Tree-graded spaces and asymptotic cones of groups

Cornelia Druţu and Mark Sapir*

with an Appendix by Denis Osin and Mark Sapir

Abstract

We introduce a concept of tree-graded metric space and we use it to show quasi-isometry invariance of certain classes of relatively hyperbolic groups, to obtain a characterization of relatively hyperbolic groups in terms of their asymptotic cones, to find geometric properties of Cayley graphs of relatively hyperbolic groups, and to construct the first example of finitely generated group with a continuum of non-π_1-equivalent asymptotic cones. Note that by a result of Kramer, Shelah, Tent and Thomas, continuum is the maximal possible number of different asymptotic cones of a finitely generated group, provided that the Continuum Hypothesis is true.

Contents

*The first author is grateful to the CNRS of France for granting her the délégation CNRS status in the 2003-2004 academic year. The research of the second author was supported in part by the NSF grant DMS 0072307 and by the US-Israeli BSF grant 1999298. We worked on the present paper during the first author's visits to Vanderbilt University, and the second author's visits to University of Lille-1 and to the IHES. The authors are grateful to the administration and faculty of these institutions for their support and hospitality.

1 Introduction

An asymptotic cone of a metric space is, roughly speaking, what one sees when one looks at the space from infinitely far away. More precisely, any asymptotic cone of a metric space (X, dist) corresponds to an ultrafilter ω, a sequence of observation points $e = (e_n)_{n\in\mathbb{N}}$ from X and a sequence of scaling constants $d = (d_n)_{n\in\mathbb{N}}$ diverging to $+\infty$. The cone $\mathrm{Con}(X; e, d)$ corresponding to e and d is the ω-limit of the sequence of spaces with basepoints $(X, \mathrm{dist}/d_n, e_n)$ (see Section 3 for precise definitions).

In particular, if X is the Cayley graph of a group G with a word metric then the asymptotic cones of X are called asymptotic cones of G.

The concept of asymptotic cone was essentially used by Gromov in [Gr$_1$] and then formally introduced by van den Dries and Wilkie [VDW].

Asymptotic cones have been used to characterize important classes of groups:

- A finitely generated group is virtually Abelian if and only if its asymptotic cones are isometric to the Euclidean space $\mathbb{R}^n$ ([Gr$_1$], [Pa]).
- A finitely generated group is virtually nilpotent if and only if its asymptotic cones are locally compact ([Gr$_1$], [VDW], [Dr$_4$]).
- A finitely generated group is hyperbolic if and only if all its asymptotic cones are $\mathbb{R}$-trees ([Gr$_3$]).

In [DP$_1$] it is shown moreover that asymptotic cones of non-elementary hyperbolic groups are all isometric to the complete homogeneous $\mathbb{R}$-tree of valence continuum. The asymptotic cones of elementary groups are isometric to either a line $\mathbb{R}$ (if the group is infinite) or to a point. In particular, every hyperbolic group has only one asymptotic cone up to isometry.

Asymptotic cones of quasi-isometric spaces are bi-Lipschitz equivalent. In particular the topology of an asymptotic cone of a finitely generated group does not depend on the choice of the generating set. This was used in [KaL_1] and [KaL_2] to prove rigidity results for fundamental groups of Haken manifolds, in [KlL] to prove rigidity for cocompact lattices in higher rank semisimple groups, and in [Dr_2] to provide an alternative proof of the rigidity for non-cocompact lattices in higher rank semisimple groups. For a survey of results on quasi-isometry invariants and their relations to asymptotic cones see [Dr_4].

The power of asymptotic cones stems from the fact that they capture both geometric and logical properties of the group, since a large subgroup of the ultrapower $\Pi G/\omega$ of the group G acts transitively by isometries on the asymptotic cone $\mathrm{Con}_\omega(G; e, d)$. Logical aspects of asymptotic cones are studied and used in the recent papers by Kramer, Shelah, Tent and Thomas [KSTT], [KT].

One of the main properties of asymptotic cones of a metric space X is that geometry of finite configurations of points in the asymptotic cone reflects the "coarse" geometry of similar finite configurations in X. This is the spirit of Gromov-Delzant's approximation statement [CDP] and of the applications of $\mathbb{R}$-trees to Rips-Sela theory of equations in hyperbolic groups and homomorphisms of hyperbolic groups [RiSe]. This was also used in Druţu's proof of hyperbolicity of groups with sub-quadratic isoperimetric inequality [Dr_3].

By a result of Gromov [Gr_3] if all asymptotic cones of a finitely presented group are simply connected then the group has polynomial isoperimetric function and linear isodiametric function. Papasoglu proved in [Pp] that groups having quadratic isoperimetric functions have simply connected asymptotic cones. In general, asymptotic cones of groups are not necessarily simply connected [Tr]. In fact, if a group G is not finitely presented then its asymptotic cones cannot all be simply connected [Gr_3, Dr_4]. A higher-dimensional version of this result is obtained by Riley [Ri]. According to the result of Gromov cited above, examples of finitely presented groups with non-simply connected asymptotic cones can be found in [Bri] and [SBR].

Although asymptotic cones can be completely described in some cases, the general perception is nevertheless that asymptotic cones are usually large and "undescribable". This might be the reason of uncharacteristically "mild" questions by Gromov [Gr_3]:

Problem 1.1. Which groups can appear as subgroups in fundamental groups of asymptotic cones of finitely generated groups?

Problem 1.2. Is it true that the fundamental group of an asymptotic cone of a group is either trivial or uncountable?

In [Gr_3], Gromov also asked the following question.

Problem 1.3. How many non-isometric asymptotic cones can a finitely generated group have?

A solution of Problem 1.1 was given by Erschler and Osin [EO]. They proved that every metric space satisfying some weak properties can be π_1- and isometrically embedded into the asymptotic cone of a finitely generated group. This implies that every countable group is a subgroup of the fundamental group of an asymptotic cone of a finitely generated group.

Notice that since asymptotic cones tend to have fundamental groups of order continuum, this result does not give information about the structure of the whole fundamental group of an asymptotic cone, or about how large the class of different asymptotic cones is: there exists a group of cardinality continuum that contains all countable groups as subgroups (for example, the group of all permutations of a countable set). One of the goals of this paper is to get more precise information about fundamental groups of asymptotic cones, and about the whole set of different asymptotic cones of a finitely generated group.

Problem 1.3 turned out to be related to the Continuum Hypothesis (i.e. the famous question of whether there exists a set of cardinality strictly between $\aleph_0$ and $2^{\aleph_0}$). Namely, in [KSTT], it is proved that if the Continuum Hypothesis is not true then any uniform lattice in $\mathrm{SL}_n(\mathbb{R})$ has $2^{2^{\aleph_0}}$ non-isometric asymptotic cones, but if the Continuum Hypothesis is true then any uniform lattice in $\mathrm{SL}_n(\mathbb{R})$ has exactly one asymptotic cone up to isometry, moreover the maximal theoretically possible number of non-isometric asymptotic cones of a finitely generated group is continuum. Recall that the Continuum Hypothesis is independent of the usual axioms of set theory (ZFC).

It is known, however, that even if the Continuum Hypothesis is true, there exist groups with more than one non-homeomorphic asymptotic cones [TV]. Nevertheless, it was not known whether there exists a group with the maximal theoretically possible number of non-isometric asymptotic cones (continuum).

In [Gr_2], Gromov introduced a useful generalization of hyperbolic groups, namely the relatively hyperbolic groups[1]. This class includes:

(1) geometrically finite Kleinian groups; these groups are hyperbolic relative to their cusp subgroups;

(2) fundamental groups of hyperbolic manifolds of finite volume (that is, non-uniform lattices in rank one semisimple groups with trivial center); these are hyperbolic relative to their cusp subgroups;

(3) hyperbolic groups; these are hyperbolic relative to the trivial subgroup or more generally to collections of quasi-convex subgroups satisfying some extra conditions;

(4) free products of groups; these are hyperbolic relative to their factors;

(5) fundamental groups of non-geometric Haken manifolds with at least one hyperbolic component; these are hyperbolic relative to the fundamental groups of the maximal graph-manifold components and to the fundamental groups of the tori and Klein bottles not contained in graph-manifold components [Bow_3];

(6) ω-residually free groups (limit groups in another terminology); these are hyperbolic relative to the collection of maximal Abelian non-cyclic subgroups [Dah_1].

There exist several characterizations of relatively hyperbolic groups which are in a sense parallel to the well known characterizations of hyperbolic groups (see [Bow_1], [Fa], [Os], [Dah_2], [Ya] and references therein). But there was no characterization in terms of asymptotic cones. Also, it was not known whether being relatively hyperbolic with respect to any kind of subgroups is a quasi-isometry invariant, except for hyperbolic groups when quasi-isometry invariance is true.

The following theorems are the main results of the paper (we formulate these results not in the most general form).

The first theorem gives more information about the possible structure of fundamental groups of asymptotic cones.

[1] These groups are also called *strongly relatively hyperbolic* in order to distinguish them from weakly relatively hyperbolic groups in the sense of Farb.

Theorem 1.4 (Theorem 7.33 and Corollary 7.32). *(1) For every countable group C, the free product of continuously many copies of C is the fundamental group of an asymptotic cone of a 2-generated group.*

(2) There exists a 2-generated group Γ such that for every finitely presented group G, the free product of continuously many copies of G is the fundamental group of an asymptotic cone of Γ.

The second theorem answers the question about the number of asymptotic cones of a finitely generated group.

Theorem 1.5 (Theorem 7.37). *Regardless of whether the Continuum Hypothesis is true or not, there exists a finitely generated group G with continuously many pairwise non-π_1-equivalent asymptotic cones.*

The third theorem shows that large classes of relatively hyperbolic groups are closed under quasi-isometry.

Theorem 1.6 (Corollary 5.19). *Let G be a finitely generated group that is hyperbolic relative to subgroups $H_1, ..., H_m$. Suppose that none of the asymptotic cones of H_i has global cut-points, $i = 1, ..., m$.*

Let G' be a group that is quasi-isometric to G. Then G' is hyperbolic relative to subgroups $H'_1, ..., H'_n$ each of which is quasi-isometric to one of $H_1, ..., H_m$.

The cardinal m of the finite collection of "parabolic" subgroups $\{H_i\}_{i\in I}$ in Theorem 1.6 is not a quasi-isometry invariant. This can be seen for instance for the fundamental groups of a finite volume hyperbolic manifold and of a finite covering of it.

There are previous results showing that some special classes of relatively hyperbolic groups are closed under quasi-isometry: the class of fundamental groups of non-geometric Haken manifolds with at least one hyperbolic component ([KaL_1], [KaL_2]) and the class of non-uniform lattices of isometries of a rank one symmetric space [Sch]. Also, in [PW] the classification up to quasi-isometry of free products of groups, and more generally of fundamental groups of graphs of groups with finite edge groups is treated.

The main ingredient in the proof of Theorem 1.6 is the following result, interesting by itself.

Theorem 1.7 (Corollary 5.5). *Let G be a finitely generated group that is hyperbolic relative to subgroups $H_1, ..., H_m$, and let S be a group with the property that none of its asymptotic cones has global cut-points. Then the image of S under any (L, C)-quasi-isometry $S \to G$ is in an M-tubular neighborhood of a coset gH_i, $g \in G, i = 1, ..., m$, where M depends on L, C, G and S.*

Note that the hypothesis of Theorem 1.7 that the group S does not have cut-points in its asymptotic cones clearly cannot be removed. An example illustrating this is the relatively hyperbolic group G itself which, if all its "parabolic" subgroups H_i are infinite proper subgroups, has cut-points in every asymptotic cone, and which is not contained in the bounded neighborhood of a left coset of a subgroup H_i.

A result similar to Theorem 1.7 is obtained in [PW, §3] for G a fundamental group of a graph of groups with finite edge groups and S a one-ended group. We should note here that the class of groups with the property that none of their asymptotic cones has global cut-points is inside the class of one-ended groups, by Stallings' Ends Theorem. Moreover, the inclusion is strict, since the asymptotic cones of any hyperbolic group are $\mathbb{R}$-trees.

Theorem 1.7 in particular gives information about which subgroups with none of their asymptotic cones having global cut-points can appear as undistorted subgroups in a relatively hyperbolic group (see Remark 8.28, (1)). The following theorem clarifies even more the question of the structure of undistorted subgroups in relatively hyperbolic groups.

Theorem 1.8 (Theorem 8.27). *Let $G = \langle S \rangle$ be a finitely generated group that is hyperbolic relative to subgroups $H_1, ..., H_n$. Let $G_1 = \langle S_1 \rangle$ be an undistorted finitely generated subgroup of G. Then G_1 is relatively hyperbolic with respect to subgroups $H_1', ..., H_m'$, where each H_i' is one of the intersections $G_1 \cap gH_jg^{-1}$, $g \in G$, $j \in \{1, 2, \dots, n\}$.*

We also obtain information about the automorphism group of a relatively hyperbolic group.

Theorem 1.9 (Corollary 8.29). *Let G be a finitely generated group that is relatively hyperbolic with respect to a subgroup H. Suppose that all asymptotic cones of H are without global cut-points. Let $\mathrm{Fix}(H)$ be the subgroup of the automorphism group of G consisting of the automorphisms that fix H as a set. Then:*

(1) $\mathrm{Inn}(G)\mathrm{Fix}(H) = \mathrm{Aut}(G)$.

(2) $\mathrm{Inn}(G) \cap \mathrm{Fix}(H) = \mathrm{Inn}_H(G)$, *where* $\mathrm{Inn}_H(G)$ *is by definition* $\{i_h \in \mathrm{Inn}(G) \mid h \in H\}$.

(3) There exists a natural homomorphism from $\mathrm{Out}(G)$ to $\mathrm{Out}(H)$ given by $\phi \mapsto i_{g_\phi}\phi|_H$, where g_ϕ is an element of G such that $i_{g_\phi}\phi \in \mathrm{Fix}(H)$, and $\psi|_H$ denotes the restriction of an automorphism $\psi \in \mathrm{Fix}(H)$ to H.

Examples of groups such that none of their asymptotic cones has global cut-points:

- Groups satisfying a law (see Corollary 6.9). We recall that a *law* is a word w in n letters $x_1, \dots, x_n$ and *a group satisfying the law* w is a group G such that $w = 1$ in G whenever $x_1, \dots, x_n$ are replaced by an arbitrary set of n elements in G. For instance Abelian groups are groups with the law $w = x_1x_2x_1^{-1}x_2^{-1}$. More generally, solvable groups are groups with a law, and so are Burnside groups. While for nilpotent groups the results of Theorems 1.6 and 1.7 are not surprising and were already known in some particular cases of relatively hyperbolic groups [Sch], for solvable non-nilpotent groups and for Burnside groups the situation is different. For instance the group Sol has asymptotic cones composed of continuously many Hawaiian earrings [Bu], so it is not *a priori* clear why such a group should have a rigid behavior with respect to quasi-isometric embeddings into relatively hyperbolic groups. Burnside groups display a similar picture.

  In the particular case of groups with a law, the constant M in Theorem 1.7 depends only on the law and not on the group S (Corollary 6.10).

- Groups not virtually cyclic with elements of infinite order in the center (see Theorem 6.3);

- Groups of isometries acting properly discontinuously and with compact quotients on products of symmetric spaces and Euclidean buildings, of rank at least two. The asymptotic cones of such groups are Euclidean buildings of rank at least two [KlL]. Most likely the same is true for such groups of isometries so that the quotients have finite volume, but the proof of this statement is not straightforward.

