.

We treat the case (ii) first. Steps of the proof are indicated in Figure 2. Let v_0 be a vertex of γ such that $\angle(v_0) = 2$ and $v_0 \prec \gamma(0)$. Denote by v_{-1} the vertex of γ that is adjacent to v_0 and $v_{-1} \prec v_0$. Introduce a coordinate system on $\mathbb{E}^2_\Delta$ by setting v_{-1} to be the base point and letting $e_1 = \overrightarrow{v_{-1}v_0}$ and $e_2 = \overrightarrow{v_{-1}w}$, where w is the unique vertex which lies inside P and is adjacent to both v_{-1} and v_0. We will write $v = (x_v, y_v)$ for the coordinates of a vertex v in basis e_1, e_2. It follows from the choice of the coordinate system and the fact that the strip P is h–invariant, that for all $k \in \mathbb{Z}$ we have

$$\gamma(k+1) = \gamma(k) + e_1 \text{ or } \gamma(k+1) = \gamma(k) + e_2, \tag{3.5}$$

and both possibilities occur an infinite number of times. In particular, the second coordinate $y_{\gamma(k)}$ of $\gamma(k)$ is a non-decreasing function of k such that

$$y_{\gamma(k)} \to -\infty \text{ as } k \to -\infty. \tag{3.6}$$

By property (3) of the map f, the distance $d_{\mathbb{E}^2_\Delta}(\gamma_0, \gamma)$ is bounded from above by $d_X(\gamma_0, \gamma) + 1$. This implies that the geodesic γ_0 also satisfies (3.5) and (3.6) (i.e. it runs parallel to γ, see Figure 2). Hence, there exists k_0, such that for $k \geqslant k_0$ the coordinate $y_{\gamma_0(-k)}$ is strictly less than y_{v_0}, which is in turn less than $y_{\gamma(0)}$, since $v_0 \prec \gamma(0)$. Therefore for each $r \geqslant d_P(\gamma_0(-k), v_0)$ the combinatorial sphere $S(\gamma_0(-k), r)$ intersects $\gamma|_{[0,\infty)}$ exactly once, where $\gamma|_{[0,\infty)}$ denotes the geodesic ray obtained by restricting the domain of γ to non-negative integers. In particular, if $S(\gamma_0(-k), r) \cap \gamma|_{[0,\infty)} = \gamma(s)$ for some $s \geqslant 0$, then $S(\gamma_0(-k), r+1) \cap \gamma|_{[0,\infty)} = \gamma(s+1)$.

Take $\widetilde{n_0}$ such that for any $n \geqslant \widetilde{n_0}$ we have $\gamma(0) = \max_{\prec}(B(-n,n) \cap \gamma)^{(0)}$. By property (3) of the map f, for such n we have $d_{\mathbb{E}^2_\Delta}(-n, \gamma(0)) \in \{n, n+1\}$. Set $n_0 = \max\{\widetilde{n_0}, k_0\}$.

For $n \geqslant n_0$ and for $t \geqslant 0$ we have $d_{\mathbb{E}^2_\Delta}(-n,\gamma(t)) = d_{\mathbb{E}^2_\Delta}(-n,\gamma(0))+t \in \{n+t, n+1+t\}$. Therefore the claim follows from property (3).

Case (i) is proven analogously. Using the notation of case (ii) we introduce the coordinate system as follows. Put $w = \gamma(0)$, $v_{-1} = \gamma(-1)$ and let v_0 be the unique vertex which lies outside of P and is adjacent to both v_{-1} and w. The rest of the proof is the same as in case (ii). □

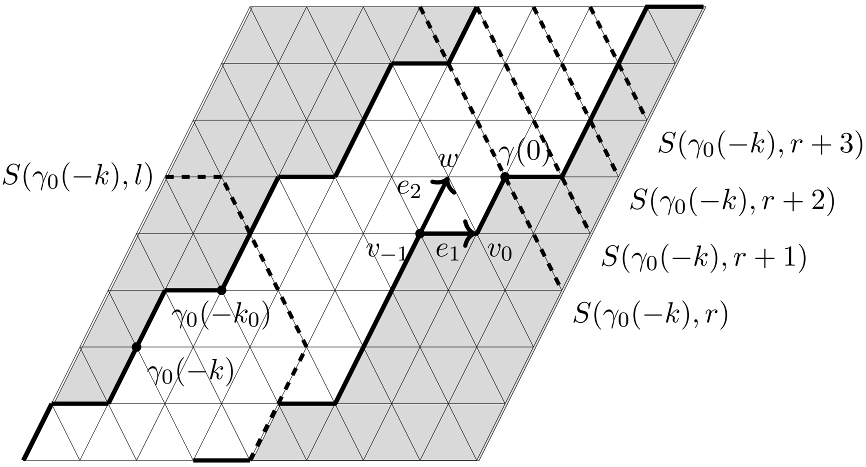

FIGURE 2. Geodesics and combinatorial spheres in the coordinate system.

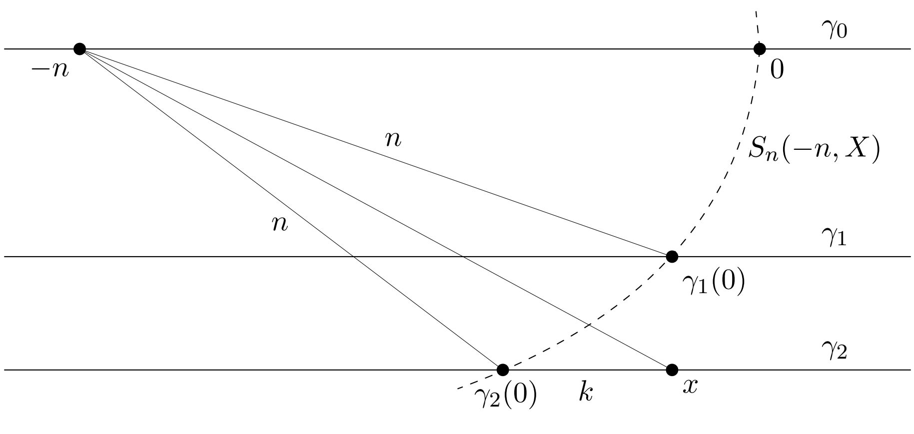

FIGURE 3. Geodesics and their origins.

We need one more metric estimate.

Lemma 3.9. *Let γ_1, γ_2 be h–invariant geodesics in X. Assume additionally, that the translation length $L(h) > 3$. Then we have the following inequality*

$$d_X(\gamma_1(0), \gamma_2(0)) \leqslant 2(L(h)+1)d_{Y(h)}(\gamma_2, \gamma_2) + 2(2L(h)+4)(L(h)+1).$$

Proof. Let L denote the translation length $L(h)$ and let D denote the distance $d_{Y(h)}(\gamma_1, \gamma_2)$. We claim that there exists a vertex $x \in \gamma_2$ such that

$$d_X(x, \gamma_1(0)) \leqslant (L+1)D + L$$

and $\gamma_2(0) \prec x$. By Lemma 3.5.(ii) there is a vertex $x \in \gamma_2$ which is at distance at most $(L+1)D$ from $\gamma_1(0)$. If $\gamma_2(0) \prec x$ does not hold, we do the following. Let α be a geodesic joining $\gamma_1(0)$ and x. Apply the isometry h^m to α, where m is the smallest integer such that $h^m(x) \succ \gamma_2(0)$. Then the concatenation of a geodesic segment joining $\gamma_2(0)$ and $h^m(x)$ with geodesic $h^m(\alpha)$ is a path joining $\gamma_2(0)$ and $h^m(\gamma_1(0))$ of length at most $(L+1)D + L$. Therefore we can switch roles of γ_1 and γ_2 and set x to be $h^m(\gamma_1(0))$. This proves the claim.

Let k denote the distance $d_X(\gamma_2(0), x)$. We will show that $k \leqslant d_X(x, \gamma_1(0)) + 1$ (see Figure 3). In order to apply Lemma 3.8, assume that $d_Y(\gamma_0, \gamma_2) \geqslant 2L + 4$, and choose n large enough, such that $\gamma_i(0) = \max_{\prec}(B_n(-n, X) \cap \gamma_i)^{(0)}$ for $i \in \{1, 2\}$ and $n \geqslant n_0$, where n_0 is the constant appearing in the formulation of Lemma 3.8 (clearly the same holds for any $n_1 \geqslant n$). By Lemma 3.8 we have

$$d_X(-n, x) \geqslant n + k - 1.$$

By the triangle inequality applied to the vertices $-n, \gamma_1(0)$ and x we get

$$d_X(-n, x) \leqslant n + d_X(\gamma_1(0), x).$$

Combining the two above inequalities gives us

$$k \leqslant d_X(\gamma_1(0), x) + 1.$$

By the triangle inequality the distance $d_X(\gamma_1(0), g_2(0))$ is at most $d_X(\gamma_1(0), x) + k$ hence we have

$$d_X(\gamma_1(0), \gamma_2(0)) \leqslant 2d_X(\gamma_1(0), x) + 1 \leqslant 2(L+1)D + 2L + 1.$$

This proves the lemma under the assumption that $d_Y(\gamma_0, \gamma_2) \geqslant 2L + 4$.

If for both γ_1 and γ_2 we have $d_Y(\gamma_0, \gamma_i) \leqslant 2L + 4$, then by directly comparing γ_i with γ_0, one can show that

$$d_X(\gamma_0(0), \gamma_i(0)) \leqslant 2d_{Y(h)}(\gamma_0, \gamma_i)(L+1)$$

for $i \in \{1, 2\}$. Using the triangle inequality one gets

$$d_X(\gamma_1(0), \gamma_2(0)) \leqslant 2d_{Y(h)}(\gamma_0, \gamma_1)(L+1) + 2d_{Y(h)}(\gamma_0, \gamma_2)(L+1) \leqslant 2(2L+4)(L+1). \quad \square$$

We are now ready to prove Theorem 3.4.

Proof of Theorem 3.4. Define the map $c\colon (Y(h)\times\mathbb{Z}, d_h)\to(\mathrm{Min}(h), d_X)$ as

$$c(\gamma,t)=\gamma(0+t),$$

where $\gamma(0+t)$ is the unique vertex of γ satisfying $\gamma(0)\prec\gamma(0+t)$ and $d_X(\gamma(0),\gamma(0+t))=t$. We will show that for every two points (γ_1,t_1) and (γ_2,t_2) of $Y(h)\times\mathbb{Z}$ we have the following inequality

$$\rho_1(d_h((\gamma_1,t_1),(\gamma_2,t_2)))\leqslant d_X(c(\gamma_1,t_1),c(\gamma_2,t_2))\leqslant\rho_2(d_h((\gamma_1,t_1),(\gamma_2,t_2))),$$

where ρ_1 and ρ_2 are non-decreasing linear functions.

We first find the function ρ_2. Without loss of generality we can assume that $|t_1|\leqslant|t_2|$ and let α be a geodesic joining $\gamma_1(t_1)$ and $\gamma_2(t_2)$. Denote $L(h)$ by L and apply the isometry h^m to α, where m is chosen such that $d_X(\gamma_1(0),h^m(\gamma_1(t_1)))\leqslant L$, and m has the smallest absolute value among such numbers. We then have $d_X(h^m(\gamma_2(t_2)),\gamma_2(0))\leqslant|t_1-t_2|+L$. Hence, by the triangle inequality we get

$$d_X(\gamma_1(t_1),\gamma_2(t_2))=d_X(h^m(\gamma_1(t_1)),h^m(\gamma_2(t_2)))\leqslant|t_1-t_2|+L+L+d_X(\gamma_1(0),\gamma_2(0)).$$

By Lemma 3.9 we obtain

$$\begin{aligned}d_X(\gamma_1(t_1),\gamma_2(t_2))&\leqslant|t_1-t_2|+2L+2(L+1)d_{Y(h)}(\gamma_2,\gamma_2)+2(2L+4)(L+1)\\&\leqslant 2(L+1)(d_{Y(h)}(\gamma_2,\gamma_2)+|t_1-t_2|)+2(2L+4)(L+1)+2L\\&=2(L+1)d_h((\gamma_1,t_1),(\gamma_2,t_2))+2(2L+4)(L+1)+2L.\end{aligned}$$

We are left with finding ρ_1. Let K denote the distance $d_{Y(h)}(\gamma_1,\gamma_2)+|t_1-t_2|$. If $d_{Y(h)}(\gamma_1,\gamma_2)\geqslant\frac{1}{10(2(L+1))}K$, then by Lemma 3.5.(i) we have

$$d_X(\gamma_1(t_1),\gamma_2(t_2))\geqslant d_{Y(h)}(\gamma_2,\gamma_2)-1\geqslant\frac{1}{10(2(L+1))}K-1. \tag{3.7}$$

If $d_{Y(h)}(\gamma_2,\gamma_2)<\frac{1}{10(2(L+1))}K$ then one has $|t_1-t_2|\geqslant\frac{1}{2}K$. In this case, using the same argument with translation by h^m and the reverse triangle inequality, we obtain

$$d_X(\gamma_1(t_1),\gamma_2(t_2))+L\geqslant|t_1-t_2|-L-d_X(\gamma_1(0),\gamma_2(0)). \tag{3.8}$$

By Lemma 3.9 we have

$$d_X(\gamma_1(0),\gamma_2(0))\leqslant 2(L+1)d_{Y(h)}(\gamma_2,\gamma_2)+2(2L+4)(L+1),$$

which combined with our assumption gives

$$d_X(\gamma_1(0),\gamma_2(0))\leqslant 2(L+1)\frac{1}{10(2(L+1))}K+2(2L+4)(L+1)=\frac{1}{10}K+2(2L+4)(L+1).$$

Putting this in the inequality (3.8) yields

$$d_X(\gamma_1(t_1),\gamma_2(t_2))+L\geqslant\frac{1}{2}K-L-\frac{1}{10}K-2(2L+4)(L+1),$$

hence finally

$$d_X(\gamma_1(t_1), \gamma_2(t_2)) \geqslant \frac{9}{20}K - 2L - 2(2L+4)(L+1). \tag{3.9}$$

In both (3.7) and (3.9), the following inequality holds

$$d_X(\gamma_1(t_1), \gamma_2(t_2)) \geqslant \frac{1}{20(L+1)}K - 2L - 2(2L+4)(L+1).$$

This proves that c is a quasi-isometric embedding. We note that by Proposition 3.2.(iii) the map c is surjective on the vertex sets, hence a quasi-isometry. □

4. Filling radius for spherical cycles

The purpose of this section is to prove that the graph of axes defined in Section 3 is quasi-isometric to a simplicial tree. Our main tool is the S^kFRC property of systolic complexes introduced in [JŚ07]. It is a coarse (i.e. quasi-isometry invariant) property that, intuitively, describes "asymptotic thinness of spheres" in a given metric space. We use numerous features of S^kFRC spaces established in [JŚ07], some of which we adjust to our setting. The crucial observation is Proposition 4.6, which says that an S^1FRC space satisfying certain homological condition is quasi-isometric to a tree. This extends a result of [JŚ07], which treats only the case of finitely presented groups.

Let (X, d) be a metric space. Given $r > 0$, the *Rips complex* $P_r(X)$ is a simplicial complex defined as follows. The vertex set of $P_r(X)$ is the set of all points in X. The subset $\{x_1, \ldots, x_n\}$ spans a simplex of $P_r(X)$ if and only if $d(x_i, x_j) \leqslant r$ for all $i, j \in \{1, \ldots, n\}$. Notice that if $R \geqslant r$ then $P_r(X)$ is naturally a subcomplex of $P_R(X)$.

In what follows we consider simplicial chains with arbitrary coefficients. For detailed definitions see [JŚ07].

A *k–spherical cycle* in a simplicial complex X is a simplicial map $f\colon S^k \to X$ from an oriented simplicial k–sphere to X. Let C_f denote the image through f of the fundamental (simplicial) k–cycle in S^k. A *filling* of a k–spherical cycle f is a simplicial $(k+1)$–chain D such that $\partial D = C_f$. Let $\operatorname{supp}(f)$ denote the image through f of the vertex set of S^k, and let $\operatorname{supp}(D)$ denote the set of vertices of all underlying simplices of D.

Definition 4.1. A metric space (X, d) *has filling radius for spherical cycles constant* (or (X, d) is S^kFRC) if for every $r > 0$ there exists $R \geqslant r$ such that any k–spherical cycle f which is null-homologous in $P_r(X)$ has a filling D in $P_R(X)$ satisfying $\operatorname{supp}(D) \subset \operatorname{supp}(f)$.

Proposition 4.2. [JŚ07, p. 16] *Let (X, d_X) be S^kFRC and let $f\colon (Y, d_Y) \to (X, d_X)$ be a coarse embedding. Then (Y, d_Y) is S^kFRC.*

Lemma 4.3. [JŚ07, Theorem 4.1, Lemma 5.3] *Let X be a systolic complex. Then X is S^kFRC for any $k \geqslant 2$.*

The following lemma describes the behaviour of property S^kFRC with respect to products. It was originally proved in [JŚ07] only for products of finitely generated groups. However, it is straightforward to check that the lemma holds for arbitrary geodesic metric spaces. The metric on a product is chosen to be the sum of metrics on the factors.

Lemma 4.4. [JŚ07, Proposition 7.2] *Let $k \in \{0, 1\}$. Assume that (X, d_X) is not S^kFRC and that there is $r > 0$ such that every k–spherical cycle $f\colon S^k \to X$ is null-homotopic in $P_r(X)$. If (Y, d_Y) is unbounded then the product $(X, d_X) \times (Y, d_Y)$ is not S^{k+1}FRC.*

The following criterion is the key tool that we use in the proof of Proposition 4.6.

Lemma 4.5. [Man05, Theorem 4.6] *Let (X, d) be a geodesic metric space. Then the following are equivalent:*

(1) X is quasi-isometric to a simplicial tree,
(2) (bottleneck property) there exists $\delta > 0$, such that for any two points $x, y \in X$ there is a midpoint $m = m(x, y)$ with $d(x, m) = d(m, y) = \frac{1}{2} d(x, y)$, and such that any path from x to y in X contains a point within distance at most δ from m.

Proposition 4.6. *Let X be a graph which is S^1FRC and assume that there exists $r > 0$, such that any 1–spherical cycle in $P_1(X^{(0)})$ is null-homologous in $P_r(X^{(0)})$. Then X is quasi-isometric to a simplicial tree.*

Note that to use Proposition 4.5 formally we need to consider a geodesic metric d_g on X – see Remark 2.10. Proposition 4.6 will be true for our standard metric d as well, since clearly $(X^{(0)}, d)$ is quasi–isometric to (X, d_g).

Proof. Let $R \geqslant r$ be such that every 1–spherical cycle $f\colon S^1 \to X$ that is null-homologous in $P_r(X)$ has a filling D in $P_R(X)$ with $\mathrm{supp}(D) \subset \mathrm{supp}(f)$.

We proceed by contradiction. Suppose that X is not quasi-isometric to a tree. Let δ be a natural number larger than $5R$. Then, by the bottleneck property (Proposition 4.5), there exist two vertices v and w, a midpoint m between them, and a path α between v and w omitting $B_\delta(m, X)$. Without loss of generality we can assume that m is a vertex. Let γ denote a geodesic between v and w that contains m.

Let $a = \lceil 3\delta/4 \rceil$ and $b = \lfloor \delta/2 \rfloor$. We define subcomplexes $A = P_R(\gamma \cap B_a(m, X))$, and $B = P_R(\alpha \cup (\gamma \setminus B_b(m, X))$ of $P_R(X)$. We claim that the following hold:

(1) A and B are path-connected,
(2) $A \cap B$ has the homotopy type of two points,
(3) $A \cup B = P_R(\alpha \cup \gamma)$.

Assertion (1) is straightforward. For (2) observe that

$$A \cap B = P_R(\gamma \cap (B_a(m, X) \setminus B_b(m, X)))$$

and that $\gamma \cap (B_a(m, X) \setminus B_b(m, X))$ consists of two geodesic segments that are separated by at least $2b > R$. The Rips complex of a geodesic segment is easily seen to be contractible. For (3) we clearly have $A \cup B \subset P_R(\alpha \cup \gamma)$. To prove the other inclusion we need to show that for any edge in $P_R(\alpha \cup \gamma)$ both of its endpoints are either in A or in B. This follows from the definition of A and B, as for any two vertices x and y with $x \in A \setminus B$ and $y \in B \setminus A$ we have $d(x, y) \geqslant a - b > R$.

Now let $\overline{\alpha}$ and $\overline{\gamma}$ be the continuous paths obtained from α and γ by connecting consecutive vertices by edges. Let $\overline{\alpha\gamma}$ be the 1–spherical cycle in $P_1(X^{(0)})$ obtained by their concatenation. By our assumption the cycle $\overline{\alpha\gamma}$ has a filling D in $P_R(\mathrm{supp}(\overline{\alpha\gamma})) = P_R(\alpha \cup \gamma) = A \cup B$ and thus $[\overline{\alpha\gamma}] = 0$ in $H_1(A \cup B)$. However, in the Mayer-Vietoris sequence for the pair A, B the boundary map

$$H_1(A \cup B) \to H_0(A \cap B)$$

sends $[\overline{\alpha\gamma}]$ to a non-zero element. This gives a contradiction and hence finishes the proof of the proposition. □

Finally we are ready to prove the main result of this section.

Corollary 4.7. *For a hyperbolic isometry h satisfying* (3.1) *and $L(h) > 3$, the graph of axes $(Y(h), d_{Y(h)})$ is quasi-isometric to a simplicial tree.*

Proof. We will show that $(Y(h), d_{Y(h)})$ satisfies the assumptions of Proposition 4.6.

First, we show that there exists an $r > 0$, such that any 1–spherical cycle in $P_1(Y(h)^{(0)})$ is null-homotopic in $P_r(Y(h)^{(0)})$. Let $f \colon S^1 \to P_1(Y(h)^{(0)})$ be such a cycle. We will show that f is null-homotopic in $P_2(Y(h)^{(0)})$ by constructing a simplicial map $p \colon \mathrm{Min}(h) \to P_1(Y(h)^{(0)})$ and a 1–spherical cycle $\tilde{f} \colon S^1 \to \mathrm{Min}(h)$, which is null-homotopic in $\mathrm{Min}(h)$, and such that $p \circ \tilde{f}$ is homotopic to f in $P_2(Y(h)^{(0)})$.

Let $\gamma_0, \ldots, \gamma_m$ be vertices of the image of f appearing in this order (i.e. γ_i and γ_{i+1} are adjacent and $\gamma_0 = \gamma_m$). For every $i \in \{0, \ldots, m-1\}$ pick a vertex $x_i \in \mathrm{Min}(h)$, such that $x_i \in \gamma_i$ and $x_i \neq x_j$ if $i \neq j$. Since γ_i and γ_{i+1} are adjacent in $Y(h)$, by definition of $Y(h)$ there exist vertices y_i and z_{i+1} in $\mathrm{Min}(h)$ such that $y_i \in \gamma_i$ and

$z_{i+1} \in \gamma_{i+1}$ and y_i, and z_{i+1} are adjacent in $\mathrm{Min}(h)$ (this can always be done, even if adjacency of γ_i and γ_{i+1} in $Y(h)$ follows from the fact that they intersect in $\mathrm{Min}(h)$). Let α_i be the path defined as the concatenation of the segment $[x_i, y_i]$ of γ_i, the edge $\{y_i, z_{i+1}\}$ and the segment $[z_{i+1}, x_{i+1}]$ of γ_{i+1}. Define $\tilde{f}$ as the concatenation of paths α_i for all $i \in \{0, \ldots, m-1\}$ (see Figure 4).

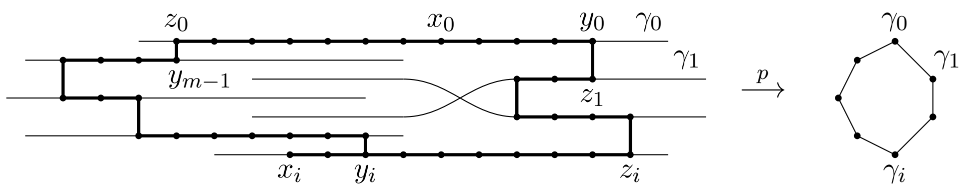

FIGURE 4. Cycles $\tilde{f}$ and f.

Define the map $p\colon \mathrm{Min}(h) \to P_1(Y(h)^{(0)})$ on vertices by $x \mapsto \gamma_x$, where γ_x is an h–invariant geodesic passing through x (such a geodesic in general is not unique, we choose one for each vertex). We claim that p is a simplicial map. Indeed, if vertices x and y form an edge in $\mathrm{Min}(h)$, then this edge connects geodesics γ_x and γ_y, hence by definition of the graph $Y(h)$ vertices γ_x and γ_y are adjacent. Since both complexes are flag, the claim follows.

The complex $\mathrm{Min}(h)$ is systolic, hence in particular it is simply connected. Thus the cycle $\tilde{f}$ is null-homotopic and so is $p \circ \tilde{f}$, since p is simplicial. It remains to prove that cycles f and $p \circ \tilde{f}$ are homotopic in $P_2(Y(h)^{(0)})$. To see that, notice that if v is a vertex in the image of $\tilde{f}$ and $v \in \gamma_i$ then its image $p(v)$ is at distance at most 1 from γ_i.

We are left with showing that the graph of axes $(Y(h), d_{Y(h)})$ is S^1FRC. We proceed as follows. By Lemma 3.1 the minimal displacement set $\mathrm{Min}(h)$ is systolic, hence S^2FRC by Lemma 4.3. By Theorem 3.4 we have a quasi-isometry

$$c\colon (Y(h) \times \mathbb{Z}, d_h) \to (\mathrm{Min}(h), d_X).$$

Therefore by Lemma 4.2 we conclude that the product $(Y(h) \times \mathbb{Z}, d_h)$ is S^2FRC. Finally, since $\mathbb{Z}$ is unbounded and since any 1–spherical cycle in $Y(h)^{(0)}$ is null-homotopic in $P_2(Y(h)^{(0)})$, Lemma 4.4 implies that $(Y(h), d_{Y(h)})$ is S^1FRC. □

For the sake of completeness we include the following well-known result.

Lemma 4.8. *Let G be a finitely generated group which acts properly by isometries on a quasi-tree (Q, d_Q). Then G is virtually free.*

Proof. Fix a finite generating set S of G and let d_S denote the associated word metric. Since the action of G is proper, the orbit map

$$(G, d_S) \to (Q, d_Q)$$

is a coarse embedding. Composing it with a quasi-isometry

$$(Q, d_Q) \to (T, d_T)$$

gives a coarse embedding of G into a tree T. Let X denote the image of this embedding with the metric restricted from T. The subspace X is quasi-connected, thus an appropriate thickening $N_R(X)$ is a connected subset of a tree, hence a tree. Clearly $N_R(X)$ is quasi-isometric to X. The composition

$$G \to X \to N_R(X)$$

is a coarse equivalence of geodesic metric spaces, hence a quasi-isometry; see e.g. [Roe03, Lemma 1.10]. This implies that G is δ–hyperbolic, and its Gromov boundary is 0–dimensional. It follows that G is virtually free. □

5. Classifying spaces for systolic groups

5.1. **Classifying space with virtually cyclic stabilisers.** In this section we gather results from Sections 2, 3 and 4 in order to prove Theorem 5.1.

Theorem 5.1. *Let G be a group acting properly on a uniformly locally finite systolic complex X of dimension d. Then there exists a model for $\underline{\underline{E}}G$ of dimension*

$$\dim \underline{\underline{E}}G = \begin{cases} d+1 & \text{if } d \leqslant 3, \\ d & \text{if } d \geqslant 4. \end{cases}$$

In the remainder of this section, let G be as in the statement of the above theorem. The model for $\underline{\underline{E}}G$ we construct is given by the cellular G–pushout of Theorem 2.2. Therefore we need to construct a model for $\underline{E}G$ and for every commensurability class of infinite virtually cyclic subgroups $[H]$, models for $\underline{E}N_G[H]$ and $E_{\mathcal{G}[H]}N_G[H]$. The first model was constructed by P. Przytycki [Prz09, Theorem 2.1], and later "refined" by V. Chepoi and the first author.

Theorem 5.2. [CO15, Theorem E] *The systolic complex X is a model for $\underline{E}G$.*

In order to construct models for the commensurators $N_G[H]$ we need a little preparation. First we show that the group G satisfies Condition (C) of Definition 2.5. Using this, in every finitely generated subgroup $K \subseteq N_G[H]$ that contains H we find a suitable normal cyclic subgroup, and show that the quotient group acts properly on a quasi-tree. This together with Propositions 2.4 and 2.7 allows us to construct the desired models.

Lemma 5.3. *The group G satisfies condition* (C) *of Definition 2.5.*

Proof. The proof is a slight modification of the one given in [Lüc09, proof of Theorem 1.1]). Take arbitrary $g, h \in G$ such that $|h| = \infty$, and assume there are $k, l \in \mathbb{Z}$ such that $g^{-1}h^k g = h^l$. We have to show that $|k| = |l|$. Since the action of G on X is proper, the element h acts as a hyperbolic isometry and by Proposition 3.2.(ii) there is an h^n–invariant geodesic $\gamma \in X$ for some $n \geqslant 1$. We get the claim by considering the following sequence of equalities for the translation length:

$$|k|L(h^n) = L(h^{nk}) = L(g^{-1}h^{nk}g) = L(h^{\pm nl}) = |l|L(h^n).$$

The first and the last of the equalities follow from the fact, that the translation length of an element can be measured on an invariant geodesic, the second one is an easy calculation and the third one is straightforward. □

Lemma 5.4. *Let K be a finitely generated subgroup of G, and $h \in K$ a hyperbolic isometry satisfying* (3.1)*, such that $\langle h \rangle$ is normal in K. Then the proper action of G on X induces a proper action of $K/\langle h \rangle$ on the graph of axes $Y(h)$.*

Proof. Since $\langle h \rangle$ is normal in K, the subcomplex $\mathrm{Min}(h)$ is invariant under K. Indeed, if $d_X(x, hx) = L(h)$, then for any $g \in K$ we have

$$d_X(gx, hgx) = d_X(x, g^{-1}hgx) = d_X(x, h^{\pm 1}x) = L(h).$$

Since h satisfies (3.1), the subcomplex $\mathrm{Min}(h)$ is spanned by the union of h–invariant geodesics. The group K acts by simplicial isometries, hence it maps h–invariant geodesics to h–invariant geodesics. This gives an action of K on the set of vertices of $Y(h)$. This action extends to the action on the graph $Y(h)$, because the adjacency relation between vertices of $Y(h)$ is preserved under simplicial isometries. The subgroup $\langle h \rangle$ acts trivially, hence there is an induced action of the quotient group $K/\langle h \rangle$.

It is left to show that the latter action is proper. For any vertex $\gamma \in Y(h)$ we show that its stabiliser $\mathrm{Stab}_{K/\langle h \rangle}(\gamma)$ is finite. Denote by π the quotient map $K \to K/\langle h \rangle$, and consider the preimage $\pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma))$. Elements of $\pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma))$ are precisely these isometries, which map geodesic γ to itself. Thus we can define a map $p : \pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma)) \to D_\infty$, where D_∞ is the infinite dihedral group, interpreted as the group of simplicial isometries of the geodesic line γ. We claim that the kernel $\ker(p)$ is finite. Indeed, the kernel consists of elements which act trivially on the whole geodesic γ, hence it is contained in the stabiliser $\mathrm{Stab}_G(x)$ of any vertex $x \in \gamma$. The group $\mathrm{Stab}_G(x)$ is finite, since the action of G on X is proper. Therefore the group $\pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma))$ is virtually cyclic, as it maps into a virtually cyclic group D_∞ with finite kernel. The infinite cyclic group $\langle h \rangle$ is contained in $\pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma))$, hence the quotient group $\pi^{-1}(\mathrm{Stab}_{K/\langle h \rangle}(\gamma))/\langle h \rangle = \mathrm{Stab}_{K/\langle h \rangle}(\gamma)$ is finite. □

Lemma 5.5. *Let K be a finitely generated subgroup of $N_G[H]$ that contains H. Then there is a short exact sequence*

$$0 \longrightarrow \langle h \rangle \longrightarrow K \longrightarrow K/\langle h \rangle \longrightarrow 0,$$

such that $h \in H$ is of infinite order and the group $K/\langle h \rangle$ is virtually free.

Proof. Choose an element of infinite order $\tilde{h} \in H$ satisfying the following two conditions:

(i) the set $\mathrm{Min}(\tilde{h})$ is the union of axes (see (3.1)),
(ii) the translation length $L(\tilde{h}) > 3$.

Both (i) and (ii) can be ensured by rising $\tilde{h}$ to a sufficiently large power. Indeed, by Proposition 3.2.(ii) there exists $n \geqslant 1$ such that $\tilde{h}^n$ satisfies condition (i). If $L(\tilde{h}^n) \leqslant 3$ then replace it with $\tilde{h}^{4n}$. The element $\tilde{h}^{4n}$ satisfies both conditions (see Proposition 3.2.(iv)). Notice that if an element satisfies conditions (i) and (ii) then, by Proposition 3.2.(iv) so does any of its powers. Since G satisfies Condition (C), by Lemma 2.6 there exists an integer $k \geqslant 1$ such that $\langle \tilde{h}^k \rangle$ is normal in K.

Put $h = \tilde{h}^k$. By Lemma 5.4 the group $K/\langle h \rangle$ acts properly by isometries on the graph of axes $(Y(h), d_{Y(h)})$, which is a quasi-tree by Corollary 4.7. Finally, Lemma 4.8 implies that the group $K/\langle h \rangle$ is virtually free. □

Lemma 5.6. *For every $[H] \in [\mathcal{VCY} \setminus \mathcal{FIN}]$ there exist*

(i) a 2–dimensional model for $E_{\mathcal{G}[H]} N_G[H]$,
(ii) a 3–dimensional model for $\underline{E} N_G[H]$.

Proof. By Proposition 2.4 it is enough to construct for every finitely generated subgroup $K \subseteq N_G[H]$, a 1–dimensional model for $E_{\mathcal{G}[H] \cap K} K$ and a 2–dimensional model for $\underline{E}K$. Notice that every finitely generated subgroup K' of G is contained in the finitely generated subgroup K that contains H (take $K = \langle K', H \rangle$). Therefore it is enough to consider only finitely generated subgroups of G that contain H.

By Lemma 5.5 for any such K there is a short exact sequence

$$0 \longrightarrow \langle h \rangle \longrightarrow K \xrightarrow{\pi} K/\langle h \rangle \longrightarrow 0,$$

where $K/\langle h \rangle$ is virtually free. The key observation is that the group $K/\langle h \rangle$ acts properly on a simplicial tree [KPS73, Theorem 1] and therefore a tree is a 1–dimensional model for $\underline{E}K/\langle h \rangle$. The claim follows then from Proposition 2.7 in the following way. First notice that for every subgroup $H \in \mathcal{G}[H]$ the image $\pi(H)$ is finite.

The preimage under π of any finite subgroup $F \in K/\langle h \rangle$ is a virtually cyclic group containing $\langle h \rangle$. In this case the intersection $\pi^{-1}(F) \cap \langle h \rangle$ is infinite, hence by definition of $\mathcal{G}[H]$ the group $\pi^{-1}(F)$ belongs to the family $\pi^{-1}(F) \cap \mathcal{G}[H]$. Thus the one point space is a 0–dimensional model for $E_{\mathcal{G}[H] \cap \pi^{-1}(F)} \pi^{-1}(F)$. This proves (i).

To prove (ii) notice that since $\pi^{-1}(F)$ is virtually cyclic, it acts on the real line with finite stabilisers [JPL06, Proposition 4]. Therefore a line is a 1–dimensional model for $\underline{E}\pi^{-1}(F)$. ☐

Proof of Theorem 5.1. By Corollary 2.3 choosing a model for $\underline{E}G$, and for every $[H] \in [\mathcal{VCY} \backslash \mathcal{FIN}]$ models for $\underline{E}N_G[H]$ and $E_{\mathcal{G}[H]}N_G[H]$, gives a model for $\underline{\underline{E}}G$ that satisfies the following inequality

$$\dim \underline{\underline{E}}G \leqslant \max\{\dim \underline{E}G, \sup_{[H]}\{\dim \underline{E}N_G[H]\} + 1, \sup_{[H]}\{\dim E_{\mathcal{G}[H]}N_G[H]\}\}.$$

If $d \geqslant 4$, then by Theorem 5.2 and Lemma 5.6 we have

$$\dim \underline{\underline{E}}G \leqslant \max\{d, 4, 2\} = d.$$

If $d \leqslant 3$, we can take the model for $\underline{E}G$ as a model for $\underline{E}N_G[H]$ instead of the one provided by Lemma 5.6, and obtain

$$\dim \underline{\underline{E}}G \leqslant \max\{d, d+1, 2\} = d + 1.$$ ☐

5.2. **Centralisers of cyclic subgroups.** As a corollary of our results we give the description of centralisers of infinite order elements in *systolic groups*, i.e. groups acting properly and cocompactly on systolic complexes, therefore confirming a conjecture of D. Wise.

Proposition 5.7. *Let G be a group that acts properly on a systolic complex X and let $h \in G$ be of infinite order. Suppose that K is a finitely generated subgroup of the centraliser $C_G(h)$ and $\langle h \rangle \subset K$. Then K is commensurable with $F_n \times \mathbb{Z}$ where F_n denotes the free group on n generators for some $n \geqslant 0$.*

Proof. The group $C_G(h)$ is contained in the commensurator $N_G[\langle h \rangle]$, hence so is K. Thus by Lemma 5.5 there is a short exact sequence

$$0 \longrightarrow \langle h^m \rangle \longrightarrow K \xrightarrow{p} VF_n \longrightarrow 0,$$

where VF_n is a virtually free group and m is some positive integer. Taking the free subgroup $F_n \subset VF_n$ gives rise to the following

$$0 \longrightarrow \mathbb{Z} \longrightarrow p^{-1}(F_n) \xrightarrow{p} F_n \longrightarrow 0.$$

Since F_n is free, the above sequence splits. Therefore, as a central extension, $p^{-1}(F_n)$ is of the form $\mathbb{Z} \times F_n$. This finishes the proof, as $[K : p^{-1}(F_n)] \leqslant [VF_n : F_n] < \infty$. ☐

Corollary 5.8. [Wis03, Conjecture 11.6] *Let G be a systolic group. Then for any element $h \in G$ of infinite order, the centraliser $C_G(h)$ is commensurable with $F_n \times \mathbb{Z}$ for some $n \geqslant 0$.*

Proof. The group G is biautomatic by [JŚ06, Theorem E], and it follows that the centraliser $C_G(h)$ is biautomatic as well [GS91a, Proposition 4.3]. In particular $C_G(h)$ is finitely generated. Thus the claim follows from Proposition 5.7. □

5.3. **Virtually abelian stabilisers.** In this section we study the family of all virtually abelian subgroups of a group G. We show that if G is systolic, then there exists a finite dimensional model for the classifying space for this family. This is due to a very special structure of abelian subgroups of systolic groups, which is in turn a consequence of the systolic Flat Torus Theorem. Our construction also carries through for certain CAT(0) groups.

Given a group G, let $\mathcal{VAB}$ denote the family of all virtually abelian subgroups of G and let $\mathcal{VAB}_{fg}$ denote the family of all finitely generated virtually abelian subgroups of G. Every subgroup in the family $\mathcal{VAB}_{fg}$ contains a finite-index free abelian subgroup of rank $n \geqslant 0$, therefore if we denote by $\mathcal{VAB}_n$ the family of all virtually abelian subgroups of rank at most n, we obtain the following filtration of the family $\mathcal{VAB}_{fg}$:

$$\mathcal{VAB}_0 \subset \mathcal{VAB}_1 \subset \mathcal{VAB}_2 \subset \ldots$$

Notice that $\mathcal{VAB}_0 = \mathcal{FIN}$ and $\mathcal{VAB}_1 = \mathcal{VCY}$. Moreover, if G is a systolic group then by Theorem 5.12.(1) it does not contain free abelian groups of rank higher than 2, and therefore the above filtration reduces to

$$\mathcal{FIN} \subset \mathcal{VCY} \subset \mathcal{VAB}_2 = \mathcal{VAB}_{fg}.$$

Moreover, in Proposition 5.14 we show that every virtually abelian subgroup of a systolic group is in fact finitely generated, and therefore for systolic groups we have $\mathcal{VAB}_{fg} = \mathcal{VAB}$. The following is the main theorem of this section.

Theorem 5.9. *Let G be a group acting properly and cocompactly on a d–dimensional systolic complex. Then there exists a model for $E_{\mathcal{VAB}}G$ of dimension* $\max\{4, d\}$.

The construction which we use is a pushout construction of [LW12] (cf. Section 2.1) applied to the inclusion of families $\mathcal{VCY} \subset \mathcal{VAB}$. More precisely we want to apply [LW12, Corollary 2.8] which requires the collection of subgroups $\mathcal{VAB} \setminus \mathcal{VCY}$ of G to satisfy the following two conditions:

(NM1): any $H \in \mathcal{VAB} \setminus \mathcal{VCY}$ is contained in a unique maximal $M \in \mathcal{VAB} \setminus \mathcal{VCY}$,
(NM2): for any maximal subgroup M of $\mathcal{VAB} \setminus \mathcal{VCY}$ we have $N_G(M) = M$.

These conditions correspond to conditions $M_{\mathcal{VCY}\subset\mathcal{VAB}}$ and $NM_{\mathcal{VCY}\subset\mathcal{VAB}}$ of [LW12, Notation 2.7]. We will keep our notation for the sake of clarity.

Lemma 5.10. *Let G be a systolic group. Then G satisfies conditions* (NM1) *and* (NM2)*.*

Assuming the lemma, we proceed with the construction of the desired model.

Proof of Theorem 5.9. Let $\mathcal{M}$ denote the complete set of representatives of conjugacy classes in G of subgroups which are maximal in $\mathcal{VAB}\setminus\mathcal{VCY}$. Since G satisfies (NM1) and (NM2), it follows from [LW12, Corollary 2.8] that a model for $E_{\mathcal{VAB}}G$ is given by the cellular G–pushout

$$\begin{array}{ccc} \coprod_{M\in\mathcal{M}} G\times_M \underline{\underline{E}}M & \xrightarrow{\ i\ } & \underline{\underline{E}}G \\ \Big\downarrow \coprod_{M\in\mathcal{M}} p_M & & \Big\downarrow \\ \coprod_{M\in\mathcal{M}} G/M & \longrightarrow & E_{\mathcal{VAB}}G, \end{array}$$

where i is an inclusion of CW–complexes and p_M is the canonical projection

$$G\times_M \underline{\underline{E}}M \to G\times_M * \cong G/M.$$

By Theorem 5.1 there exists a d–dimensional model for $\underline{\underline{E}}G$ as long as $d \geqslant 4$. It follows from [LW12, Theorem 5.13.(iii)] that there exists a 3–dimensional model for $\underline{\underline{E}}M$ (since M contains a finite-index subgroup isomorphic to $\mathbb{Z}^2$) and it is in fact a model of the lowest possible dimension. The existence of the map i follows from the universal property of the classifying space $\underline{\underline{E}}G$. To ensure that i is injective, we replace it with an inclusion into the mapping cylinder (cf. Corollary 2.3). Finally, we have that G/M has dimension 0. Therefore applying the above pushout to these models gives us a model for $E_{\mathcal{VAB}}G$ of dimension $\max\{0, 4, d\}$. □

It remains to prove Lemma 5.10. The main tool that we use in the proof is the systolic Flat Torus Theorem of T. Elsner. Before stating the theorem we need to recall some terminology. For details we refer the reader to [Els09a].

Definition 5.11. Let $\mathbb{E}^2_\Delta$ denote the equilaterally triangulated Euclidean plane. A *flat* in a systolic complex X is a simplicial map $F\colon \mathbb{E}^2_\Delta \to X$ which is an isometric embedding. We will identify F with its image and treat it as a subcomplex of X.

We say that two flats are *equivalent* if they are at finite Hausdorff distance. This gives an equivalence relation on the set of all flats which we call a *flat equivalence.* Let $Th(F)$ denote the subcomplex of X spanned by all the flats that are equivalent to F. We call $Th(F)$ the *thickening* of F. Any two equivalent flats are in fact at

Hausdorff distance 1 [Els09a, Theorem 5.4]. Therefore for any F' that is equivalent to F, the inclusion $F' \hookrightarrow Th(F)$ is a quasi-isometry.

If $H \subset G$ is a free abelian subgroup of a systolic group G we define minimal displacement set of H as follows

$$\mathrm{Min}(H) = \bigcap_{h \in H \setminus \{e\}} \mathrm{Min}(h).$$

Theorem 5.12 (Flat Torus Theorem)**.** [Els09a, Theorem 6.1] *Let G be a systolic group and let $H \subset G$ be a free abelian subgroup of rank at least 2. Then:*

(1) the group H is isomorphic to $\mathbb{Z}^2$,
(2) there exists an H–invariant flat F, unique up to flat equivalence,
(3) we have $\mathrm{Min}(H) = Th(F)$ for an H–invariant flat F.

Proof of Lemma 5.10. (NM1) First we show that any rank 2 virtually abelian subgroup $H \subset G$ is contained in a maximal one. This is equivalent to the statement that any ascending chain of rank 2 virtually abelian subgroups $H_1 \subset H_2 \subset \ldots$ stabilises, i.e. we have $H_i = H_{i+1}$ for i sufficiently large.

Suppose $H_1 \subset H_2 \subset \ldots$ is such a chain and let A be a finite-index subgroup of H_1 isomorphic to $\mathbb{Z}^2$. Since for every i the group H_i contains a finite-index subgroup isomorphic to $\mathbb{Z}^2$, it follows that the index of A in H_i is finite. We will show that this index is bounded from above by a constant which is independent of i.

By [Els09a, Corollary 6.2] the group H_i preserves the thickening of an A_i–invariant flat F_i where A_i is a certain finite-index subgroup of H_i. Since A and A_i are finite-index subgroups of H_i, so is their intersection $A \cap A_i$. By Theorem 5.12.(2) there exists an A–invariant flat F. Note that $A \cap A_i \cong \mathbb{Z}^2$ and both F and F_i are $A \cap A_i$–invariant. Therefore, again by Theorem 5.12.(2), we have $Th(F_i) = Th(F)$. Therefore any H_i preserves $Th(F)$.

Now, since G acts properly and cocompactly, for any integer $R > 0$ there exists an integer N_R such that for every vertex $v \in X$ the cardinality of the set $\{g \in G \mid d(gv, v) \leqslant R\}$ is at most N_R. Since A acts cocompactly on F and since $Th(F)$ is quasi-isometric to F, there is an integer $R > 0$ such that for any vertex $w \in F$, the orbit of a combinatorial ball $B_R(w, X)$ under A covers the thickening $Th(F)$. Fix a vertex $v \in F$. For any $h \in H_i$ there exists $a \in A$ such that $d(v, ahv) \leqslant R$. It follows that the index of A in H_i is bounded by N_R.

Now we prove the uniqueness. Assume that H_1 and H_2 are maximal subgroups in $\mathcal{VAB} \setminus \mathcal{VCY}$ that contain H. By [Els09a, Corollary 6.2.(2)] there are flats F_1 and F_2 such that $H_1 = \mathrm{Stab}_G(Th(F_1))$ and $H_2 = \mathrm{Stab}_G(Th(F_2))$. By [Els09a, Corollary 6.2.(1)] there exists a flat F, unique up to flat equivalence, such that H preserves

$Th(F)$. Since H is contained in both H_1 and H_2, the thickenings $Th(F_1)$ and $Th(F_2)$ are both H–invariant. Hence we have $Th(F_1) = Th(F_2) = Th(F)$ and therefore $H_1 = H_2$.

(NM2) Let $H \in \mathcal{VAB} \setminus \mathcal{VCY}$ be a maximal subgroup and let A' be a finite-index subgroup of H that is isomorphic to $\mathbb{Z}^2$. Define the subgroup $A \subset H$ as the intersection of all subgroups of H of index $[H : A']$. Since H is finitely generated, there is finitely many of subgroups of this kind. Therefore $A \subset H$ is a rank 2 free abelian subgroup of finite index, and by construction A is a characteristic subgroup of H. It follows that the group $N_G(H)$ normalises A, and hence it preserves the subcomplex $\mathrm{Min}(A)$.

The action of $N_G(H)$ on $\mathrm{Min}(A)$ is proper and therefore the induced action of $N_G(H)/A$ on $\mathrm{Min}(A)/A$ is proper. By Theorem 5.12.(3) we have $\mathrm{Min}(A) = Th(F)$ where F is an A–invariant flat. This implies that the action of A on $\mathrm{Min}(A)$ is cocompact. Since the quotient $N_G(H)/A$ acts properly on a compact space $\mathrm{Min}(A)/A$, it follows that $N_G(H)/A$ is a finite group. Therefore $N_G(H)$ is a rank 2 virtually abelian group and hence we have $N_G(H) = H$ by the maximality of H. □

The methods used above apply also to a certain class of CAT(0) groups, namely the groups acting geometrically on CAT(0) spaces that do not contain flats of dimension greater than 2. For details about CAT(0) spaces and groups we refer the reader to [BH99].

Corollary 5.13. *Let G be a group acting properly and cocompactly by isometries on a complete* CAT(0) *space X of topological dimension $d > 0$. Furthermore, assume that for $n > 2$ there is no isometric embedding $\mathbb{E}^n \to X$ where $\mathbb{E}^n$ is the Euclidean space. Then there exists a model for $E_{\mathcal{VAB}}G$ of dimension* $\max\{4, d+1\}$.

Among the CAT(0) spaces satisfying the assumptions of the above corollary there are CAT(0) spaces of dimension 2, e.g. CAT(0) square complexes, and rank–2 symmetric spaces. In particular, the corollary applies to lattices in rank–2 symmetric spaces, thus answering a special case of a question by J.-F. Lafont [Cha08, Problem 46.7]. On the other hand, our approach fails if X contains flats of dimension bigger than 2. The construction of models for $E_{\mathcal{VAB}}G$ in this case would require techniques significantly different from ours.

Proof of Corollary 5.13. By [Lüc09] there exists a model for $\underline{\underline{E}}G$ of dimension at most $d+1$. Since X does not contain isometrically embedded $\mathbb{E}^n$ for $n > 2$ it follows from the Flat Torus Theorem [BH99, Theorem II.7.1] that G does not contain free abelian subgroups of rank bigger than 2. This together with the fact that every virtually abelian subgroup of G is finitely generated [BH99, Corollary II.7.6] implies that the family $\mathcal{VAB}$ reduces to $\mathcal{VAB}_2$. It remains to show that conditions (NM1) and (NM2) are satisfied. The proof of this is analogous to the proof of Lemma 5.10. The "existence" part of (NM1) follows from [BH99, Theorem II.7.5]. Both the "uniqueness" part of (NM1) and condition (NM2) follow from [BH99, Corollary II.7.2]. □

We finish this section with the aforementioned proposition.

Proposition 5.14. *Let G be a group acting properly and cocompactly on a finite dimensional systolic complex X. Then every virtually abelian subgroup of G is finitely generated.*

Proof. It is enough to prove that every abelian subgroup A of G is finitely generated. Since G acts properly and cocompactly on X, there is a uniform bound on the order of finite subgroups of G. Therefore the torsion subgroup of A must be finitely generated. Now let $A' \subset A$ be the torsion-free part and let $\mathrm{rk}(A')$ denote its rank. By Theorem 5.12.(1) we have $\mathrm{rk}(A') \leqslant 2$.

If $\mathrm{rk}(A') = 1$ then we claim that $A' \cong \mathbb{Z}$. To show this, by the classification of torsion-free abelian groups of rank 1, it is enough to show that for any $a \in A'$ there are only finitely many positive integers n, such that there exists $b \in A'$ with $a = nb$. Suppose we have $a = nb$ for some b and n. Since both a and b are hyperbolic isometries of X, we can compare their translation lengths. If b has an axis, then it is straightforward to see that $L(nb) = n \cdot L(b)$. If b has no axis, then by [Els09b, Theorem 1.3] it has a "thick axis" of thickness $k \leqslant \dim X$, and using [Els09b, Fact 3.7] one easily checks that $L(nb) \geqslant \lfloor \frac{n}{k} \rfloor \cdot L(b)$. Therefore, in both cases the following holds:

$$L(a) = L(nb) \geqslant \left\lfloor \frac{n}{k} \right\rfloor \cdot L(b) \geqslant \left\lfloor \frac{n}{k} \right\rfloor . \tag{5.1}$$

Since $k \leqslant \dim X$, for a fixed element a there are only finitely many positive integers n satisfying (5.1). Therefore we get that $A' \cong \mathbb{Z}$.

If $\mathrm{rk}(A') = 2$ then proceeding as in the proof of Lemma 5.10.(NM2) we obtain that A' acts properly and cocompactly on a thickening of an A''–invariant flat where $A'' \subset A'$ is a subgroup isomorphic to $\mathbb{Z}^2$. Therefore A'/A'' is finite and thus A' is finitely generated. □

Remark 5.15. All results in this section hold under the following weakened assumptions. Instead of a cocompact action we assume that X is uniformly locally finite (which is automatically true if the action is cocompact) and that there is a uniform bound on the order of finite subgroups of G.

In the CAT(0) case, instead of a cocompact action we assume that X is proper, the action is via semisimple isometries and the set of translation lengths of hyperbolic elements is discrete at 0.

6. Graphical small cancellation complexes

In this section we begin the study of graphical small cancellation complexes. Our goal is to show that for any group G acting properly on graphical small cancellation complex, there is a (canonical) systolic complex on which G acts properly, and use the latter to construct low-dimensional models for various classifying spaces for G. This requires substantial preparations, including notation and terminology.

We begin with introducing combinatorial and graphical 2–complexes. Then we state and prove a version of the so-called Lyndon-van Kampen Lemma, and use it to establish certain combinatorial properties of graphical small cancellation complexes. In Section 7 we define the dual complex of a graphical small cancellation complex and show that these two are G–homotopy equivalent, where G is any group that acts on a graphical small cancellation complex. Finally we give the construction of classifying spaces for families $\mathcal{FIN}, \mathcal{VCY}$ and $\mathcal{VAB}$ for graphical small cancellation groups.

6.1. **Combinatorial 2–complexes.** The purpose of this section is to give the basic definitions and to establish terminology regarding combinatorial 2–complexes. In our exposition we mainly follow [MW02].

A map $X \to Y$ of CW–complexes is *combinatorial* if its restriction to every open cell of X is a homeomorphism onto an open cell of Y. A CW–complex is *combinatorial* if the attaching map of every n–cell is combinatorial for a suitable subdivision of the sphere S^{n-1}. An *immersion* is a combinatorial map that is locally injective.

Unless otherwise stated, all combinatorial CW–complexes that we consider are 2–dimensional and all the attaching maps are immersions. We will refer to them simply as "2–complexes". Consequently all the maps between 2–complexes are assumed to be combinatorial.

Notice that according to the above definition, a *graph* may contain loops and multiple edges, as opposed to graphs considered in Sections 2–5.

Example 6.1 (Presentation complex)**.** Let $\langle S|R\rangle$ be a group presentation. The *presentation complex* is a 2–complex that has a single 0–cell, a directed labeled 1–cell for each generator $s \in S$, and a 2–cell attached along the closed combinatorial path corresponding to each relator $r \in R$.

A *polygon* is a 2–disc with the cell structure that consists of n vertices, n edges and a single 2–cell. For any 2–cell C of 2–complex X there exists a map $R \to X$, where R is a polygon and the attaching map for C factors as $S^1 \to \partial R \to X$. In the remainder of this section by a *cell* we will mean a map $R \to X$ where R is a polygon. An *open cell* is the image in X of the single 2–cell of R.

A *path* in X is a combinatorial map $P \to X$ where P is either a subdivision of the interval or a single vertex. If P is a vertex, we call path $P \to X$ a *trivial* path. If the target space is clear from the context, we will refer to the path $P \to X$ as "the path P". The *interior* of the path is the path minus its endpoints. Let P^{-1} denote the path P traversed in the opposite direction. Given paths $P_1 \to X$ and $P_2 \to X$ such that the terminal point of P_1 is equal to the initial point of P_2, their *concatenation* is an obvious path $P_1P_2 \to X$ whose domain is the union of P_1 and P_2 along these points. A *cycle* is a map $C \to X$, where C is a subdivision of the circle S^1. The cycle $C \to X$ is *non-trivial* if its image is not a tree. Therefore a homotopically non-trivial cycle is non-trivial, but the converse is not necessarily true. A path or cycle is *simple* if it is injective on vertices. Notice that a simple cycle (of length at least 3) is non-trivial. A *length* of a path P or a cycle C denoted by $|P|$ or $|C|$ respectively is the number of 1–cells in the domain. A subpath $Q \to X$ of a path $P \to X$ (or a cycle) is a path that factors as $Q \to P \to X$ such that $Q \to P$ is an injective map. Notice that the length of a subpath does not exceed the length of the path.

A *disc diagram* is a contractible finite 2–complex D with a specified embedding into the real plane. We call D *nonsingular* if it is homeomorphic to the 2–disc. Otherwise D is called *singular*. The *area* of D is the number of 2–cells. The boundary cycle ∂D is the attaching map of the 2–cell that contains the point $\{\infty\}$, when we regard $S^2 = \mathbb{R}^2 \cup \{\infty\}$. A *boundary path* is any path $P \to D$ that factors as $P \to \partial D \to D$. An *interior path* is a path such that none of its vertices, except for possibly endpoints, lie on the boundary of D.

If X is a 2–complex a *disc diagram in* X is a map $D \to X$.

The following definition is crucial in small cancellation theory.

Definition 6.2. A *piece* in a disc diagram D is a path $P \to D$ for which there exist two different lifts to 2–cells of D, i.e. there is a commutative diagram

$$\begin{array}{ccc} P & \longrightarrow & R_i \\ \downarrow & & \downarrow \\ R_j & \longrightarrow & D \end{array}$$

where $R_i \to D$ and $R_j \to D$ are 2–cells and the maps $P \to R_i$ and $P \to R_j$ are different (although it might be that $R_i = R_j$).

Now we turn to graphical complexes.

Definition 6.3. Let $\mathbf{\Gamma} \to \Theta$ be an immersion of graphs and assume that Θ is connected. For convenience we will write $\mathbf{\Gamma}$ as the union of its connected components

$$\mathbf{\Gamma} = \bigsqcup_{i \in I} \Gamma_i,$$

and refer to the connected graphs Γ_i as *relators*.

A *thickened graphical complex* X is a 2–complex with 1–skeleton Θ and a 2–cell attached along every immersed cycle in $\mathbf{\Gamma}$, i.e. if a cycle $C \to \mathbf{\Gamma}$ is immersed, then in X there is a 2–cell attached along the composition $C \to \mathbf{\Gamma} \to \Theta$.

The term "thickened" comes from the fact, that for any connected component Γ_i, we have a "thick cell" $Th(\Gamma_i)$ which is formed by gluing 2–cells along all immersed cycles in Γ_i. As long as Γ_i is not a tree, there is infinitely many 2–cells in $Th(\Gamma_i)$. This definition may seem odd, however, it allows us to avoid certain technical complications in the proof of a version of the Lyndon-van Kampen Lemma in Section 6.2.

Definition 6.4. Let X be a thickened graphical complex. A *piece* in X is a path $P \to X$ for which there exist two different lifts to $\mathbf{\Gamma}$, i.e. there is a commutative diagram

$$\begin{array}{ccc} P & \longrightarrow & \Gamma_i \\ \downarrow & & \downarrow \\ \Gamma_j & \longrightarrow & X \end{array}$$

such that the maps $P \to \Gamma_i$ and $P \to \Gamma_j$ are different.

6.2. **The Lyndon-van Kampen Lemma.**

Definition 6.5. Let X be a thickened graphical complex. A disc diagram $D \to X$ is *reduced* if for every piece $P \to D$ the composition $P \to D \to X$ is a piece in X.

Observe that the definitions of a *piece* in D and in X are different (cf. Definition 6.2 and Definition 6.4). We use the same name as it will always be clear out of context what piece we consider.