- Fundamental groups of closed graph-manifolds. This follows from results in [KaL$_3$] as well as from Corollary 6.9 in the particular case when the manifold has the *Sol*-geometry.

The main tool in this paper are tree-graded spaces.

Definition 1.10 (tree-graded spaces). Let $\mathbb{F}$ be a complete geodesic metric space and let $\mathcal{P}$ be a collection of closed geodesic subsets of $\mathbb{F}$ (called *pieces*) such that the following two properties are satisfied:

(T_1) Every two different pieces have at most one common point.

(T_2) Every simple geodesic triangle (a simple loop composed of three geodesics) in $\mathbb{F}$ is contained in one piece.

Then we say that the space $\mathbb{F}$ is *tree-graded with respect to* $\mathcal{P}$.

The main interest in the notion of tree-graded space resides in the following characterization of relatively hyperbolic groups of which the converse part is proven in Section 8 and the direct part in the Appendix written by D. Osin and M. Sapir.

Theorem 1.11 (Theorem 8.5). *A finitely generated group G is relatively hyperbolic with respect to finitely generated subgroups $H_1, ..., H_n$ if and only if every asymptotic cone* $\mathrm{Con}_\omega(G; e, d)$ *is tree-graded with respect to ω-limits of sequences of cosets of the subgroups H_i.*

Section 2 contains many general properties of tree-graded spaces.

In particular, by Lemma 2.28 any complete homogeneous geodesic metric space with global cut-points is tree-graded with respect to a certain uniquely defined collection of pieces without cut-points.

We prove in Proposition 2.15, that the property (T_2) in the definition of tree-graded spaces can be replaced by the property (T_2'):

> For every topological arc $\mathfrak{c} : [0, d] \to \mathbb{F}$ and $t \in [0, d]$, let $\mathfrak{c}[t-a, t+b]$ be a maximal sub-arc of $\mathfrak{c}$ containing $\mathfrak{c}(t)$ and contained in one piece. Then every other topological arc with the same endpoints as $\mathfrak{c}$ must contain the points $\mathfrak{c}(t-a)$ and $\mathfrak{c}(t+b)$.

Moreover, when (T_2) is replaced by (T_2') the condition that the pieces are geodesic is no longer needed. Thus, if we do not ask that the whole space be geodesic either, tree-graded spaces can be considered in a purely topological setting.

Notice that there are similarities in the study of asymptotic cones of groups and that of boundaries of groups. Boundaries of groups do not necessarily have a natural metric, and rarely are geodesic spaces, but they have a natural topology and they are also, in many interesting cases, homogeneous spaces with respect to actions by homeomorphisms. Thus, if the boundary of a group is homogeneous and has a global cut-point then most likely it is tree-graded (in the topological sense) with respect to pieces that do not have cut-points. Such a study of boundaries of groups with global cut-points appeared, for example, in the work of Bowditch [Bow_2] on the Bestvina-Mess conjecture. Bowditch developed a general theory appropriate for the study of topological homogeneous spaces with global cut-points that is related to the study of tree-graded spaces that we do in this paper. Results related to Bowditch's work in this general setting can be found in [AN].

As a byproduct of the arguments in Sections 4 and 8, we obtain many facts about the geometry of Cayley graphs of relatively hyperbolic groups. Recall that given a finitely generated group $G = \langle S \rangle$ and a finite collection $H_1, ..., H_n$ of subgroups of it, one can consider the standard Cayley graph $\mathrm{Cayley}(G, S)$ and the modified Cayley graph $\mathrm{Cayley}(G, S \cup \mathcal{H})$, where

$\mathcal{H} = \bigsqcup_{i=1}^n (H_i \setminus \{e\})$. The standard definition of relative hyperbolicity of a group G with respect to subgroups $H_1, ..., H_n$ is given in terms of the modified Cayley graph $\mathrm{Cayley}(G, S \cup \mathcal{H})$. Theorem 1.11 and the results of Section 4 allow us to define the relative hyperbolicity of G with respect to $H_1, ..., H_n$ in terms of $\mathrm{Cayley}(G, S)$ only. This is an important ingredient in our rigidity results.

An important part in studying tree-graded spaces is played by *saturations* of geodesics. If G is relatively hyperbolic with respect to $H_1, ..., H_n$, $\mathfrak{g}$ is a geodesic in $\mathrm{Cayley}(G, S)$ and M is a positive number, then the *M-saturation of* $\mathfrak{g}$ is the union of $\mathfrak{g}$ and all left cosets of H_i whose M-tubular neighborhoods intersect $\mathfrak{g}$. We show that in the study of relatively hyperbolic groups, saturations play the same role as the geodesics in the study of hyperbolic groups.

More precisely, we use Bowditch's characterization of hyperbolic graphs [Bow_1], and show that tubular neighborhoods of saturations of geodesics can play the role of "lines" in that characterization. In particular, we show that for every geodesic triangle $[A, B, C]$ in $\mathrm{Cayley}(G, S)$ the M-tubular neighborhoods of the saturations of its sides (for some M depending on G and S) have a common point which is at a bounded distance from the sides of the triangle or a common left coset which is at a bounded distance from the sides. Notice that this fact, in a slightly different form, was proved independently in the recent preprint [CR] for a completely different purpose.

We also obtain the following analog for relatively hyperbolic groups of the Morse Lemma for hyperbolic spaces. Recall that the Morse lemma states that every quasi-geodesic in a hyperbolic space is at a bounded distance from a geodesic joining its endpoints. We again do not write the statements in the whole generality.

Notations: Throughout the whole paper, $\mathcal{N}_\delta(A)$ denotes the δ-tubular neighborhood of a subset A in a metric space X, that is $\{x \in X \mid \mathrm{dist}(x, A) < \delta\}$. We denote by $\overline{\mathcal{N}}_\delta(A)$ its closure, that is $\{x \mid \mathrm{dist}(x, A) \leq \delta\}$. In the particular case when $A = \{x\}$ we also use the notations $B(x, \delta)$ and $\overline{B(x, \delta)}$ for the tubular neighborhood and its closure.

Theorem 1.12 (Morse property for relatively hyperbolic groups). *Let $G = \langle S \rangle$ be a group that is hyperbolic relative to the collection of subgroups $H_1, \ldots, H_m$. Then there exists a positive constant M depending only on S such that the following holds. Let $\mathfrak{g}$ be a geodesic in* $\mathrm{Cayley}(G, S)$. *Let $\mathfrak{q} : [0, \ell] \to \mathrm{Cayley}(G, S)$ be an (L, C)-quasi-geodesic joining the endpoints of $\mathfrak{g}$. Then:*

(1) $\mathfrak{q}$ is contained in the τ-tubular neighborhood of the M-saturation of $\mathfrak{g}$, where τ depends only on L, C.

(2) Let gH_i and $g'H_j$ be two left cosets contained in the M-saturation of $\mathfrak{g}$. Let $\mathfrak{q}'$ be a sub-quasi-geodesic of $\mathfrak{q}$ of endpoints $a \in \mathcal{N}_\tau(gH_i)$ and $b \in \mathcal{N}_\tau(g'H_j)$ which intersects $\mathcal{N}_\tau(gH_i)$ and $\mathcal{N}_\tau(g'H_j)$ in sets of bounded diameter. Then a and b belong to the κ-tubular neighborhood of $\mathfrak{g}$, where κ depends only on L, C, τ.

(3) In the Cayley graph $\mathrm{Cayley}(G, S \cup \mathcal{H})$*, $\mathfrak{q}$ is at Hausdorff distance at most $\varkappa$ from any geodesic in the same Cayley graph joining its endpoints, where $\varkappa$ depends only on S.*

The proof of this theorem and more facts about the geometry of relatively hyperbolic groups are contained in Lemmas 4.25, 4.26, 4.28 and Proposition 8.25.

Theorem 1.11 and statements about tree-graded spaces from Section 2 imply that for relatively hyperbolic groups, Problem 1.2 has a positive answer.

Corollary 1.13. *The fundamental group of an asymptotic cone of a relatively hyperbolic group G is either trivial or of order continuum.*

Proof. Suppose that the fundamental group of an asymptotic cone of the group G is non-trivial. By Theorem 1.11, the asymptotic cone of G is tree-graded with respect to a set of pieces that are isometric copies of asymptotic cones of the parabolic subgroups H_i with the induced metric. The induced metric on each H_i is equivalent to the natural word metric by quasi-convexity (see Lemma 4.15). Moreover, in that set, every piece appears together with continuously many copies.

The argument in the first part of the proof of Proposition 2.20 shows that at least one of the pieces has non-trivial fundamental group Γ.

The argument in the second part of the proof of Proposition 2.20 implies that the fundamental group of the asymptotic cone of G contains the free product of continuously many copies of Γ. □

The following statement is another straightforward consequence of Theorem 1.11.

Corollary 1.14. *If a group G is hyperbolic relative to $\{H_1, \dots, H_m\}$ and if each H_i is hyperbolic relative to a collection of subgroups $\{H_i^1, \dots, H_i^{n_i}\}$ then G is hyperbolic relative to $\{H_i^j \mid i \in \{1, \dots, m\}, j \in \{1, \dots, n_i\}$.*

Thus one can consider a "descending process", finding smaller and smaller subgroups of G with respect to which G is relatively hyperbolic. If that process stops then the subgroups in the last collection are not hyperbolic relative to any proper subgroups anymore. The somewhat parallel operation in the asymptotic cones of G is the following. Each time when a piece has a cut-point, we use Lemma 2.28 and represent the piece as a tree-graded space, then we replace the piece by the set of pieces in that tree-graded space. The latter operation stops when the pieces have no more cut-points. This justifies once more the open Problem 1.17 in Section 1.1. Moreover, an affirmative answer to this question would imply that there always exists a collection of minimal subgroups (consequently non-relatively hyperbolic subgroups) relative to which the initial group is hyperbolic.

We note that in the alternative geometric definition of relatively hyperbolic groups given in Theorem 1.11 we do not need the hypothesis that H_i are finitely generated. This is implied by the quasi-convexity of the groups H_i seen as sets in $\mathrm{Cayley}(G, S)$ (Lemma 4.15). Moreover, this geometric definition makes sense when G is replaced by a geodesic metric space X and the collection of cosets of the subgroups H_i is replaced by a collection $\mathcal{A}$ of subsets of X. A similar generalization can be considered for Farb's definition of relative hyperbolicity (including the BCP condition). Thus, both definitions allow to speak of geodesic spaces hyperbolic relative to families of subsets. Such spaces, completely unrelated to groups, do appear naturally. For instance the complements of unions of disjoint open horoballs in rank one symmetric spaces are hyperbolic with respect to the boundary horospheres. Also, the free product of two metric spaces with basepoints (X, x_0) and (Y, y_0), as defined in [PW, §1], is hyperbolic with respect to all the isometric copies of X and Y. It might be interesting for instance to study actions of groups on such spaces, hyperbolic with respect to collections of subsets. To some extent, this is already done in the proof of our Theorem 5.10, where a particular case of action of a group by quasi-isometries on an asymptotically tree-graded (=relatively hyperbolic) space is studied.

Bowditch's characterization of hyperbolic graphs can be easily generalized to arbitrary geodesic metric spaces. So one can expect that an analog of Theorem 1.11 is true for arbitrary geodesic metric spaces.

1.1 Open questions

Problem 1.15. Is it possible to drop the condition that all the asymptotic cones of the "parabolic subgroups" H_i do not have global cut-points from the formulation of Theorem 1.6?

An obvious candidate to a counterexample would be for instance the pair of groups $G = A * A * A * A$, where $A = \mathbb{Z}^2$, and $G' = (A * A * A * A) \rtimes \mathbb{Z}/4\mathbb{Z}$, where $\mathbb{Z}/4\mathbb{Z}$ permutes the factors. The group G is relatively hyperbolic with respect to $A * A * 1 * 1$ and $1 * 1 * A * A$. It is easy to check that the group G' is not relatively hyperbolic with respect to any isomorphic copy of $A * A$. Unfortunately this example does not work. Indeed, G' is quasi-isometric to $A * A$ by [PW], so G' is relatively hyperbolic with respect to a subgroup that is quasi-isometric to $A * A$, namely itself. Moreover, it is most likely that G' is hyperbolic relative to a proper subgroup isomorphic to $A * \mathbb{Z}$ which is also quasi-isometric to $A * A$ by [PW].

Problem 1.16. Corollary 5.21 shows the following. Let G be a group, asymptotically tree-graded as a metric space with respect to a family of subsets $\mathcal{A}$ satisfying the following conditions:

(1) $\mathcal{A}$ is uniformly asymptotically without cut-points, where all sets $A \in \mathcal{A}$ are endowed with the metric induced from G (see Definition 5.2 for the notion of collection of metric spaces uniformly asymptotically without cut-points);

(2) every $A \in \mathcal{A}$ has infinite diameter;

(3) For a fixed $x_0 \in G$ and every $R > 0$ the ball $B(x_0, R)$ intersects only finitely many $A \in \mathcal{A}$.

Then the group G is relatively hyperbolic with respect to subgroups $H_1, ..., H_m$ such that every H_i is quasi-isometric to some $A \in \mathcal{A}$.

Can one remove some of the conditions (1), (2), (3) from this statement?

Problem 1.17. If a group G has the property that all its asymptotic cones have cut-points, consequently by Lemma 2.28 all its asymptotic cones are tree-graded spaces, does this imply that the group G is relatively hyperbolic with respect to a collection of proper subgroups $\{H_1, \ldots, H_m\}$? In case the answer to the previous question is affirmative, there is the following more specific question: if one considers in each asymptotic cone of G the minimal pieces provided by Lemma 2.28, (b), that is pieces without cut-points, is there a (finite) collection of subgroups $\{H_1, \ldots, H_m\}$ or a collection of subsets $\mathcal{A}$ such that all pieces in all asymptotic cones are ultralimits of left cosets of H_i or of subsets in $\mathcal{A}$?

Problem 1.18. A group $G = \langle S \rangle$ is weakly hyperbolic relative to subgroups $H_1, ..., H_n$ if the Cayley graph $\mathrm{Cayley}(G, S \cup \mathcal{H})$ is hyperbolic. It would be interesting to investigate the behavior of weak relatively hyperbolic groups up to quasi-isometry. In particular, it would be interesting to find out if an analog of Theorem 1.6 holds. The arguments used in this paper for the (strong) relative hyperbolicity no longer work. This can be seen on the example of $\mathbb{Z}^n$. That group is weakly hyperbolic relative to $\mathbb{Z}^{n-1}$. But a quasi-isometry $\mathfrak{q} : \mathbb{Z}^n \to \mathbb{Z}^n$ can transform left cosets of $\mathbb{Z}^{n-1}$ into polyhedral or even more complicated surfaces (see [KlL, Introduction] for examples). Nevertheless it is not a real counter-example to a theorem similar to Theorem 1.6 for weak hyperbolic groups, as every group quasi-isometric to $\mathbb{Z}^n$ is virtually $\mathbb{Z}^n$.

1.2 Plan of the paper

In Section 2, we establish some basic properties of tree-graded spaces. In particular, we show that tree-graded spaces behave "nicely" with respect to homeomorphisms.

In Section 3, we establish general properties of asymptotic cones and their ultralimits. In particular, we show that the ultralimit of asymptotic cones of a metric space X is an asymptotic cone of X itself.

In Section 4, we give an "internal" characterization of *asymptotically tree-graded metric spaces*, i.e. pairs of a metric space X and a collection of subsets $\mathcal{A}$, such that every asymptotic cone $\mathrm{Con}_\omega(X; e, d)$ is tree-graded with respect to ω-limits of sequences of sets from $\mathcal{A}$.

In Section 5, we show that being asymptotically tree-graded with respect to a family of subsets is a quasi-isometry invariant provided the ω-limits of subsets from that family do not have cut-points.

In Section 6, we show that asymptotic cones of groups that are not virtually cyclic but have infinite cyclic subgroups in the center, as well as asymptotic cones of groups satisfying non-trivial laws, have no cut-points.

In Section 7, we modify a construction from the paper [EO] to prove, in particular, Theorems 1.4 and 1.5.

In Section 8 and in the Appendix (written by D. Osin and M. Sapir), we prove the characterization of relatively hyperbolic groups in terms of their asymptotic cones given in Theorem 1.11. Theorem 1.8 about undistorted subgroups of relatively hyperbolic groups is also proved in Section 8.

Acknowledgement. The authors are grateful to Anna Erschler and Denis Osin for many helpful conversations.

2 Tree-graded spaces

2.1 Properties of tree-graded spaces

The next several lemmas establish some useful properties of tree-graded spaces. Until Proposition 2.15, $\mathbb{F}$ is a tree-graded space with respect to $\mathcal{P}$.

Lemma 2.1. *If all pieces in $\mathcal{P}$ are $\mathbb{R}$-trees then $\mathbb{F}$ is an $\mathbb{R}$-tree.*

Proof. It is an immediate consequence of (T_2). □

Lemma 2.2. *Let M be a piece and x a point outside M. If y and z are points in M such that there exist geodesics $[x, y]$ and $[x, z]$, joining them to x which intersect M only in y and z, respectively, then $y = z$.*

Proof. Suppose that $y \neq z$. We join y and z by a geodesic $[y, z]$ in M. Let x' be the farthest from x intersection point of the geodesics $[x, y]$ and $[x, z]$. The triangle $x'yz$ is simple because by the assumption $[x, y] \cup [x, z]$ intersects with $[y, z]$ only in y and z. Therefore that triangle is contained in one piece M' by (T_2). Since $M \cap M'$ contains $[y, z]$, $M = M'$ by (T_1), so $x' \in M$, a contradiction since x' belongs both to $[x, y]$ and to $[x, z]$ but cannot coincide with both y and z at the same time. □

Lemma 2.3. *Every simple quadrangle (i.e. a simple loop composed of four geodesics) in $\mathbb{F}$ is contained in one piece.*

Proof. Let A_1, A_2, A_3 and A_4 be the vertices of the quadrangle. We suppose that each vertex is not on a geodesic joining its neighbors, otherwise we have a geodesic triangle and the statement is trivial. Let $\mathfrak{g}$ be a geodesic joining A_1 and A_3. Let P be its last intersection point with

$[A_1, A_2] \cup [A_1, A_4]$. Suppose that $P \in [A_1, A_2]$ (the other case is symmetric). Let Q be the first intersection point of $\mathfrak{g}$ with $[A_2, A_3] \cup [A_3, A_4]$. We replace the arc of $\mathfrak{g}$ between A_1 and P with the arc of $[A_1, A_2]$ between these two points, and the arc of g between Q and A_3 with the corresponding arc of $[A_2, A_3] \cup [A_3, A_4]$. Then $\mathfrak{g}$ thus modified cuts the quadrangle into two simple triangles having in common the geodesic $[P, Q]$. Both triangles are in the same piece by (T_2), and so is the quadrangle. □

Lemma 2.4. *(1) Each piece is a convex subset of* $\mathbb{F}$.

(2) For every point $x \in \mathbb{F}$ *and every piece* $M \in \mathcal{P}$*, there exists a unique point* $y \in M$ *such that* $\mathrm{dist}(x, M) = \mathrm{dist}(x, y)$*. Moreover, every geodesic joining* x *with a point of* M *contains* y*.*

Proof. (1) Suppose that there exists a geodesic $\mathfrak{g}$ joining two points of M and not contained in M. Let z be a point in $\mathfrak{g} \setminus M$. Then z is on a sub-arc $\mathfrak{g}'$ of $\mathfrak{g}$ intersecting M only in its endpoints, a, b. Lemma 2.2 implies $a = b = z \in M$, a contradiction.

(2) Let $y_n \in M$ be such that $\lim_{n\to\infty}\mathrm{dist}(x, y_n)\mathrm{dist}(x, M)$. Since M is closed, we may suppose that every geodesic $[x, y_n]$ intersects M only in y_n. It follows by Lemma 2.2 that $y_1 = y_2 = \ldots = y$.

Let $z \in M$ and let $\mathfrak{g}$ be a geodesic joining z with x. Let z' be the last point on $\mathfrak{g}$ contained in M. Then $z' = y$, by Lemma 2.2. □

Definition 2.5. We call the point y in part (2) of Lemma 2.4 *the projection of x onto the piece M*.

Lemma 2.6. *Let M be a piece and x a point outside it with* $\mathrm{dist}(x, M) = \delta$*, and let y be the projection of x onto M. Then the projection of every point $z \in \overline{B}(x, \delta)$ onto M is equal to y.*

Proof. Notice that by part (2) of Lemma 2.4 $\overline{B}(x, \delta) \cap M = \{y\}$. Suppose that the projection z' of $z \in B(x, \delta)$ onto M is different from y. Then $z \neq y$, hence z does not belong to M.

Consider a geodesic quadrangle with vertices x, z, z' and y. By the definition of projection, the interiors of $[z, z'] \cup [x, y]$ and $[y, z']$ do not intersect.

If there is a common point p of $[x, y]$ and $[z, z']$ then we get a contradiction with Lemma 2.2, so $[x, y]$ and $[z, z']$ are disjoint. In particular $[z, z'] \cup [z', y] \cup [y, x]$ is a topological arc. Since $z \in \overline{B}(x, \delta) \setminus \{y\}$, the side $[x, z]$ of this quadrangle does not intersect M. By part (1) of Lemma 2.4 it follows that $[x, z]$ does not intersect $[y, z']$.

We can replace if necessary z with the last intersection point of $[z, x]$ with $[z, z']$ and x with the last intersection point of the geodesics $[x, y]$ and $[x, z]$. We get a simple geodesic quadrangle $xzz'y$ in which the side $[x, z]$ possibly reduces to a point. By Lemma 2.3, it belongs to one piece. Since it has $[y, z']$ in common with M, that piece is M by (T_1). But this contradicts the fact that $[x, z] \cap M = \emptyset$. □

Corollary 2.7. *Every continuous path in* $\mathbb{F}$ *which intersects a piece M in at most one point, projects onto M in a unique point.*

Proof. If the path does not intersect the piece, it suffices to cover it with balls of radius less than the distance from the path to the piece and use Lemma 2.6.

If the path intersects M in a point x, we may suppose that x is one of its ends and that the interior of the path does not pass through x. Let z be another point on the path and let y be its projection onto M. By the previous argument every point t on the path, $t \neq x$, has the same projection y onto M. Let $\lim_{n\to\infty} t_n = x$, $t_n \neq x$. Then $\lim_{n\to\infty}\mathrm{dist}(t_n, M) = \lim_{n\to\infty}\mathrm{dist}(t_n, y) = 0$. Therefore $x = y$. □

Corollary 2.8. *(1) Every topological arc in $\mathbb{F}$ joining two points in a piece is contained in the piece.*

(2) Every non-empty intersection between a topological arc in $\mathbb{F}$ and a piece is a point or a sub-arc.

Proof. (1) If there exists a topological arc $\mathfrak{p}$ in $\mathbb{F}$ joining two points of a piece M and not contained in M, then a point z in $\mathfrak{p} \setminus M$ is on a sub-arc $\mathfrak{p}'$ of $\mathfrak{p}$ intersecting M only in its endpoints, a, b. Corollary 2.7 implies that both a and b are projections of z into M, contradiction.

(2) immediately follows from (1). □

Corollary 2.9. *Let A be a connected subset (possibly a point) in $\mathbb{F}$ which intersects a piece M in at most one point.*

(1) The subset A projects into M in a unique point x.

(2) Every path joining a point in A with a point in M contains x.

Notation: Let $x \in \mathbb{F}$. We denote by T_x the set of points $y \in \mathbb{F}$ which can be joined to x by a topological arc intersecting every piece in at most one point.

Lemma 2.10. *Let $x \in \mathbb{F}$ and $y \in T_x$, $y \neq x$. Then every topological arc with endpoints x, y intersects each piece in at most one point. In particular the arc is contained in T_x.*

Proof. Suppose, by contradiction, that there exists a topological arc $\mathfrak{p}$ in $\mathbb{F}$ connecting x, y and intersecting a piece M in more than one point. By Corollary 2.8, $M \cap \mathfrak{p}$ is a topological arc with endpoints $a \neq b$. By definition, there also exists an arc $\mathfrak{q}$ connecting x and y and touching each piece in at most one point.

Now consider the two paths connecting x and M. The first path $\mathfrak{p}'$ is a part of $\mathfrak{p}$ connecting x and a. The second path $\mathfrak{q}'$ is the composition of the path $\mathfrak{q}$ and a portion of $\mathfrak{p}^{-1}$ connecting y and b. By Corollary 2.9, the path $\mathfrak{q}'$ must pass through the point a. Since the portion $[y, b]$ of $\mathfrak{p}^{-1}$ does not contain a, the path $\mathfrak{q}$ must contain a. But then there exists a part $\mathfrak{q}''$ of $\mathfrak{q}'$ connecting a and b and intersecting M in exactly two points. This contradicts part (1) of Corollary 2.9, as a point in $\mathfrak{q}'' \setminus \{a, b\}$ would project onto M in both a and b. □

Lemma 2.11. *Let $x \in \mathbb{F}$ and $y \in T_x$. Then $T_x = T_y$.*

Proof. It suffices to prove $T_y \subset T_x$. Let $z \in T_y$. By Lemma 2.10, any geodesics connecting y with x or z intersects every piece in at most one point. Let t be the farthest from y intersection point between two geodesics $\mathfrak{p} = [y, x]$ and $\mathfrak{q} = [y, z]$. Then $\mathfrak{r} = [x, t] \cup [t, z]$ is a topological arc. The arc $\mathfrak{r}$ intersects every piece in at most one point. Indeed, if $\mathfrak{r}$ intersects a piece M in two points a, b then it intersects it in a subarc by Corollary 2.8, so at least one of the two segments $[x, t], [t, z]$ intersects M in an arc, contradiction. Thus $z \in T_x$. □

Lemma 2.12. *Let $x \in \mathbb{F}$.*

(1) Every topological arc joining two distinct points in T_x is contained in T_x.

(2) The subset T_x is a real tree.

Proof. (1) is an immediate consequence of the two previous lemmas.

(2) First we prove that for every $y, z \in T_x$ there exists a unique geodesic joining y and z, also contained in T_x. Since $\mathbb{F}$ is a geodesic space, there exists a geodesic in $\mathbb{F}$ joining x and y. By the first part of the lemma, this geodesic is contained in T_x. Suppose there are two distinct geodesics $\mathfrak{g}, \mathfrak{g}'$ in T_x joining y and z. A point on $\mathfrak{g}$ which is not on $\mathfrak{g}'$ is contained in a simple bigon composed of a sub-arc of $\mathfrak{g}$ and a sub-arc of $\mathfrak{g}'$. This bigon, by (T_2), is contained in a piece. This contradicts Lemma 2.10.

Now consider a geodesic triangle yzt in T_x. Deleting, if necessary, a common sub-arc we can suppose that $[y,z] \cap [y,t] = \{y\}$. If $y \notin [z,t]$ then let z' be the nearest to y point of $[y,z] \cap [z,t]$ and let t' be the nearest to y point of $[y,t] \cap [z,t]$. The triangle $yz't'$ is simple, therefore it is contained in one piece by (T_2). This again contradicts Lemma 2.10. Thus $y \in [z,t]$. □

Convention: We assume that a 1-point metric space has a cut-point.

Lemma 2.13. *Let A be a path connected subset of $\mathbb{F}$ without a cut-point. Then A is contained in a piece. In particular every simple loop is contained in a piece.*

Proof. By our convention, A contains at least two points. Fix a point $x \in A$. The set A cannot be contained in the real tree T_x, because otherwise it would have a cut-point. Therefore, a topological arc joining in A the point x and some $y \in A$ intersects a piece M in a sub-arc $\mathfrak{p}$. Suppose that $A \not\subset M$. Let $z \in A \setminus M$ and let z' be the projection of z onto M. Corollary 2.9 implies that every continuous path joining z to any point α of $\mathfrak{p}$ contains z'. In particular $z' \in A$, and z and α are in two distinct connected components of $\mathbb{F} \setminus \{z'\}$. Thus, z' is a cut-point of A, a contradiction. □

Proposition 2.14. *Let $\mathbb{F}$ and $\mathbb{F}'$ be two tree-graded spaces with respect to the sets of pieces $\mathcal{P}$ and $\mathcal{P}'$, respectively. Let $\Psi \colon \mathbb{F} \to \mathbb{F}'$ be a homeomorphism. Suppose that all pieces in $\mathcal{P}$ and $\mathcal{P}'$ do not have cut-points. Then Ψ sends any piece from $\mathcal{P}$ onto a piece from $\mathcal{P}'$, and $\Psi(T_x) = T_{\Psi(x)}$ for every $x \in \mathbb{F}$.*

Proof. Indeed, for every piece M in $\mathbb{F}$, $\Psi(M)$ is a path connected subset of $\mathbb{F}'$ without cut-points. Therefore $\Psi(M)$ is inside a piece M' of $\mathbb{F}'$ by Lemma 2.13. Applying the same argument to Ψ^{-1}, we have that $\Psi^{-1}(M')$ is contained in a piece M''. Then $M \subseteq \Psi^{-1}(M') \subseteq M''$, hence $M = M''$ and $\Psi(M) = M'$. □

Proposition 2.15. *Condition (T_2) in the definition of tree-graded spaces can be replaced by each of the following conditions:*

(T_2') *For every topological arc $\mathfrak{c} : [0,d] \to \mathbb{F}$ and $t \in [0,d]$, let $\mathfrak{c}[t-a, t+b]$ be a maximal sub-arc of $\mathfrak{c}$ containing $\mathfrak{c}(t)$ and contained in one piece. Then every other topological arc with the same endpoints as $\mathfrak{c}$ must contain the points $\mathfrak{c}(t-a)$ and $\mathfrak{c}(t+b)$.*

(T_2'') *Every simple loop in $\mathbb{F}$ is contained in one piece.*

Proof. Obviously (T_1) and (T_2') imply (T_2). Therefore it is enough to establish the implications $(T_1)\&(T_2'') \Rightarrow (T_2')$ and $(T_1)\&(T_2) \Rightarrow (T_2'')$. The second of these implications is given by Lemma 2.13.

Suppose that (T_1) and (T_2'') hold for some space $\mathbb{F}$ with respect to some set of pieces $\mathcal{P}$.