Lemma 6.6 (Lyndon-van Kampen Lemma)**.** *Let X be a thickened graphical complex and let $C \to X$ be a closed homotopically trivial path. Then*

(1) there exists a (possibly singular) disc diagram $D \to X$ such that the path C factors as $C \to \partial D \to X$, and $C \to \partial D$ is an isomorphism,

(2) if a diagram $D \to X$ is not reduced, then there exists a diagram $D' \to X$ with smaller area and the same boundary cycle in the sense that there is a commutative diagram:

$$\begin{array}{ccc} \partial D' & \xrightarrow{\cong} & \partial D \\ & \searrow & \downarrow \\ & & X, \end{array}$$

(3) any minimal area diagram $D \to X$ such that C factors as $C \xrightarrow{\cong} \partial D \to X$ is reduced.

Proof. (1) Since C is null-homotopic, there exists a disc diagram $D \to X$ such that the map $C \to X$ factors as $C \to D \to X$ and $C \to D$ is the boundary cycle of D (see [ECH$^+$92, Section 2.2] for a proof).

(2) Since $D \to X$ is not reduced, there is a piece $P \to D$ such that $P \to D \to X$ is not a piece. Let $R_1 \to D$ and $R_2 \to D$ be the 2–cells such that P factors through both of them.

We will first treat the case when $R_1 = R_2$. Let $p_1, p_2 \colon P \to R_1$ denote the two different maps. Since $P \to D \to X$ is not a piece, the map $\partial R_1/(p_1 \sim p_2) \to X$ lifts to $\mathbf{\Gamma}$ (the graph $\partial R_1/(p_1 \sim p_2)$ is the quotient of the boundary of R_1, obtained by identifying the images $p_1(P)$ and $p_2(P)$ pointwise). Assume that P is maximal, i.e. it is not a proper subpath of a piece $P' \to D$. The attaching map for R_1 can be written as the concatenation $PS_1P^{-1}S_2 \to D$, such that S_1 and S_2 are closed paths, see Figure 5. Either S_1 or S_2 bounds a (possibly singular) subdiagram D' of D, assume that it is S_1. Remove from D the open cell R_1 together with the path P (retaining its initial vertex) and the subdiagram D' bounded by S_1. Call the resulting complex D'' (formally D'' is not a diagram as it is not contractible). Observe that D'' has a hole, whose boundary cycle is precisely S_2.

The lift $S_2 \to \mathbf{\Gamma}$ (given by $S_2 \to \partial R_1/(p_1 \sim p_2) \to \mathbf{\Gamma}$) is immersed everywhere, except for possibly at its initial vertex. Write $S_2 \to D$ as $Q_1SQ_2 \to D$ where Q_1 and Q_2 are the maximal paths such that lifts $Q_1 \to \mathbf{\Gamma}$ and $Q_2^{-1} \to \mathbf{\Gamma}$ are the same.

Now consider the quotient $\tilde{D}$ of D'' obtained by identifying the domains of Q_1 and Q_2. The boundary cycle of a hole is now equal to S, and by construction S lifts to an immersed cycle $S \to \mathbf{\Gamma}$. Therefore we can glue to $\tilde{D}$ a 2–cell $\tilde{R}$ determined by $S \to \mathbf{\Gamma}$. The area of the resulting diagram $\tilde{D} \cup \tilde{R}$ is smaller than the one of D.

Now suppose that $P \to D$ factors through two distinct cells $R_1 \to D$ and $R_2 \to D$. Assume that P is maximal and consider the lift $\partial R_1 \cup_P \partial R_2 \to \mathbf{\Gamma}$. Let $S_1 \to D$ and $S_2 \to D$ be paths such that the concatenations $PS_1 \to D$ and $PS_2 \to D$ are attaching maps for R_1 and R_2 respectively. Consider the lift $S_1S_2^{-1} \to \mathbf{\Gamma}$ of the closed path

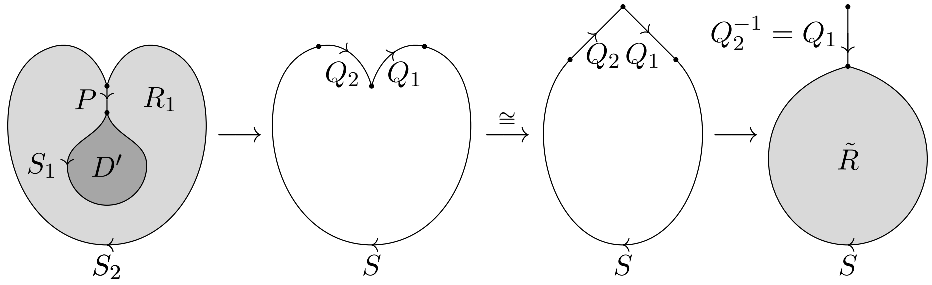

FIGURE 5. Replacing the open cell R_1 and the subdiagram $D' \cup \mathrm{Int}P$ with the 2–cell determined by the cycle S.

$S_1S_2^{-1} \to D$. If the paths $S_1 \to \mathbf{\Gamma}$ and $S_2 \to \mathbf{\Gamma}$ are equal, then we cut out from D open cells R_1 and R_2 together with the interior of the image of the path P and we "sew up" the resulting hole. For a proof of this see [MW02, Lemma 2.16].

We may therefore assume that $S_1 \to \mathbf{\Gamma}$ and $S_2 \to \mathbf{\Gamma}$ are not equal. Write $S_1 \to \mathbf{\Gamma}$ as the concatenation $J_1S_1'T_1 \to \mathbf{\Gamma}$ and $S_2 \to \mathbf{\Gamma}$ as $J_2S_2'T_2 \to \mathbf{\Gamma}$, such that $J_1 = J_2$ and $T_1 = T_2$ and both pairs are chosen to be maximal among paths having this property, see Figure 6. The subpaths S_1' and S_2' cannot be trivial since S_1 is not equal to S_2, and therefore the concatenation $S_1'S_2'^{-1} \to \mathbf{\Gamma}$ is a closed immersed path.

Remove from D open cells R_1 and R_2 together with the interior of the image of the path P, and consider the quotient D' of D obtained by identifying domains of paths J_1 and J_2 and of paths T_1 and T_2 respectively. The resulting diagram D' has a hole, whose boundary cycle lifts to the closed immersed path $S_1'S_2'^{-1} \to \mathbf{\Gamma}$. Hence we can attach the 2–cell $\tilde{R}$ along this path, thus removing the hole. This establishes (2) as the area of the resulting diagram is smaller than the area of D.

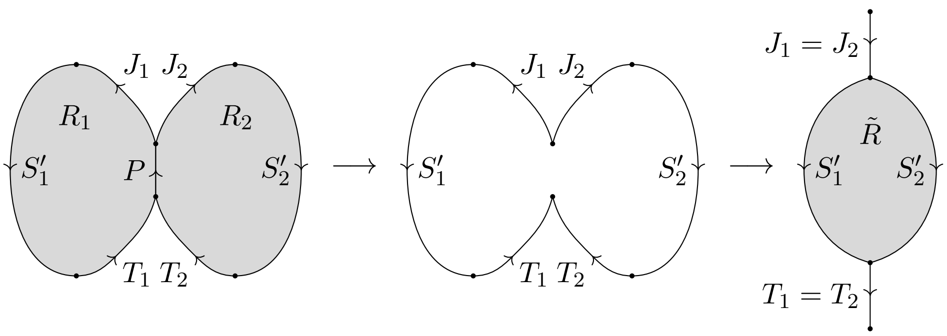

FIGURE 6. Replacement procedure.

(3) Let $D \to X$ be a minimal area diagram and suppose that it is not reduced. Then applying (2) gives a diagram of lower area and with the same boundary cycle, which contradicts the minimality of D. $\square$

6.3. **Properties of graphical small cancellation complexes.** In this section we define $C(p)$ and $C'(\lambda)$ small cancellation conditions and prove basic results about relators in graphical small cancellation complexes.

Definition 6.7. Let X be a thickened graphical complex, and let p be a positive integer and λ a positive real number. We say that X satisfies the

- $C(p)$ *small cancellation condition* if no non-trivial cycle $C \to X$ that factors as $C \to \Gamma_i \to X$ is the concatenation of less than p pieces.
- $C'(\lambda)$ *small cancellation condition* if for every piece $P \to X$ that is a subpath of a simple cycle $C \to \Gamma_i \to X$ we have $|P| < \lambda \cdot |C|$.

We abbreviate the $C(p)$ small cancellation condition to the "$C(p)$ condition" and call X a "$C(p)$ thickened graphical complex" (we use the same abbreviations in the $C'(\lambda)$ case). Mostly we will be concerned with the $C(p)$ condition for $p \geqslant 6$. Notice that if $p \geqslant q$ then the $C(p)$ condition implies the $C(q)$ condition. Therefore some results will be stated and proven in the $C(6)$ case only.

If D is a disc diagram we define small cancellation conditions in a very similar way, except that a *piece* is understood in the sense of Definition 6.2. For clarity we include the definition.

Definition 6.8. Let D be a disc diagram. We say that D satisfies the

- $C(p)$ small cancellation condition if no boundary cycle ∂R of a 2–cell R is the concatenation of less than p pieces.
- $C'(\lambda)$ small condition if for every piece P that factors as $P \to \partial R \to D$ for some 2–cell R, we have $|P| < \lambda \cdot |\partial R|$.

One can show that the $C'(\lambda)$ condition implies the $C(\lfloor \frac{1}{\lambda} \rfloor + 1)$ condition. This follows from the fact that it is enough to check the $C(p)$ condition on simple cycles.

Proposition 6.9. *If X is a $C(p)$ (respectively $C'(\lambda)$) thickened graphical complex and $D \to X$ is a reduced disc diagram, then D is $C(p)$ (respectively $C'(\lambda)$) diagram.*

Proof. The assertion follows immediately from the definitions of a reduced map and a piece. □

The next lemma is the crucial tool in small cancellation theory. It describes the possible shapes of the $C(6)$ disc diagrams. Before stating the lemma we need the following definition.

Let D denote a disc diagram. A *spur* is an edge of a boundary path of D that has a vertex of valence 1. In this case the boundary path is not immersed. Let $i \geqslant 0$ be an integer. A 2–cell $R \to D$ is called an *i–shell* if its boundary cycle ∂R is the concatenation $P_1 \cdots P_i Q$, such that every P_j is a simple interior path (and hence a piece), and Q is a simple boundary path of D. We call Q the *outer path* of ∂R.

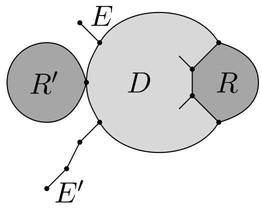

FIGURE 7. Diagram D with spurs E and E', a 0–shell R' and a 3–shell R.

Theorem 6.10 (Greendlinger's Lemma)**.** [MW02, Theorem 9.4] *Let D be a $C(6)$ disc diagram. Then one of the following holds:*

(1) D is a single 0–cell or it has exactly one 2–cell,
(2) D has at least two spurs or/and i–shells with $i \leqslant 3$.

The statement of Theorem 6.10 is actually weaker than the quoted Theorem 9.4 of [MW02], which distinguishes two further subcases of case (2). We present the simplified statement for the sake of clarity, as it is sufficient for our purposes.

Lemma 6.11. *Let X be a simply connected $C(6)$ thickened graphical complex. Then the following hold:*

(i) For every relator Γ_i, the map $\Gamma_i \to X$ is an embedding.
(ii) The intersection of (the images of) any two relators is either empty or it is a finite tree.
(iii) If three relators pairwise intersect then they triply intersect and the intersection is a finite tree.

Proof. (i) Assume conversely that $\Gamma_i \to X$ is not an embedding. Therefore there exist two distinct vertices of Γ_i which are mapped to a single vertex of X. Let $P \to \Gamma_i$ be a path joining these vertices. We can assume that P is an immersion. By construction P is non-closed and the projection $P \to \Gamma_i \to X$ is closed. Since X is simply connected, the path $P \to X$ is homotopically trivial. By Lemma 6.6.(1) there exists a disc diagram $D \to X$ with the boundary cycle $P \to X$. Assume that D is chosen such that the area of D is minimal among all examples of paths P of this type. Hence by Lemma 6.6.(3) diagram D is reduced, and therefore by Proposition 6.9 it satisfies $C(6)$ condition.

Thus one of the assertions of Theorem 6.10 applies to D. Clearly D cannot be trivial as in that case the path P would be trivial. It also does not contain spurs since $P \to X$ is an immersion. Thus it consists of either a single 2–cell or it contains at least one i–shell R with $i \leqslant 3$ and outer path Q, such that the endpoint of $P \to X$ is not contained in the interior of Q.

In the case when D consists of a single 2–cell, its boundary path $P \to D \to X$ lifts to a closed path in some Γ_j. This lift cannot be equal to the path $P \to \Gamma_i$ we started with, since by the assumption $P \to \Gamma_i$ is not a closed path. Hence $P \to X$ is a piece and since it is a non-trivial closed path, this violates the $C(6)$ hypothesis.

Now suppose R is an i–shell with $i \leqslant 3$ and the interior of its outer path $Q \to X$ avoids the endpoint of $P \to X$, see Figure 8. We claim that $Q \to X$ is a piece. If it is not the case, then the lift $Q \to \Gamma_i$ determined by the path $P \to \Gamma_i$ extends to a lift $\partial R \to \Gamma_i$.

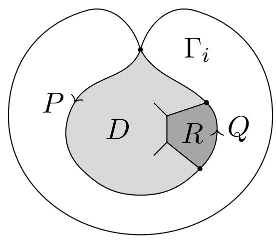

FIGURE 8. Diagram D with an i–shell R.

Thus we can remove from D the open cell R together with the interior of the path Q, and obtain a lower area diagram D' whose boundary path P' lifts to a non-closed path in Γ_i. If the resulting path $P' \to \Gamma_i$ is not immersed, we can fold the boundary of D' until all back-tracks are removed. The obtained diagram D' contradicts the minimality of D and hence proves the claim.

Given that $Q \to X$ is a piece, observe that the cycle $\partial R \to X$ is the concatenation of at most 4 pieces as $R \to D$ is an i-shell with $i \leqslant 3$. This contradicts the $C(6)$ hypothesis and hence establishes (i).

(ii) Given a relator Γ_i recall that a thick cell $Th(\Gamma_i)$ is a 2–complex obtained by gluing 2–cells along all immersed cycles in Γ_i. We shall argue by contradiction. Let Γ_1 and Γ_2 be two relators that meet along maximal disjoint connected subgraphs U and V and let

$$C = Th(\Gamma_1) \cup_{U \sqcup V} Th(\Gamma_2)$$

be a 2–complex obtained by gluing $Th(\Gamma_1)$ and $Th(\Gamma_2)$ along U and V. Note that there is an immersion $C \to X$ and consider the closed immersed path $P \to C \to X$ such that $P \to C$ is a generator for the fundamental group $\pi_1(C) \cong \mathbb{Z}$. Let $D \to X$ be a disc diagram whose boundary cycle is P and assume that the area of D is minimal among all examples of this type. Hence D is a non-trivial diagram without spurs and the map $D \to X$ is reduced. By Theorem 6.10 there is an i–shell R in D with $i \leqslant 3$ (if D consists of a single 2–cell we treat this cell as a 0–shell). Let Q denote the outer path of R in D.

We claim that any edge of Q is a piece in X. To show this assume the contrary, that there is an edge $E \to Q$ that is not a piece. Without loss of generality we can assume that the image of E in C (determined by the path $P \to C$) is contained in the relator Γ_1. Since $E \to X$ is not a piece, there exists a lift of the boundary ∂R to Γ_i extending the lift $E \to \Gamma_1$. Therefore as in the case (i) above, we can remove from D the open cell R together with the interior $\mathrm{Int}(Q)$ and obtain a lower area diagram D' whose boundary path P' is obtained from P by pushing the subpath Q through R. The paths P and P' are homotopic in C and therefore P' is a generator for $\pi_1(C)$. Thus D' is a lower area counterexample which contradicts the minimality of D and hence proves the claim (if D consists of a single 2–cell R, then Q is equal to the entire boundary ∂R and therefore pushing Q through R collapses D to a trivial cycle, hence contradicting the fact that U and V are disjoint).

Hence the path $Q \to X$ is the concatenation of n pieces, where n is a positive integer. Since R is an i–shell with $i \leqslant 3$, the $C(6)$ hypothesis implies that $n \geqslant 3$. The only situation when this can happen (up to changing roles of Γ_1 and Γ_2) is when the path $Q \to X$ travels in Γ_1 then passes to Γ_2 through the subgraph U and it comes back to Γ_1 through the subgraph V. More precisely the path Q has a subpath that is the concatenation $U'WV'$ where U' and V' are paths in Γ_1 which are not entirely contained in $\Gamma_1 \cap \Gamma_2$ and W is path in Γ_2 such that its initial vertex belongs to the subgraph U and its terminal vertex belongs to the subgraph V, see Figure 9.

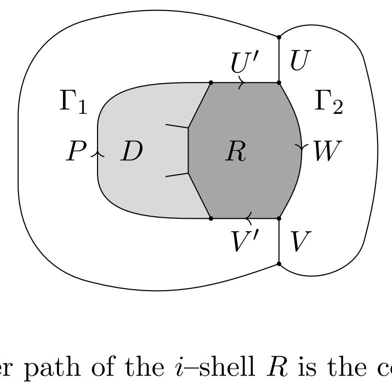

FIGURE 9. The outer path of the i–shell R is the concatenation $U'WV'$.

Notice that if both endpoints of W belong to one component of $\Gamma_1 \cap \Gamma_2$, say U (but W is not entirely contained in U), then we have a contradiction, as taking any path

$W' \to U$ connecting endpoints of W gives a cycle $WW' \to X$ that is a concatenation of 2 pieces.

Hence assume that we are in the situation shown in Figure 9. We have two cases to consider:

a) The path $U'WV' \to X$ is not closed. Let Γ_3 denote the relator containing a lift of the cycle ∂R and let $\overline{U}$ and $\overline{V}$ be the maximal connected components of $\Gamma_1 \cap \Gamma_3$ that contain paths U' and V' respectively. We claim that the intersection $\overline{U} \cap \overline{V}$ is empty. Assume it is not the case and pick a path $T \to \overline{U} \cup \overline{V} \subset \Gamma_1 \cap \Gamma_3$ joining the endpoint of V' to the origin of U'. The concatenation $U'WV'T \to X$ is then a closed path that is the concatenation of two pieces: W (lifts to Γ_2 and Γ_3) and $V'TU'$ (lifts to Γ_1 and Γ_3). This is a contradiction provided that the cycle $U'WV'T \to X$ is non-trivial, i.e. its image is not a tree. However, if it was trivial then T would be equal to $(U'WV')^{-1}$ and hence there would be a path in $\Gamma_1 \cap \Gamma_2$ joining subgraphs U and V, contradicting the assumption that U and V are disjoint. Consequently, the intersection $\overline{U} \cap \overline{V}$ is empty.

Thus we can replace C with $C' = Th(\Gamma_1) \cup_{\overline{U} \sqcup \overline{V}} Th(\Gamma_3)$ and D with $D' = D \setminus (\mathrm{Int}(R) \cup \mathrm{Int}(Q))$ and the path P with the path P' obtained by pushing the subpath Q through R. After removing possible back-tracks (in order for P' to be an immersion), we get a lower area counterexample.

b) The path $U'WV' \to X$ is closed. Let Γ_3 be the same as in case a) above. Then we get a contradiction as $U'WV' \to X$ is the concatenation of two pieces: $V'U'$ and W. Notice that the terminal vertex of V' and the initial vertex of U' lift to the same vertex of Γ_3 for otherwise $\Gamma_3 \to X$ is not an embedding what contradicts (i).

This shows that the intersection $\Gamma_1 \cap \Gamma_2$ is connected. Notice that there is no simple cycles in $\Gamma_1 \cap \Gamma_2$ as any simple cycle $C \to \Gamma_1 \cap \Gamma_2$ would be a piece itself. Therefore $\Gamma_1 \cap \Gamma_2$ is a tree.

(iii) Let Γ_1, Γ_2 and Γ_3 be relators that pairwise intersect but do not triply intersect, and let U_{12}, U_{13} and U_{23} be respective maximal connected subgraphs along which they intersect. Notice that the U_{ij} are disjoint from each other as otherwise there would be a triple intersection. Let C denote the union of $Th(\Gamma_1)$, $Th(\Gamma_2)$ and $Th(\Gamma_3)$ along subgraphs U_{12}, U_{13} and U_{23} and let $P \to C \to X$ be an immersed path such that $P \to C$ is a generator for $\pi_1(C)$. Let $D \to X$ be a disc diagram for P and suppose that the area of D is minimal among all examples as above. Thus $D \to X$ is reduced. Proceeding as in the proof of (ii) we conclude that D contains an i–shell R with $i \leqslant 3$, with outer path Q, such that every edge of Q is a piece. Since R is an i–shell with

$i \leqslant 3$, the $C(6)$ hypothesis implies that Q has a subpath $V_1V_2V_3$ such that V_1, V_2 and V_3 are non-trivial paths in Γ_1, Γ_2 and Γ_3 respectively, neither of them being contained entirely in an appropriate double intersection, and such that the origin of V_2 belongs to U_{12} and the endpoint of V_2 belongs to U_{23}, see Figure 10.

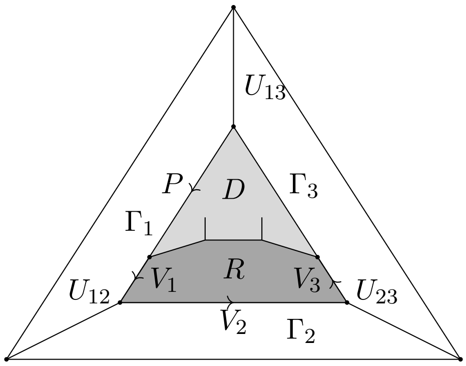

FIGURE 10. The outer path of the i–shell R is the concatenation $V_1V_2V_3$.

Similarly as in (ii) we consider two cases:

a) The path $V_1V_2V_3 \to X$ is not closed. Let Γ_4 denote the relator containing a lift of the cycle ∂R. We claim that the triple intersection $\Gamma_1 \cap \Gamma_3 \cap \Gamma_4$ is empty. Assume conversely that there exists a vertex $v \in \Gamma_1 \cap \Gamma_3 \cap \Gamma_4$. Choose paths $T_1 \to \Gamma_1 \cap \Gamma_4 \to X$ joining v to the initial vertex of V_1 and $T_2 \to \Gamma_3 \cap \Gamma_4 \to X$ joining the terminal vertex of V_3 to v. These paths exist because by (ii) the intersections $\Gamma_1 \cap \Gamma_4$ and $\Gamma_3 \cap \Gamma_4$ are connected. The concatenation $T_1V_1V_2V_3T_2 \to X$ is a non-trivial closed path that is the concatenation of three pieces: T_1V_1, V_2 and V_3T_2. This contradicts the $C(6)$ hypothesis and hence proves the claim.

   Now let $\overline{V_1}$ be the connected component of $\Gamma_1 \cap \Gamma_4$ that contains path V_1 and let $\overline{V_3}$ be the connected component of $\Gamma_3 \cap \Gamma_4$ that contains path V_3. We replace C with C' which is the union of $Th(\Gamma_1)$, $Th(\Gamma_3)$ and $Th(\Gamma_4)$ along the subgraphs U_{13}, $\overline{V_1}$ and $\overline{V_3}$. We replace D with $D' = D \setminus (\mathrm{Int}(R) \cup \mathrm{Int}(Q))$ and the path P with the path P' obtained by pushing the subpath Q through R. This gives a lower area counterexample.

b) The path $V_1V_2V_3 \to X$ is closed. Then it is the concatenation of 3 pieces, hence we get a contradiction with the $C(6)$ hypothesis.

It remains to show that the intersection $\Gamma_1 \cap \Gamma_2 \cap \Gamma_3$ is a tree. First we show that it is connected. Assume the converse and let u and v be vertices lying in different connected components of $\Gamma_1 \cap \Gamma_2 \cap \Gamma_3$. Since double intersections are connected we can pick paths $P \to \Gamma_1 \cap \Gamma_2$ and $Q \to \Gamma_1 \cap \Gamma_3$ both joining u to v. The concatenation

PQ^{-1} is then a closed path which is non-trivial since u and v lie in different connected components of $\Gamma_1 \cap \Gamma_2 \cap \Gamma_3$. Since PQ^{-1} is the concatenation of two pieces we get a contradiction. Therefore $\Gamma_1 \cap \Gamma_2 \cap \Gamma_3$ is connected. The proof that it is a tree is the same as in case (ii) above. □

Lemma 6.12. *Let X be a simply connected $C(6)$ thickened graphical complex and consider a finite collection of relators $\{\Gamma_i \to X\}_{i \in \{0,\ldots,n\}}$. If for every $i, j \in \{0, \ldots, n\}$ the intersection $\Gamma_i \cap \Gamma_j$ is non-empty then the intersection $\bigcap_{i \in \{0,\ldots,n\}} \Gamma_i$ is a non-empty tree.*

Proof. Consider first the intersection $\Gamma_0 \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$. This intersection is connected by Lemma 6.11.(ii)(iii). We claim that $\Gamma_0 \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$ is a tree.

Assuming the claim we proceed with the proof of the lemma. By Lemma 6.11.(ii)(iii) all intersections $\{\Gamma_0 \cap \Gamma_i\}_{i \in \{1,\ldots,n\}}$ are pairwise intersecting, non-empty subtrees of a tree $\Gamma_0 \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$. Therefore by the Helly property of trees the intersection $\bigcap_{i \in \{0,\ldots,n\}} \Gamma_i$ is a non-empty tree.

It remains to prove the claim. Assume conversely that there is a non-trivial simple cycle $C \to \Gamma_0 \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$. Let Γ_j be any relator different from Γ_0 through which C passes and let P_j be a maximal subpath of C that lifts to Γ_j. Choose P_i and P_k to be the paths that lift to Γ_i and Γ_k respectively, such that the concatenation $P_i P_j P_k$ is a maximal subpath C with these properties (for Γ_j fixed), see Figure 11. If $P_i P_j P_k = C$ then we get a contradiction with the fact that X is a $C(6)$ complex (the same happens if already $P_i P_j$ or $P_j P_k$ is equal to C).

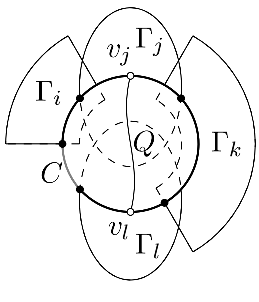

FIGURE 11. Cycle C partially covered by relators.

Hence assume that it is not the case and let P_l be the subpath of C that lifts to Γ_l and appears right after P_k. Choose vertices $v_j \in P_j$ and $v_l \in P_l$, such that $v_l \notin P_k \cup P_i$.

Since the intersection $\Gamma_j \cap \Gamma_0 \cap \Gamma_l$ is non-empty, there is a path $Q \to \Gamma_0 \cap (\Gamma_j \cup \Gamma_l)$ joining v_j to v_l. If Q is equal to the subpath of C from v_j to v_l which contains P_k then we get a contradiction with the choice of P_k, as in such case P_l appears right after P_j and covers a larger portion of C. Similarly if Q is equal to the subpath of C from v_j to v_l which contains P_i then we get a contradiction with the choice of P_i. Otherwise the concatenation of the subpath of C from v_j to v_l containing P_k with Q is a non-trivial cycle in Γ_0 that is covered by the images of three relators Γ_j, Γ_k and Γ_l. This contradicts the $C(6)$ condition and therefore finishes the proof of the claim. □

7. Dual of a $C(p)$ complex is p–systolic

Let X be a simply connected $C(p)$ thickened graphical complex X and suppose that $p \geqslant 6$. The purpose of this section is to construct a p–systolic simplicial complex $W(X)$ such that any group acting on X, acts naturally on $W(X)$. Furthermore, after replacing X with a "non-thickened" graphical complex X' we show that X' and $W(X)$ are G–homotopy equivalent. This replacement is necessary, as in general the thickened complex contains non-trivial 2–spheres, whereas systolic complexes are contractible. Roughly speaking, the non-thickened graphical complex has the same 1–skeleton as the thickened one, but instead of thick cells, it has topological cones glued along relators. The non-thickened complex, combinatorially being equivalent to the thickened one, has better topological properties (in particular it is contractible).

7.1. **Equivariant nerve theorem.** Our main tool in showing that X and $W(X)$ are G–homotopy equivalent is the Equivariant Nerve Theorem. This theorem is formulated in the abstract language of G–posets, therefore we begin by recalling some terminology.

A G–poset is a partially ordered set with an order-preserving action of a group G. A *geometric realisation* of a poset X is a simplicial complex $|X|$ whose n–simplices are chains $x_0 \leqslant y_1 \leqslant \ldots \leqslant x_n$ in X. If X is a G–poset then its geometric realisation $|X|$ is naturally a G–simplicial complex. All topological notions applied to a poset X are to be understood as corresponding notions applied to its geometric realisation $|X|$. For an element y of a poset Y define the subposet

$$Y_{\leqslant y} = \{x \in Y \mid x \leqslant y\}.$$

The following theorem is an equivariant analogue of the celebrated Quillen's "Theorem A".

Theorem 7.1 ([TW91, Theorem 1])**.** *Let G be a group and let $f\colon X \to Y$ be a G–map between G–posets X and Y. If for every $y \in Y$, the preimage $f^{-1}(Y_{\leqslant y})$ is G_y–contractible then f is a G–homotopy equivalence.*

Let $(X, \leqslant)$ be a poset. We say that a subset $U \subset X$ is closed with respect to $\leqslant$, if for any $x \in U$ and any y such that $y \leqslant x$, we have $y \in U$. A *cover* of a poset $(X, \leqslant)$ is a family $\mathcal{U}$ of subsets of X, such that every $U \in \mathcal{U}$ is closed with respect to $\leqslant$, and $\bigcup_{U \in \mathcal{U}} U = X$.