Let $\mathfrak{c} : [0,d] \to \mathbb{F}$ be a topological arc, $t \in [0,d]$, and a, b as in (T_2'). If $\mathfrak{c}' : [0,d'] \to \mathbb{F}$ is another topological arc with the same endpoints as c, then $K = \mathfrak{c}^{-1}(\mathfrak{c}'[0,d'])$ is a compact set

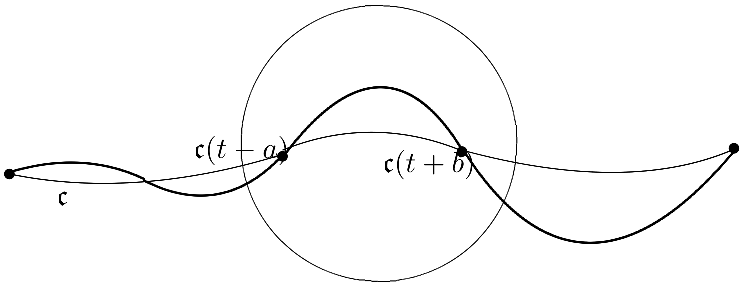

Figure 1: Property (T_2').

containing 0 and d. Suppose that, say, $t-a \notin K$. Let α be the supremum of $K \cap [0, t-a]$ and β be the infimum of $K \cap [t-a, d]$. Then $\alpha < t-a < \beta$. Since $\alpha, \beta \in K$, there exist $\alpha', \beta' \in [0, d']$ such that $\mathfrak{c}'(\alpha') = \mathfrak{c}(\alpha), \mathfrak{c}'(\beta') = \mathfrak{c}(\beta)$. The restriction of $\mathfrak{c}$ to $[\alpha, \beta]$ and the restriction of $\mathfrak{c}'$ to $[\alpha', \beta']$ form a simple loop which is contained in one piece by (T_2''). In particular $\mathfrak{c}([\alpha, \beta])$ is contained in one piece. Since $[t-a, t+b]$ is the maximal interval containing t such that the restriction of $\mathfrak{c}$ to that interval is contained in one piece, it follows that $b+a \neq 0$. Therefore the intersection of the intervals $[\alpha, \beta]$ and $[t-a, t+b]$ has a non-empty interior. Hence the pieces containing $\mathfrak{c}([\alpha, \beta])$ and $\mathfrak{c}([t-a, t+b])$ must coincide by property (T_1). But this contradicts the maximality of the interval $[t-a, t+b]$. $\square$

Remark 2.16. If a collection of subsets $\mathcal{P}$ of a geodesic metric space X satisfy (T_1) and (T_2''), and each set in $\mathcal{P}$ is path connected then each set in $\mathcal{P}$ is a geodesic subspace. Thus if one replaces property (T_2) by the stronger property (T_2'') in Definition 1.10 then one can weaken the condition on $\mathcal{P}$.

Proof. Let $M \in \mathcal{P}$, let x, y be two points in M and let $\mathfrak{r}$ be a topological arc joining x and y in M. Suppose that a geodesic $\mathfrak{g}$ connecting x and y in X is not contained in M. Let $z \in \mathfrak{g} \setminus M$. There exists a simple non-trivial bigon with one side a sub-arc in $\mathfrak{r}$ and the other a sub-arc in $\mathfrak{g}$ containing z. Property (T_2'') implies that this bigon is contained in a piece, and property (T_1) implies that this piece is M. Hence z is in M, a contradiction. $\square$

Lemma 2.17. *For every $x \in \mathbb{F}$, T_x is a closed subset of $\mathbb{F}$.*

Proof. Let (y_n) be a sequence in T_x converging to a point y. Suppose that the geodesic $[x, y]$ intersects a piece M in a maximal non-trivial sub-arc $[\alpha, \beta]$. We can assume that the geodesic $[y_n, y]$ intersects $[y_n, x]$ only in y_n. Otherwise we can replace y_n with the farthest from it intersection point between these two geodesics. By property (T_2') the arc $[x, y_n] \cup [y_n, y]$ must contain $[\alpha, \beta]$. Since $y_n \in T_x$, it follows by Lemma 2.10 that $[\alpha, \beta] \subset [y_n, y]$ and so $\mathrm{dist}(y_n, y) \geq \mathrm{dist}(\alpha, \beta) > 0$. This contradicts $\mathrm{dist}(y_n, y) \to 0$. We conclude that $[x, y]$ intersects every piece in at most one point and that $y \in T_x$. $\square$

Lemma 2.18. *The projection of $\mathbb{F}$ onto any of the pieces is a metric retraction.*

Proof. Let M be a piece, x, y two points in $\mathbb{F}$ and $[x, y]$ a geodesic joining them. If $[x, y] \cap M = \emptyset$ then $[x, y]$ projects onto one point z, by Corollary 2.7, and $d(x, y) \geq d(z, z) = 0$.

If $[x, y] \cap M = [\alpha, \beta]$ then α is the projection of x onto M and β is the projection of y onto M, by Corollary 2.7. Obviously $d(x, y) \geq d(\alpha, \beta)$. $\square$

Lemma 2.19. *Let $\mathfrak{p}\colon [0,l] \to \mathbb{F}$ be a path in a tree-graded space $\mathbb{F}$. Let $U_{\mathfrak{p}}$ be the union of open subintervals $(a,b) \subset [0,l]$ such that the restriction of $\mathfrak{p}$ onto (a,b) belongs to one piece (we include the trees T_x into the set of pieces). Then $U_{\mathfrak{p}}$ is an open and dense subset of $[0,l]$.*

Proof. Suppose that $U_{\mathfrak{p}}$ is not dense. Then there exists a non-trivial interval (c,d) in the complement $[0,l] \setminus U_{\mathfrak{p}}$. Suppose that the restriction $\mathfrak{p}'$ of $\mathfrak{p}$ on (c,d) intersects a piece P in two points $y = \mathfrak{p}(t_1), z = \mathfrak{p}(t_2)$. We can assume that y is not in the image of $(t_1, t_2]$ under $\mathfrak{p}$. Since $y \notin U_{\mathfrak{p}}$ there is a non-empty interval (t_1, t_3) such that the restriction of $\mathfrak{p}$ onto that interval does not intersect P. Let $t > t_1$ be the smallest number in $(t_1, t_2]$ such that $z' = \mathfrak{p}(t)$ is in P. Then $z' \neq y$. Applying Corollary 2.9 to the restriction of $\mathfrak{p}$ onto $[t_1, t]$, we get a contradiction. This means that $\mathfrak{p}'$ intersects every piece in at most one point. Therefore $\mathfrak{p}'$ is contained in a tree T_x for some x, a contradiction. □

Proposition 2.20. *Let $\mathbb{F}$ be a tree-graded space with the set of pieces $\mathcal{P}$. If the pieces in $\mathcal{P}$ are locally uniformly contractible then $\pi_1(\mathbb{F})$ is the free product of $\pi_1(M)$, $M \in \mathcal{P}$.*

Proof. We include all trees T_x into $\mathcal{P}$. Fix a base point x in $\mathbb{F}$ and for every piece $M_i \in \mathcal{P}$ let y_i be the projection of x onto M_i, and let $\mathfrak{g}_i$ be a geodesic connecting x and y_i. We identify $\pi_1(M_i)$ with the subgroup $G_i = \mathfrak{g}_i \pi_1(M_i, y_i) \mathfrak{g}_i^{-1}$ of $\pi_1(\mathbb{F}, x)$. Consider an arbitrary loop $\mathfrak{p}\colon [0,l] \to \mathbb{F}$ in $\mathbb{F}$ based at x. Let $\mathfrak{p}'$ be the image of $\mathfrak{p}$. Let $\mathcal{P}_{\mathfrak{p}}$ be the set of pieces from $\mathcal{P}$ which are intersected by $\mathfrak{p}'$ in more than one point. By Lemma 2.19 the set $\mathcal{P}_{\mathfrak{p}}$ is countable.

Let $M \in \mathcal{P}_{\mathfrak{p}}$. The projection $\mathfrak{p}_M$ of $\mathfrak{p}'$ onto M is a loop containing the intersection $\mathfrak{p}' \cap M$. Let us prove that $\mathfrak{p}_M = \mathfrak{p}' \cap M$. If there exists a point $z \in \mathfrak{p}_M \setminus \mathfrak{p}'$ then z is a projection of some point $y \in \mathfrak{p}' \setminus M$ onto M. By Corollary 2.9, a subpath of $\mathfrak{p}$ joining y with a point in $\mathfrak{p}' \cap P$ must contain z, a contradiction.

Therefore $\mathfrak{p}'$ is a union of at most countably many loops $\mathfrak{p}_i$, $i \in \mathbb{N}$, contained in pieces from $\mathcal{P}_{\mathfrak{p}}$. By uniform local contractibility of the pieces, all but finitely many loops $\mathfrak{p}_i$ are contractible inside the corresponding pieces. Consequently, in the fundamental group $\pi_1(\mathbb{F})$, $\mathfrak{p}$ is a product of finitely many loops from G_i. Hence $\pi_1(\mathbb{F}, x)$ is generated by the subgroups G_i.

It remains to prove that for every finite sequence of loops $\mathfrak{p}_i \in G_i$, $i = 1, ..., k$, if $M_i \neq M_j$ for $i \neq j$, and if the loops $\mathfrak{p}_i$ are not null-homotopic in M_i, then the loop $\mathfrak{p}_1\mathfrak{p}_2 ... \mathfrak{p}_n$ is not null-homotopic in $\mathbb{F}$. Suppose that $\mathfrak{p}$ is null-homotopic, and that $\gamma : t \to \mathfrak{p}(t)$ is the homotopy, $\mathfrak{p}(0) = \mathfrak{p}$, $\mathfrak{p}(1)$ is a point. Let π_i be the projection of $\mathbb{F}$ onto M_i. Lemma 2.18 implies that $\pi_i \circ \gamma : t \to \mathfrak{p}_i'(t)$ is a homotopy which continuously deforms $\mathfrak{p}_i'$ in M_i into a point. Hence each of the loops $\mathfrak{p}_i$ is null-homotopic, a contradiction. □

2.2 Modifying the set of pieces

Lemma 2.21 (gluing pieces together). *Let $\mathbb{F}$ be a space which is tree-graded with respect to $\mathcal{P}\{M_k \mid k \in K\}$.*

(1) Let $Y = \bigcup_{k \in F} M_k$ be a finite connected union of pieces. Then $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}' = \{M_k \mid k \in K \setminus F\} \cup \{Y\}$.

(2) Let $\mathfrak{c}$ be a topological arc in $\mathbb{F}$ (possibly a point) and let $Y(\mathfrak{c})$ be a set of the form $\mathfrak{c} \cup \bigcup_{j \in J} M_j$, where J is a subset of K such that every M_j with $j \in J$ has a non-empty intersection with $\mathfrak{c}$, and J contains all $i \in K$ such that $M_i \cap \mathfrak{c}$ is a non-trivial arc.

Then $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}' = \{M_k \mid k \in K \setminus J\} \cup \{Y(\mathfrak{c})\}$.

(3) Let $\{\mathfrak{c}_i\,;\, i \in F\}$ be a finite collection of topological arcs in $\mathbb{F}$ and let $Y(\mathfrak{c}_i) = \mathfrak{c}_i \cup \bigcup_{j\in J_i} M_j$ be sets defined as in (2). If $Y = \bigcup_{i\in F} Y(\mathfrak{c}_i)$ is connected then $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}' = \{M_k \mid k \in K \setminus \bigcup_{i\in F} J_i\} \cup \{Y\}$.

Remark 2.22. In particular all properties on projections on pieces obtained till now hold for sets Y defined as in (1)-(3). We shall call sets of the form $Y(\mathfrak{c})$ *sets of type* Y.

Proof. **(1)** We first prove that Y is convex. Every $y, y' \in Y$ can be joined by a topological arc $\mathfrak{c} : [0, d] \to Y$. By Corollary 2.8, we may write $\mathfrak{c}[0, d] = \bigcup_{k\in F'} [\mathfrak{c}[0, d] \cap M_k]$, where $F' \subset F$ and $\mathfrak{c}[0, d] \cap M_k$ is a point or an arc. Property (T_1) implies that every two such arcs have at most one point in common. Therefore there exists a finite sequence $t_0 = 0 < t_1 < t_2 < \cdots < t_{n-1} < t_n = d$ such that $\mathfrak{c}[t_i, t_{i+1}] = \mathfrak{c}[0, d] \cap M_{k(i)}$, $k(i) \in F'$, for every $i \in \{0, 1, \ldots n-1\}$. Property (T_2') implies that every geodesic between y and y' must contain $\mathfrak{c}(t_1), \mathfrak{c}(t_2), \ldots \mathfrak{c}(t_{n-1})$. Hence every such geodesic is of the form $[y, \mathfrak{c}(t_1)] \cup [\mathfrak{c}(t_1), \mathfrak{c}(t_2)] \cup \cdots \cup [\mathfrak{c}(t_{n-1}), y]$, so by Corollary 2.8 it is contained in Y.

For every $k \in K \setminus F$, $M_k \cap Y$, if non-empty, is a convex set composed of finitely many points. Hence it is a point. This and the previous discussion imply that $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}'$.

(2) In order to prove that Y is convex, let $\mathfrak{g}$ be a geodesic joining two points $x, y \in Y$. We show that $\mathfrak{g}$ is inside Y.

Case I. Suppose that $x, y \in \mathfrak{c}$. Consider a point $z = \mathfrak{g}(t)$ in $\mathfrak{g}$. Take the maximal interval $[t-a, t+b]$ such that $\mathfrak{g}([t-a, t+b])$ is contained in one piece M. If $a+b \neq 0$ then by property (T_2') the path $\mathfrak{c}$ must pass through $\mathfrak{g}(t-a)$ and $\mathfrak{g}(t+b)$. By part (1) of Corollary 2.8 the (non-trivial) subarc of $\mathfrak{c}$ joining $\mathfrak{g}(t-a)$ and $\mathfrak{g}(t+b)$ is contained in M. Then M is one of the pieces contained in Y. Therefore $z \in Y$. If $a+b = 0$ then again by (T_2') the curve $\mathfrak{c}$ must pass through z, so $z \in Y$. We conclude that in both cases $z \in Y$.

Case II. Suppose that $x \in \mathfrak{c}$ and $y \in M \setminus \mathfrak{c}$, where M is a piece in Y. By the definition of Y, M has a non-trivial intersection with $\mathfrak{c}$. If $x \in M$, we can use the convexity of M (Corollary 2.8). So suppose that $x \notin M$.

Let α be the projection of x onto M. By Corollary 2.9, part (2), $\alpha \in \mathfrak{c}$. Then the sub-arc $\mathfrak{c}'$ of $\mathfrak{c}$ with endpoints x and α forms together with the geodesic $[\alpha, y] \subseteq M$ a topological arc. Property (T_2') implies that $\alpha \in \mathfrak{g}$. Corollary 2.8, part (1), implies that the portion of $\mathfrak{g}$ between α and y is contained in Y. For the remaining part of $\mathfrak{g}$ we apply the result in Case I of the proof (since both endpoints of that part of $\mathfrak{g}$ belong to $\mathfrak{c}$).

Case III. Suppose that $x \in M_1 \setminus \mathfrak{c}$ and that $y \in M_2 \setminus \mathfrak{c}$. Let α be the projection of x onto M_2. As before, we obtain that $\alpha \in \mathfrak{c}$, $\alpha \in \mathfrak{g}$ and that the portion of $\mathfrak{g}$ between α and y is contained in M_2, hence in Y. For the remaining part of $\mathfrak{g}$ we apply the result of Case II.

(3) We argue by induction on the size k of the set F. The statement is true for $k = 1$ by part (2) of this Proposition. Suppose it is true for some $k \geq 1$. Let us prove it for $k+1$. We have two cases.