The *nerve* $N(\mathcal{U})$ of a cover $\mathcal{U}$ is a simplicial complex whose vertex set is $\mathcal{U}$, and vertices $U_0, \ldots, U_n$ span an n–simplex of $N(\mathcal{U})$ if and only if $\bigcap_{0 \leqslant i \leqslant n} U_i \neq \emptyset$. If G acts on X and for any element $U \in \mathcal{U}$ and any $g \in G$ we have $gU \in \mathcal{U}$ then we say that $\mathcal{U}$ is G*–cover*. In this case the G–action on X induces the G–action on $N(\mathcal{U})$. In particular, element $g \in G$ stabilises a simplex σ of $N(\mathcal{U})$ if and only if g leaves the intersection $\bigcap_{U \in \sigma} U \subset X$ invariant.

Definition 7.2. A G–cover $\mathcal{U}$ of a G–poset X is G*–contractible* if for any simplex σ of $N(\mathcal{U})$, the subposet $\bigcap_{U \in \sigma} U$ is a G_σ–contractible subposet of X, where G_σ denotes the G–stabiliser of σ.

The following result and its proof are immediate equivariant analogues of [Smi11, Theorem 4.5.2]. To the best of our knowledge there is no proof of this theorem in the literature.

Theorem 7.3 (Equivariant Nerve Theorem)**.** *Let G be a group and let X be a G–poset. Let $\mathcal{U}$ be a G–contractible cover of X. Then $N(\mathcal{U})$ is G–homotopy equivalent to $|X|$.*

Proof. We work with the face poset $N'(\mathcal{U})$ of $N(\mathcal{U})$, with the reversed inclusion order. More precisely, the elements of $N'(\mathcal{U})$ are simplices of $N(\mathcal{U})$, i.e. tuples $\{U_i\}_{i \in I}$ such that $\bigcap_{i \in I} U_i \neq \emptyset$ and $\{U_i\}_{i \in I} \leqslant \{U_j\}_{j \in J}$ in $N'(\mathcal{U})$ if and only if $J \subseteq I$.

The geometric realisation of $N'(\mathcal{U})$ is homeomorphic to $N(\mathcal{U})$ and the G–action on $N(\mathcal{U})$ induces a G–action on $N'(\mathcal{U})$. We define the map $f\colon X \to N'(\mathcal{U})$ as

$$f(x) = \{U \in \mathcal{U} \mid x \in U\}.$$

This is a map of posets since if $y \leqslant x$ then $y \in U$ whenever $x \in U$, by closedness of U. It is straightforward to check that f is a G–map.

Let $U_I = \{U_i\}_{i \in I}$ be an element of $N'(\mathcal{U})$. If $x \in f^{-1}(N'(\mathcal{U})_{\leqslant U_I})$ then $x \in U$, for every $U \in U_I$. Therefore

$$f^{-1}(N'(\mathcal{U})_{\leqslant U_I}) = \bigcap_{U \in U_I} U = \bigcap_{i \in I} U_i.$$

Since the cover $\mathcal{U}$ is G–contractible, each preimage $f^{-1}(N'(\mathcal{U})_{\leqslant U_I})$ is G_{U_I}–contractible. Therefore, by Theorem 7.1, the map f is a G–homotopy equivalence. □

In the remainder of this section we show how to apply Theorem 7.3 to the case of $C(p)$ graphical complexes. For this we need to introduce the "non-thickened" graphical complex.

Let Γ be a finite graph. A *cone* on Γ is the quotient space

$$C(\Gamma) = \Gamma \times [0,1]/\Gamma \times \{1\}.$$

Definition 7.4. Let $\varphi\colon \mathbf{\Gamma} \to \Theta$ be an immersion of graphs and assume that Θ is connected. Write $\mathbf{\Gamma}$ as the union of its connected components $\mathbf{\Gamma} = \bigsqcup \Gamma_i$ and let φ_i denote the composition $\Gamma_i \to \mathbf{\Gamma} \to \Theta$. Therefore $\varphi = \sqcup \varphi_i$.

A *graphical complex* X is a 2–complex obtained by gluing a cone $C(\Gamma_i)$ along each $\varphi_i\colon \Gamma_i \to \Theta$:

$$X = \Theta \cup_\varphi \bigsqcup_{i \in I} C(\Gamma_i).$$

For a map $\Gamma_i \to X$ a *cone-cell* is the corresponding map $C(\Gamma_i) \to X$.

Notice that X is not a 2–complex in the sense of Section 6.1. However, one can put a structure of a combinatorial 2–complex on X (or even a simplicial complex) by appropriately subdividing every cone. For most of our purposes though, it will be enough to treat entire cone-cells as "2–cells". Consequently we would like to treat the graph Θ as the 1–skeleton of X. In particular, any path $P \to X$ necessarily factors as $P \to \Theta \to X$.

Remark 7.5. To an immersion of graphs $\varphi\colon \mathbf{\Gamma} \to \Theta$ we assigned two complexes: a thickened graphical complex (see Definition 6.3) and a graphical complex (see Definition 7.4). Let us denote them by $Th(X)$ and X respectively. We emphasise that both constructions depend only on the map $\varphi\colon \mathbf{\Gamma} \to \Theta$ and therefore one construction determines another.

Moreover, notice that the fundamental groups of $Th(X)$ and X are isomorphic. Indeed one can construct a map $Th(X) \to X$ which is the identity on 1–skeleton, and which sends 2–cells of $Th(X)$ to the cone-cells of X. After a suitable subdivision this map becomes combinatorial, and one can easily show that it induces an isomorphism on fundamental groups.

We now proceed with the definitions of small cancellation conditions for a graphical complex. Notice that both definitions of a piece (Definition 6.4) and of small cancellation conditions (Definition 6.7) for a thickened graphical complex depend only on the map $\mathbf{\Gamma} \to \Theta$. Therefore we can use the exact same definitions for a graphical complex. For the sake of completeness we include the following (tautological) definition.

Definition 7.6. Let X be a graphical complex and let $Th(X)$ denote the corresponding thickened graphical complex. A path $P \to X$ is a *piece* if the corresponding path $P \to Th(X)$ is a piece. Consequently we say that X satisfies $C(p)$ or $C'(\lambda)$ condition if $Th(X)$ does so.

Remark 7.7. Definition 7.6 together with the fact that $\pi_1(Th(X)) \cong \pi_1(X)$ implies that Lemma 6.11 and Lemma 6.12 hold for a $C(6)$ graphical complex X as well.

From now on a *graphical complex* will always be understood in the sense of Definition 7.4. We proceed with the definition of the aforementioned simplicial complex $W(X)$.

Definition 7.8. Let X be a simply connected $C(p)$ graphical complex for $p \geqslant 6$. Assume that X is the union of its cone-cells, i.e. that every edge and vertex of X is in the image of $C(\Gamma_i) \to X$ for some relator Γ_i. Notice that by Lemma 6.11.(i) every map $\Gamma_i \to X$ is an embedding, and therefore we can identify a cone-cell $C(\Gamma_i) \to X$ with its image. Let

$$\mathbf{U} = \{C(\Gamma_i) \mid \Gamma_i \subset \mathbf{\Gamma}\}$$

be the covering of X by its cone-cells. Define the simplicial complex $W(X)$ to be the nerve of the covering $\mathbf{U}$. This complex was introduced by D. Wise in the classical $C(p)$ setting [Wis03], therefore we will refer to $W(X)$ as the *Wise complex.*

Notice that any cellular G–action on X (i.e. cellular on 1–skeleton and maps cone-cells to cone-cells) induces a simplicial G–action on $W(X)$. Our goal is to show that in fact X and $W(X)$ are G–homotopy equivalent. To show this, we will present X as a realisation of a certain G–poset, and we will find a G–cover of this poset whose nerve will be isomorphic to $W(X)$. The claim will then follow from Theorem 7.3.

Remark 7.9. We remark that the assumption in Definition 7.8 is not very restrictive. Indeed, if X contains such "free edges", i.e. edges not contained in any cone-cell, one can consider a new complex X' obtained by gluing to X a cone over every free edge (this cone is homeomorphic to the triangle in this case). The complex X' satisfies the assumptions of Definition 7.8 and any cellular G–action on X induces a cellular G–action on X'. It is straightforward to check that the quotient map $X' \to X$ which retracts every cone over the free edge onto this edge is a G–homotopy equivalence.

Let X be as in Definition 7.8. We define an associated poset $\mathcal{X}$ as follows. Elements of $\mathcal{X}$ are cone-cells, edges, and vertices of X ordered by inclusion. The geometric realisation $|\mathcal{X}|$ of the poset $\mathcal{X}$ is homeomorphic to X, and if G acts on X then there is an induced action on $\mathcal{X}$, and the homeomorphism is equivariant.

Let $\mathcal{U}$ be the cover of $\mathcal{X}$ given by

$$\mathcal{U} = \{\mathcal{X}_{\leqslant c} \mid c \text{ is a cone-cell of } X\}.$$

By construction the cover $\mathcal{U}$ is closed with respect to $\leqslant$ and it is straightforward to check that it is a G–cover of $\mathcal{X}$. Observe that the geometric realisation of any element $\mathcal{X}_{\leqslant c}$ of $\mathcal{U}$ is homeomorphic to the cone-cell c. Therefore the nerve $N(\mathcal{U})$ is isomorphic to the complex $W(X)$.

Lemma 7.10. *The G–cover $\mathcal{U}$ is G–contractible.*

Proof. For any $\sigma \in N(\mathcal{U})$ the geometric realisation of the intersection $\bigcap_{U\in\sigma} U$ is a tree by Lemma 6.12, hence it is G_σ–contractible. □

The above discussion together with Lemma 7.10 and Theorem 7.3 gives the following.

Theorem 7.11. *Let X be simply connected $C(6)$ graphical G–complex satisfying the assumptions of Definition 7.8. Then X is G–homotopy equivalent to the simplicial complex $W(X)$.*

7.2. **Graphical small cancellation groups are systolic.** In this section we show that if X satisfies the $C(p)$ small cancellation condition then the complex $W(X)$ is p–systolic, and we use the latter to construct models for the classifying spaces $\underline{\underline{E}}G$ and $E_{\mathcal{VAB}}G$ for a group G acting properly on X.

Theorem 7.12. *Suppose $p \geqslant 6$ and let X be a simply connected $C(p)$ graphical complex. Then its Wise complex $W(X)$ is p–systolic.*

Proof. The idea as well as the strategy of the proof come from D. Wise who proved this theorem for classical $C(6)$ complexes (cf. Theorem 10.6 in [Wis03]). We need to show that $W(X)$ is simply connected, flag and that links of vertices of $W(X)$ are p–large.

Simple connectedness of $W(X)$ follows from Theorem 7.11. To show that $W(X)$ is flag, suppose that $v_0, \ldots, v_n$ are vertices of $W(X)$ which are pairwise adjacent. We claim that these vertices span an n–simplex of $W(X)$. Let $C(\Gamma_0), \ldots, C(\Gamma_n)$ be the corresponding cone-cells in X. By our assumption we have $\Gamma_i \cap \Gamma_j \neq \emptyset$ for all $0 \leqslant i, j \leqslant n$ (cone-cells can intersect only at the relators). Thus by Lemma 6.12 the intersection $\bigcap_{i=1}^n \Gamma_i$ is non-empty, and therefore the vertices $v_0, \ldots, v_n$ span an n–simplex of $W(X)$.

It remains to show that for any vertex $v \in W(X)$ the link $W(X)_v$ is p–large. Let $(v_1, \ldots, v_n)$ be a cycle in $W(X)_v$ of length less than p. This corresponds to a sequence $C(\Gamma_1), \ldots, C(\Gamma_n)$ of cone-cells such that $\Gamma_i \cap \Gamma_{i+1} \neq \emptyset$ and all Γ_i intersect a fixed relator Γ.

The intersection $\Gamma \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$ is a connected graph (cf. proof of Lemma 6.12). We claim that it is a tree. Indeed, any non-trivial cycle in $\Gamma \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$ is a concatenation of at most n pieces, which contradicts the $C(p)$ hypothesis as $n < p$. Now choose vertices $u_i \in \Gamma \cap \Gamma_i \cap \Gamma_{i+1}$ and let $P_i \to \Gamma \cap \Gamma_i$ be a non-backtracking path joining u_{i-1} to u_i. The concatenation of paths $P_1 P_2 \cdots P_n$ is a cycle in the tree $\Gamma \cap (\Gamma_1 \cup \ldots \cup \Gamma_n)$. It is straightforward to check that there are two nonconsecutive paths P_j and P_k that intersect, see Figure 12 on the right. Therefore the cone-cells $C(\Gamma_j)$ and $C(\Gamma_k)$ intersect and this gives a diagonal (v_j, v_k) in a cycle $(v_1, \ldots, v_n)$ in $W(X)_v$. □

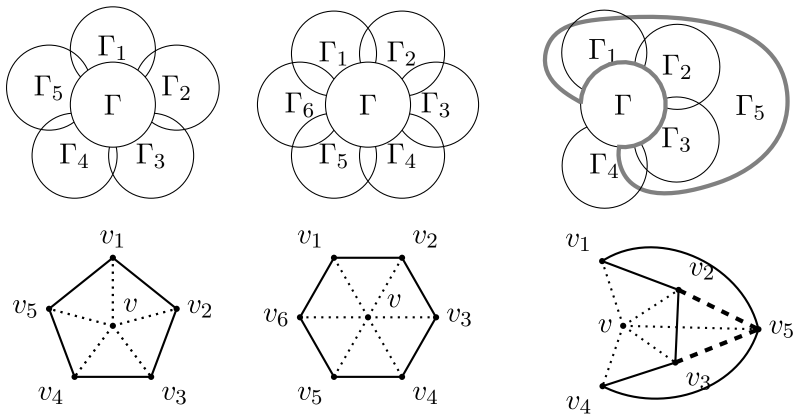

FIGURE 12. Vertex links of $W(X)$ and the corresponding subcomplexes of a $C(6)$ complex X. On the left there is an illegal configuration leading to a cycle of length 5 without diagonals. In the middle and on the right legal configurations are shown.

Definition 7.13. We say that a graphical complex X is (uniformly) locally finite, if after subdividing each cone-cell $C(\Gamma_i)$ into triangles spanned by the edges of Γ_i and the apex of the cone $C(\Gamma_i)$, the resulting complex is a (uniformly) locally finite simplicial complex.

It follows directly from the construction that X is uniformly locally finite if and only if $W(X)$ is so. Consequently, since X and $W(X)$ are G–homotopy equivalent, the G–action on X is proper if and only if the G–action on $W(X)$ is proper. Finally, the G–action on X is cocompact if and only if the G–action on $W(X)$ is so. These observations lead to the following corollary, which is interesting in its own right.

Corollary 7.14. *Let G be a group acting properly and cocompactly on a simply connected $C(p)$ graphical complex for $p \geqslant 6$. Then G acts properly and cocompactly on a systolic complex, i.e. G is a systolic group.*

Being systolic implies many properties including e.g. biautomaticity [JŚ06, Theorem 13.1]. For further results see e.g. [JŚ06, JŚ07, Prz09, Osa15, OŚ15] and references therein.

We now state and prove the main theorem of this section.

Theorem 7.15. *Let a group G act properly on a simply connected uniformly locally finite $C(6)$ graphical complex X. Then:*

(1) the complex X is a model for $\underline{E}G$,

(2) there exists a 3–dimensional model for $\underline{\underline{E}}G$,

(3) there exists a 4–dimensional model for $E_{\mathcal{VAB}}G$, provided the action is additionally cocompact.

Proof. (1) By Theorem 7.11 the group G acts properly on a uniformly locally finite systolic complex $W(X)$, and hence by Theorem 5.2 the complex $W(X)$ is a model for $\underline{E}G$. Therefore X is a model for $\underline{E}G$ as well, since X and $W(X)$ are G–homotopy equivalent.

(2) By Corollary 2.3 it is enough to find for every $[H] \in [\mathcal{VCY} \setminus \mathcal{FIN}]$ a 2–dimensional models for $\underline{E}N_G[H]$ and $E_{\mathcal{G}[H]}N_G[H]$. By (1) the complex X may serve as a model for $\underline{E}N_G[H]$. Notice that G acts properly on a systolic complex $W(X)$, hence by Lemma 5.6.(i) there exists a 2–dimensional model for $E_{\mathcal{G}[H]}N_G[H]$.

(3) Since G acts properly and cocompactly on a systolic complex, by Lemma 5.10 it satisfies conditions (NM1) and (NM2) (cf. Section 5.3). Therefore proceeding exactly as in the proof of Theorem 5.9, we obtain a model for $E_{\mathcal{VAB}}G$ of dimension $\max\{4, d\}$ where d is the dimension of a model for $\underline{\underline{E}}G$. By (2) the latter can be chosen to be at most 3, hence the claim. □

8. Examples

In this section we provide few classes of examples of groups to which our theory applies. When relevant, we mention that our constructions give new bounds on dimensions of classifying spaces.

8.1. **Graphical small cancellation presentations.** A *graphical presentation* $\mathcal{P} = \langle S \mid \varphi \rangle$ is a graph

$$\mathbf{\Gamma} = \bigsqcup_{i \in I} \Gamma_i,$$

and an immersion

$$\varphi \colon \mathbf{\Gamma} \to R_S,$$

where every Γ_i is finite and connected, and R_S is a rose, i.e. a wedge of circles with edges (cycles) labelled by a set S. Alternatively, the map $\varphi \colon \mathbf{\Gamma} \to R_S$, called a *labelling*, may be thought of as an assignment: to every edge of $\mathbf{\Gamma}$ we assign a direction (orientation) and an element of S.

A graphical presentation $\mathcal{P}$ defines a group

$$G = G(\mathcal{P}) = \pi_1(R_S)/\left\langle\left\langle \varphi_*(\pi_1(\Gamma_i))_{i\in I}\right\rangle\right\rangle .$$

In other words G is the quotient of the free group $F(S)$ by the normal closure of the group generated by all words (over $S \cup S^{-1}$) read along cycles in $\mathbf{\Gamma}$ (where an oriented edge labelled by $s \in S$ is identified with the edge of the opposite orientation and the label s^{-1}). A *piece* is a path P labelled by S such that there exist two immersions $p_1 \colon P \to \mathbf{\Gamma}$ and $p_2 \colon P \to \mathbf{\Gamma}$, and there is no automorphism $\Phi \colon \mathbf{\Gamma} \to \mathbf{\Gamma}$ such that $p_1 = \Phi \circ p_2$. The presentation $\mathcal{P}$ satisfies the $C(p)$ *small cancellation condition*, for $p \geqslant 6$, if no cycle in $\mathbf{\Gamma}$ is covered by less than p pieces; see eg. [Gru15] for a systematic treatment.

Consider the following graphical complex (see Definition 7.4):

$$X = R_S \cup_\varphi \bigsqcup_{i\in I} C(\Gamma_i).$$

The fundamental group of X is isomorphic to G. In the universal cover $\widetilde{X}$ of X there might be multiple copies of cones $C(\Gamma_i)$ whose attaching maps differ by lifts of $\mathrm{Aut}(\Gamma_i)$. After identifying all such copies, we obtain the complex $\widetilde{X}^*$. The group G acts geometrically, but not necessarily freely on $\widetilde{X}^*$. If $\mathcal{P}$ is a $C(p)$ graphical small cancellation presentation then the complex $\widetilde{X}^*$ is a $C(p)$ small cancellation complex. Moreover, the complex $\widetilde{X}^*$ satisfies the assumptions of Definition 7.8 as long as the map $\varphi \colon \mathbf{\Gamma} \to R_S$ is surjective. This happens precisely when the presentation $\mathcal{P}$ has no free generators.

Graphical small cancellation presentations provide a powerful tool for constructing groups with often unexpected properties, see e.g. [Osa14]. For such groups with torsion our result concerning the model for $\underline{E}G$ is new. If a $C(6)$ graphical small cancellation group is torsion-free then it admits a model for $\underline{\underline{E}}G$ of dimension at most three, by the work of D. Degrijse [Deg17, Corollary 3]. There are however $C(6)$ graphical small cancellation groups to which Degrijse's result does not apply. In such cases our constructions of low-dimensional $\underline{\underline{E}}G$ and $E_{\mathcal{VAB}}G$ are the only general tools available.

8.2. **Groups acting on $\mathcal{VH}$–complexes.** Not all groups acting geometrically on graphical small cancellation complexes possess graphical small cancellation presentations. The simplest example is $\mathbb{Z}^2$. It acts simply transitively on a tessellation of the plane by regular hexagons (a simple example of a $C(6)$ complex), but possesses no graphical $C(6)$ presentation. The following class of examples is more interesting.

The notion of *$\mathcal{VH}$–complexes* was introduced by D. Wise [Wis96]. Recall that a square complex X, i.e. a combinatorial 2–complex whose cells are squares, is a $\mathcal{VH}$–complex if the following holds. The edges of X can be partitioned into two classes $\mathcal{V}$ and $\mathcal{H}$ called *vertical* and *horizontal* edges respectively, such that every square has two opposite vertical and two opposite horizontal edges.

Let X be a simply connected $\mathcal{VH}$–complex that is a CAT(0) space with respect to the standard piecewise Euclidean metric. We now show how to turn X into a simply connected $C(6)$ graphical small cancellation complex. Subdivide every square into 24 triangles as shown in Figure 13 on the left. More precisely, the subdivision is invariant

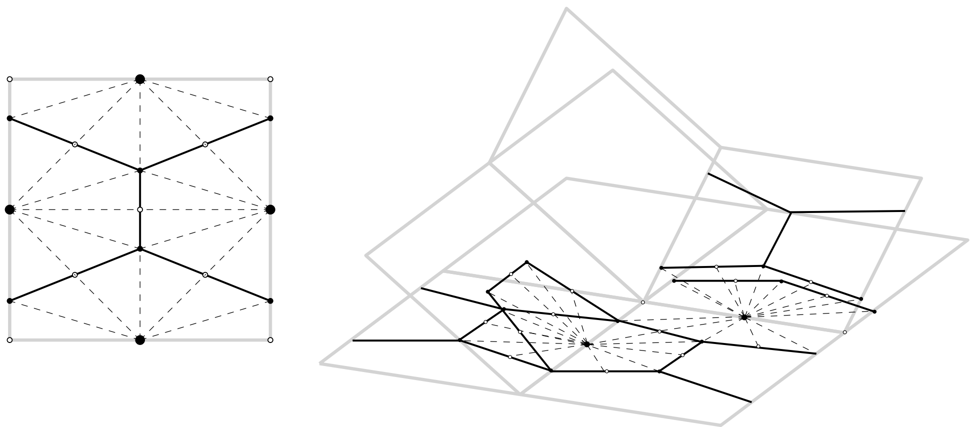

FIGURE 13. The subdivision of a $\mathcal{VH}$–square into 24 triangles (left), and the $C(6)$ graphical small cancellation complex structure on a CAT(0) $\mathcal{VH}$–complex. Two cones on relators are highlighted: one with a vertical and one with a horizontal apex.

with respect to $\mathcal{VH}$–isometries of the square, the vertical edges are subdivided into four sub-edges each, and the horizontal edges are subdivided into two sub-edges each. This defines a triangulation of X. Call the vertices of this triangulation, being mid-points of vertical and horizontal edges *vertical* and *horizontal apexes*, respectively. Consider links of apexes. Such a link is a graph of girth 12. Two such links intersect in a subgraph (possibly empty) of diameter at most 2; see Figure 13 on the right. Therefore, the complex X has a structure of the union of cones on links of apexes (relators). This defines the $C(6)$ graphical small cancellation complex X^*. It is clear that every $\mathcal{VH}$–automorphism, i.e. an automorphism respecting types of edges of X induces an automorphism of X^*.

Theorem 8.1. *Let X be a simply connected $\mathcal{VH}$–complex. Then the complex X^* is a $C(6)$ graphical small cancellation complex. In particular, every group of $\mathcal{VH}$–automorphisms of X acts by automorphisms on X^*. One action is proper and/or cocompact if and only if the other is so.*

T. Elsner and P. Przytycki [EP13] showed that a group acting properly or geometrically on a simply connected $\mathcal{VH}$–complex acts, respectively, properly or geometrically on a 3–dimensional systolic complex. Theorem 8.1 together with Theorem 7.12 provide a higher dimensional systolic complex in such a case. Nevertheless, the theorem above equips $\mathcal{VH}$–groups with a new 2–dimensional structure, extending in a way the Elsner-Przytycki result.

Of course, $\mathcal{VH}$–complexes carry a natural CAT(0) metric so that constructions of the corresponding low-dimensional models for $\underline{E}G$ and $\underline{\underline{E}}G$ are available by [Lüc09]. Our results provide a 4–dimensional model for $E_{\mathcal{VAB}}G$ for groups acting geometrically on such complexes.

8.3. **Lattices in $\widetilde{A}_2$–buildings.** Here we present another example of a group acting properly on a graphical small cancellation complex. An $\widetilde{A}_2$–building is a building with apartments isomorphic to the equilaterally triangulated plane $\mathbb{E}^2_\Delta$, see Definition 5.11.

Consider such a building Y. Let Y' be its barycentric subdivision. Define a *dual graph* Θ of Y as follows. Vertices of Θ are edges of Y and triangles of Y. There is an edge in Θ between every edge of Y and a triangle of Y containing this edge; see Figure 14.

The link in Y' of any vertex of Y is a 12–large graph (a subdivision of a spherical building) that may be considered as a subgraph of Θ. The complex Y' is thus obtained by attaching cones on such links to the graph Θ. Two such cones intersect in a set of diameter at most 2. Therefore Y' may be seen as a $C(6)$ graphical small cancellation complex. Lattices in $\mathrm{Isom}(Y)$ act naturally on Y'. Notice that such lattices may be very different from groups in the previous example because they may have Kazhdan's property (T).

Each $\widetilde{A}_2$–building possesses a natural structure of a systolic 2–dimensional complex or even a CAT(0) complex. Our results provide a 4–dimensional model for $E_{\mathcal{VAB}}G$ for lattices in its isometry group.

In fact, by exactly the same construction as above one equips any 2–dimensional p–systolic complex with a structure of a $C(p)$ graphical complex.

8.4. **A 3–dimensional systolic example.** As the last example we present a non-hyperbolic group G acting geometrically on a 3–dimensional systolic pseudomanifold X which does not admit a G–invariant CAT(0) metric.

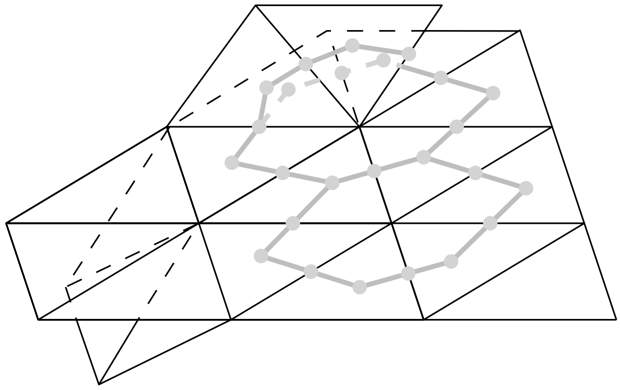

FIGURE 14. A part of an $\widetilde{A}_2$–building together with a part of its dual graph (thick gray).

We start with a simplex of groups $\mathcal{G}_{54}$, introduced by J. Świątkowski in [Świ06] (for some details on complexes of groups we refer the reader to [JŚ06, Świ06]). Let T_{54} be a 6–large triangulation of the flat 2–torus consisting of 54 equilateral triangles; see Figure 15 on the right (with the opposite sides of the hexagon and the appropriate vertices identified). Let G_{54} be a group of automorphisms of T_{54} generated by reflections with respect to edges of triangles. For the G_{54}–action on T_{54} the stabilisers of triangles are trivial, the stabilisers of edges are isomorphic to $\mathbb{Z}_2$, and the stabilisers of vertices are isomorphic to the dihedral group D_3. The quotient T_{54}/G_{54} is a single triangle.

The 3–simplex of groups $\mathcal{G}_{54}$ is defined as follows. The group of the 3–simplex is trivial, the triangle groups are $\mathbb{Z}_2$, the edge groups are D_3, the vertex groups are G_{54} and the inclusion maps correspond to inclusions of respective stabilisers in the G_{54}–action on T_{54}; see Figure 15 on the left.

Since $\mathcal{G}_{54}$ is a locally 6–large simplex of groups (see [JŚ06, Section 6]) it is developable by [JŚ06, Theorem 6.1]. Its fundamental group $\overline{G} = \pi_1(\mathcal{G}_{54})$ acts geometrically (with the corresponding stabilisers of faces) on an infinite 3–dimensional systolic pseudomanifold X, whose vertex links are all isomorphic to the torus T_{54}. The quotient of this action is a 3–simplex.

Using Świątkowski's construction we now define a new simplex of groups $\mathcal{G}^*_{54}$, whose fundamental group acts on the barycentric subdivision X' of the pseudomanifold X, transitively on 3–simplices. It is obtained by assigning appropriate groups to faces of a simplex of the barycentric subdivision of the 3–simplex underlying $\mathcal{G}_{54}$. Let G^*_{54} be a group of isometries of the barycentric subdivision T'_{54} of the torus T_{54} generated by reflections with respect to all edges. That is, besides the elements of G_{54} we consider also reflections with respect to lines like, for example, the dashed ones in Figure 15.

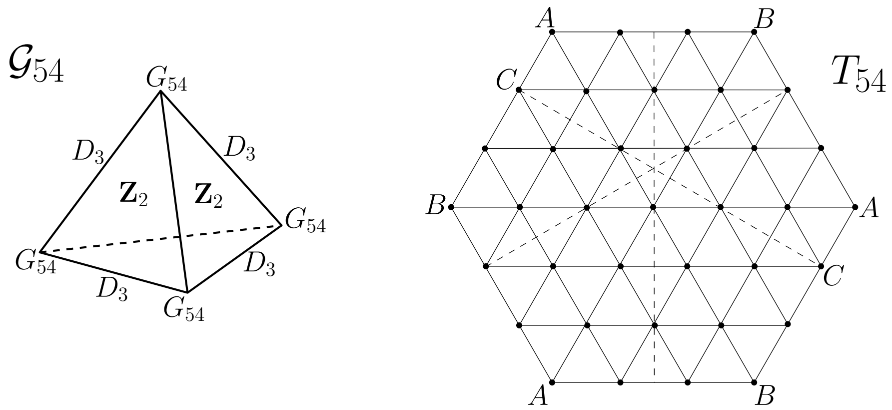

FIGURE 15. The simplex of groups $\mathcal{G}_{54}$ (left), and the local development T_{54} (the development of the triangle of groups in the link of a vertex).

Observe that in this case the stabiliser of a triangle in T'_{54} is trivial, stabilisers of edges are $\mathbb{Z}_2$, and the stabilisers of vertices are as follows: the stabiliser of a barycentre of a triangle of T_{54} is D_3; the stabiliser of a barycentre of an edge of T_{54} is $\mathbb{Z}_2^2$; the stabiliser of a vertex of T_{54} is D_6. The 3–simplex of groups $\mathcal{G}^*_{54}$ is now defined as follows. We consider a 3–simplex in the barycentric subdivision of a tetrahedron P underlying $\mathcal{G}_{54}$, see Figure 16. The 3–simplex group is trivial. The triangle faces groups are $\mathbb{Z}_2$. The assignment of the edge and vertex groups is shown in Figure 16.

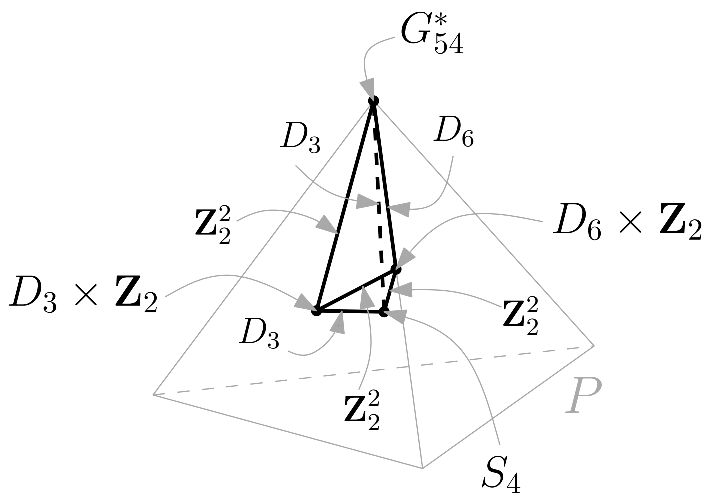

FIGURE 16. The simplex of groups $\mathcal{G}^*_{54}$.

The fundamental group G of $\mathcal{G}^*_{54}$ acts on X' with the corresponding stabilisers of cells and with the quotient being a 3–simplex in the barycentric subdivision of X.

Proposition 8.2. *The complex X does not admit a G–invariant* CAT(0) *metric.*

Proof. Suppose such a metric exists. By the high transitivity of the G–action every edge of X has the same length. It follows that all triangles in X are equilateral. Hence, by the CAT(0) property, angles between edges in triangles are at most $\frac{\pi}{3}$, that is, the angle length of every edge in the link of a vertex of X does not exceed $\frac{\pi}{3}$. Every such link is isomorphic to the barycentric subdivision T'_{54} of T_{54} and the vertex group G^*_{54} acts transitively on edges. Therefore, all the edges in T_{54} have the same length. Consider now the straight line connecting the vertices labeled C in Figure 15. This is a homotopically non-trivial loop in the link of length strictly less that 2π. It follows from the fact that, by the CAT(1) property of the link, every segment of this line contained in a single triangle has length smaller then the length $\frac{\pi}{3}$ of edges of this triangle. This contradicts the fact that the metric is CAT(0). □

It is relatively easy to observe that X contains flats and hence the group G is not hyperbolic [Wie08]. We believe that G acts geometrically on a high dimensional CAT(0) cube complex. It seems that methods developed in the current article provide the only way of constructing low-dimensional models for the classifying spaces $\underline{\underline{E}}G$ and $E_{\mathcal{VAB}}G$. There are other examples of non-hyperbolic systolic groups (of high dimension) to which our theory applies.

References

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR1744486

[Cha08] *Guido's book of conjectures*, Monographies de L'Enseignement Mathématique [Monographs of L'Enseignement Mathématique], vol. 40, L'Enseignement Mathématique, Geneva, 2008. A gift to Guido Mislin on the occasion of his retirement from ETHZ June 2006; Collected by Indira Chatterji. MR2499538

[Che00] Victor Chepoi, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179. MR1748966 (2001a:57004)

[CO15] Victor Chepoi and Damian Osajda, *Dismantlability of weakly systolic complexes and applications*, Trans. Amer. Math. Soc. **367** (2015), no. 2, 1247–1272, DOI 10.1090/S0002-9947-2014-06137-0. MR3280043

[Deg17] Dieter Degrijse, *A cohomological characterization of locally virtually cyclic groups*, Adv. Math. **305** (2017), 935–952, DOI 10.1016/j.aim.2016.10.014. MR3570151

[DP14] Dieter Degrijse and Nansen Petrosyan, *Geometric dimension of groups for the family of virtually cyclic subgroups*, J. Topol. **7** (2014), no. 3, 697–726, DOI 10.1112/jtopol/jtt045. MR3252961

[Els09a] Tomasz Elsner, *Flats and the flat torus theorem in systolic spaces*, Geom. Topol. **13** (2009), no. 2, 661–698, DOI 10.2140/gt.2009.13.661. MR2469526 (2009m:20065)

[Els09b] ———, *Isometries of systolic spaces*, Fund. Math. **204** (2009), no. 1, 39–55, DOI 10.4064/fm204-1-3. MR2507689 (2010g:51005)

[EP13] Tomasz Elsner and Piotr Przytycki, *Square complexes and simplicial nonpositive curvature*, Proc. Amer. Math. Soc. **141** (2013), no. 9, 2997–3004, DOI 10.1090/S0002-9939-2013-11568-6. MR3068952

[ECH+92] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston, *Word processing in groups*, Jones and Bartlett Publishers, Boston, MA, 1992. MR1161694

[FJ87] Martin Farber and Robert E. Jamison, *On local convexity in graphs*, Discrete Math. **66** (1987), no. 3, 231–247, DOI 10.1016/0012-365X(87)90099-9. MR900046 (89e:05167)

[GS91a] S. M. Gersten and H. B. Short, *Rational subgroups of biautomatic groups*, Ann. of Math. (2) **134** (1991), no. 1, 125–158, DOI 10.2307/2944334. MR1114609 (92g:20092)

[GS91b] ———, *Small cancellation theory and automatic groups. II*, Invent. Math. **105** (1991), no. 3, 641–662, DOI 10.1007/BF01232283. MR1117155

[Gro03] M. Gromov, *Random walk in random groups*, Geom. Funct. Anal. **13** (2003), no. 1, 73–146, DOI 10.1007/s000390300002.1978492 (2004j:20088a)

[Gru15] D. Gruber, *Groups with graphical $C(6)$ and $C(7)$ small cancellation presentations*, Trans. Amer. Math. Soc. **367** (2015), no. 3, 2051–2078, DOI 10.1090/S0002-9947-2014-06198-9. MR3286508

[Hag03] F. Haglund, *Complexes simpliciaux hyperboliques de grande dimension*, Prepublication Orsay **71** (2003), preprint, available at `http://www.math.u-psud.fr/~haglund/cpl_hyp_gde_dim.pdf`.

[JŚ06] Tadeusz Januszkiewicz and Jacek Świątkowski, *Simplicial nonpositive curvature*, Publ. Math. Inst. Hautes Études Sci. **104** (2006), 1–85. MR2264834 (2007j:53044)

[JŚ07] ———, *Filling invariants of systolic complexes and groups*, Geom. Topol. **11** (2007), 727–758. MR2302501 (2008d:20079)

[JPL06] Daniel Juan-Pineda and Ian J. Leary, *On classifying spaces for the family of virtually cyclic subgroups*, Recent developments in algebraic topology, Contemp. Math., vol. 407, Amer. Math. Soc., Providence, RI, 2006, pp. 135–145, DOI 10.1090/conm/407/07674. MR2248975 (2007d:19001)

[KPS73] A. Karrass, A. Pietrowski, and D. Solitar, *Finite and infinite cyclic extensions of free groups*, J. Austral. Math. Soc. **16** (1973), 458–466. Collection of articles dedicated to the memory of Hanna Neumann, IV. MR0349850

[Lüc05] Wolfgang Lück, *Survey on classifying spaces for families of subgroups*, Infinite groups: geometric, combinatorial and dynamical aspects, Progr. Math., vol. 248, Birkhäuser, Basel, 2005, pp. 269–322. MR2195456

[Lüc09] ———, *On the classifying space of the family of virtually cyclic subgroups for CAT(0)-groups*, Münster J. of Math. **2** (2009), 201–214.

[LW12] Wolfgang Lück and Michael Weiermann, *On the classifying space of the family of virtually cyclic subgroups*, Pure Appl. Math. Q. **8** (2012), no. 2, 497–555, DOI 10.4310/PAMQ.2012.v8.n2.a6. MR2900176

[Man05] Jason Fox Manning, *Geometry of pseudocharacters*, Geom. Topol. **9** (2005), 1147–1185 (electronic), DOI 10.2140/gt.2005.9.1147. MR2174263 (2006j:57002)

[MW02] Jonathan P. McCammond and Daniel T. Wise, *Fans and ladders in small cancellation theory*, Proc. London Math. Soc. (3) **84** (2002), no. 3, 599–644, DOI 10.1112/S0024611502013424. MR1888425 (2003b:20047)

[Osa14] Damian Osajda, *Small cancellation labellings of some infinite graphs and applications* (2014), preprint, available at `arXiv:1406.5015`.

[Osa15] ———, *Normal subgroups of SimpHAtic groups* (2015), preprint, available at `arXiv:1501.00951`.

[OŚ15] Damian Osajda and Jacek Świątkowski, *On asymptotically hereditarily aspherical groups*, Proc. London Math. Soc. **111** (2015), no. 1, 93–126, DOI 10.1112/plms/pdv021.

[Prz09] Piotr Przytycki, *$\underline{E}G$ for systolic groups*, Comment. Math. Helv. **84** (2009), no. 1, 159–169, DOI 10.4171/CMH/156. MR2466079 (2009k:20100)

[Roe03] John Roe, *Lectures on coarse geometry*, University Lecture Series, vol. 31, American Mathematical Society, Providence, RI, 2003. MR2007488 (2004g:53050)

[Smi11] Stephen D. Smith, *Subgroup complexes*, Mathematical Surveys and Monographs, vol. 179, American Mathematical Society, Providence, RI, 2011. MR2850680 (2012k:20028)

[SC83] V. P. Soltan and V. D. Chepoĭ, *Conditions for invariance of set diameters under d-convexification in a graph*, Kibernetika (Kiev) **6** (1983), 14–18, DOI 10.1007/BF01068561 (Russian, with English summary); English transl., Cybernetics **19** (1983), no. 6, 750–756 (1984). MR765117 (86k:05102)

[Świ06] Jacek Świątkowski, *Simplicial nonpositive curvature* (2006), lecture notes, available at `http://www.math.uni.wroc.pl/~swiatkow/montreal/notes.pdf`.

[TW91] J. Thévenaz and P. J. Webb, *Homotopy equivalence of posets with a group action*, J. Combin. Theory Ser. A **56** (1991), no. 2, 173–181, DOI 10.1016/0097-3165(91)90030-K. MR1092846 (92k:20049)

[Wie08] Jacek Wieszaczewski, *Sympleksy grup o torusowych grupach wierzchołkowych*, M.Sc. thesis, Wrocław University, 2008 (Polish).

[Wis96] Daniel T. Wise, *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*, ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)–Princeton University. MR2694733

[Wis03] ———, *Sixtolic complexes and their fundamental groups* (2003), unpublished manuscript.

Instytut Matematyczny, Uniwersytet Wrocławski, pl. Grunwaldzki 2/4, 50–384 Wrocław, Poland

Institute of Mathematics, Polish Academy of Sciences, Śniadeckich 8, 00–656 Warszawa, Poland

E-mail address: `dosaj@math.uni.wroc.pl`

Institut for Matematiske Fag, Københavns Universitet, Universitetsparken 5, 2100 København Ø, Denmark

E-mail address: `t.prytula@math.ku.dk`

HYPERBOLIC ISOMETRIES AND BOUNDARIES OF SYSTOLIC COMPLEXES

TOMASZ PRYTUŁA

Abstract. Given a group G acting geometrically on a systolic complex X and a hyperbolic isometry $h \in G$, we study the associated action of h on the systolic boundary ∂X. We show that h has a canonical pair of fixed points on the boundary and that it acts trivially on the boundary if and only if it is virtually central. The key tool that we use to study the action of h on ∂X is the notion of a K–displacement set of h, which generalises the classical minimal displacement set of h. We also prove that systolic complexes equipped with a geometric action of a group are almost extendable.

1. Introduction

A systolic complex is a simply connected simplicial complex whose vertex links satisfy a certain combinatorial condition called 6–*largeness*. The condition of 6–largeness serves as an upper bound for the combinatorial curvature, and thus systolic complexes may be seen as combinatorial analogues of metric spaces of nonpositive curvature, the so-called CAT(0) spaces. Systolic complexes were first introduced in [Che00] under the name of *bridged complexes*, although their 1–skeleta had appeared much earlier in metric graph theory (see e.g., [SC83]). In this article we are interested in systolic complexes that are equipped with a geometric action of a group. Any such group is called a *systolic group*. The theory of systolic complexes and groups, as developed in [JŚ06], is to a large extent parallel to the theory of CAT(0) spaces and groups. In particular, over the last fifteen years many of the nonpositive curvature-like properties of systolic complexes have been established (see [Che00, JŚ06, Els09a, Els09b, Els09c, OP16] and references therein). On the other hand, a combinatorial approach led to constructions of examples of systolic groups whose behaviour is very different from the classical nonpositively curved groups [JŚ06].

An important invariant of a CAT(0) space X is its *boundary at infinity* ∂X. The boundary is a topological space which, as a set, consists of equivalence classes of geodesic rays in X, such that asymptotic rays are equivalent. One topologises it in a

Date: April 4, 2017.

2010 *Mathematics Subject Classification.* 20F67 (Primary), 20F65, 20F69 (Secondary).

Key words and phrases. Systolic complex, boundary at infinity, hyperbolic isometry.

way that two geodesic rays are 'close' if they fellow travel 'long time'. Any G–action by isometries on X gives rise to a G–action by homeomorphisms on ∂X. It turns out that many algebraic properties of a group are reflected in topological properties of the boundary and in the action itself.

In this article we study this correspondence in the setting of systolic complexes. The boundary for systolic complexes was constructed in [OP09]. The construction is similar to the one for CAT(0) spaces, however it is much more technical. The points of the systolic boundary are also represented by geodesic rays in (the 1–skeleton of) a systolic complex X, but not every geodesic ray in X gives a point in the boundary: in order to ensure good properties of the boundary, a choice of a certain subclass of geodesics was necessary. This is mainly due to the fact that arbitrary geodesics in a systolic complex do not satisfy any form of the Fellow Traveller Property (indeed, two geodesics of length D with the same endpoints may get $\frac{D}{2}$ apart). In [OP09] the authors introduce *good geodesics* and *good geodesic rays*, and define the systolic boundary ∂X as a set of equivalence classes of good geodesic rays in X. The topology on ∂X is defined analogously as in the CAT(0) case. Both good geodesics and good geodesic rays are preserved by simplicial automorphisms of X, and therefore any simplicial G–action on X induces a G–action (by homeomorphisms) on ∂X. Intuitively, a good geodesic ray is a geodesic ray which, whenever contained in a flat F, follows the CAT(0) geodesic in F. In particular good geodesic rays have the desired metric properties, similar to those of geodesic rays in CAT(0) spaces.

An isometry (i.e., a simplicial automorphism) h of a systolic complex X is *hyperbolic* if it does not fix any simplex of X. Note that if G acts geometrically on X then every infinite order element of G is a hyperbolic isometry of X. The main point of this article is to study the associated action of h on the systolic boundary ∂X. We start by determining when this action is trivial (i.e., when h acts as the identity on ∂X). Denote by $C_G(h)$ the centraliser of h in G. The following is the systolic analogue of a result of K. Ruane for CAT(0) spaces [Rua01].

Theorem 1 (Theorem 4.5)**.** *Let G be a group acting geometrically on a systolic complex X, and let $h \in G$ be a hyperbolic isometry. Then h acts trivially on the boundary ∂X if and only if the centraliser $C_G(h)$ has finite index in G.*

The canonical object used to study the action of h on X is the *minimal displacement set* of h, which is a subcomplex of X spanned by all the vertices which are moved by h the minimal (combinatorial) distance. This distance is called the *translation length* of h and it is denoted by $L(h)$. Due to a 'coarse nature' of ∂X, in order to study the action of h on ∂X it is convenient to replace the minimal displacement set by its coarse equivalent – the *K–displacement set* of h, for some $K \geqslant L(h)$, which is a subcomplex of X spanned by all the vertices that are moved by h the distance at most K. The K–displacement set has all the desired (from our point of view) features of the minimal displacement set, while it has the advantage of being more flexible as one can let K vary.

The proof of Theorem 1 is based on the interplay between K–displacement sets of h (for different values of K) and the centraliser $C_G(h)$. In particular, the 'if' direction essentially boils down to showing the following two facts:

(1) Any point in ∂X represented by a geodesic ray that lies inside some K–displacement set of h is fixed by h.
(2) The centraliser $C_G(h)$ acts cocompactly on any K–displacement set of h.

The 'only if' direction is more involved. In this case we are given the information about the action on the boundary, and we need to extract the information about the action on the complex. For this we need X to satisfy the following property. We say that X is *almost extendable* if there exists a constant $E \geqslant 0$ such that for every pair of vertices x, y in X there is a good geodesic ray issuing from x and passing within distance E from y. The following theorem is also of independent interest.

Theorem 2 (Theorem 3.2)**.** *Let X be a noncompact systolic complex, on which a group G acts geometrically. Then X is almost extendable.*

The proof of this theorem relies on the study of topology at infinity of systolic complexes. It is similar to the proof of an analogous theorem for CAT(0) spaces [Ont05]. The main difference is that our proof uses the notion of connectedness at infinity, whereas the one in [Ont05] uses cohomology with compact supports. The key fact is that in the setting above, the complex X is not 1–connected at infinity (see [Osa07]).

In the second part of the article we consider arbitrary hyperbolic isometries of X (not necessarily the virtually central ones). One can still ask whether such an isometry h has any fixed points in ∂X. In the setting of CAT(0) spaces, a hyperbolic isometry h has an *axis*, that is, an h–invariant geodesic line, and this axis determines two fixed points of h in ∂X. In our situation h also has a kind of axis (see [Els09b]), but unfortunately this axis does not have to determine an h–invariant good geodesic. In fact, an h–invariant good geodesic may not exist. However, we do prove that h has a pair of fixed points in ∂X.

Theorem 3 (Proposition 6.2). *Let G be a group acting geometrically on a systolic complex X and let $h \in G$ be a hyperbolic isometry. Then:*

(1) *there exist points $h^{-\infty}$ and $h^{+\infty}$ in the boundary ∂X which are fixed by h,*
(2) *for any vertex $x \in X$ we have $(h^n \cdot x)_n \to h^{+\infty}$ and $(h^{-n} \cdot x)_n \to h^{-\infty}$ as $n \to \infty$ in the compactification $\overline{X} = X \cup \partial X$.*

The second statement shows that $h^{+\infty}$ and $h^{-\infty}$ are in a certain sense the canonical fixed points of h. To find $h^{+\infty}$ and $h^{-\infty}$ we construct an 'almost h–invariant' good geodesic in X, by which we mean a geodesic that is contained in some K–displacement set of h. This requires analysing the construction of good geodesics. We go through the steps of the construction and show that given any two vertices x and y in the minimal displacement set of h, a good geodesic between x and y is contained in a K–displacement set where K is independent of distance between x and y. Then we construct a bi-infinite good geodesic as a limit of finite good geodesics contained in the K–displacement set.

In order to prove the second part of Theorem 3 we also study good geodesics contained in the flats of X. In particular, we give a simple criterion for a geodesic contained in a flat to be a good geodesic and we show that any geodesic that is good in the flat is also good in the complex X.

We believe that the results presented in this article may be used in the further study of systolic groups via their boundaries. Theorems 1 and 3 are the first steps in analysing the dynamics of the action of h on ∂X, which in the CAT(0) setting plays the key role in e.g., [PS09], where the topology of ∂X is related to splittings of G over 2–ended subgroups. Theorem 2 seems to be of a more general nature; its CAT(0) counterpart has been widely used in the study of CAT(0) groups and boundaries.

Organisation. The article consists of an introductory Section 2, where we give background on systolic complexes and boundaries, and of the two main parts. In the first part, which occupies Sections 3 and 4, we prove Theorem 2 and after establishing basic facts about K–displacement sets we give a proof of Theorem 1. In the second part (Sections 5 and 6) we first sketch the construction of good geodesics, and then we prove Theorem 3.

Acknowledgements. I would like to thank Piotr Przytycki for suggesting the problem and for many helpful discussions. I would also like to thank Damian Osajda for helpful discussions and comments, in particular leading to the proof of Theorem 2. Finally, I thank Dieter Degrijse and Damian Osajda for a careful proofreading of the manuscript and help in improving the exposition.

The author was supported by the Danish National Research Foundation through the Centre for Symmetry and Deformation (DNRF92).

2. Systolic complexes and their boundaries

In this section we give some background on systolic complexes and their boundaries. We also fix the terminology and notation that is used throughout the article.

2.1. **Systolic simplicial complexes.** Let X be a simplicial complex. We assume that X is finite dimensional and uniformly locally finite, i.e., there is a uniform bound on the degree of vertices in X. We equip X with the CW–topology, and always treat it as a topological space (we do not make a distinction between an abstract simplicial complex and its geometric realisation). Let $X^{(k)}$ denote the k–skeleton of X. In particular $X^{(0)}$ is the vertex set of X. For any subset $A \subset X^{(0)}$, a subcomplex *spanned* by A is the largest subcomplex of X that has A as its vertex set. We denote this subcomplex by $\mathrm{span}(A)$. A map $f\colon X \to Y$ of simplicial complexes is *simplicial* if $f(X^{(0)}) \subset Y^{(0)}$ and whenever vertices $x_0, x_1, \ldots, x_n$ span a simplex of X then their images $f(x_0), f(x_1), \ldots, f(x_n)$ span a simplex of Y. Note that a simplicial map is continuous, and in particular a simplicial automorphism is a homeomorphism of X.

In this article we will be particularly interested in metric aspects of simplicial complexes.

Definition 2.1. We endow the vertex set $X^{(0)}$ with a metric, where the distance $d(x, y)$ between vertices x and y is defined to be the combinatorial distance in the 1–skeleton, i.e., the minimal number of edges of an edge–path joining x and y.

For two subcomplexes $A, B \subset X$, we define the distance $d(A, B)$ to be the minimal distance between vertices $a \in A$ and $b \in B$.

Whenever we refer to the *metric* on X we mean the metric on $X^{(0)}$ defined above. Consequently, a *geodesic* in X is a sequence of vertices $(v_0, v_1, \ldots, v_n)$ such that for any $0 \leqslant i, j \leqslant n$ we have $d(v_i, v_j) = |j - i|$. Analogously we define a geodesic that is infinite in one or both ends. In the first case we call it a *geodesic ray*, in the second case: a *geodesic line* or a *bi-infinite geodesic*. Observe that a simplicial map is 1–Lipschitz and any simplicial automorphism is an isometry of X.

We now briefly recall the notions needed to define systolic complexes. We say that X is *flag* if every set of vertices of X pairwise connected by edges spans a simplex of X. A flag simplicial complex is completely determined by its 1–skeleton $X^{(1)}$ or, equivalently, by its vertex set $X^{(0)}$ and the metric d defined above. For a vertex $v \in X$ the *link* of v is a subcomplex $\mathrm{Lk}(v, X)$ of X, that consists of all the simplices of X that do not contain v, but together with v span a simplex of X. A *cycle* in X is the image of a simplicial map $f\colon S^1 \to X$ from the triangulation of the 1–sphere to X. A cycle is *embedded* if f is injective. Let γ be an embedded cycle. The *length* of γ, denoted by $|\gamma|$ is the number of edges of γ. A *diagonal* of γ is an edge in X that connects two nonconsecutive vertices of γ.

We are ready to define systolic complexes. Our main reference for the theory of systolic complexes is [JŚ06].

Definition 2.2. Given a natural number $k \geqslant 6$, a simplicial complex X is *k–large* if every embedded cycle γ in X with $4 \leqslant |\gamma| < k$ has a diagonal. We say that X is *∞–large* if it is k–large for every $k \geqslant 6$.

Definition 2.3. A simplicial complex X is *k–systolic* if it is simply connected and if for every vertex $v \in X$ the link $\mathrm{Lk}(v, X)$ is flag and k–large. If $k = 6$ then we abbreviate 6–*systolic* to *systolic.*

A k–systolic complex is flag and k–large [JŚ06, Proposition 1.4 and Fact 1.2(4)]. Note that if $k \leqslant m$ then 'm–systolic' implies 'k–systolic'. In this article we will be interested in the (most general) case of $k = 6$. This case is of particular importance in the theory, as for $k \geqslant 7$ one shows that k–systolic complexes are δ–hyperbolic [JŚ06, Theorem 2.1].

The condition of k–largeness, when applied to the link of a vertex $v \in X$, serves as a kind of upper bound for the curvature around v. In particular complexes with 6–large links are called complexes of *simplicial nonpositive curvature* (SNPC). Consequently, systolic complexes can be thought of as simplicial analogues of CAT(0) metric spaces.

Definition 2.4. Let $v \in X$ be a vertex and let n be a positive integer. Define the *ball of radius n centred at v* by $B_n(v, X) = \mathrm{span}\{x \in X^{(0)} \mid d(x, v) \leqslant n\} \subset X$. Define the *sphere of radius n centred at v* by $S_n(v, X) = \mathrm{span}\{x \in X^{(0)} \mid d(x, v) = n\}$. For a subcomplex $A \subset X$ define the *ball of radius n around A* by

$$B_n(A, X) = \bigcup_{v \in A^{(0)}} B_n(v, X).$$

We also refer to $B_n(A, X)$ as an *n–neighbourhood of A* in X.

A subcomplex $A \subset X$ is *convex* if for every two vertices $x, y \in A$ any geodesic between x and y in X is contained in A. Note that since geodesics in X are not necessarily unique, a subcomplex $A \subset X$ can be isometrically embedded and not convex.

Proposition 2.5. *Let X be a systolic complex. Then the following hold:*

(1) *For any convex subcomplex $A \subset X$ the ball $B_n(A, X)$ is convex and contractible* [JŚ06, Corollary 7.5]. *In particular for any vertex $v \in X$ the ball $B_n(v, X)$ is convex and contractible.*
(2) *The complex X is contractible* [JŚ06, Theorem 4.1(1)].

We finish this section with some terminology regarding group actions on simplicial complexes. Let G be a (discrete) group acting on a simplicial complex X. We assume that G acts via simplicial automorphisms. We say that the action is:

- *proper* if for every vertex $v \in X$ the stabiliser G_v is finite,
- *cocompact* if there exists a compact subset $K \subset X$ such that $G \cdot K = X$,
- *geometric* if it is proper and cocompact.

A group is called *systolic* if it acts geometrically on a systolic complex.

2.2. **Boundaries of systolic complexes.** Given a (noncompact) systolic complex X one can define the *boundary at infinity* (or shortly the *boundary*) ∂X of X. Analogously to the cases of δ–hyperbolic and CAT(0) spaces, the boundary for systolic complexes is given by a set of equivalence classes of geodesic rays, such that asymptotic rays are equivalent. In this section we give the definition of the boundary and briefly discuss its key features that are needed in this article. For more details we refer the reader to [OP09].

The main difference from δ–hyperbolic and CAT(0) cases is that, instead of arbitrary geodesic rays, to define the boundary one uses a canonically defined subcollection of the so-called *good geodesic rays*. To define good geodesic rays one first defines *good geodesics*. The actual definition of good geodesics (which is quite involved) is needed only in Section 6, and therefore we give this definition in Section 5.

In order to follow the arguments in Sections 3 and 4 it is enough to know that a good geodesic is a certain geodesic in X, and that the subclass of good geodesics has the following properties:

(1) for any two vertices there exists a (not necessarily unique) good geodesic joining these vertices,
(2) any subgeodesic of a good geodesic is a good geodesic,
(3) any simplicial automorphism of X maps good geodesics to good geodesics.

A *good geodesic ray* is a geodesic ray, such that any of its finite subgeodesics is a good geodesic. By (3) any simplicial automorphism of X maps good geodesic rays to good geodesic rays.

Let $\mathcal{R}$ denote the set of all good geodesic rays in X. For a vertex $O \in X$ let $\mathcal{R}_O$ denote the set of all good geodesic rays starting at O.

Definition 2.6. Let X be a systolic complex. Define the *boundary of* X to be the set $\partial X = \mathcal{R}/\sim$ where for rays $\eta = (v_0, v_1, \ldots)$ and $\xi = (w_0, w_1, \ldots)$ we have $\eta \sim \xi$ if and only if there exists $K \geqslant 0$, such that for every $i \geqslant 0$ we have $d(v_i, w_i) \leqslant K$.

Define the *boundary of* X *with respect to the basepoint* O to be the set $\partial_O X = \mathcal{R}_O/\sim$, where $\sim$ is the same equivalence relation as above. In both cases let $[\eta]$ denote the equivalence class of η.

For any vertex $O \in X$ there is a bijection $\partial X \to \partial_O X$ [OP09, Corollary 3.10]. In particular this means that for every geodesic ray $\eta \subset X$ and for every vertex $O \in X$ there is a ray $\xi \subset X$ starting at O such that $[\eta] = [\xi]$ in ∂X. This fact will be used many times in this article. The set $\overline{X} = X \cup \partial_O X$ can be equipped with a topology that extends the standard topology on X, and turns $\overline{X}$ into a compact topological space [OP09, Propositions 4.4 and 5.3]. For any two vertices $O, O' \in X$ there is a homeomorphism between $X \cup \partial_O X$ and $X \cup \partial_{O'} X$ [OP09, Lemma 5.5]. Any simplicial action of a group on X extends to an action by homeomorphisms on $\overline{X}$ [OP09, Theorem A(4)].