Case I. Suppose that there exist $i, j \in F, i \neq j$, such that the intersection $\mathfrak{c}_i \cap Y(\mathfrak{c}_j)$ is not empty. According to part (2) of the Proposition and Corollary 2.8, part (2), the intersection is a sub-arc, and $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}_j' = \{M_k \mid k \in K \setminus J_j\} \cup \{Y(\mathfrak{c}_j)\}$. Let $Y'(\mathfrak{c}_i) = Y(\mathfrak{c}_i) \cup Y(\mathfrak{c}_j)$. Then $Y'(\mathfrak{c}_i)$ is a set defined as in part (2) of the Proposition but with $\mathcal{P}$ replaced by $\mathcal{P}_j'$. Thus we can write $Y = Y'(\mathfrak{c}_i) \cup \bigcup_{s\in F\setminus\{i,j\}} Y(\mathfrak{c}_s)$ and use the induction hypothesis.

Case II. For every $i, j \in F, i \neq j$, we have $\mathfrak{c}_i \cap Y(\mathfrak{c}_j) = \emptyset$.

Then there are no pieces that appear in both $Y(\mathfrak{c}_i)$ and $Y(\mathfrak{c}_j)$ for $i \neq j \in F$. Hence by (T_1), for every $k \in J_i$, $l \in J_j$, $M_k \cap M_l$ consists of at most one point. By part (2) of the Proposition and Corollary 2.9 that point must be equal to the projection of $\mathfrak{c}_i$ onto $Y(\mathfrak{c}_j)$. Therefore $Y(\mathfrak{c}_i) \cap Y(\mathfrak{c}_j)$ is either empty or one point. This implies that $\mathbb{F}$ is tree-graded with respect to $\mathcal{P}'' = \{M_k \mid k \in K \setminus \bigcup_{i \in F} J_i\} \cup \{Y(\mathfrak{c}_i) \mid i \in F\}$. It remains to apply part (1) of the Proposition. □

Definition 2.23. Let (M_1, x_1), $(M_2, x_2), \ldots, (M_k, x_k)$ be finitely many metric spaces with fixed basepoints. The *bouquet* of these spaces, denoted by $\bigvee_{i=1}^{k}(M_i, x_i)$, is the metric space obtained from the disjoint union of all M_i by identifying all the points x_i. We call the point x thus obtained *the cut-point of the bouquet*. The metric on $\bigvee_{i=1}^{k}(M_i, x_i)$ is induced by the metrics on M_i in the obvious way.

Clearly each M_i is a closed subset of the bouquet $\bigvee_{i=1}^{k}(M_i, x_i)$. It is also clear that the bouquet is a geodesic metric space if and only if all M_i are geodesic metric spaces.

Lemma 2.24 (cutting pieces by cut-points). *Let $\mathbb{F}$ be a space which is tree-graded with respect to $\mathcal{P}\{M_k \mid k \in K\}$. Let $I \subset K$ be such that for every $i \in I$ the piece M_i is the bouquet of finitely many subsets of it, $\{M_i^j\}_{j \in F_i}$, and its cut-point is x_i.*

Then $\mathbb{F}$ is tree-graded with respect to the set

$$\mathcal{P}' = \{M_k \mid k \in K \setminus I\} \cup \{M_i^j \mid j \in F_i,\, i \in I\}.$$

Proof. Since $M_i^j \cap M_k \subset M_i \cap M_k$ for $i \in I$, $k \in K \setminus I$, and $M_i^j \cap M_t^s \subset M_i \cap M_t$ for $i \neq t$, $i, t \in I$, property (T_1) for $(\mathbb{F}, \mathcal{P}')$ is an immediate consequence of property (T_1) for $(\mathbb{F}, \mathcal{P})$.

Let Δ be a simple geodesic triangle. Property (T_2) for $(\mathbb{F}, \mathcal{P})$ implies that either $\Delta \subset M_k$ for some $k \in K \setminus I$ or $\Delta \subset M_i$ for some $i \in I$. We place ourselves in the second case. Assume that Δ has a point in $M_i^{j_1}$ and a point in $M_i^{j_2}$, with $j_1 \neq j_2$. Then x_i is a cut-point for Δ. This contradicts the fact that Δ is a simple loop. We conclude that there exists $j \in F_i$ such that M_i^j contains Δ. Thus $\mathcal{P}'$ satisfies (T_2). □

Lemma 2.11 implies that two trees T_x and T_y are either disjoint or coincident. We consider $\{T_i \mid i \in I\}$ the collection of all the trees $\{T_x \mid x \in \mathbb{F}\}$.

Remark 2.25. The set $\mathcal{P}' = \mathcal{P} \cup \{T_i \mid i \in I\}$ also satisfies properties (T_1) and (T_2). Therefore all the properties and arguments done for $\mathbb{F}$ and $\mathcal{P}$ up to now also hold for $\mathbb{F}$ and $\mathcal{P}'$. In this case, $T_x = \{x\}$ for every $x \in \mathbb{F}$. The disadvantage of this point of view is that trees T_x always have cut-points.

2.3 Geodesics in tree-graded spaces

Notation: For every path $\mathfrak{p}$ in a metric space X, we denote the start of $\mathfrak{p}$ by $\mathfrak{p}_-$ and the end of $\mathfrak{p}$ by $\mathfrak{p}_+$.

Lemma 2.26. *Let $\mathfrak{g} = \mathfrak{g}_1\mathfrak{g}_2 \ldots \mathfrak{g}_{2m}$ be a curve in a tree-graded space $\mathbb{F}$ which is a composition of geodesics. Suppose that all geodesics $\mathfrak{g}_{2k}$ with $k \in \{1, \ldots, m-1\}$ are non-trivial and for every $k \in \{1, \ldots, m\}$ the geodesic g_{2k} is contained in a piece M_k while for every $k \in \{0, 1, \ldots, m-1\}$ the geodesic $\mathfrak{g}_{2k+1}$ intersects M_k and M_{k+1} only in its respective endpoints. In addition assume that if $\mathfrak{g}_{2k+1}$ is empty then $M_k \neq M_{k+1}$. Then $\mathfrak{g}$ is a geodesic.*

Proof. Suppose that $\mathfrak{g}$ is not simple. By (T_2''), any simple loop formed by a portion of $\mathfrak{g}$ has to be contained in one piece M. On the other hand the loop must contain the whole neighborhood of one vertex $(\mathfrak{g}_i)_+ = (\mathfrak{g}_{i+1})_-$ in $\mathfrak{g}$. Let k be such that $\{\mathfrak{g}_i, \mathfrak{g}_{i+1}\} = \{\mathfrak{g}_{2k}, \mathfrak{g}_{2k\pm 1}\}$. The intersection of M and M_k contains a sub-arc of $\mathfrak{g}_{2k}$, whence $M = M_k$. At the same time, M contains a subarc of $\mathfrak{g}_{2k\pm 1}$ or (if $\mathfrak{g}_{2k\pm 1}$ is empty) of $\mathfrak{g}_{2k-2}$. In all cases we immediately get a contradiction.

Therefore $\mathfrak{g}$ is simple and has two distinct endpoints x, y. We apply (T_2') to the topological arc $\mathfrak{g}$ and to any geodesic $\mathfrak{r}$ joining x, y. We obtain that $\mathfrak{r}$ contains all the endpoints of all geodesics $\mathfrak{g}_i$. Therefore the length of $\mathfrak{g}$ coincides with the length of $\mathfrak{r}$. □

Corollary 2.27. *Let M and M' be two distinct pieces in a tree-graded space $\mathbb{F}$. Suppose that M' projects onto M in x and M projects on M' in y. Let A be a set in $\mathbb{F}$ that projects onto M' in $z \neq y$. Then A projects onto M in x and* $\mathrm{dist}(A, M) \geq \mathrm{dist}(M', M)$.

Proof. Let $a \in A$ and let $[a, z]$, $[z, y]$ and $[y, x]$ be geodesics. Then $\mathfrak{g}_a = [a, z] \cup [z, y] \cup [y, x]$ is a geodesic, according to Lemma 2.26. It cannot intersect M in a sub-geodesic, because $[z, y] \cup [y, x]$ intersects M in x. Hence $\mathfrak{g}_a \cap M = \{x\}$ and x is the projection of a onto M. Also $\mathrm{dist}(a, x) \geq \mathrm{dist}(y, x)$. □

2.4 Cut-points and tree-graded spaces

Property (T_2') implies that any tree-graded space containing more than one piece has a global cut-point. Here we shall show that any geodesic metric space with cut-points can be represented in a unique way as a tree-graded space with cut-point-free pieces.

Lemma 2.28. *Let X be a complete geodesic metric space containing at least two points and let $\mathcal{C}$ be a non-empty set of global cut-points in X.*

(a) There exists a collection $\mathcal{P}$ of subsets M of X, $M \neq X$, such that X is tree-graded with respect to $\mathcal{P}$ and the intersection of any two distinct sets from $\mathcal{P}$ is in $\mathcal{C}$ or is empty.

(b) If $\mathcal{C} = X$ then all pieces M in $\mathcal{P}$ are either points or are without cut-point and $\mathcal{P}$ is the unique collection of subsets without cut-point with respect to which X is tree graded. In particular this is true if X is a homogeneous space with a cut-point.

(c) Let X be a homogeneous space with a cut-point and $\mathcal{C} = X$. Let $\mathcal{P}$ be the collection of subsets from part (a). Then for every $M \in \mathcal{P}$ every $x \in M$ is the projection of a point $y \in X \setminus M$ onto M.

Proof. (a) We define $\mathcal{P}$ as the set of all maximal path connected subsets M with the property that either $|M| = 1$ or cut-points of M do not belong to $\mathcal{C}$. The existence of maximal subsets with this property immediately follows from Zorn's lemma.

Any $M \in \mathcal{P}$ is closed. Indeed, let $\bar{M}$ be the closure of M in X and suppose that $\bar{M} \neq M$. Let $a \in \bar{M} \setminus M$. There exists a sequence of points (a_n) in M converging to a. Let M' be the union of M and geodesics $[a, a_n]$, $n = 1, 2, ...$ (one geodesic for each n). By construction, the set M' is path connected. We prove that cut-points of M' do not belong to $\mathcal{C}$. This will contradict the maximality of M.

Let $c \in \mathcal{C} \cap M'$, $x, y \in M' \setminus \{c\}$. We want to connect x and y with a path avoiding c. If $x, y \in M \setminus \{c\}$ then we are done.

Suppose that $x \in M \setminus \{c\}$ and $y \in [a_n, a]$ for some n. The point x can be connected by some path $\mathfrak{p}_k \subseteq M$ avoiding c with a_k for every $k \in \mathbb{N}$.

If $c \notin [a_n, y]$ then the path $\mathfrak{p}_n \cup [a_n, y] \subseteq M'$ avoids c and we are done.

If $c \in [a_n, y]$ then $\mathrm{dist}(c, a) > \mathrm{dist}(y, a)$. In particular c is not in $[a, a_m]$ for m large enough. Then we join y with x by a path $[y, a] \cup [a, a_m] \cup \mathfrak{p}_m$ avoiding c.

It remains to consider the case when $x \in [a_m, a]$ and $y \in [a_n, a]$ for some m, n. If $c \notin [a_m, x]$ then we can replace x with a_m and use the previous argument. Likewise if $c \notin [a_n, y]$. If $c \in [a_m, x] \cap [a_n, y]$ then we join x and y in $X \setminus \{c\}$ by $[x, a] \cup [a, y]$.

Let M_1, M_2 be distinct sets from $\mathcal{P}$, $c \in \mathcal{C}$. Suppose that $M_1 \cap M_2$ contains a point x that is different from c. Then any point $z_i \in M_i$, $z_i \neq c$, $i = 1, 2$, can be joined with x by a path in M_i avoiding c. Hence z_1 and z_2 can be joined in $M_1 \cup M_2$ by a path avoiding c. Consequently if $M_1 \cap M_2$ contains more than one point or contains a point not from $\mathcal{C}$, we get a contradiction with the maximality of M_i. Thus $\mathcal{P}$ satisfies (T_1) and the intersection of any two sets from $\mathcal{P}$ is in $\mathcal{C}$ or empty.

To prove (T_2'') notice that every simple loop is path connected and does not have cut-points, hence it is contained in some M.

The fact that each piece $M \in \mathcal{P}$ is a geodesic subset follows from Remark 2.16.

(b) The first part of this statement follows immediately from part (a). The second part is obtained by applying Proposition 2.14 to the identity map of X.

(c) Let $M \in \mathcal{P}$. Since $M \neq X$ it follows that one point $x_0 \in M$ is the projection on M of a point $y_0 \in X \setminus M$. If M is a point this ends the proof. Suppose in the sequel that M has at least two points. Let $[y_0, x_0]$ be a geodesic joining y_0 and x_0 and let $[x_0, z_0]$ be a geodesic in M. By the definition of the projection, $[y_0, x_0] \cap [x_0, z_0] = \{x_0\}$. Let x be an arbitrary point in M. We consider an isometry g such that $g(x_0) = x$. Let $[y, x]$ and $[x, z]$ be the respective images of $[y_0, x_0]$ and $[x_0, z_0]$ under g. If $g(M) = M$ then x is the projection of y on M. Suppose $g(M) \neq M$. Then $g(M) \cap M = \{x\}$, hence $[x, z] \subset g(M)$ intersects M in x. Corollary 2.9 implies that z projects on M in x. □

Remarks 2.29. (1) In general not every point in $\mathcal{C}$ is the intersection point of two distinct pieces. An example is an $\mathbb{R}$-tree and $\mathcal{P}$ is the set of all one-point sets.

(2) Statement (b) implies that every asymptotic cone of a group which has a cut-point is tree-graded with respect to a uniquely determined collection of closed geodesic subsets without cut-point.

3 Ultralimits and asymptotic cones

3.1 Preliminaries

Most of the interesting examples of tree-graded spaces that we know are asymptotic cones of groups. In this section, we start with giving the definitions of ultralimit, asymptotic cone and related objects (most of these definitions are well known). We show that the collection of asymptotic cones of a space is closed under ultralimits. We also show that simple geodesic triangles in ultralimits and asymptotic cones can be approximated by ultralimits of polygons with certain properties. As a consequence we show that the family of tree-graded spaces is also closed under ultralimits. These results play a central part in the theorems obtained in Sections 4 and 7.

Convention: In the sequel I will denote an arbitrary countable set.

Definition 3.1 (ultrafilter). A (non-principal[2]) ultrafilter ω over I is a set of subsets of I satisfying the following conditions:

[2]We shall only use non-principal ultrafilters in this paper, so the word non-principal will be omitted.

1. If $A, B \in \omega$ then $A \cap B \in \omega$;
2. If $A \in \omega$, $A \subseteq B \subseteq I$, then $B \in \omega$;
3. For every $A \subseteq I$ either $A \in \omega$ or $I \setminus A \in \omega$;
4. No finite subset of I is in ω.

Equivalently ω is a finitely additive measure on the class $\mathcal{P}(I)$ of subsets of I such that each subset has measure either 0 or 1 and all finite sets have measure 0. If some statement $P(n)$ holds for all n from a set X belonging to an ultrafilter ω, we say that $P(n)$ holds *ω-almost surely.*

Remark 3.2. By definition ω has the property that $\omega(\sqcup_{i=1}^m A_i) = 1$ (here $\sqcup$ stands for disjoint union) implies that there exists $i_0 \in \{1, 2, \ldots, m\}$ such that $\omega(A_{i_0}) = 1$ and $\omega(A_i) = 0$ for every $i \neq i_0$. This can be reformulated as follows: let $P_1(n)$, $P_2(n)$, $\ldots$, $P_m(n)$ be properties such that for any $n \in I$ no two of them can be true simultaneously. If the disjunction of these properties holds ω-almost surely then there exists $i \in \{1, 2, \ldots, m\}$ such that ω-almost surely $P_i(n)$ holds and all $P_j(n)$ with $j \neq i$ do not hold.