In this article we will mostly be concerned with the induced action on the boundary, not on the entire $\overline{X}$. Moreover, we will be interested in the action on the boundary seen as a set, not as a topological space. For this, we can use a slightly simpler definition.

Definition 2.7. Suppose that a group G acts simplicially on X. We define an action of G on the set ∂X as follows. Let $[\eta] \in \partial X$ where $\eta = (v_0, v_1, \ldots)$. Then define $g \cdot [\eta] = [g \cdot \eta]$ where $g \cdot \eta = (g \cdot v_0, g \cdot v_1, \ldots)$. It is straightforward to check that this is well defined and it defines an action of G on ∂X.

One can also verify, that via the bijection $\partial X \to \partial_O X$ the action described above agrees with the action on $\partial_O X$ defined in [OP09].

We conclude this section with certain metric properties of good geodesics. The following is a crucial property, which can be seen as a coarse version of CAT(0) inequality for good geodesics.

Theorem 2.8. [OP09, Corollary 3.4] *Let* $(v_0, v_1, v_2, \ldots, v_n)$ *and* $(w_0, w_1, w_2, \ldots, w_m)$ *be good geodesics in a systolic complex* X *such that* $v_0 = w_0$. *Then for any* $0 \leqslant c \leqslant 1$ *we have*

$$d(v_{\lfloor cn \rfloor}, w_{\lfloor cm \rfloor}) \leqslant c \cdot d(v_n, w_m) + D,$$

where D *is a universal constant.*

This leads to the following corollary, which will be very useful to us.

Corollary 2.9. *Let* $\eta = (v_0, v_1, v_2, \ldots)$ *and* $\xi = (w_0, w_1, w_2, \ldots)$ *be good geodesic rays in a systolic complex* X*. If* $[\eta] = [\xi]$ *in* ∂X *then for every* $i \geqslant 0$ *we have*

$$d(v_i, w_i) \leqslant d(v_0, w_0) + 2D + 1,$$

where D *is the constant appearing in Theorem 2.8.*

Proof. Since $[\eta] = [\xi]$, there is a constant $K \geqslant 0$ such that for all i we have $d(v_i, w_i) \leqslant K$. Fix $i \geqslant 0$, pick $n > K$ and let $z = (z_0 = v_0, z_1, z_2, \ldots, w_{ni})$ be a good geodesic joining v_0 and w_{ni}. By Theorem 2.8 applied to $(v_0, v_1, v_2, \ldots, v_{ni})$ and z we have

$$d(v_i, z_i) \leqslant \frac{1}{n} d(v_{ni}, w_{ni}) + D \leqslant \frac{K}{n} + D \leqslant 1 + D. \tag{2.1}$$

Applying Theorem 2.8 to z and $(w_0, w_1, w_2, \ldots, w_{ni})$ (with the direction reversed) we obtain that $d(z_i, w_i) \leqslant d(v_0, w_0) + D$. This, together with (2.1) and the triangle inequality gives the claim. □

3. Almost extendability of systolic complexes

In this section we study a property of metric spaces called the *almost extendability*. This property can be defined for arbitrary geodesic metric spaces. The definition we present is adjusted to the setting of systolic complexes.

Definition 3.1. A systolic complex X is *almost extendable*, if there exists a constant $E \geqslant 0$ such that for any two vertices x and y of X, there is a good geodesic ray starting at y and passing within distance E from x.

It is easy to construct systolic complexes (in fact, trees) that are not almost extendable. For example, let T denote the half-line $\mathbb{R}_+$ with the interval of length n attached to every integer $n \in \mathbb{R}_+$. The standard triangulation turns T into a systolic complex in which every combinatorial geodesic is a good geodesic. One can easily see that T is not almost extendable. When we equip a systolic complex with a geometric action of a group then the situation changes.

Theorem 3.2. *Let* X *be a noncompact systolic complex, on which a group* G *acts geometrically. Then* X *is almost extendable.*

The analogous theorem is true in the CAT(0) setting [Ont05, Theorem B], and it is an exercise in the setting of δ–hyperbolic groups (see [Ont05]). Our proof is similar to the one for CAT(0) spaces, however, it can be seen as more direct. The main difference is that instead of cohomology with compact supports, our proof uses the notion of connectedness at infinity. We begin by recalling this notion.

Definition 3.3. Let Y be a topological space and let $n \geqslant -1$ be an integer. We say that Y is *n–connected at infinity* if for every $-1 \leqslant k \leqslant n$ the following condition holds: for every compact set $K \subset Y$ there exists a compact set $L \subset Y$ such that $K \subset L$ and every map $S^k = \partial B^{k+1} \to Y \setminus L$ extends to a map $B^{k+1} \to Y \setminus K$.

For $k = -1$ we define $S^{-1} = \emptyset$ and $B^0 = \{*\}$. In particular Y is (-1)–connected at infinity if and only if it is not compact.

Note that if Y is a simplicial complex then (in view of the Simplicial Approximation Theorem) in the above definition it is enough to consider only simplicial maps.

The following theorem of D. Osajda is the crucial ingredient in the proof of Theorem 3.2.

Theorem 3.4. [Osa07, Theorem 3.2] *Let X be a noncompact systolic complex, on which a group G acts geometrically. Then X is not 1–connected at infinity.*

Proof of Theorem 3.2. First we show that it is enough to prove the following claim.

Claim 1. *Let p be a fixed vertex. Then there exists a constant E' such that, for any $g \in G$ there is a good geodesic ray starting at p and passing within E' from $g \cdot p$.*

Indeed, let x and y be arbitrary vertices of X. By cocompactness there exists $R > 0$ and elements $g_1, g_2 \in G$ such that we have $d(g_1 \cdot p, x) \leqslant R$ and $d(g_2 \cdot p, y) \leqslant R$. By Claim 1 there exists a good geodesic ray η starting at p and passing within E' from $g_1^{-1} g_2 \cdot p$. Then the ray $g_1 \cdot \eta$ starts at $g_1 \cdot p$ and passes within E' from $g_2 \cdot p$, and hence it passes within $E' + R$ from y. Now let ξ be a good geodesic ray starting at x and such that $[\xi] = [g_1 \cdot \eta]$. Write $g_1 \cdot \eta = (g_1 \cdot p = v_0, v_1, \ldots)$ and $\xi = (x = w_0, w_1, \ldots)$. Then by Corollary 2.9 for every $i \geqslant 0$ we have

$$d(v_i, w_i) \leqslant d(g_1 \cdot p, x) + 2D + 1 \leqslant R + 2D + 1.$$

Since $g_1 \cdot \eta$ passes within $E' + R$ from y, we have that ξ passes within $E' + R + R + 2D + 1$ from y. Therefore Claim 1 implies the theorem (with constant $E = E' + R + R + 2D + 1$).

The rest of the proof is devoted to proving Claim 1. We need a little preparation. In what follows, for a good geodesic or a good geodesic ray η we will denote its vertices by $\eta(i)$, for $i \geqslant 0$, i.e., $\eta = (\eta(0), \eta(1), \eta(2), \ldots)$. In other words, the geodesic η may be seen as a map $\mathbb{N} \to X$. We still treat η as a subset of X; the above notation is introduced only to simplify the exposition.

For a good geodesic η let l_η denote the supremum of natural numbers l, such that η can be extended to a good geodesic on the interval $[0, l] = \{0, 1, \ldots, l\} \subset \mathbb{N}$. Note that l_η does not have to be attained, in that case we write $l_\eta = \infty$. Observe that if $l_\eta < \infty$ then there is an extension of η to the interval $[0, l_\eta]$. If $l_\eta = \infty$ then by the fact that $X \cup \partial X$ is compact, there is an extension of η to the interval $[0, \infty)$, i.e., the geodesic η can be extended to a good geodesic ray. For vertices $x, y \in X$ let $[\![x, y]\!]$ denote a good geodesic between these two vertices. Note that such a geodesic in general is not unique.

Now we begin the proof of Claim 1. We proceed by contradiction. Assume that Claim 1 does not hold, then we have the following:

$(*)$ For every $r > 0$ there exists $g_r \in G$ such that for every vertex $x \in B_r(g_r \cdot p, X)$ we have $l_{[\![p,x]\!]} < \infty$ for every good geodesic $[\![p, x]\!]$.

Claim 2. *For every $r > 0$ we have*

$$\sup\{l_{[\![p,x]\!]} \mid [\![p, x]\!] \text{ where } x \in B_r(g_r \cdot p, X)\} < \infty.$$

(The supremum is taken over all possible good geodesics that start at p and end at a vertex of $B_r(g_r \cdot p, X)$.)

To prove Claim 2, assume conversely that there exists a sequence of good geodesics $([\![p, x_i]\!])_i$ with $x_i \in B_r(g_r \cdot p, X)$, such that $l_{[\![p,x_i]\!]} \to \infty$ as $n \to \infty$. Let η_i denote a good geodesic extending $[\![p, x_i]\!]$ to the interval $[0, l_{[\![p,x_i]\!]}]$ (we choose one for each i). Using a diagonal argument, out of the sequence $(\eta_i)_i$ one constructs an infinite geodesic ray ξ that issues from p, and such that for any interval $[0, l]$ we have $\xi\big|_{[0,l]} = \eta_i$ for some $i = i_l$ (cf. [OP09, Proposition 5.3]). In particular, the ray ξ intersects the ball $B_r(g_r \cdot p, X)$, which contradicts $(*)$.

Claim 3. *For every $r > 0$ there exists $r' > r$ such that for every vertex $y \in X \setminus B_{r'}(p, X)$, every good geodesic $[\![p, y]\!]$ misses the ball $B_r(g_r \cdot p, X)$, i.e., we have*

$$[\![p, y]\!] \cap B_r(g_r \cdot p, X) = \emptyset.$$

Note that we have $p \notin B_r(g_r \cdot p, X)$, for otherwise we would get a contradiction with $(*)$ as there always is a geodesic ray issuing from p (since X is noncompact). Let $r' = \sup\{l_{[\![p,x]\!]} \mid [\![p, x]\!] \text{ where } x \in B_r(g_r \cdot p, X)\}$. Then the claim follows from the definition of r'.

Claim 4. *The complex X is 1–connected at infinity.*

First observe that since X is noncompact, it is (-1)–connected at infinity. Let $K \subset X$ be a compact subset. Take $M > 0$ such that $K \subset B_M(p, X)$ and consider the ball $B_{M+D+2}(p, X)$, where D is the constant appearing in Theorem 2.8 (the reason why we need to pass to the larger ball will become clear later on).

Pick $r > M + D + 2$. By Claim 3 (after 'translating its statement by g_r^{-1}') there exists $r' > r$ such that every good geodesic joining a vertex $y \in X \setminus B_{r'}(g_r^{-1} \cdot p, X)$

with $g_r^{-1} \cdot p$, misses the ball $B_r(p, X)$. Set $L = B_{r'}(g_r^{-1} \cdot p, X)$. By construction we have $K \subset L$, and $g_r^{-1} \cdot p \in L \setminus K$.

Suppose $f \colon S^0 \to X \setminus L$ is a simplicial map. Let v_1 and v_2 be the two vertices in the image of f. For $i \in \{0, 1\}$ let η_i be a good geodesic joining $g_r^{-1} \cdot p$ with v_i. Both η_i miss the ball $B_r(p, X)$ (and hence they miss K) and therefore their union defines a map $F \colon B^1 \to X \setminus K$ that extends f. This shows that X is 0–connected at infinity.

Now let $f \colon S^1 \to X \setminus L$ be a simplicial map. Let $(v_0, v_1, \ldots, v_n, v_{n+1} = v_0)$ be the vertices of the image of f appearing in this order, i.e., for all i vertices v_i and v_{i+1} are adjacent. For every $i \in \{0, \ldots, n\}$ let η_i be a good geodesic joining $g_r^{-1} \cdot p$ and v_i. Observe that no η_i intersects the ball $B_r(p, X)$. We will use η_i's to construct the required extension of f to the disk B^2.

For any i consider the cycle $\alpha_i \subset X$ which is the union

$$\alpha_i = \eta_i \cup \eta_{i+1} \cup [v_i, v_{i+1}].$$

We will show that α_i can be contracted to a point in its $(D + 2)$–neighbourhood. First note that either η_i and η_{i+1} have the same length, or their lengths differ by 1.

In the first case put $k = d(g_r^{-1} \cdot p, v_i) = d(g_r^{-1} \cdot p, v_{i+1})$. Since η_i and η_{i+1} start at the same vertex and end at vertices that are adjacent, it follows from Theorem 2.8 that for any $j \in \{0, 1, \ldots, k\}$ we have

$$d(\eta_i(j), \eta_{i+1}(j)) \leqslant D + 1.$$

For any $j \in \{0, \ldots, k\}$ let β_j^i be a geodesic between $\eta_i(j)$ and $\eta_{i+1}(j)$ (note that β_0^i is the vertex $g_r^{-1} \cdot p$ and β_k^i is the edge $[v_i, v_{i+1}]$). Now for every $j \in \{0, \ldots, k - 1\}$ consider a cycle γ_j^i defined as

$$\gamma_j^i = \beta_j^i \cup [\eta_{i+1}(j), \eta_{i+1}(j + 1)] \cup \beta_{j+1}^i \cup [\eta_i(j), \eta_i(j + 1)].$$

By construction γ_j^i is contained in the ball $B_{D+2}(\eta_i(j), X)$ and therefore it can be contracted inside $B_{D+2}(\eta_i(j), X)$, as balls in X are contractible (see Proposition 2.5.(1)). These contractions of γ_j^i for all $j \in \{0, \ldots, k\}$ form a contraction of α_i inside the ball $B_{D+2}(\eta_i, X)$ around the geodesic η_i. (Formally, by a *contraction* we mean a simplicial map from a simplicial disk $f \colon B^2 \to B_{D+2}(\eta_i, X)$ such that f maps the boundary ∂B^2 isomorphically onto α_i.)

In the second case, assume that η_{i+1} is longer than η_i, i.e., we have $d(g_r^{-1} \cdot p, v_i) = k$ and $d(g_r^{-1} \cdot p, v_{i+1}) = k + 1$. In this case the concatenation $\eta_i \cup [v_i, v_{i+1}]$ is a geodesic. Then it follows from [JŚ06, Lemma 7.7] that $\eta_{i+1}(k)$ and v_i are adjacent, and therefore $\eta_{i+1}(k), v_i$ and v_{i+1} span a 2–simplex. Now η_i and $\eta_{i+1}\big|_{[0,k]}$ are of the same length and their endpoints v_i and $\eta_{i+1}(k)$ are adjacent. Proceeding as in the first case we obtain a contraction of the cycle

$$\eta_i \cup \eta_{i+1}\big|_{[0,k]} \cup [v_i, \eta_{i+1}(k)]$$

inside the ball $B_{D+2}(\eta_i, X)$. Adding the 2–simplex $[v_i, \eta_{i+1}(k), v_{i+1}]$ we obtain the desired contraction of $\alpha_i = \eta_i \cup \eta_{i+1} \cup [v_i, v_{i+1}]$.

Finally, contractions of α_i for all $i \in \{0, \ldots, n\}$ form the contraction of $(v_0, \ldots v_n, v_0)$ that is performed in the $(D+2)$–neighbourhood of the union of all η_i's. Since every η_i misses the ball $B_r(p, X)$, the $(D+2)$–neighbourhood of η_i misses the ball $B_M(p, X)$ as $r > M + D + 2$, and hence it misses K as $K \subset B_M(p, X)$. We conclude that the constructed contraction of $(v_0, \ldots, v_n, v_0)$ defines the extension of f that misses K. This finishes the proof of Claim 4.

This gives a contradiction with Theorem 3.4 and hence proves Claim 1. □

4. Isometries acting trivially on the boundary

In this section, given a group G acting geometrically on a systolic complex X, we investigate which elements of G act trivially on the boundary ∂X. Before proving the main theorem which characterises such elements in terms of their centralisers in G, we introduce the terminology and briefly discuss the tools needed in the proof.

4.1. **Hyperbolic isometries and their K–displacement sets.** Let h be an isometry (i.e., a simplicial automorphism) of a systolic complex X. We say that h is *hyperbolic* if it does not fix any simplex of X. If h is hyperbolic, then any of its powers is hyperbolic as well ([Els09b]). To such h one associates the *displacement function* $d_h \colon X^{(0)} \to \mathbb{N}$ defined as $d_h(x) = d(x, h \cdot x)$. The minimum of d_h (which is always attained) is called the *translation length* of h and is denoted by $L(h)$.

Definition 4.1. Let h be a hyperbolic isometry of a systolic complex X. The *minimal displacement set* $\mathrm{Min}(h)$ is the subcomplex of X spanned by all the vertices of X which are moved by h the minimal distance, i.e.:

$$\mathrm{Min}(h) = \mathrm{span}\{x \in X^{(0)} \mid d(x, h \cdot x) = L(h)\}.$$

More generally, for a natural number $K \geqslant L(h)$ define the *K–displacement set* as

$$\mathrm{Disp}_K(h) = \mathrm{span}\{x \in X^{(0)} \mid d(x, h \cdot x) \leqslant K\}.$$

Clearly we have $\mathrm{Disp}_K(h) \subset \mathrm{Disp}_{K'}(h)$ for $K \leqslant K'$ and $\mathrm{Disp}_{L(h)}(h) = \mathrm{Min}(h)$.

Let us mention that $\mathrm{Min}(h)$ is a systolic complex on its own, and its inclusion into X is an isometric embedding ([Els09b]). We do not know whether the same is true for

$\mathrm{Disp}_K(h)$ for $K > L(h)$. In this article we are interested only in the coarse-geometric behaviour of $\mathrm{Disp}_K(h)$.

Observe that if $x \in X$ is a vertex such that $d(x, \mathrm{Disp}_K(h)) \leqslant C$ for some $C \geqslant 0$, then by the triangle inequality we have $d(x, h \cdot x) \leqslant K + 2C$. This means that $B_C(\mathrm{Disp}_K(h), X) \subseteq \mathrm{Disp}_{K+2C}(h)$. In the presence of a geometric action of a group, the (partial) converse also holds.

Lemma 4.2. *Let G be a group acting geometrically on a systolic complex X and suppose that $h \in G$ is a hyperbolic isometry. Pick $K \geqslant L(h)$. Then there exists $C > 0$ such that for all $K' \leqslant K$ we have $\mathrm{Disp}_K(h) \subset B_C(\mathrm{Disp}_{K'}(h), X)$.*

The lemma is an easy consequence of the following theorem.

Theorem 4.3. *Let G act geometrically on a systolic complex X, and let $h \in G$ be a hyperbolic isometry. Then for any $K \geqslant L(h)$ the centraliser $C_G(h) \subset G$ acts geometrically on the subcomplex $\mathrm{Disp}_K(h) \subset X$.*

Proof. The proof is a verbatim translation of K. Ruane's proof of a similar result for CAT(0) spaces [Rua01, Theorem 3.2]. The original proof treats only the case where $\mathrm{Disp}_K(h) = \mathrm{Min}(h)$, however it is straightforward to check that it carries through for any $\mathrm{Disp}_K(h)$. We include the proof for the sake of completeness.

First we check that $C_G(h)$ leaves $\mathrm{Disp}_K(h)$ invariant. Let $x \in \mathrm{Disp}_K(h)$ be a vertex and take $g \in C_G(h)$. Then we have

$$d(g \cdot x, hg \cdot x) = d(g \cdot x, gh \cdot x) = d(x, h \cdot x) \leqslant K$$

and thus $g \cdot x \in \mathrm{Disp}_K(h)$. Observe that the action of $C_G(h)$ on $\mathrm{Disp}_K(h)$ is proper, since the action of G on X is proper. We only need to check cocompactness. We proceed by contradiction. Assume that there is no compact subset of $\mathrm{Disp}_K(h)$ whose $C_G(h)$–translates cover $\mathrm{Disp}_K(h)$, and pick a vertex $x_0 \in \mathrm{Disp}_K(h)$. Then there exists a sequence of vertices $(x_n)_{n=1}^{\infty}$ of $\mathrm{Disp}_K(h)$ such that $d(C_G(h) \cdot x_0, x_n) \to \infty$ as $n \to \infty$. Let $D \subset X$ be a compact set containing x_0 such that $G \cdot D = X$, and let $(g_n)_{n=1}^{\infty}$ be a sequence of elements of G such that $g_n \cdot x_n \in D$. We can assume (by passing to a subsequence if necessary) that $g_n \neq g_m$ for $n \neq m$. Indeed, we have

$$d(x_0, g_n^{-1} \cdot x_0) \geqslant d(x_0, x_n) - \mathrm{diam} D \to \infty \text{ as } n \to \infty.$$

Now consider the family of elements $\{g_n h g_n^{-1}\}_{n \geqslant 1}$ of G. We claim that the displacement functions $d_{g_n h g_n^{-1}}$ are uniformly bounded on D. Let $y \in D$ be a vertex.

We have

$$\begin{aligned}
d_{g_nhg_n^{-1}}(y) &= d(y, g_nhg_n^{-1}\cdot y)\\
&\leqslant d(y, g_n\cdot x_n) + d(g_n\cdot x_n, g_nhg_n^{-1}(g_n\cdot x_n))\\
&\quad + d(g_nhg_n^{-1}(g_n\cdot x_n), g_nhg_n^{-1}\cdot y)\\
&= d(y, g_n\cdot x_n) + d(g_n\cdot x_n, g_nh\cdot x_n) + d(g_nh\cdot x_n, g_nhg_n^{-1}\cdot y)\\
&\leqslant \operatorname{diam}(D) + K + \operatorname{diam}(D).
\end{aligned}$$

Since G acts properly, it must be $g_nhg_n^{-1} = g_mhg_m^{-1}$ for $n \neq m$ (after passing to a subsequence). Therefore for all $n \neq m$ we have that $g_m^{-1}g_n \in C_G(h)$. Now for any $n \neq 1$ we get

$$d(x_n, g_n^{-1}g_1\cdot x_0) \leqslant d(x_n, g_n^{-1}g_1\cdot x_1) + d(g_n^{-1}g_1\cdot x_1, g_n^{-1}g_1\cdot x_0) \leqslant \operatorname{diam}(D) + d(x_0, x_1).$$

This gives a contradiction since $g_n^{-1}g_1 \in C_G(h)$, and by the choice of x_n we have that $d(x_n, C_G(h)\cdot x_0) \to \infty$ as $n \to \infty$. □

We are ready now to prove Lemma 4.2.

Proof of Lemma 4.2. Pick any $K' \leqslant K$ and let $x_0 \in \operatorname{Disp}_{K'}(h)$ be a vertex. By Theorem 4.3 the centraliser $C_G(h)$ acts cocompactly on $\operatorname{Disp}_K(h)$. Hence there exists $R > 0$ such that $\operatorname{Disp}_K(h) \subset C_G(h)\cdot B_R(x_0, X)$. Since $C_G(h)\cdot x_0 \subset \operatorname{Disp}_{K'}(h)$, taking R as C proves the lemma. □

Remark 4.4. We believe that in Lemma 4.2 one can obtain a concrete distance estimate, i.e., in the formula $\operatorname{Disp}_{K'}(h) \subset B_C(\operatorname{Disp}_K(h), X)$ one can express C as an explicit function of K and K'. However, for our purposes, the existence of any constant C is sufficient.

4.2. **Trivial action on the boundary.** In this section we characterise hyperbolic isometries that act trivially on the boundary as being virtually central. More precisely, we show the following.

Theorem 4.5. *Let G be a group acting geometrically on a systolic complex X, and let $h \in G$ be a hyperbolic isometry. Then h acts trivially on the boundary ∂X if and only if the centraliser $C_G(h)$ has finite index in G.*

The theorem is a systolic analogue of a theorem of K. Ruane for CAT(0) spaces [Rua01, Theorem 3.4]. In a certain way, our situation is more restrictive. Namely, by [OP16, Corollary 5.8] the centraliser $C_G(h)$ is commensurable with the product $F_n \times \mathbb{Z}$, where F_n is the free group on n generators for some $n \geqslant 0$. It follows that either of the assertions of Theorem 4.5 holds true if and only if the group G itself is commensurable with $F_n \times \mathbb{Z}$.

Proof of Theorem 4.5. **"if" direction.** By Theorem 4.3 the centraliser $C_G(h)$ acts cocompactly on the minimal displacement set $\mathrm{Min}(h) \subset X$. Since the index $[G\colon C_G(h)]$ is finite, it follows that the action of $C_G(h)$ on X is cocompact as well. Therefore there exists a constant $K \geqslant 0$ such that for any vertex $x \in X$, there is a vertex $y \in \mathrm{Min}(h)$ with $d(x, y) \leqslant K$. Hence, by the triangle inequality, for any $x \in X$ we have $d(x, h \cdot x) \leqslant L(h) + 2K$.

Now let $\eta = (v_0, v_1, v_2, \ldots)$ be a good geodesic ray in X. For any $i \geqslant 0$ we have $d(v_i, h \cdot v_i) \leqslant L(h) + 2K$ and hence $[\eta] = [h \cdot \eta]$ in ∂X.

"only if" direction. Choose a vertex $y \in X$ and let $x \in X$ be an arbitrary vertex. By Theorem 3.2 there exists a good geodesic ray $\eta = (v_0, v_1, v_2, \ldots)$ such that $v_0 = y$ and for some $i \geqslant 0$ we have $d(v_i, x) \leqslant E$, where E is a constant independent of x and y.

The isometry h acts trivially on the boundary, so we have $[\eta] = [h \cdot \eta]$. Applying Corollary 2.9 we obtain

$$d(v_i, h \cdot v_i) \leqslant d(v_0, h \cdot v_0) + 2D + 1,$$

where D is the constant appearing in Theorem 2.8. In other words, we have $v_i \in \mathrm{Disp}_K(h)$ where $K = d(v_0, h \cdot v_0) + 2D + 1$. Since $d(v_i, x) \leqslant E$, the triangle inequality implies that $x \in \mathrm{Disp}_{K+2E}(h)$ (see the discussion after Definition 4.1). Because x was arbitrary, we have $X = \mathrm{Disp}_{K+2E}(h)$. By Theorem 4.3 the centraliser $C_G(h)$ acts cocompactly on X and so it has finite index in G. □

We obtain the following corollary.

Corollary 4.6. *Let G be a torsion-free group, acting geometrically on a systolic complex X. Then G acts trivially on ∂X if and only if $G \cong \mathbb{Z}$ or $G \cong \mathbb{Z}^2$.*

Proof. If G is isomorphic to either $\mathbb{Z}$ or $\mathbb{Z}^2$ then every element of G is central, and by Theorem 4.5 it acts trivially on ∂X.

Now assume that G acts trivially on ∂X. Since G is torsion-free, all of its elements are hyperbolic, and therefore by Theorem 4.5 every element is virtually central. Pick $h \in G$. By [OP16, Corollary 5.8] the centraliser $C_G(h)$ contains a finite-index subgroup $H \cong F_n \times \mathbb{Z}$, where F_n is the free group on n generators for some $n \geqslant 0$. We must have $n \leqslant 1$ for otherwise no non-trivial element of F_n would be virtually central in G. This means that H is isomorphic to either $\mathbb{Z}$ or $\mathbb{Z}^2$. Now since H has finite index in $C_G(h)$ and $C_G(h)$ has finite index in G we get that H has finite index in G.

If $H \cong \mathbb{Z}$ then G is a virtually cyclic torsion-free group, and hence it must be isomorphic to $\mathbb{Z}$. If $H \cong \mathbb{Z}^2$ then G is a torsion-free group that contains $\mathbb{Z}^2$ as a finite-index subgroup. Such G must be isomorphic to either $\mathbb{Z}^2$ or to the fundamental group of the Klein bottle. Since the latter contains elements that are not virtually central, we conclude that $G \cong \mathbb{Z}^2$. □

5. Good geodesics

The proofs in Section 6 require going through the construction of good geodesics (unlike proofs in previous sections, where only certain 'formal' properties were needed). In this section we give a sketch of this construction.

In order to define good geodesics we first describe the construction of *Euclidean geodesics.* This construction is fairly involved, and hence it is divided into a few steps. Our exposition is based on [OP09, Sections 7–9] (we refer the reader there for the proofs of various statements discussed below). Throughout this section let X be a systolic complex. Some notions appearing in the construction are presented in Figures 1 and 2 in Section 6 (in the special case of $X = \mathbb{E}^2_\Delta$).

5.1. **Directed geodesics.** Let $x, y \in X$ be two vertices and put $n = d(x, y)$. A *directed geodesic from x to y* is a sequence of simplices $(\sigma_i)_{i=0}^n$ such that $\sigma_0 = x$, $\sigma_n = y$ and the following two conditions are satisfied:

(1) any two consecutive simplices σ_i and σ_{i+1} are disjoint and together they span a simplex of X,
(2) for any three consecutive simplices $\sigma_{i-1}, \sigma_i, \sigma_{i+1}$we have

$$\mathrm{Res}(\sigma_{i-1}, X) \cap B_1(\sigma_{i+1}, X) = \sigma_i,$$

where $\mathrm{Res}(\sigma_{i-1}, X)$ is the union of all simplices of X that contain σ_{i-1}.

A directed geodesic from x to y always exists and it is unique. One can show that $\sigma_i \subset S_i(x, X) \cap S_{n-i}(y, X)$, and therefore any sequence of vertices $(v_i)_{i=0}^n$ such that $v_i \in \sigma_i$ is a geodesic. Finally, as the name suggests, directed geodesics in general are not symmetric – usually a directed geodesic from x to y is not equal to a directed geodesic from y to x.

5.2. **Layers.** The intersection $S_i(x, X) \cap S_{n-i}(y, X)$ is called the *layer i between x and y* and it is denoted by L_i. For any i the layer L_i is convex and ∞–large. (Layers in fact can be defined in the same way for any two convex subcomplexes V and W such that for every $v \in V$ and $w \in W$ one has $d(v, w) = n$ for some fixed $n > 0$.)

Convention 5.1. Suppose that $(\sigma_i)_{i=0}^n$ is a directed geodesic from x to y and $(\tau_i)_{i=0}^n$ is a directed geodesic from y to x. We introduce the following convention: despite $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$ go in the opposite directions, we index simplices of $(\tau_i)_{i=0}^n$ in the same direction as for $(\sigma_i)_{i=0}^n$, i.e., $\tau_0 = x, \tau_1, \ldots, \tau_{n-1}, \tau_n = y$.

Observe that both σ_i and τ_i are contained in the layer L_i. Define the *thickness of layer* L_i (with respect to $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$) to be the maximal distance between vertices of σ_i and τ_i (since layers are convex, this distance is always realised inside L_i). The layer is *thin* if its thickness is at most 1, and it is *thick* otherwise.

A pair of indices (j, k) such that $0 < j < k < n$ and $j < k - 1$ is called a *thick interval* if layers L_j and L_k are thin, and for every i such that $j < i < k$ the layer L_i is thick. If for some i we have $j < i < k$ then we say that i belongs to the interval (j, k).

5.3. **Characteristic surfaces.** Let (j, k) be a thick interval and let $s_i \in \sigma_i$ and $t_i \in \tau_i$ be vertices such that for any $j \leqslant i \leqslant k$ the distance between s_i and t_i is equal to the thickness of layer L_i. Consider the sequence of vertices

$$(s_j, s_{j+1}, \ldots, s_{k-1}, s_k, t_k, t_{k-1}, \ldots, t_{j+1}, t_j, s_j).$$

Observe that any two consecutive vertices in the above sequence are adjacent and therefore this sequence defines a closed loop which we denote by γ. In fact γ is always an embedded loop (this amounts to saying that $s_j \neq t_j$ and $s_k \neq t_k$).

Let $S\colon \Delta \to X$ be a minimal surface spanned by γ, i.e., a simplicial map from a triangulation of a 2–disk Δ such that:

(1) the boundary of Δ is mapped isomorphically to γ,
(2) the disk Δ consists of the least possible number of triangles (among all disks Δ' for which there exists a simplicial map $S'\colon \Delta' \to X$ satisfying (1)).

We call $S\colon \Delta \to X$ a *characteristic surface* (for the thick interval (j, k)) and we call Δ a *characteristic disk*. It is a standard fact that a minimal disk is always systolic, i.e., every of its internal vertices is incident to at least 6 triangles.

The cycle γ does not have to be unique, and hence there could be many characteristic surfaces. For any two characteristic surfaces $S\colon \Delta \to X$ and $S'\colon \Delta' \to X$ the disks Δ and Δ' are isomorphic. We can thus identify all such disks and denote *the* characteristic disk by Δ. Now for any simplex $\rho \in \Delta$ the images $S(\rho)$ for all possible characteristic surfaces span a simplex of X, which we denote $\mathcal{S}(\rho)$. This assignment, called the *characteristic mapping*, respects inclusions, i.e., if $\rho_1 \subseteq \rho_2$ then $\mathcal{S}(\rho_1) \subseteq \mathcal{S}(\rho_2)$.

5.4. **Geometry of characteristic disks.** For any i such that $j \leqslant i \leqslant k$, let v_i and w_i be vertices of Δ that are preimages of s_i and t_i respectively, for some characteristic surface $S\colon \Delta \to X$. In fact vertices v_i and w_i are uniquely defined and the sequence

$$(v_j, v_{j+1}, \ldots, v_{k-1}, v_k, w_k, w_{k-1}, \ldots, w_{j+1}, w_j, v_j)$$

constitutes the boundary of the disk Δ.

Denote by $\mathbb{E}^2_\Delta$ the equilateral triangulation of the Euclidean plane. Clearly $\mathbb{E}^2_\Delta$ viewed as a simplicial complex is systolic. One shows that the characteristic disk Δ can be isometrically embedded in $\mathbb{E}^2_\Delta$ (such disk is called *flat*). Moreover, after embedding $\Delta \subset \mathbb{E}^2_\Delta$, the edges $[v_j, w_j]$ and $[v_k, w_k]$ are parallel, and consecutive layers in Δ between them are contained in straight lines of $\mathbb{E}^2_\Delta$ (treated as subcomplexes of $\mathbb{E}^2_\Delta$), that are parallel to the lines containing $[v_j, w_j]$ and $[v_k, w_k]$. In particular for any i the vertices v_i and w_i lie on a straight line inside $\mathbb{E}^2_\Delta$. The subpath of this line between v_i and w_i is the unique geodesic between v_i and w_i in Δ, which we denote by $v_i w_i$. The geodesic $v_i w_i$ is in fact equal to the entire layer i in Δ (between the edges $[v_j, w_j]$ and $[v_k, w_k]$).

Finally, for any characteristic surface $S\colon \Delta \to X$ (and hence for a characteristic mapping $\mathcal{S}\colon \Delta \to X$ as well) the image of the geodesic $v_i w_i$ is contained in the layer i in X. Also, any characteristic surface $S\colon \Delta \to X$ restricted to $v_i w_i$ is an isometric embedding.

5.5. **Euclidean diagonals.** Given the characteristic disk Δ, for every $i \in \{j, \ldots, k\}$ let v_i' and w_i' be points on the unique geodesic between v_i and w_i in Δ, that are at distance $\frac{1}{2}$ from v_i and w_i respectively. In particular $v_j' = w_j'$ and $v_k' = w_k'$. Consider a piecewise linear loop defined as the concatenation of straight segments between consecutive points in the sequence

$$(v_j' = w_j', v_{j+1}', \ldots v_{k-1}', v_k' = w_k', w_{k-1}', \ldots, w_{j+1}', w_j' = v_j'),$$

and let Δ' be a polygonal domain inside Δ enclosed by this loop. We call Δ' the *modified characteristic disk.* We endow Δ' with a path metric induced from the Euclidean metric on $\mathbb{E}^2_\Delta \cong \mathbb{E}^2$. Observe that Δ' is simply connected, and therefore this path metric is in fact a CAT(0) metric. (The disk Δ' does not have to be a convex subset of $\mathbb{E}^2_\Delta$ with respect to the Euclidean metric on $\mathbb{E}^2_\Delta$, and therefore a CAT(0) geodesic inside Δ' does not have to be a straight line.)

The *Euclidean diagonal* of Δ is a sequence of simplices $(\rho_i)_{j+1}^{k-1}$ of Δ defined as follows. Let α be a CAT(0) geodesic in Δ' between points $v_j' = w_j'$ and $v_k' = w_k'$. For every i such that $j < i < k$ choose vertices on the geodesic $v_i w_i$, that are closest to the point of intersection $\alpha \cap v_i w_i$. For any i, it is either a single vertex, in that case we put ρ_i to be that vertex, or in the case when α goes through the barycentre of some edge of $v_i w_i$, then we put ρ_i to be this edge. One can show that the Euclidean diagonal for Δ satisfies the following two conditions:

(1) for any i such that $j < i < k-1$ simplices ρ_i and ρ_{i+1} span a simplex,
(2) vertices v_j, w_j, σ_{j+1} span a simplex and vertices v_k, w_k, ρ_{k-1} span a simplex (in particular ρ_{j+1} and ρ_{k-1} are necessarily vertices).

5.6. **Euclidean geodesics.** We are ready now to define Euclidean geodesics.

Definition 5.2. The *Euclidean geodesic* between vertices x and y in X, such that $d(x, y) = n$ is the sequence of simplices $(\delta_i)_{i=0}^n$ defined as follows. For any i such that $0 < i < n$, if the layer L_i is thin then set

$$\delta_i = \mathrm{span}\{\sigma_i, \tau_i\},$$

where σ_i and τ_i are the simplices of the directed geodesics between x and y that are contained in layer L_i. For any i such that the layer L_i is thick, consider the thick interval (j, k) that contains i and put

$$\delta_i = \mathcal{S}(\rho_i),$$

where ρ_i is the simplex of Euclidean diagonal that is contained in the layer i in the characteristic disk for (j, k), and $\mathcal{S}$ denotes the characteristic mapping. Finally let $\delta_0 = x$ and $\delta_n = y$.

By definition, consecutive simplices of the Euclidean geodesic $(\delta_i)_{i=0}^n$ are contained in consecutive layers between x and y. Unlike for directed geodesics, not every two consecutive simplices δ_i, δ_{i+1} span a simplex of X. However, the following holds.

Proposition 5.3. [OP09, Remark 3.1] *Suppose $(\delta_i)_{i=0}^n$ is a Euclidean geodesic between vertices x and y. Then there exists a sequence of vertices $(v_i)_{i=0}^n$ such that $v_i \in \delta_i$ and $(v_i)_{i=0}^n$ is a geodesic.*

Also note that Euclidean geodesics are symmetric with respect to their endpoints. We now present two theorems describing the crucial properties of Euclidean geodesics. The first one, roughly speaking, says that Euclidean geodesics are coarsely closed under taking subsegments. The second one is a coarse form of a CAT(0) inequality for Euclidean geodesics.

Theorem 5.4. [OP09, Theorem B] *Let $(\delta_i)_{i=0}^n$ be a Euclidean geodesic between vertices x and y. Take $j, k \in \{0, \ldots, n\}$ with $j < k$ and let $(r_i)_{i=j}^k$ be a geodesic such that $r_i \in \delta_i$ for $i \in \{j, \ldots, k\}$. Let $(\delta_i^{j,k})_{i=j}^k$ denote the Euclidean geodesic between vertices r_j and r_k. Then for every $i \in \{j, \ldots, k\}$ for any vertices $v_i \in \delta_i$ and $u_i \in \delta_i^{j,k}$ we have*

$$d(v_i, u_i) \leqslant C,$$

where $C > 0$ is a universal constant.

Theorem 5.5. [OP09, Theorem C] *Let x, y and $\widetilde{y}$ be vertices of X with $d(x, y) = n$ and $d(x, \widetilde{y}) = m$. Let $(v_i)_{i=0}^n$ and $(\widetilde{v}_i)_{i=0}^m$ be geodesics such that for all appropriate i we* *have* $v_i \in \delta_i$ *and* $\widetilde{v}_i \in \widetilde{\delta}_i$*, where* $(\delta_i)_{i=0}^n$ *and* $(\widetilde{\delta}_i)_{i=0}^m$ *are the Euclidean geodesics between* x *and* y *and between* x *and* $\widetilde{y}$ *respectively. Then for any* $0 \leqslant c \leqslant 1$ *we have*

$$d(v_{\lfloor cn \rfloor}, \widetilde{v}_{\lfloor cm \rfloor}) \leqslant c \cdot d(v_n, \widetilde{v}_m) + C,$$

where $C > 0$ *is a universal constant.*

5.7. **Good geodesics.** From now on let $C > 0$ be a fixed constant which satisfies the assertions of both Theorem 5.4 and Theorem 5.5. In particular this means that $C \geqslant 200$ ([OP09, p. 2877]). Having an explicit lower bound will be needed in Section 6.

Theorem 5.4 presents a model behaviour, which motivates the definition of good geodesics.

Definition 5.6 (C'–good geodesic)**.** Let $(v_i)_{i=0}^n$ be a geodesic in X. For $j, k \in \{0, \ldots, n\}$ let $(\delta_i^{j,k})_{i=j}^k$ denote the Euclidean geodesic between vertices v_j and v_k. We say that $(v_i)_{i=0}^n$ is a C'*–good geodesic* if for every two vertices v_j and v_k, for every $i \in \{j, \ldots, k\}$ for any vertex $u_i \in \delta_i^{j,k}$ we have

$$d(v_i, u_i) \leqslant C',$$

where $C' > 0$ is a positive constant. An infinite geodesic is a C'–good geodesic if every of its finite subgeodesics is a C'–good geodesic. Observe that for a C'–good geodesic any of its subgeodesics is a C'–good geodesic as well.

Definition 5.7 (Good geodesic)**.** A geodesic $(v_i)_i$ (finite or infinite) is a *good geodesic* if it is a C–good geodesic.

In particular, by Theorem 5.4 any geodesic arising from Proposition 5.3 is a good geodesic. Consequently, any two vertices of X can be joined by a good geodesic (cf. [OP09, Corollary 3.3]).

We finish this section with the following two remarks.

Remark 5.8. By going through the steps of the construction, one observes that the directed, Euclidean, and good geodesics are preserved by simplicial automorphisms of X.

Remark 5.9. The main goal of the construction outlined in this section is to establish Theorem 2.8. This theorem plays the key role in showing various properties of the boundary in [OP09]. Theorem 2.8 follows easily from Definition 5.7, Proposition 5.3 and Theorem 5.5. In particular, the constant D appearing in Theorem 2.8 may be taken to be $3C$.

6. Fixed points on the boundary

The purpose of this section is to show that every hyperbolic isometry h of a systolic complex X fixes a pair of points on the boundary. These two points, denoted by $h^{+\infty}$ and $h^{-\infty}$ are the canonical fixed points of h, in the sense that for any vertex $x \in X$ we have $(h^n \cdot x)_n \to h^{+\infty}$ and $(h^{-n} \cdot x)_n \to h^{-\infty}$ as $n \to \infty$ in $\overline{X} = X \cup \partial X$.

To obtain $h^{+\infty}$ and $h^{-\infty}$ we show that there exists a bi-infinite good geodesic γ such that γ and $h \cdot \gamma$ are asymptotic. In principal, one could expect a stronger result, namely the existence of an h–invariant good geodesic. However, there are examples of systolic complexes where a hyperbolic isometry has no invariant geodesics at all [Els09b, Example 1.2]. It is true though, that for every hyperbolic isometry h there is a geodesic γ, such that $h \cdot \gamma$ and γ are Hausdorff 1–close [Els09b, Theorem 1.3]. Unfortunately, in our construction the distance between $h \cdot \gamma$ and γ depends on $L(h)$.

Theorem 6.1. *Let X be a systolic complex on which a group G acts geometrically. Let $h \in G$ be a hyperbolic isometry. Then either there exists a bi-infinite good geodesic which is contained in $\mathrm{Disp}_K(h)$ for some $K = K(h)$ and $\mathrm{Disp}_K(h)$ is h–cocompact, or there exists an h–invariant good geodesic.*

We now state and prove the main result of this section assuming Theorem 6.1.

Proposition 6.2. *Let G be a group acting geometrically on a systolic complex X and let $h \in G$ be a hyperbolic isometry. Then:*

1. *there exist points $h^{-\infty}$ and $h^{+\infty}$ in the boundary ∂X which are fixed by h,*
2. *for any vertex $x \in X$ we have $(h^n \cdot x)_n \to h^{+\infty}$ and $(h^{-n} \cdot x)_n \to h^{-\infty}$ as $n \to \infty$ in the compactification $\overline{X} = X \cup \partial_O X$, where $O \in X$ is some base vertex.*

Since for any two vertices $O, O' \in X$ there is a homeomorphism between $X \cup \partial_O X$ and $X \cup \partial_{O'} X$ (see Subsection 2.2), the choice of O does not really matter. In order to simplify the argument we will choose O during the proof.

Proof. To show that a sequence $(x_n)_{n=0}^{\infty}$ converges to a point $[\xi] \subset \partial_O X$ in $\overline{X}$, it is enough to find a sequence $(v_n)_{n=0}^{\infty} \subset \xi$, such that $d(v_n, O) \to \infty$ as $n \to \infty$ and such that $d(v_n, x_n)$ is uniformly bounded (see [OP09, Definition 4.1]).

By Theorem 6.1 there either exists a bi-infinite good geodesic $\gamma \subset \mathrm{Disp}_K(h)$ and $\mathrm{Disp}_K(h)$ is h–cocompact, or there exists an h–invariant good geodesic γ. We first give the proof assuming that there exists a good geodesic γ which is h–invariant.

Choose a vertex $O \in \gamma$ and parametrise vertices of γ by integers such that $\gamma(0) = O$ and $h \cdot \gamma(0) = \gamma(L(h))$. Then γ splits into two good geodesic rays $\gamma\big|_{[0,+\infty]}$ and $\gamma\big|_{[0,-\infty]}$

starting at O. Define $h^{+\infty} = [\gamma\big|_{[0,+\infty]}]$ and $h^{-\infty} = [\gamma\big|_{[0,-\infty]}]$. Since γ is h–invariant, for every $i \geqslant 0$ we have

$$d(\gamma\big|_{[0,\pm\infty]}(i), h \cdot \gamma\big|_{[0,\pm\infty]}(i)) = L(h),$$

and thus both $h^{+\infty}$ and $h^{-\infty}$ are fixed by h.

Let $x \in X$ be an arbitrary vertex and let $M = d(x, O)$. For any $n \geqslant 0$ we have:

(1) $h^n \cdot O \in \gamma\big|_{[0,+\infty]}$,
(2) $d(h^n \cdot O, O) \to \infty$ as $n \to \infty$,
(3) $d(h^n \cdot x, h^n \cdot O) = M$.

This implies that $h^n \cdot x \to [\gamma\big|_{[0,+\infty]}] = h^{+\infty}$ and $h^{-n} \cdot x \to [\gamma\big|_{[0,-\infty]}] = h^{-\infty}$.

In the case where $\gamma \subset \mathrm{Disp}_K(h)$ is a bi-infinite good geodesic and $\mathrm{Disp}_K(h)$ is h–cocompact we proceed similarly as above. Choosing a vertex $O \in \gamma$ splits γ into two good geodesic rays. We denote them by $\gamma\big|_{[0,+\infty]}$ and $\gamma\big|_{[0,-\infty]}$, even though we did not specify how we choose an orientation of γ (it will become clear from the proof). Consequently, let $h^{+\infty} = [\gamma\big|_{[0,+\infty]}]$ and $h^{-\infty} = [\gamma\big|_{[0,-\infty]}]$. Both $\gamma\big|_{[0,+\infty]}$ and $\gamma\big|_{[0,-\infty]}$ are contained in $\mathrm{Disp}_K(h)$ and therefore for every $i \geqslant 0$ we have

$$d(\gamma\big|_{[0,\pm\infty]}(i), h \cdot \gamma\big|_{[0,\pm\infty]}(i)) \leqslant K,$$

and hence h fixes both $h^{+\infty}$ and $h^{-\infty}$.

Let $x \in X$ be an arbitrary vertex and let $M = d(x, O)$. Since $\mathrm{Disp}_K(h)$ is h–cocompact and γ is bi-infinite it follows that there exists $R > 0$ such that

$$\mathrm{Disp}_K(h) \subset B_R(\gamma, X).$$

Consider the sequence $(h^n \cdot O)_n$. We do not necessarily have $(h^n \cdot O)_n \subset \gamma\big|_{[0,\infty]}$, but for any $n \geqslant 0$ there exists a vertex $v_n \in \gamma\big|_{[0,\infty]}$ with $d(h^n \cdot O, v_n) \leqslant R$. Then, by the triangle inequality we have $d(O, v_n) \to \infty$ as $n \to \infty$ since $d(O, h^n \cdot O) \to \infty$ as $n \to \infty$. Finally, for any $n \geqslant 0$ we have

$$d(h^n \cdot x, v_n) \leqslant d(h^n \cdot x, h^n \cdot O) + d(h^n \cdot O, v_n) \leqslant M + R,$$

and therefore $h^n \cdot x \to [\gamma\big|_{[0,+\infty]}] = h^{+\infty}$ and $h^{-n} \cdot x \to [\gamma\big|_{[0,-\infty]}] = h^{-\infty}$. □

It remains to prove Theorem 6.1. Since the proof is fairly long we outline it first.

Outline of the proof of Theorem 6.1. Observe that the action of h on X preserves $\mathrm{Min}(h)$. We consider two cases: when $\mathrm{Min}(h)$ is h–cocompact and when it is not h–cocompact. These two cases will lead to the two respective claims of the theorem.

In the first case we show in Lemma 6.3 that for any two vertices $x, y \in \mathrm{Min}(h)$ the Euclidean geodesic between these vertices is contained in $\mathrm{Disp}_K(h)$ for some $K > 0$. This is achieved by showing that both directed geodesics between these vertices, and characteristic disks spanned by those geodesics, belong to $\mathrm{Disp}_K(h)$. The main tool in the proof of Lemma 6.3 is the Fellow Traveller Property of directed geodesics. Then we construct the desired good geodesic, roughly, as a limit of good geodesics between vertices $h^{-n} \cdot x$ and $h^n \cdot x$ for a fixed vertex $x \in \mathrm{Min}(h)$.

In the second case, since $C_G(h)$ acts geometrically on $\mathrm{Min}(h)$, we deduce that there is a hyperbolic isometry g that commutes with h, such that $\langle g, h\rangle \cong \mathbb{Z}^2$. Thus by the systolic Flat Torus Theorem there is a flat $F \subset X$ on which the subgroup $\langle g, h\rangle$ acts by translations. Then using Lemma 6.4 we find an h–invariant C'–good geodesic inside F (treated as a systolic complex on its own), for a certain $C' > 0$. By Lemma 6.5 any C'–good geodesic in F is a $(C' + 10)$–good geodesic in X. In the above procedure we are able to choose C' so that $C' + 10$ is less than C and therefore the constructed $(C' + 10)$–good geodesic is a good geodesic in X.

Before giving the proof of Theorem 6.1 we state and prove the three lemmas mentioned above.

Lemma 6.3. *Consider two vertices $x, y \in \mathrm{Min}(h) \subset X$ and let $(\delta_i)_{i=0}^n$ be the Euclidean geodesic between x and y. Then we have $(\delta_i)_{i=0}^n \subset \mathrm{Disp}_K(h)$, where $K = 9 \cdot L(h) + 6$.*

Proof. Let $(\sigma_i)_{i=0}^n$ be a directed geodesic from x to y (i.e., $\sigma_0 = x$ and $\sigma_n = y$). Then, by the Fellow Traveller Property [JS06, Proposition 11.2] applied to directed geodesics $(\sigma_i)_{i=0}^n$ and $(h \cdot \sigma_i)_{i=0}^n$, for each $i \in \{0, \dots, n\}$ for any vertex $s \in \sigma_i$ we have

$$d(s, h \cdot s) \leqslant 3 \cdot \max\{d(x, h \cdot x), d(y, h \cdot y)\} + 1 = 3 \cdot L(h) + 1,$$

since $x, y \in \mathrm{Min}(h)$. Put $K' = 3 \cdot L(h) + 1$. By the above inequality we get that $(\sigma_i)_{i=0}^n \subset \mathrm{Disp}_{K'}(h)$.

Now let $(\sigma_i)_{i=0}^n$ be as above, and let $(\tau_i)_{i=0}^n$ be the directed geodesic from y to x (see Convention 5.1). Clearly, by the argument above, we also have $(\tau_i)_{i=0}^n \subset \mathrm{Disp}_{K'}(h)$. If for some $i \in \{0, \dots, n\}$ the layer L_i is thin, then by definition $\delta_i = \mathrm{span}\{\sigma_i, \tau_i\}$ and therefore it is contained in $\mathrm{Disp}_{K'}(h)$, since both σ_i and τ_i are so. If the layer L_i is thick then we proceed as follows.

Take any vertex $z \in \delta_i$. We claim that there exist vertices $s_i \in \sigma_i$ and $t_i \in \tau_i$ and a geodesic α between s_i and t_i, such that z lies on α. Indeed, consider a thick interval that contains i, let Δ be an appropriate characteristic disk and let $v_i w_i$ be a geodesic in Δ that is the layer i in Δ. Any characteristic surface $S\colon \Delta \to X$ restricted to $v_i w_i$ is an isometric embedding. Moreover, any vertex of δ_i lies in the image $S(v_i w_i)$ for some such surface. Take a surface $S\colon \Delta \to X$ such that $z \in S(v_i w_i)$ and put $s_i = S(v_i)$, $t_i = S(w_i)$ and let $\alpha = S(v_i w_i)$. This proves the claim.

Put $m = d(s_i, t_i)$ and let $(\rho_j)_{j=0}^m$ be a directed geodesic from s_i to t_i (actually, here the direction is not important). Let $(u_j)_{j=0}^m$ be a geodesic such that $u_j \in \rho_j$ for every $j \in \{0, \ldots, m\}$ (in particular $u_0 = s_i$ and $u_m = t_i$).

Since the layer L_i is convex, both α and $(u_j)_{u=0}^m$ are contained in L_i. Since L_i is ∞–large, any two geodesics with the same endpoints are Hausdorff 1–close [JŚ06, Lemma 2.3]. Therefore α and $(u_j)_{j=0}^m$ are 1–close. Finally, by the Fellow Traveller Property (applied to $(\rho_j)_{j=0}^m$ and $(h \cdot \rho_j)_{j=0}^m$) for any $j \in \{0, \ldots, m\}$ we have

$$d(u_j, h \cdot u_j) \leqslant 3 \cdot \max\{d(s_i, h \cdot s_i), d(t_i, h \cdot t_i)\} + 1 \leqslant 3 \cdot K' + 1,$$

since both s_i and t_i belong to $\mathrm{Disp}_{K'}(h)$. Because α and $(u_j)_{j=0}^m$ are 1–close and $z \in \alpha$, the above inequality implies that $d(z, h \cdot z) \leqslant 3 \cdot K' + 1 + 2$ and hence $z \in \mathrm{Disp}_{3K'+3}(h)$. Since $z \in \delta_i$ was arbitrary we obtain that $\delta_i \subset \mathrm{Disp}_{3K'+3}(h)$ for any i such that L_i is thick. This, together with the assertion that $\delta_i \subset \mathrm{Disp}_{K'}(h)$ for any i such that L_i is thin, finishes the proof of the lemma, as $3 \cdot K' + 3 = 9 \cdot L(h) + 6$. □

In the next lemma we study the equilaterally triangulated Euclidean plane $\mathbb{E}^2_\Delta$. We view it simultaneously as a systolic simplicial complex and as a CAT(0) metric space (cf. Subsections 5.4 and 5.5). We denote the CAT(0) distance between points of $\mathbb{E}^2_\Delta$ by $d_{\mathbb{E}^2}$ in order to distinguish it from the standard (combinatorial) distance d.

Lemma 6.4. *Let h be a hyperbolic isometry of $\mathbb{E}^2_\Delta$ and let $\gamma \subset \mathbb{E}^2_\Delta$ be an h–invariant geodesic. Suppose that β is a* CAT(0) *geodesic (i.e., a straight line) such that γ and β are Hausdorff K–close with respect to the* CAT(0) *distance, for some $K > 0$. Then γ is a $(\frac{4K}{\sqrt{3}} + 1)$–good geodesic.*

Proof. We first give the idea of the proof. We observe that for any thick interval in $\mathbb{E}^2_\Delta$ there is a unique characteristic surface. After identifying the characteristic disk Δ with its image, the CAT(0) geodesic α in the modified characteristic disk Δ' is uniformly close to β, and this distance depends only on K. This will imply the lemma as the simplices of the Euclidean geodesic are 1–close to α, and γ is K–close to β. The case of thin layers will follow easily from the methods used to prove the case of thick intervals.

We now begin the proof. Let x and y be any two vertices of γ. Let $(\sigma_i)_{i=0}^n$ be the directed geodesic from x to y and let $(\tau_i)_{i=0}^n$ be the directed geodesic going in the opposite direction (again, simplices of $(\tau_i)_{i=0}^n$ are indexed in the same direction as for $(\sigma_i)_{i=0}^n$, see Convention 5.1). One checks that these geodesics have the form shown in Figure 1.

For $k \in \{0, 1, \ldots, n\}$ let l_k denote the infinite line in $\mathbb{E}^2_\Delta$ that contains σ_k and τ_k. In particular, for $k \in \{1, \ldots, n-1\}$ the line l_k contains the layer L_k between x and y.

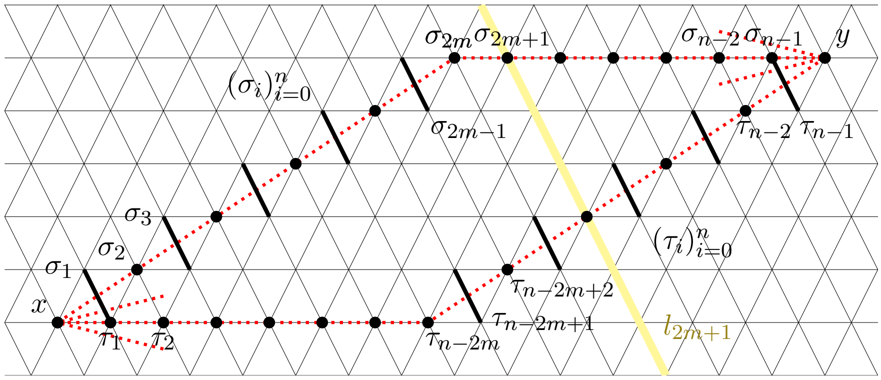

FIGURE 1. Generic form and position of the directed geodesics $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$ in $\mathbb{E}^2_\Delta$.

Note that the geodesic $(\sigma_i)_{i=0}^n$ splits into two parts: the part $(\sigma_i)_{i=0}^{2m}$, where σ_i is a vertex for even i and an edge for odd i, and the part $(\sigma_i)_{i=2m+1}^n$, which consists entirely of vertices. Now observe that if $2m \in \{0, 2, n-1, n\}$ then for every $k \in \{1, \ldots, n-1\}$ the layer L_k with respect to $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$ is thin. If $2m = n-2$ then for $k \in \{1, 2m+1\}$ the layer L_k is thin, and for $k \in \{2, \ldots, 2m\}$ the layer L_k is thin for k even, and thick of thickness 2 for k odd. The above five cases ($2m \in \{0, 2, n-2, n-1, n\}$) will be dealt with at the end.

Now assume that $2 < 2m < n-2$. In this case for $k \in \{1, 2, n-1, n-2\}$ layers L_k are thin and $(2, n-2)$ is a thick interval (i.e., layers $L_3, L_4, \ldots, L_{n-3}$ are thick). Let $S\colon \Delta \to \mathbb{E}^2_\Delta$ be a characteristic surface for the interval $(2, n-2)$. The image of S is presented in Figure 2. Observe that S is the unique characteristic surface for this interval, and that it is an isometric embedding. Therefore we can identify Δ with $S(\Delta)$ and treat it as a subcomplex of $\mathbb{E}^2_\Delta$.

Let α denote the CAT(0) geodesic in the modified characteristic disc $\Delta' \subset \Delta \subset \mathbb{E}^2_\Delta$ between the midpoints of edges $[\sigma_2, \tau_2]$ and $[\sigma_{n-2}, \tau_{n-2}]$. Denote these midpoints by v_2 and v_{n-2} respectively. Let α' be a CAT(0) geodesic joining x and y. Both α and α' are shown in Figure 2. Note that α' does not appear in the construction of the Euclidean diagonal; it is an auxiliary geodesic that will be used to estimate distances between α and β.