Definition 3.3 (ω-limit). Let ω be an ultrafilter over I. For every sequence of points $(x_n)_{n\in I}$ in a topological space X, its ω-limit $\lim_\omega x_n$ is a point x in X such that for every neighborhood U of x the relation $x_n \in U$ holds ω-almost surely.

Remark 3.4. If ω-limit $\lim_\omega x_n$ exists then it is unique, provided the space X is Hausdorff. Every sequence of elements in a compact space has an ω-limit [Bou].

Definition 3.5 (ultraproduct). For every sequence of sets $(X_n)_{n\in I}$ the *ultraproduct* $\Pi X_n/\omega$ corresponding to an ultrafilter ω consists of equivalence classes of sequences $(x_n)_{n\in I}$, $x_n \in X_n$, where two sequences (x_n) and (y_n) are identified if $x_n = y_n$ ω-almost surely. The equivalence class of a sequence (x_n) in $\Pi X_n/\omega$ is denoted by $(x_n)^\omega$. In particular, if all X_n are equal to the same X, the ultraproduct is called the *ultrapower* of X and is denoted by X^ω.

Recall that if G_n, $n \geq 1$, are groups then $\Pi G_n/\omega$ is again a group with the operation $(x_n)^\omega (y_n)^\omega = (x_n y_n)^\omega$.

Definition 3.6 (ω-limit of metric spaces). Let (X_n, dist_n), $n \in I$, be a sequence of metric spaces and let ω be an ultrafilter over I. Consider the ultraproduct $\Pi X_n/\omega$ and an *observation point* $e = (e_n)^\omega$ in $\Pi X_n/\omega$. For every two points $x = (x_n)^\omega, y = (y_n)^\omega$ in $\Pi X_n/\omega$ let

$$D(x, y) = \lim_\omega \mathrm{dist}_n(x_n, y_n)\,.$$

The function D is a pseudo-metric on $\Pi X_n/\omega$ (i.e. it satisfies the triangle inequality and the property $D(x, x) = 0$, but for some $x \neq y$, the number $D(x, y)$ can be 0 or ∞). Let $\Pi_e X_n/\omega$ be the subset of $\Pi X_n/\omega$ consisting of elements which are finite distance from e with respect to D. The *ω-limit* $\lim^\omega (X_n)_e$ *of the metric spaces* (X_n, dist_n) *relative to the observation point* e is the metric space obtained from $\Pi_e X_n/\omega$ by identifying all pairs of points x, y with $D(x, y) = 0$. The equivalence class of a sequence (x_n) in $\lim^\omega (X_n)_e$ is denoted by $\lim^\omega (x_n)$.

Remark 3.7 (changing the observation point). It is easy to see that if $e, e' \in \Pi X_n/\omega$ and $D(e, e') < \infty$ then $\lim^\omega (X_n)_e = \lim^\omega (X_n)_{e'}$.

Definition 3.8 (asymptotic cone). Let (X, dist) be a metric space, ω be an ultrafilter over a set I, $e = (e_n)^\omega$ be an observation point. Consider a sequence of numbers $d = (d_n)_{n\in I}$ called *scaling constants* satisfying $\lim_\omega d_n = \infty$. The ω-limit $\lim^\omega (X, \frac{\mathrm{dist}}{d_n})_e$ is called an *asymptotic cone of X*. It is denoted by $\mathrm{Con}^\omega(X; e, d)$ (see [Gr$_1$], [Gr$_3$], [VDW]).

Every finitely generated group $G = \langle X \rangle$ can be considered a metric space where the distance between two elements a, b is the length of the shortest group word in X representing $a^{-1}b$. Let G_n, $n \in I$, be the metric space G with metric $\frac{\mathrm{dist}}{d_n}$ for some sequence of scaling constants $(d_n)_{n\in I}$. Then for any observation point e, the set $\Pi_e G_n/\omega$ denoted by G_e^ω is a subgroup of the ultrapower G^ω. Clearly for different observation points e (but the same scaling constants) all these groups are conjugate in the ultrapower G^ω. In particular, these groups are isomorphic.

Definition 3.9. For a sequence $(A_n), n \in I$, of subsets of (X, dist) we denote by $\lim^\omega(A_n)$ the subset of $\mathrm{Con}^\omega(X; e, d)$ that consists of all the elements $\lim^\omega(x_n)$ such that $x_n \in A_n$ ω-almost surely. Notice that if $\lim_\omega \frac{\mathrm{dist}(e_n, A_n)}{d_n} = +\infty$ then the set $\lim^\omega(A_n)$ is empty.

Remark 3.10. It is proved in [VDW] that any asymptotic cone of a metric space is complete. The same proof gives that $\lim^\omega(A_n)$ is always a closed subset of the asymptotic cone $\mathrm{Con}^\omega(X; e, d)$.

Definition 3.11 (quasi-isometries). A *quasi-isometric embedding* of a metric space (X, dist_X) into a metric space (Y, dist_Y) is a map $\mathfrak{q}\colon X \to Y$ such that

$$\frac{1}{L}\mathrm{dist}_X(x, x') - C \le \mathrm{dist}_Y(\mathfrak{q}(x), \mathfrak{q}(x')) \le L\mathrm{dist}_X(x, x') + C, \text{ for all } x, x' \in X.$$

In particular if (X, dist_X) is an interval of the real line $\mathbb{R}$ then $\mathfrak{q}$ is called a *quasi-geodesic* or an (L, C)-*quasi-geodesic*.

A *quasi-isometry* is a quasi-isometric embedding $\mathfrak{q}\colon X \to Y$ such that there exists a quasi-isometric embedding $\mathfrak{q}'\colon Y \to X$ with the property that $\mathfrak{q} \circ \mathfrak{q}'$ and $\mathfrak{q}' \circ \mathfrak{q}$ are at finite distance from the identity maps.

Remark 3.12 (quasi-injectivity). Although a quasi-isometric embedding is not necessarily injective, a weaker version of injectivity holds: If $\mathfrak{q}$ is an (L, C)-quasi-isometric embedding then $\mathrm{dist}(x, y) > LC$ implies $\mathrm{dist}(\mathfrak{q}(x), \mathfrak{q}(y)) > 0$.

Definition 3.13 (Lipschitz maps). Let $L \ge 1$. A map $\mathfrak{q}\colon (X, \mathrm{dist}_X) \to (Y, \mathrm{dist}_Y)$ is called *Lipschitz* if

$$\mathrm{dist}_Y(\mathfrak{q}(x), \mathfrak{q}(x')) \le L\mathrm{dist}_X(x, x')$$

for every $x, x' \in X$. The map $\mathfrak{q}$ is called *bi-Lipschitz* if it also satisfies

$$\mathrm{dist}_Y(\mathfrak{q}(x), \mathfrak{q}(x')) \ge \frac{1}{L}\mathrm{dist}_X(x, x')\,.$$

Remark 3.14. Let (X_n) and (Y_n) be sequences of metric spaces, $e_n \in X_n$, $e'_n \in Y_n$ $(n \in I)$. Then it is easy to see that any sequence $\mathfrak{q}_n\colon X_n \to Y_n$ of (L_n, C_n)-quasi-isometries with $\mathfrak{q}_n(e_n) = e'_n$, $n \in I$, induces an (L, C)-quasi-isometry $\mathfrak{q}\colon \lim^\omega(X_n)_e \to \lim^\omega(Y_n)_{e'}$ where $e = (e_n)^\omega$, $e' = (e'_n)^\omega$, and $L = \lim_\omega L_n$, $C = \lim_\omega C_n$ provided $L < \infty$, $C < \infty$. Moreover, the ω-limit of the images $\mathfrak{q}_n(X_n)$ coincides with the image of $\mathfrak{q}$.

Remark 3.15. Let $\mathfrak{q}_n\colon [0, \ell_n] \to X$ be a sequence of (L, C)-quasi-geodesics in a geodesic metric space (X, dist). Then the ω-limit $\lim^\omega(\mathfrak{q}_n([0, \ell_n]))$ in any asymptotic cone $\mathrm{Con}^\omega(X, e, d)$ is either empty, or a bi-Lipschitz arc or a bi-Lipschitz ray or a bi-Lipschitz line. This immediately follows from Remark 3.14.

Remark 3.16. Any quasi-isometric embedding $\mathfrak{q}$ of (X, dist_X) into (Y, dist_Y) induces a bi-Lipschitz embedding of $\mathrm{Con}^\omega(X; e, d)$ into $\mathrm{Con}^\omega(Y; (q(e_n)), d)$ for every ω, e and d [Gr_3].

If G is a finitely generated group then the asymptotic cones corresponding to different observation points are isometric [Gr3]. Thus when we consider an asymptotic cone of a finitely generated group, we shall always assume that the observation point e is $(1)^\omega$.

Remark 3.17. Notice [Gr3] that the group G_e^ω acts on $\mathrm{Con}^\omega(G; e, d)$ by isometries:

$$(g_n)^\omega \lim^\omega(x_n) = \lim^\omega(g_n x_n).$$

This action is transitive, so, in particular, every asymptotic cone of a group is homogeneous.

More generally if a group G acts by isometries on a metric space (X, dist) and there exists a bounded subset $B \subset X$ such that $X = GB$ then all asymptotic cones of X are homogeneous metric spaces.

Definition 3.18 (asymptotic properties). We say that a space *has a certain property asymptotically* if each of its asymptotic cones has this property. For example, a space may be asymptotically CAT(0), asymptotically without cut-point etc.

Definition 3.19 (asymptotically tree-graded spaces). Let (X, dist) be a metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X. In every asymptotic cone $\mathrm{Con}^\omega(X; e, d)$, we consider the collection of subsets

$$\mathcal{A}_\omega = \left\{ \lim^\omega(A_{i_n}) \mid (i_n)^\omega \in I^\omega \text{ such that the sequence } \left(\frac{\mathrm{dist}(e_n, A_{i_n})}{d_n}\right) \text{ is bounded} \right\}.$$

We say that X is *asymptotically tree-graded with respect to* $\mathcal{A}$ if every asymptotic cone $\mathrm{Con}^\omega(X; e, d)$ is tree-graded with respect to $\mathcal{A}_\omega$.

Corollary 4.30 will show that there is no need to vary the ultrafilter in Definition 3.19: if a space is tree-graded with respect to a collection of subsets for one ultrafilter, it is tree-graded for any other with respect to the same collection of subsets.

3.2 Ultralimits of asymptotic cones are asymptotic cones

Definition 3.20 (an ultraproduct of ultrafilters). Let ω be an ultrafilter over I and let $\mu = (\mu_n)_{n \in I}$ be a sequence of ultrafilters over I. We consider each μ_n as a measure on the set $\{n\} \times I$ and ω as a measure on I.

For every subset $A \subseteq I \times I$ we set $\omega\mu(A)$ equal to the ω-measure of the set of all $n \in I$ such that $\mu_n(A \cap (\{n\} \times I)) = 1$.

In other words

$$\omega\mu(A) = \int \mu_n \left(A \cap (\{n\} \times I)\right) \, d\omega(n).$$

Notice that this is a generalization of the standard notion of product of ultrafilters (see [Sh, Definition 3.2 in Chapter VI]).

Lemma 3.21. *(cf [Sh, Lemma 3.6 in Chapter VI]) $\omega\mu$ is an ultrafilter over $I \times I$.*

Proof. It suffices to prove that $\omega\mu$ is finitely additive and that it takes the zero value on finite sets.

Let A and B be two disjoint subsets of $I \times I$. Then for every $n \in I$ the sets $A \cap (\{n\} \times I)$ and $B \cap (\{n\} \times I)$ are disjoint. Hence (by the additivity of μ_n) for every $n \in I$

$$\mu_n((A \cup B) \cap (\{n\} \times I)) = \mu_n(A \cap (\{n\} \times I)) + \mu_n(B \cap (\{n\} \times I)).$$

Therefore (by the additivity of ω)

$$\omega\mu(A \sqcup B) = \omega\mu(A) + \omega\mu(B).$$

Let now A be a finite subset of $I\times I$. Then the set of numbers n for which $\mu_n(A\cap(\{n\}\times I)) = 1$ is empty. So $\omega\mu(A) = 0$ by definition. $\square$

Lemma 3.22 (double ultralimit of sequences). *Let ω, μ_n, $n \in I$, be as in Definition 3.20. Let $r_k^{(n)}$ be an uniformly bounded double indexed sequence of real numbers, $k, n \in I$. Then*

$$\lim_{\omega\mu} r_k^{(n)} = \lim_\omega \lim_{\mu_n} r_k^{(n)} \tag{1}$$

(the internal limit is taken with respect to k).

Proof. Let $r = \lim_{\omega\mu} r_k^{(n)}$. It follows that, for every $\varepsilon > 0$,

$$\omega\mu\left\{(n,k) \mid r_k^{(n)} \in (r-\varepsilon, r+\varepsilon)\right\} = 1 \Leftrightarrow$$

$$\omega\left\{n \in I \mid \mu_n\left\{k \mid r_k^{(n)} \in (r-\varepsilon, r+\varepsilon)\right\} = 1\right\} = 1\,.$$

It follows that

$$\omega\left\{n \in I \mid \lim_{\mu_n} r_k^{(n)} \in [r-\varepsilon, r+\varepsilon]\right\} = 1\,,$$

which implies that

$$\lim_\omega \lim_{\mu_n} r_k^{(n)} \in [r-\varepsilon, r+\varepsilon]\,.$$

Since this is true for every $\varepsilon > 0$ we conclude that $\lim_\omega \lim_{\mu_n} r_k^{(n)} = r$. $\square$

Lemma 3.22 immediately implies:

Proposition 3.23 (double ultralimit of metric spaces). *Let ω and μ be as in Definition 3.20. Let $\left(X_k^{(n)}, \mathrm{dist}_k^{(n)}\right)$ be a double indexed sequence of metric spaces, $k, n \in I$, and let e be a double indexed sequence of points $e_k^{(n)} \in X_k^{(n)}$. We denote by $e^{(n)}$ the sequence $\left(e_k^{(n)}\right)_{k\in I}$.*

The map

$$\lim_{\omega\mu}\left(x_k^{(n)}\right) \mapsto \lim_\omega\left(\lim_{\mu_n}\left(x_k^{(n)}\right)\right)\,, \tag{2}$$

is an isometry from $\lim^{\omega\mu}\left(X_k^{(n)}\right)_e$ onto $\lim^\omega\left(\lim^{\mu_n}\left(X_k^{(n)}\right)_{e^{(n)}}\right)_{e'}$, where $e'_n = \lim^{\mu_n}(e^{(n)})$

Corollary 3.24 (ultralimits of cones are cones). *Let X be a metric space. Let ω and μ be as above. For every $n \in I$ let $e^{(n)} = \left(e_k^{(n)}\right)_{k\in I}$ be an observation point, $d^{(n)} = \left(d_k^{(n)}\right)_{k\in I}$ be a sequence of scaling constants satisfying $\lim_{\mu_n} d_k^{(n)} = \infty$ for every $n \in I$. Let $\mathrm{Con}^{\mu_n}\left(X; e^{(n)}, d^{(n)}\right)$ be the corresponding asymptotic cone of X. Then the map*

$$\lim_{\omega\mu}\left(x_k^{(n)}\right) \mapsto \lim_\omega\left(\lim_{\mu_n}\left(x_k^{(n)}\right)\right), \tag{3}$$

is an isometry from $\mathrm{Con}^{\omega\mu}(X; e, d)$ onto

$$\lim^\omega\left(\mathrm{Con}^{\mu_n}\left(X; e^{(n)}, d^{(n)}\right)\right)_{(\lim^{\mu_n}(e^{(n)}))}\,,$$

where $e = \left(e_k^{(n)}\right)_{(n,k)\in I\times I}$ and $d = \left(d_k^{(n)}\right)_{(n,k)\in I\times I}$.