Introduce a coordinate system on $\mathbb{E}^2_\Delta$ by setting x as the base vertex, and vectors $\overrightarrow{x\tau_1}$ and $\overrightarrow{xw}$, where w is the vertex of σ_1 that does not belong to τ_1, as the base vectors (see Figure 2). Note that both α and α' are contained in the sector of $\mathbb{E}^2_\Delta$ bounded by x and half-lines emanating from x in the directions of $\overrightarrow{x\tau_1}$ and $\overrightarrow{xw}$. Moreover, the Euclidean angle between α and $\overrightarrow{x\tau_1}$ and the Euclidean angle between α' and $\overrightarrow{x\tau_1}$ are

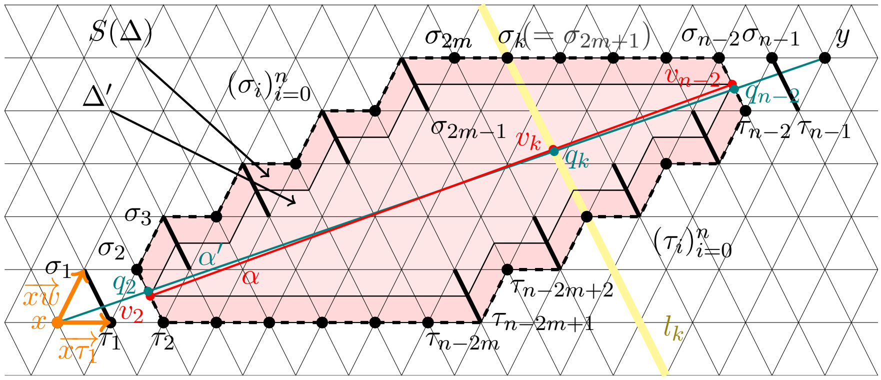

FIGURE 2. Image of the characteristic surface $S\colon \Delta \to \mathbb{E}^2_\Delta$ and the modified characteristic disc Δ'. Geodesics α and α' and the coordinate system. The CAT(0) distance between v_k and q_k is at most $\frac{1}{2}$.

both between $0°$ and $30°$. In particular α' intersects the interior of edges $[\sigma_2, \tau_2]$ and $[\sigma_{n-2}, \tau_{n-2}]$. Call the intersection points q_2 and q_{n-2} respectively.

Note that the CAT(0) distances between v_2 and q_2 and between v_{n-2} and q_{n-2} are less than $\frac{1}{2}$. For $k \in \{2, \ldots, n-2\}$ let q_k and v_k denote the intersection points of α and α' respectively with the layer L_k (see Figure 2). Since the edges $[\sigma_2, \tau_2]$ and $[\sigma_{n-2}, \tau_{n-2}]$ are parallel to all the layers, by elementary Euclidean geometry, for any $k \in \{2, \ldots, n-2\}$ we have:

$$d_{\mathbb{E}^2}(q_k, v_k) \leqslant \max\{d_{\mathbb{E}^2}(q_2, v_2), d_{\mathbb{E}^2}(q_{n-2}, v_{n-2})\} \leqslant \frac{1}{2}. \tag{6.1}$$

Now consider the CAT(0) geodesic β. First we determine the position of β with respect to the base vectors $\overrightarrow{x\tau_1}$ and $\overrightarrow{xw}$. Since both $\overrightarrow{xw}$ and $\overrightarrow{x\tau_1}$ coordinates of y are greater than the coordinates of x, and since γ is h–invariant, it follows that there are vertices $z \in \gamma$ such that (x, y, z) lie on γ in this order and both coordinates of z are arbitrarily large. Since β stays K–close to γ, it follows that the Euclidean angle directed counter-clockwise from $\overrightarrow{x\tau_1}$ to β is between $0°$ and $60°$. Examples of possible geodesics β and γ are shown in Figure 3.

It follows that β intersects every line l_k at a single point, which we denote by p_k. For any $k \in \{1, \ldots, n-1\}$ let t_k be a vertex of γ that belongs to layer L_k (and hence to the line l_k), and let $t_0 = x$ and $t_n = y$. By assumption, geodesic β passes within the CAT(0) distance K from any t_k. Then, since the directed angle from $\overrightarrow{x\tau_1}$ to β is between $0°$ and $60°$, we claim that for any $k \in \{0, \ldots, n\}$ we have:

$$d_{\mathbb{E}^2}(p_k, t_k) \leqslant \frac{2K}{\sqrt{3}}. \tag{6.2}$$

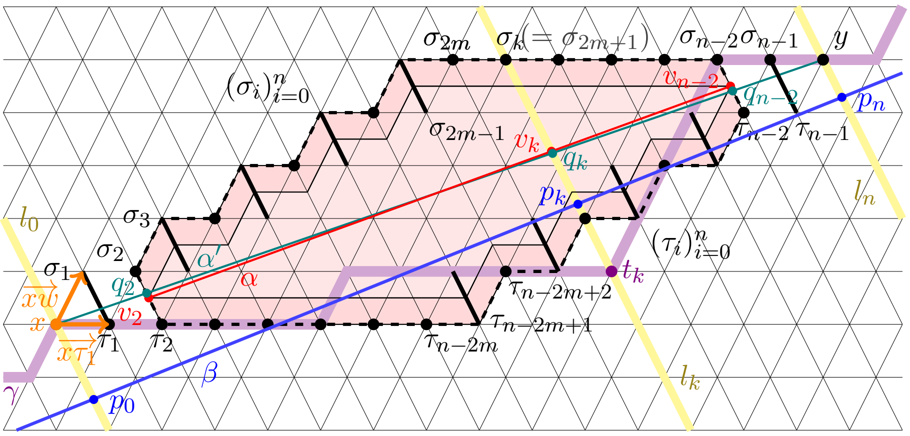

FIGURE 3. Examples of (parts of) geodesics β and γ and their position with respect to α and α'. Distances between vertices v_k, q_k, p_k, t_k are measured on the line l_k.

To obtain the constant $\frac{2K}{\sqrt{3}}$ $(=\frac{K}{\cos 30^\circ})$ one checks the extreme cases where β is parallel to either $\overrightarrow{x\tau_1}$ or $\overrightarrow{xw}$. In particular we have $d_{\mathbb{E}^2}(p_0, x) \leqslant \frac{2K}{\sqrt{3}}$ and $d_{\mathbb{E}^2}(p_n, y) \leqslant \frac{2K}{\sqrt{3}}$. Now since α' connects x and y and since β passes through p_0 and p_n it follows that for $k \in \{0, \ldots, n\}$ we have:

$$d_{\mathbb{E}^2}(p_k, q_k) \leqslant \frac{2K}{\sqrt{3}}. \tag{6.3}$$

Denote by $(\delta_i)_{i=0}^n$ the Euclidean geodesic between x and y. By definition, for thin layers we have $\delta_1 = \sigma_1$, $\delta_2 = [\sigma_2, \tau_2]$, $\delta_{n-2} = [\sigma_{n-2}, \tau_{n-2}]$ and $\delta_{n-1} = \tau_{n-1}$. For thick layers we have $\delta_i = u_i$ where $u_i \in L_i$ is a vertex that is at distance less than $\frac{1}{2}$ from v_i or $\delta_i = [u_i, u_i']$ if $u_i, u_i' \in L_i$ and $d_{\mathbb{E}^2}(u_i, v_i) = d_{\mathbb{E}^2}(u_i', q_k) = \frac{1}{2}$. In any case, for $k \in \{3, \ldots, n-3\}$ for any vertex $u_k \in \delta_k$ we have:

$$d_{\mathbb{E}^2}(u_k, v_k) \leqslant \frac{1}{2}. \tag{6.4}$$

Finally, by combining (6.1), (6.2), (6.3) and (6.4) for any $k \in \{3, \ldots, n-3\}$ for any $u_k \in \delta_k$ we have:

$$\begin{aligned} d_{\mathbb{E}^2}(t_k, u_k) &\leqslant d_{\mathbb{E}^2}(t_k, p_k) + d_{\mathbb{E}^2}(p_k, q_k) + d_{\mathbb{E}^2}(q_k, v_k) + d_{\mathbb{E}^2}(v_k, u_k) \\ &\leqslant \frac{2K}{\sqrt{3}} + \frac{2K}{\sqrt{3}} + \frac{1}{2} + \frac{1}{2} \\ &= \frac{4K}{\sqrt{3}} + 1. \end{aligned} \tag{6.5}$$

Since all the distances above are measured on the line l_k, the same estimate holds for the standard (combinatorial) distance.

For $k \in \{2, n-2\}$ we have $\delta_k = [\sigma_k, \tau_k]$ and one observes that $d_{\mathbb{E}^2}(q_k, \sigma_k) \leqslant 1$ and $d_{\mathbb{E}^2}(q_k, \tau_k) \leqslant 1$. Combining this with (6.2) and (6.3) we obtain the same estimate as in (6.5). Now for $k \in \{1, n-1\}$, let q_k be the point of intersection $\alpha' \cap \delta_k$. By a direct observation we see that for any vertex $u_k \in \delta_k$ we have $d_{\mathbb{E}^2}(v_k, u_k) \leqslant 1$, and thus again we obtain the same estimate as in (6.5).

We conclude that for any $k \in \{1, \ldots, n-1\}$ for any $u_k \in \delta_k$ we have

$$d(t_k, u_k) \leqslant \frac{4K}{\sqrt{3}} + 1.$$

Since x and y were arbitrary, this proves that γ is a $(\frac{4K}{\sqrt{3}} + 1)$–good geodesic.

For the remaining cases, where geodesics $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$ are close to each other ($2m \in \{0, 2, n-2, n-1, n\}$), we proceed analogously. One observes that the auxiliary CAT(0) geodesic α' joining x and y passes at the CAT(0) distance at most 1 from all the vertices of $(\delta_i)_{i=0}^n$ in all the appropriate layers. The rest of the argument goes the same as in the first case (i.e., one combines the above observation with (6.2) and (6.3) and obtains the same estimate as in (6.5)). □

A *flat* in a systolic complex X is an isometric embedding $F\colon \mathbb{E}^2_\Delta \hookrightarrow X$.

Lemma 6.5. *Let $F\colon \mathbb{E}^2_\Delta \hookrightarrow X$ be a flat, and suppose that $\gamma \subset \mathbb{E}^2_\Delta$ is a C'–good geodesic (where $\mathbb{E}^2_\Delta$ is treated as a systolic complex on its own). Then $F(\gamma)$ is a $(C' + 10)$–good geodesic in X.*

Proof. Pick any two vertices $x, y \in \mathbb{E}^2_\Delta$, let $(\delta_i)_{i=0}^m \subset \mathbb{E}^2_\Delta$ be the Euclidean geodesic between x and y in $\mathbb{E}^2_\Delta$, and let $(\widetilde{\delta_i})_{i=0}^m \subset X$ be the Euclidean geodesic between $F(x)$ and $F(y)$ in X. To prove the lemma it is enough to prove the following claim.

Claim. *For any $i \in \{0, 1, \ldots, m\}$ for any two vertices $z_i \in \delta_i$ and $\widetilde{z}_i \in \widetilde{\delta_i}$ we have*

$$d(F(z_i), \widetilde{z}_i) \leqslant 10.$$

The rest of the argument is devoted to proving the claim. Let $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$ be directed geodesics in $\mathbb{E}^2_\Delta$ going respectively from x to y and from y to x, and let $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$ be the corresponding directed geodesics between $F(x)$ and $F(y)$ in X (see Convention 5.1). As in Lemma 6.4, we will first deal with the case when there is a single thick interval for $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$, i.e., with the notation from Lemma 6.4, for $(\sigma_i)_{i=0}^n$ we assume that $2 < 2m < n-2$.

Consider the thick interval $(2, n-2)$ and for every $i \in \{2, 3, \ldots, n-2\}$ choose vertices $s_i \in \sigma_i$ and $t_i \in \tau_i$ that realise the thickness of layer i in $\mathbb{E}^2_\Delta$. Let $S\colon \Delta \to \mathbb{E}^2_\Delta$ be a characteristic surface for the cycle

$$\alpha = (s_2, s_3, \ldots, s_{n-2}, t_{n-2}, t_{n-1}, \ldots, t_2, s_2).$$

Observe that S is the unique characteristic surface for the interval $(2, n-2)$, and that it is an isometric embedding. Therefore we will identify the characteristic disk Δ with its image $S(\Delta) \subset \mathbb{E}^2_\Delta$.

The map $F\big|_\Delta \colon \Delta \to X$ does not have to be a characteristic surface for geodesics $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$, e.g., not all of the vertices of $F(\alpha)$ belong to the appropriate simplices of $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$. A priori we do not even know whether $(2, n-2)$ is a thick interval for these geodesics.

We will now show how to modify F to obtain a characteristic surface for $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$. The idea is that the images $F((\sigma_i)_{i=0}^n)$ and $F((\tau_i)_{i=0}^n))$ are 1–close to the geodesics $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$ respectively, and therefore a small perturbation of the map F would give the desired characteristic surface. We will now make this idea precise. For every $i \in \{2, 4, 6, \dots, 2m\}$ choose any vertex $u_i \in \widetilde{\sigma}_i$ and for every $i \in \{n-2m, n-2m+2, n-2m+4, \dots, n-2\}$ choose any vertex $w_i \in \widetilde{\tau}_i$. Then by [Els09c, Lemma 3.9] (applied simultaneously to geodesics $(\sigma_i)_{i=0}^n$ and $(\tau_i)_{i=0}^n$) there exists a flat $F' \colon \mathbb{E}^2_\Delta \to X$ such that:

(1) for $i \in \{2, 4, \dots, 2m\}$ we have $F'(\sigma_i) = u_i$,
(2) for $i \in \{n-2m, n-2m+2, \dots, n-2\}$ we have $F'(\tau_i) = w_i$,
(3) for every vertex $x \in \mathbb{E}^2_\Delta$ not considered in (1) and (2), we have $F'(x) = F(x)$,
(4) for every $0 \leqslant i \leqslant n$ we have $F'(\sigma_i) = \mathrm{Im}(F') \cap \widetilde{\sigma}_i$ and $F'(\tau_i) = \mathrm{Im}(F') \cap \widetilde{\tau}_i$.

Images of F and F' are shown in Figure 4.

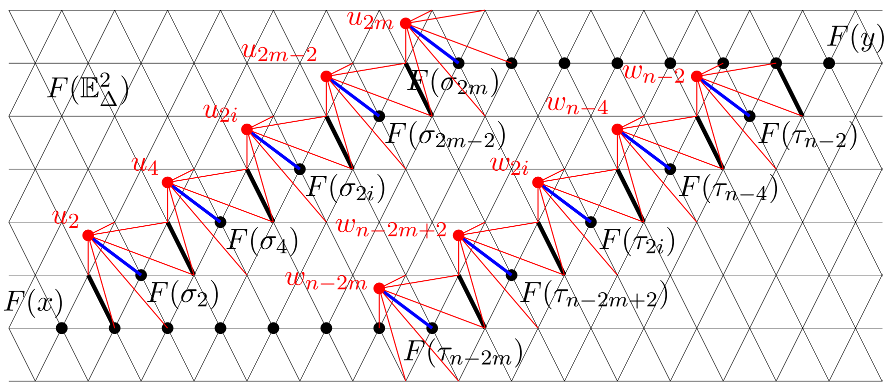

FIGURE 4. Images of flats F and F'. The part of $F'(\mathbb{E}^2_\Delta)$ which differs from $F(\mathbb{E}^2_\Delta)$ is red.

Observe that for $i \in \{2, 4, \dots, 2m\}$ vertices $F(\sigma_i)$ and $F'(\sigma_i) = u_i$ are connected. Indeed, since F and F' agree on all the neighbours of σ_i in $\mathbb{E}^2_\Delta$, taking a pair of neighbours $v_1, v_2 \in \mathbb{E}^2_\Delta$ of σ_i that are not adjacent gives a 4–cycle

$$(F(v_1), F(\sigma_i), F(v_2), F'(\sigma_i)).$$

Since X is 6–large, this cycle has a diagonal. It cannot be $[F(v_1), F(v_2)]$ since F is an isometric embedding, and v_1 and v_2 are not adjacent. Thus it must be $[F(\sigma_i), F'(\sigma_i)]$.

Analogously, for any $i \in \{n-2m, n-2m+2, \ldots, n-2\}$ vertices $F(\tau_i)$ and $F'(\tau_i) = w_i$ are connected. The above assertions, together with property (3) of the map F', imply that for any vertex $x \in \mathbb{E}^2_\Delta$ we have

$$d(F(x), F'(x)) \leqslant 1.$$

We claim that for any $2 \leqslant i \leqslant n-2$ the vertices $F'(s_i) \in \widetilde{\sigma}_i$ and $F'(t_i) \in \widetilde{\tau}_i$ realise the thickness of layer i for $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$ in X, and so layers $L_3, L_4, \ldots, L_{n-3}$ are thick, and layers L_2 and L_{n-2} are thin. This follows essentially from the fact that $F' \colon \mathbb{E}^2_\Delta \to X$ is an isometric embedding. Similarly, one can show that layers L_1 and L_{n-1} in X are thin. Finally, we conclude that

$$F'\big|_\Delta \colon \Delta \to X$$

is a characteristic surface for the thick interval $(2, n-2)$ for geodesics $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$ in X. Let $(\rho_i)_{i=3}^{n-3}$ be the Euclidean diagonal for $\Delta \subset \mathbb{E}^2_\Delta$ and denote by $(\delta_i)_{i=0}^m$ the Euclidean geodesic between vertices x and y in $\mathbb{E}^2_\Delta$. We have $\delta_i = \operatorname{span}\{\sigma_i, \tau_i\}$ for $i \in \{1, 2, n-2, n-1\}$. For all remaining i we have $\delta_i = \rho_i$ (since $\Delta \cong S(\Delta) \subset \mathbb{E}^2_\Delta$ is the unique characteristic surface for the thick interval $(2, n-2)$). For any $i \in \{0, 1, \ldots, n\}$, for any vertex $z_i \in \delta_i$ we have

$$d(F(z_i), F'(z_i)) \leqslant 1. \tag{6.6}$$

For any $i \in \{0, 1, \ldots, n\}$, for any vertices $z_i \in \delta_i$ and $\widetilde{z}_i \in \widetilde{\delta}_i$ we claim that

$$d(F'(z_i), \widetilde{z}_i) \leqslant 1. \tag{6.7}$$

This follows for $i \in \{1, 2, n-2, n-1\}$ from the property (4) of the map $F' \colon \mathbb{E}^2_\Delta \to X$. Namely we have that $F'(\sigma_i) \subset \widetilde{\sigma}_i$ and $F'(\tau_i) \subset \widetilde{\tau}_i$, and by definition $\delta_i = \operatorname{span}\{\sigma_i, \tau_i\}$ and $\widetilde{\delta}_i = \operatorname{span}\{\widetilde{\sigma}_i, \widetilde{\tau}_i\}$. For $i \in \{3, 4, \ldots, n-3\}$ by definition of a Euclidean geodesic and the fact that $F'\big|_\Delta$ is a characteristic surface for $(\widetilde{\sigma}_i)_{i=0}^n$ and $(\widetilde{\tau}_i)_{i=0}^n$ we obtain that $F'(\rho_i) \subset \widetilde{\delta}_i$.

Finally, combining (6.6) and (6.7), for any $i \in \{0, 1, \ldots, n\}$ for any two vertices $z_i \in \delta_i$ and $\widetilde{z}_i \in \widetilde{\delta}_i$ we have

$$d(F(z_i), \widetilde{z}_i) \leqslant 2.$$

This finishes the proof of the claim under the assumption that $2 < 2m < n-2$.

Now assume $2m \in \{0, 2, n-2, n-1, n\}$. In this case any layer $L_i \subset \mathbb{E}^2_\Delta$ has thickness at most 2. By [Els09c, Proposition 3.8] (which is a weaker formulation of [Els09c, Lemma 3.9] used above) for any vertices $s_i \in \sigma_i$ and $u_i \in \widetilde{\sigma}_i$ we have $d(F(s_i), u_i) \leqslant 2$. The same estimate holds for vertices of τ_i and $\widetilde{\tau}_i$. It follows from the triangle inequality, that for any $i \in \{0, 1, \ldots, n\}$ the thickness of the layer i in X is at most 6.

Observe that by definition of the Euclidean geodesic, any simplex δ_i lies between simplices σ_i and τ_i in the layer $L_i \subset \mathbb{E}^2_\Delta$. More precisely, the distance between any vertex $z_i \in \delta_i$ and any vertex $u_i \in \sigma_i$ is less than the thickness of L_i. Clearly the same estimate holds for vertices of $\widetilde{\delta}_i$ and $\widetilde{\sigma}_i$, if one replaces thickness of L_i by thickness of layer i in X. From these considerations we conclude that for any $i \in \{0, 1, \ldots, n\}$, for any $z_i \in \delta_i$ and $\widetilde{z}_i \in \widetilde{\delta}_i$ and for any vertices $s_i \in \sigma_i$ and $u_i \in \widetilde{\sigma}_i$ we have:

$$d(F(z_i), \widetilde{z}_i) \leqslant d(F(z_i), F(s_i)) + d(F(s_i), u_i) + d(u_i, \widetilde{z}_i) \leqslant 2 + 2 + 6 = 10.$$

This estimate is by no means optimal. □

We are ready now to prove Theorem 6.1.

Proof of Theorem 6.1. **Case 1:** $\mathrm{Min}(h)$ **is** h**–cocompact.** Let K be the constant appearing in Lemma 6.3. Since $\mathrm{Min}(h)$ is h–cocompact, by Lemma 4.2 the subcomplex $\mathrm{Disp}_K(h)$ is h–cocompact as well. Pick a vertex $x \in \mathrm{Min}(h) \subset \mathrm{Disp}_K(h)$, and for any $n \geqslant 0$ consider vertices $h^{-n} \cdot x, h^n \cdot x \in \mathrm{Min}(h)$. Note that $d(h^{-n} \cdot x, h^n \cdot x)$ is not necessarily equal to $2n \cdot L(h)$, but we can assume that it is even (by passing to a subsequence of the form $n_i = ik$ for some $k \geqslant 1$ if necessary, see [Els09b, Theorem 1.1]). Put $m_n = \frac{1}{2} \cdot d(h^{-n} \cdot x, h^n \cdot x)$ and let $(\delta^n_i)^{m_n}_{i=-m_n}$ be the Euclidean geodesic between $h^{-n} \cdot x$ and $h^n \cdot x$. By Lemma 6.3 we have $(\delta^n_i)^{m_n}_{i=-m_n} \subset \mathrm{Disp}_K(h)$. Since $\mathrm{Disp}_K(h)$ is h–cocompact, there exists $R > 0$ such that for every n the geodesic $(\delta^n_i)^{m_n}_{i=-m_n}$ intersects the ball $B_R(x, X)$. Let i_n be an integer such that $\delta^n_{i_n}$ is a simplex of $(\delta^n_i)^{m_n}_{i=-m_n}$ that intersects $B_R(x, X)$ (such i_n is not unique in general, we choose one for each n). By replacing R with $R + 1$ we can assume that $\delta^n_{i_n} \subset B_R(x, X)$.

Since the ball $B_R(x, X)$ contains only finitely many simplices, there are infinitely many n such that $\delta^n_{i_n}$ is equal to a fixed simplex of $B_R(x, X)$. Denote this simplex by $\widetilde{\delta}_0$. Now since the sphere $S_1(\widetilde{\delta}_0, X)$ is finite, among geodesics $(\delta^n_i)^{m_n}_{i=-m_n}$ for which $\delta^n_{i_n} = \widetilde{\delta}_0$ there are infinitely many such that $\delta^n_{i_n+1}$ is equal to a fixed simplex $\widetilde{\delta}_1$ and $\delta^n_{i_n-1}$ is equal to a fixed simplex $\widetilde{\delta}_{-1}$. By continuing this procedure for spheres $S_k(\widetilde{\delta}_0, X)$ for $k > 1$, we obtain a bi-infinite sequence of simplices

$$(\widetilde{\delta}_i)^{\infty}_{i=-\infty} \subset \mathrm{Disp}_K(h),$$

such that any of its finite subsequences is a Euclidean geodesic.

By Proposition 5.3 for any finite subsequence, say $(\widetilde{\delta}_i)^m_{i=-m}$, there exists a geodesic $\gamma_m = (v_i)^m_{i=-m}$ such that $v_i \in \widetilde{\delta}_i$. By Theorem 5.4 any γ_m is a good geodesic. By a diagonal argument, from the sequence $(\gamma_m)^{\infty}_{m=0}$ we can extract a bi-infinite geodesic $\gamma = (v_i)^{\infty}_{i=-\infty}$, which is a good geodesic, as any of its finite subgeodesics is contained in a good geodesic γ_m for some $m > 0$. Since for every $i \in \mathbb{N}$ we have $v_i \in \widetilde{\delta_i} \subset \mathrm{Disp}_K(h)$, we conclude that $\gamma \subset \mathrm{Disp}_K(h)$.

Case 2. $\mathrm{Min}(h)$ **is not** h**–cocompact.** By [OP16, Corollary 5.8] the centraliser $C_G(h)$ is commensurable with the product $F_n \times \mathbb{Z}$, such that the subgroup $\langle h \rangle \subset C_G(h)$ is commensurable with the '$\mathbb{Z}$' factor of the latter. By Theorem 4.3 the group $C_G(h)$ acts cocompactly on $\mathrm{Min}(h)$. Since $\mathrm{Min}(h)$ is not h–cocompact, we conclude that $n \geqslant 1$, and so there exists an element $g \in C_G(h)$ such that $\langle g, h \rangle \cong \mathbb{Z}^2$. By the Flat Torus Theorem ([Els09a, Theorem 6.1]) there exists a flat $F \colon \mathbb{E}^2_\Delta \to X$ whose image is preserved by the action of $\langle g, h \rangle$. We will now construct an h–invariant geodesic $\gamma \subset F(\mathbb{E}^2_\Delta)$ which satisfies the assumptions of Lemma 6.4.

Take any vertex $x \in F(\mathbb{E}^2_\Delta)$ and consider a CAT(0) geodesic γ' in $F(\mathbb{E}^2_\Delta)$ that passes through vertices x and $h \cdot x$. The isometry h acts on $F(\mathbb{E}^2_\Delta) \cong \mathbb{E}^2$ as a translation along γ' by distance equal to the CAT(0) length of segment $\gamma'\big|_{[x,h\cdot x]}$. Let α be any (combinatorial) geodesic between x and $h \cdot x$ that is Hausdorff 1–close to $\gamma'\big|_{[x,h\cdot x]}$. (To obtain such α one proceeds similarly as when defining the Euclidean diagonal in a characteristic disk in Subsection 5.5.) Define γ as

$$\gamma = \bigcup_{n \in \mathbb{Z}} (h^n \cdot \alpha).$$

By definition γ is an h–invariant geodesic, that is Hausdorff 1–close to a CAT(0) geodesic γ' in $F(\mathbb{E}^2_\Delta)$. By Lemma 6.4 we get that γ is a $(\frac{4}{\sqrt{3}} + 1)$–good geodesic in $F(\mathbb{E}^2_\Delta)$. Lemma 6.5 implies that γ is a $(\frac{4}{\sqrt{3}} + 11)$–good geodesic in X. This implies that γ is a good geodesic in X since we have $\frac{4}{\sqrt{3}} + 11 < C$, where C is the constant appearing in Definition 5.7 (cf. the discussion at the beginning of Subsection 5.7). □

References

[Che00] Victor Chepoi, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179. MR1748966 (2001a:57004)

[Els09a] Tomasz Elsner, *Flats and the flat torus theorem in systolic spaces*, Geom. Topol. **13** (2009), no. 2, 661–698, DOI 10.2140/gt.2009.13.661. MR2469526 (2009m:20065)

[Els09b] ———, *Isometries of systolic spaces*, Fund. Math. **204** (2009), no. 1, 39–55, DOI 10.4064/fm204-1-3. MR2507689 (2010g:51005)

[Els09c] ———, *Systolic spaces with isolated flats* (2009), preprint, available at `http://www.math.uni.wroc.pl/~elsner/papers/isolated-flats.pdf`.

[JŚ06] Tadeusz Januszkiewicz and Jacek Świątkowski, *Simplicial nonpositive curvature*, Publ. Math. Inst. Hautes Études Sci. **104** (2006), 1–85. MR2264834 (2007j:53044)

[Ont05] Pedro Ontaneda, *Cocompact CAT(0) spaces are almost geodesically complete*, Topology **44** (2005), no. 1, 47–62, DOI 10.1016/j.top.2004.01.010. MR2104000

[Osa07] Damian Osajda, *Connectedness at infinity of systolic complexes and groups*, Groups Geom. Dyn. **1** (2007), no. 2, 183–203, DOI 10.4171/GGD/9. MR2319456

[OP16] Damian Osajda and Tomasz Prytuła, *Classifying spaces for families of subgroups for systolic groups* (2016), preprint, available at `https://arxiv.org/abs/1604.08478`.

[OP09] Damian Osajda and Piotr Przytycki, *Boundaries of systolic groups*, Geom. Topol. **13** (2009), no. 5, 2807–2880, DOI 10.2140/gt.2009.13.2807. MR2546621

[PS09] Panos Papasoglu and Eric Swenson, *Boundaries and JSJ decompositions of CAT(0)-groups*, Geom. Funct. Anal. **19** (2009), no. 2, 559–590, DOI 10.1007/s00039-009-0012-8. MR2545250

[Rua01] Kim E. Ruane, *Dynamics of the action of a* CAT(0) *group on the boundary*, Geom. Dedicata **84** (2001), no. 1-3, 81–99, DOI 10.1023/A:1010301824765. MR1825346

[SC83] V. P. Soltan and V. D. Chepoĭ, *Conditions for invariance of set diameters under d-convexification in a graph*, Kibernetika (Kiev) **6** (1983), 14–18, DOI 10.1007/BF01068561 (Russian, with English summary); English transl., Cybernetics **19** (1983), no. 6, 750–756 (1984). MR765117 (86k:05102)

Institut for Matematiske Fag, Københavns Universitet, Universitetsparken 5, 2100 København Ø, Denmark

E-mail address: `t.prytula@math.ku.dk`