Proof. Let us prove that $\lim_{\omega\mu} d_k^{(n)} = \infty$. Let $M > 0$. For every $n \in I$ we have that $\lim_{\mu_n} d_k^{(n)} = \infty$, whence $\mu_n \left\{k \in I \mid d_k^{(n)} > M\right\} = 1$. It follows that $\left\{n \in I \mid \mu_n \left\{k \in I \mid d_k^{(n)} > M\right\} = 1\right\} = I$, therefore its ω-measure is 1. We conclude that $\omega\mu \left\{(n,k) \mid d_k^{(n)} > M\right\} = 1$.

It remains to apply Proposition 3.23 to the sequence of metric spaces $\left(X, \frac{1}{d_k^{(n)}}\mathrm{dist}\right)$ and to e. □

3.3 Another definition of asymptotic cones

In [Gr$_3$], [VDW] and some other papers, a more restrictive definition of asymptotic cones is used. In that definition, the set I is equal to $\mathbb{N}$ and the scaling constant d_n must be equal to n for every n. We shall call these asymptotic cones *restrictive.*

It is easy to see that every restrictive asymptotic cone is an asymptotic cone in our sense. The converse statement can well be false although we do not have any explicit examples.

Also for every ultrafilter ω over I and every sequence of scaling constants $d = (d_n)_{n\in I}$, there exists an ultrafilter μ over $\mathbb{N}$ such that the asymptotic cone $\mathrm{Con}^\omega(G; e, d)$ contains an isometric copy of the restrictive asymptotic cone $\mathrm{Con}^\mu(G; e, (n))$. Indeed, let ϕ be a map $I \to \mathbb{N}$ such that $\phi(i) = [d_i]$. Now define the ultrafilter μ on $\mathbb{N}$ by $\mu(A) = \omega(\phi^{-1}(A))$ for every set $A \subseteq \mathbb{N}$. The embedding $\mathrm{Con}^\mu(G; e, (n)) \to \mathrm{Con}^\omega(G; e, d)$ is defined by $(x_n)^\mu_{n\in\mathbb{N}} \mapsto (x_{\phi(i)})^\omega_{i\in I}$.

Remark 3.25. In the particular case when the sets $\{i \in I \mid [d_i] = k\}$ are of uniformly bounded (finite) size, this embedding is a surjective isometry [Ri].

The restrictive definition of asymptotic cones is, in our opinion, less natural because the ω-limit of restrictive asymptotic cones is not canonically represented as a restrictive asymptotic cone (see Corollary 3.24). Conceivably, it may even not be a restrictive asymptotic cone in general. The next statement shows that it is a restrictive asymptotic cone in some particular cases.

Proposition 3.26. *Let ν_n, $n \in \mathbb{N}$ be a sequence of ultrafilters over $\mathbb{N}$. Let (I_n) be sequence of pairwise disjoint subsets of $\mathbb{N}$ such that $\nu_n(I_n) = 1$. Let $C_n = \mathrm{Con}^{\nu_n}(X; e^{(n)}, (n))$, $n \in \mathbb{N}$, be a restrictive asymptotic cone of a metric space X. Then the ω-limit of asymptotic cones C_n is a restrictive asymptotic cone.*

Proof. Let μ_n be the restriction of ν_n onto I_n, $n \in \mathbb{N}$. Then C_n is isometric to $\mathrm{Con}^{\mu_n}(X; e^{(n)}, d^{(n)})$ where $d^{(n)}$ is the sequence of all numbers from I_n in the increasing order. By Corollary 3.24, $\lim^\omega (C_n)_{\lim^{\mu_n}(e^{(n)})}$ is the asymptotic cone $\mathrm{Con}^{\omega\mu}(X; e, d)$ where $e = \left(e_k^{(n)}\right)_{(n,k)\in\mathbb{N}\times\mathbb{N}}$ and $d = \left(d_k^{(n)}\right)_{(n,k)\in\mathbb{N}\times\mathbb{N}}$. For every natural number a the set of pairs (n,k) such that $d_k^{(n)} = a$ contains at most one element because the subsets $I_n \subseteq \mathbb{N}$ are disjoint. It remains to apply Remark 3.25. □

3.4 Simple triangles in ultralimits of metric spaces

Definition 3.27 (k-gons). We say that a metric space P is a geodesic (quasi-geodesic) k-gon if it is a union of k geodesics (quasi-geodesics) $\mathfrak{q}_1, ..., \mathfrak{q}_k$ such that $(\mathfrak{q}_i)_+ = (\mathfrak{q}_{i+1})_-$ for every $i = 1, ..., k$ (here $k+1$ is identified with 1).

For every $i = 1, ..., k$, we denote the polygonal curve $P \setminus (\mathfrak{q}_{i-1} \cup \mathfrak{q}_i)$ by $\mathcal{O}_{x_i}(P)$, where $x_i = (\mathfrak{q}_{i-1})_+ = (\mathfrak{q}_i)_-$. When there is no possibility of confusion we simply denote it by $\mathcal{O}_{x_i}$.

Lemma 3.28. *(1) Let P_n, $n \in \mathbb{N}$, be a sequence of geodesic k-gons in metric spaces (X_n, dist_n). Let ω be an ultrafilter over $\mathbb{N}$, such that $\lim^\omega(P_n) = P$, where P is a simple geodesic k-gon in the metric space $\lim^\omega(X_n)_e$ with metric* dist*. Let $\mathcal{V}_n$ be the set of vertices of P_n in the clockwise order. Let D_n be the supremum over all points x contained in two distinct edges of P_n of the distances* $\mathrm{dist}\,(x, \mathcal{V}_n)$ *. Then $\lim_\omega D_n = 0$.*

(2) Let P be a simple k-gon in (X, dist). For every $\delta > 0$ we define $D_\delta = D_\delta(P)$ to be the supremum over all k-gons P_δ in X that are at Hausdorff distance at most δ from P and over all points x contained in two distinct edges of P_δ of the distances $\mathrm{dist}\,(x, \mathcal{V}_\delta)$*, where $\mathcal{V}_\delta$ is the set of vertices of P_δ. Then $\lim_{\delta\to 0} D_\delta = 0$.*

Proof. (1) Since the ω-limit of the diameters of P_n is the diameter of P, it follows that the diameters of P_n are uniformly bounded ω-almost surely. In particular D_n is uniformly bounded ω-almost surely, therefore its ω-limit exists and it is finite. Suppose that $\lim_\omega D_n = 2D > 0$. Then ω-almost surely there exists x_n contained in two distinct edges of P_n such that $\mathrm{dist}_n\,(x_n, \mathcal{V}_n) > D$. Without loss of generality we may suppose that $x_n \in [A_n, B_n] \cap [B_n, C_n]$ for every n, where $[A_n, B_n], [B_n, C_n]$ are two consecutive edges of P_n such that $\lim^\omega([A_n, B_n]) = [A, B]$, $\lim^\omega([B_n, C_n]) = [B, C]$, where $[A, B], [B, C]$ are two consecutive edges of P. Then $\lim^\omega(x_n) \in [A, B] \cap [B, C]$, which by simplicity of P implies that $\lim^\omega(x_n) = B$. On the other hand we have that $\mathrm{dist}_n\,(x_n, \mathcal{V}_n) > D$, which implies that $\mathrm{dist}(\lim^\omega(x_n), B) \geq D$. We have obtained a contradiction.

(2) Assume that $\lim_{\delta\to 0} D_\delta = 2D > 0$. It follows that there exists a sequence (P_n) of k-gons endowed with metrics such that their Hausdorff distance to P tends to zero and such that there exists x_n contained in two distinct edges of P_n and at distance at least D of the vertices of P_n. According to [KaL$_1$], it follows that $\lim^\omega(P_n) = P$ for every ultrafilter ω. On the other hand $D_n > D$ for all $n \in \mathbb{N}$. We thus obtain a contradiction of (1). □

Proposition 3.29 (limits of simple polygons). *Consider an ultrafilter ω over $\mathbb{N}$ and a sequence of metric spaces, (X_n, dist_n), $n \in \mathbb{N}$. Let $e \in \Pi X_n/\omega$ be an observation point. For every simple geodesic triangle Δ in $\lim^\omega(X_n)_e$, for every sufficiently small $\varepsilon > 0$ there exists $k_0 = k_0(\varepsilon)$ and a simple geodesic triangle Δ_ε with the properties:*

- *(a) The Hausdorff distance between Δ and Δ_ε does not exceed ε;*
- *(b) Δ_ε contains the midpoints of the edges of Δ;*
- *(c) The triangle Δ_ε can be written as $\lim^\omega(P_n^\varepsilon)$, where each P_n^ε is a geodesic k-gon in X_n, $k \leq k_0$, P_n^ε is simple and the lengths of all edges of P_n^ε are $O(1)$ ω-almost surely.*

Proof. Let A, B, C be the vertices of Δ, in the clockwise order, and let M_{AB}, M_{BC} and M_{AC} be the midpoints of $[A, B], [B, C]$ and $[A, C]$, respectively.

We construct Δ_ε in several steps.

Step I. Constructing not necessarily simple geodesic triangles Δ_ε.

For every small $\delta > 0$ we divide each of the halves of edges of Δ determined by a vertex and a midpoint into finitely many segments of length at least δ and at most 2δ. Let $\mathcal{V}$ be the set of endpoints of all these segments, endowed with the natural cyclic order. We call $\mathcal{V}$ a *δ-partition of* Δ. We assume that $\{A, B, C, M_{AB}, M_{BC}, M_{AC}\} \subset \mathcal{V}$. Every $t \in \mathcal{V}$ can be written as $t = \lim^\omega(t_n)$, hence $\mathcal{V} = \lim^\omega(\mathcal{V}_n)$, where each $\mathcal{V}_n$ is endowed with a cyclic order. Let P_n be a geodesic k-gon with vertices $\mathcal{V}_n$, where $k = |\mathcal{V}|$. The limit set $\Delta_\delta = \lim^\omega(P_n)$ is a geodesic triangle with vertices A, B, C and at Hausdorff distance at most δ from Δ.

Notation: Let E, F be two points on an edge of Δ_δ. We denote the part of the geodesic side of Δ_δ between E and F in Δ_δ by $[E, F]_\delta$. If E, F are two points on an edge of Δ, we denote the part of the side of Δ between E and F by $[E, F]$. This is to avoid confusion between different geodesics joining two such points.

Step II. Making Δ_ε simple.

For every $\delta > 0$, we consider $D_\delta = D_\delta(\Delta)$ given by Lemma 3.28. Let

$$\alpha(\Delta) \inf \{\mathrm{dist}\,(x, \mathcal{O}_x(\Delta)) \mid x \in \{A, B, C\}\} \ .$$

By Lemma 3.28 we have $\lim_{\delta\to 0} D_\delta = 0$. Therefore, for δ small enough we have

$$2D_\delta + 4\delta < \alpha(\Delta) \text{ and } D_\delta + 2\delta \le \frac{1}{10} \min\{\mathrm{dist}(A, B), \mathrm{dist}(B, C), \mathrm{dist}(C, A)\} \ . \tag{4}$$

Fix a δ satisfying (4), a δ-partition $\mathcal{V}$ of Δ, and a corresponding triangle $\Delta_\delta = \lim^\omega(P_n)$.

Let A_1 and A_2 be the nearest to A points of $\mathcal{V} \setminus \mathcal{N}_{D_\delta+\delta}(A)$ on the edges $[A, B]$ and $[A, C]$, respectively. For an appropriate choice of Δ_δ, we may suppose that $\mathrm{dist}(A, A_1) = \mathrm{dist}(A, A_2)$. We note that $\mathrm{dist}(A, A_1) \in [D_\delta + \delta, D_\delta + 2\delta]$. Similarly we take $B_1 \in [B, C] \cap \mathcal{V}$, $B_2 \in [B, A] \cap \mathcal{V}$ and $C_1 \in [C, A] \cap \mathcal{V}$, $C_2 \in [C, B] \cap \mathcal{V}$ with $\mathrm{dist}(B, B_1) = \mathrm{dist}(B, B_2) \in [D_\delta + \delta, D_\delta + 2\delta]$ and $\mathrm{dist}(C, C_1) = \mathrm{dist}(C, C_2) \in [D_\delta + \delta, D_\delta + 2\delta]$.

Suppose that $[A_1, B_2]_\delta$ and $[B_1, C_2]_\delta$ have a point E in common. The definition of D_δ implies that $E \in \mathcal{N}_{D_\delta}(\{A, B, C\})$. On the other hand $E \in [A_1, B_2]_\delta$ implies $E \notin \mathcal{N}_{D_\delta}(\{A, B\})$ and $E \in [B_1, C_2]_\delta$ implies $E \notin \mathcal{N}_{D_\delta}(\{B, C\})$, a contradiction.

We conclude, by repeating the previous argument, that the segments $[A_1, B_2]_\delta$, $[B_1, C_2]_\delta$ and $[C_1, A_2]_\delta$ are pairwise disjoint. Since $\mathrm{dist}(A, A_1)$, $\mathrm{dist}(B, B_2) \le D_\delta + 2\delta \le \frac{1}{10}\mathrm{dist}(A, B)$, it follows that M_{AB} is contained in $[A_1, B_2]_\delta$. Likewise, M_{BC} and M_{AC} are contained in $[B_1, C_2]_\delta$ and $[C_1, A_2]_\delta$, respectively.

Let d_A be the supremum of $\mathrm{dist}(E, A)$ for all E satisfying two conditions: $E \in [A_1, A]_\delta$ and $\mathrm{dist}(A_2, E) + \mathrm{dist}(E, A) = \mathrm{dist}(A_2, A)$. Since these two conditions define a closed set, it follows that there exists $A' \in [A_1, A]_\delta$ such that $\mathrm{dist}(A_2, A') + \mathrm{dist}(A', A) = \mathrm{dist}(A_2, A)$ and $\mathrm{dist}(A, A') = d_A$. Obviously $A' \notin \{A_1, A_2\}$. In other words, A' is the farthest from A point in $[A_1, A]_\delta$ which is contained in a geodesic joining A_2 and A. Hence A' has the property that every geodesic joining it with A_2 intersects $[A_1, A']_\delta$ only in A'. Similarly we find points $B' \in [B_1, B]_\delta$ and $C' \in [C_1, C]_\delta$.

Recall that $\Delta_\delta = \lim^\omega(P_n)$. Let P_n^A be a sequence of polygonal lines in P_n with endpoints A'_n, B_n^2, having as limit $[A', B_2]_\delta$. Likewise let P_n^B and P_n^C be sequences of polygonal lines in P_n, with endpoints B'_n, C_n^2 and C'_n, A_n^2, having as limits $[B', C_2]_\delta$ and $[C', A_2]_\delta$, respectively. We consider the new sequence of polygons $P'_n = P_n^A \cup [B_n^2, B'_n] \cup P_n^B \cup [C_n^2, C'_n] \cup P_n^C \cup [A_n^2, A'_n]$. The limit set $\lim^\omega(P'_n)$ is $[A', B_2]_\delta \cup \mathfrak{g}_{B_2B'} \cup [B', C_2]_\delta \cup \mathfrak{g}_{C_2C'} \cup [C', A_2]_\delta \cup \mathfrak{g}_{A_2A'}$ where $\mathfrak{g}_{B_2B'} = \lim^\omega([B_n^2, B'_n])$ is a geodesic and likewise for $\mathfrak{g}_{C_2C'}, \mathfrak{g}_{A_2A'}$.

We have $\mathrm{dist}(C', A) = \mathrm{dist}(C', A_2) + \mathrm{dist}(A_2, A) = \mathrm{dist}(C', A_2) + \mathrm{dist}(A_2, A') + \mathrm{dist}(A', A)$. It follows that by joining the pairs of points (C', A_2), (A_2, A') and (A', A) by geodesics we obtain a geodesic from C' to A. In particular $[C', A_2]_\delta \cup \mathfrak{g}_{A_2A'}$ is a geodesic. Likewise, $[A', B_2]_\delta \cup \mathfrak{g}_{B_2B'}$ and $[B', C_2]_\delta \cup \mathfrak{g}_{C_2C'}$ are geodesics. Therefore $\lim^\omega(P'_n)$ is a geodesic triangle Δ'_δ with vertices A', B', C'. By construction the Hausdorff distance between Δ'_δ and Δ_δ is at most $D_\delta + 2\delta$, hence the Hausdorff distance between Δ'_δ and Δ is at most $D_\delta + 3\delta$.

Suppose that two edges of Δ'_δ have a common point E. Suppose the two edges are $[A', B_2]_\delta \cup \mathfrak{g}_{B_2B'}$ and $[B', C_2]_\delta \cup \mathfrak{g}_{C_2C'}$. If $E \in [A', A_1]_\delta$ then $\mathrm{dist}(A, E) \le D_\delta + 2\delta$. On the other hand $E \in [B', C_2]_\delta \cup \mathfrak{g}_{C_2C'}$ implies $E \in \mathcal{N}_{D_\delta+2\delta}([B, C])$. Hence $\mathrm{dist}(A, [B, C]) \le 2D_\delta + 4\delta < \alpha(\Delta)$, a contradiction.

If $E \in \mathfrak{g}_{C_2C'}$ then $\mathrm{dist}(C, E) \le D_\delta + 2\delta$ which together with $E \in [A', B_2]_\delta \cup \mathfrak{g}_{B_2B'} \subset \mathcal{N}_{D_\delta+2\delta}([A, B])$ implies $\mathrm{dist}(C, [A, B]) \le 2D_\delta + 4\delta < \alpha(\Delta)$, a contradiction.

If $E \in [A_1, B_2]_\delta$ then $E \notin [B_1, C_2]_\delta$. Also since $\mathrm{dist}(B, E) \ge \mathrm{dist}(B, B_2) = \mathrm{dist}(B, B_1)$ it follows that $E \notin [B', B_1]_\delta$, a contradiction.

If $E \in \mathfrak{g}_{B_2B'}$ then an argument similar to the previous gives $E \notin [B_1, C_2]_\delta$. We conclude that $E \in [B', B_1]_\delta$. By the choice of B' we have $E = B'$.

We conclude that Δ'_δ is a simple geodesic triangle, containing the midpoints of the edges of Δ, at Hausdorff distance at most $D_\delta + 3\delta$ from Δ, and $\Delta'_\delta = \lim^\omega(P'_n)$, where P'_n is a geodesic m-gon, with $m \le k + 3$.

Step III. Making polygons simple.

Let D_n be the supremum over all points x contained in two distinct edges of P'_n of the distances from x to the vertices of P'_n. Applying Lemma 3.28, (1), to (P'_n) and to $\Delta'_\delta = \lim^\omega(P'_n)$ we obtain that D_n tends to zero as $n \to \infty$. Let v_n be a vertex of P'_n. We consider the farthest point v'_n in the ball $B(v_n, 2D_n)$ contained in both edges of endpoint the vertex v_n. Cut the bigon of vertices v_n, v'_n from the polygon, and repeat this operation for every vertex v_n of P'_n. As a result, we obtain a new polygon P''_n which is simple and at Hausdorff distance at most $2D_n$ from P'_n. It follows that $\lim^\omega(P''_n) = \lim^\omega(P'_n) = \Delta'_\delta$. □

Theorem 3.30 (being tree-graded is closed under ultralimits). *For every $n \in \mathbb{N}$ let $\mathbb{F}_n$ be a complete geodesic metric space which is tree-graded with respect to a collection $\mathcal{P}_n$ of closed geodesic subsets of $\mathbb{F}_n$. Let ω be an ultrafilter over $\mathbb{N}$ and let $e \in \Pi\mathbb{F}_n/\omega$ be an observation point. The ultralimit $\lim^\omega(\mathbb{F}_n)_e$ is tree-graded with respect to the collection of limit sets*

$$\mathcal{P}_\omega = \{\lim^\omega(M_n) \mid M_n \in \mathcal{P}_n \,,\ \mathrm{dist}(e_n, M_n) \text{ bounded uniformly in } n\} \,.$$

Proof. **Property** (T_1). Let $\lim^\omega(M_n)\,, \lim^\omega(M'_n) \in \mathcal{P}_\omega$ be such that there exist two distinct points x_ω, y_ω in $\lim^\omega(M_n) \cap \lim^\omega(M'_n)$. It follows that $x_\omega = \lim^\omega(x_n) = \lim^\omega(x'_n)$ and $y_\omega \lim^\omega(y_n) = \lim^\omega(y'_n)$, where $x_n, y_n \in M_n$, $x'_n, y'_n \in M'_n$, $\mathrm{dist}(x_n, x'_n) = o(1)$, $\mathrm{dist}(y_n, y'_n) = o(1)$, while $\mathrm{dist}(x_n, y_n) = O(1)$, $\mathrm{dist}(x'_n, y'_n) = O(1)$.

By contradiction suppose that $M_n \neq M'_n$ ω-almost surely. Then property (T_2) of the space $\mathbb{F}_n$ and Corollary 2.9 imply that M_n projects into M'_n in a unique point z_n and that $z_n \in [x_n, x'_n] \cap [y_n, y'_n]$. It follows that $\mathrm{dist}(x_n, z_n) = o(1)$ and $\mathrm{dist}(y_n, z_n) = o(1)$, therefore that $\mathrm{dist}(x_n, y_n) = o(1)$. This contradiction implies that $M_n = M'_n$ ω-almost surely, so $\lim^\omega(M_n) = \lim^\omega(M'_n)$.

Property (T_2)**.** Let Δ be a simple geodesic triangle in $\lim^\omega(\mathbb{F}_n)_e$. Consider an arbitrary sufficiently small $\varepsilon > 0$ and apply Proposition 3.29. We obtain a simple geodesic triangle Δ_ε satisfying properties (a), (b), (c) in the conclusion of the Proposition. In particular $\Delta_\varepsilon = \lim^\omega(P^\varepsilon_n)$, where P^ε_n is a simple geodesic polygon in $\mathbb{F}_n$. Property (T''_2) applied to $\mathbb{F}_n$ implies that P^ε_n is contained in one piece M_n. Consequently $\Delta_\varepsilon \subset \lim^\omega(M_n)$. Property (b) of Δ_ε implies that $\lim^\omega(M_n)$ contains the three distinct middle points of the edges of Δ. This and property (T_1) already proven imply that all triangles Δ_ε are contained in the same $\lim^\omega(M_n)$. Property (a) and the fact that $\lim^\omega(M_n)$ is closed imply that $\Delta \subset \lim^\omega(M_n)$. □

Definition 3.31. Let P be a polygon with quasi-geodesic edges and with set of vertices $\mathcal{V}$. Points in $P \setminus \mathcal{V}$ are called *interior points of* P. Let $p \in P$. The *inscribed radius in* p with respect to P is either the distance from p to the set $\mathcal{O}_p$, if p is a vertex, or the distance from p to the set $P \setminus \mathfrak{q}$ if p is contained in the interior of the edge $\mathfrak{q}$.

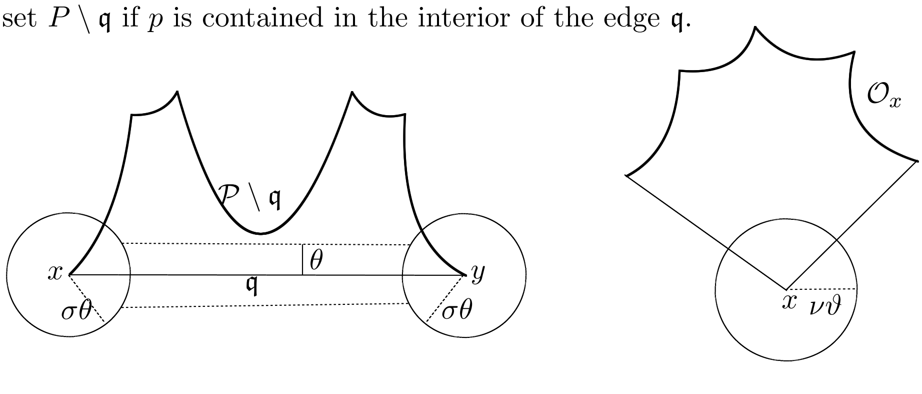

Figure 2: Properties (F_1) and (F_2).

Definition 3.32 (fat polygons). Let $\vartheta > 0$, $\sigma \geq 1$ and $\nu \geq 4\sigma$. We call a k-gon P with quasi-geodesic edges (ϑ, σ, ν)*-fat* if the following properties hold:

(F_1) (**large comparison angles, large inscribed radii in interior points**) for every edge $\mathfrak{q}$ with endpoints $\{x, y\}$ we have

$$\operatorname{dist}(\mathfrak{q} \setminus \mathcal{N}_{\sigma\vartheta}(\{x, y\}), P \setminus \mathfrak{q}) \geq \vartheta;$$

(F_2) (**large edges, large inscribed radii in vertices**) for every vertex x we have

$$\operatorname{dist}(x, \mathcal{O}_x) \geq \nu\vartheta.$$

Remark 3.33. Property (F_1) implies that in each of the vertices x, y certain comparison angles are at least $\frac{1}{\sigma}$ and that in the interior points of P outside $\mathcal{N}_{\sigma\vartheta}(\mathcal{V})$ the inscribed radius is at least ϑ.

Property (F_2) ensures that for every edge $\mathfrak{q}$ the set $\mathfrak{q} \setminus \mathcal{N}_{\sigma\vartheta}(\{x, y\})$ has diameter at least $2\sigma\vartheta$, in particular it is never empty. It also ensures that the inscribed radius in every vertex is at least $\nu\vartheta$.

Proposition 3.34 (triangles in an asymptotic cone are ω-limits of fat polygons). *For every simple geodesic triangle Δ in $\mathrm{Con}^{\omega}(X; e, d)$, for every sufficiently small $\varepsilon > 0$ there exists $k_0 = k_0(\varepsilon)$ and a simple geodesic triangle Δ_ε with the properties:*

(a) The Hausdorff distance between Δ and Δ_ε does not exceed ε;

(b) Δ_ε contains the midpoints of the edges of Δ;

(c) For every $\vartheta > 0$ and $\nu \geq 8$, the triangle Δ_ε can be written as $\lim^{\omega}(P_n^{\varepsilon})$, where each P_n^{ε} is a geodesic k-gon in X, $k \leq k_0$, and P_n^{ε} is $(\vartheta, 2, \nu)$-fat ω-almost surely.

Proof. Proposition 3.29 applied to $\left(X\,,\,\frac{1}{d_n}\operatorname{dist}\right)$, ω, e and Δ implies that for every $\varepsilon > 0$ there exists $k_0 = k_0(\varepsilon)$ and Δ_ε satisfying (a) and (b) and such that $\Delta_\varepsilon = \lim^{\omega}(P_n)$, where P_n are simple geodesic k-gons in X, $3 \leq k \leq k_0$, such that the lengths of all edges in P_n are $O(d_n)$

ω-almost surely. Remark 3.2 implies that there exists $m \in \{3, \dots, k_0\}$ such that P_n have m edges ω-almost surely. Let $\vartheta > 0$ and $\nu \geq 8$. We modify the sequence of polygons (P_n) so that their limit set stays the same while the polygons become $(\vartheta, 2, \nu)$-fat.

Let $\mathcal{V}_n = \{v_1^n, v_2^n, \dots, v_m^n\}$ be the set of vertices of P_n in the clockwise order. We denote the limit set $\lim^\omega(\mathcal{V}_n)$ by $\mathcal{V}$, and we endow it with the clockwise order on Δ_ε. There exists $\varrho > 0$ such that for every $v \in \mathcal{V}$, the distance between v and $\mathcal{O}_v(\Delta_\varepsilon)$ is at least 2ϱ, where $\mathcal{O}_v(\Delta_\varepsilon)$ is taken in Δ_ε considered as a polygon with vertices $\mathcal{V}$. It follows that ω-almost surely for every $i \in \{1, 2, \dots, m\}$ we have $\operatorname{dist}\left(v_i^n, \mathcal{O}_{v_i^n}(P_n)\right) \geq \varrho d_n$. In particular, ω-almost surely all the edges of P_n have length at least ϱd_n.

Convention: In what follows we use the notation $[v_i^n, v_{i+1}^n]$ for a generic edge of P_n, where $i+1$ is taken modulo m.

Let ϵ_n be the supremum of distances $\operatorname{dist}(x, \mathcal{V}_n)$ for all $x \in [v_i^n, v_{i+1}^n] \cap \mathcal{N}_{\nu\vartheta}\left([v_j^n, v_{j+1}^n]\right)$, $i \neq j$, $i, j \in \{1, 2, \dots, m\}$. Suppose that $\lim_\omega \frac{\epsilon_n}{d_n} = 2\eta > 0$. Then there exist $x_n \in [v_i^n, v_{i+1}^n] \cap \mathcal{N}_{\nu\vartheta}\left([v_j^n, v_{j+1}^n]\right)$, $i \neq j$, $i, j \in \{1, 2, \dots, m\}$, with $\operatorname{dist}(x_n, \mathcal{V}_n) \geq \eta d_n$ ω-almost surely. Taking the ω-limit, we get a contradiction with the fact that Δ_ε is simple. Therefore $\lim_\omega \frac{\epsilon_n}{d_n} = 0$.

Notation: We denote by $\mathfrak{N}$ the set of all $n \in \mathbb{N}$ such that for every $i \in \{1, 2, \dots, m\}$ we have $\operatorname{dist}\left(v_i^n, \mathcal{O}_{v_i^n}\right) \geq \varrho d_n$ and such that $\varrho d_n \geq 2\epsilon_n + 2 + (2\nu + 1)\vartheta$. Obviously $\mathfrak{N} \in \omega$.

Let $[v_{i-1}^n, v_i^n]$ and $[v_i^n, v_{i+1}^n]$ be two consecutive edges of P_n. Let $\bar{v}_i^n$ be the farthest point of v_i^n in $[v_{i-1}^n, v_i^n] \cap \mathcal{N}_{\epsilon_n+1}(v_i^n)$ contained in the $\nu\vartheta$-tubular neighborhood of a different edge $\mathfrak{p}$ of P_n. The edge $\mathfrak{p}$ has to be at a distance at most $\epsilon_n + 1 + \nu\vartheta$ from v_i^n. It follows that for every $n \in \mathfrak{N}$ the edge $\mathfrak{p}$ must be $[v_i^n, v_{i+1}^n]$. Therefore $\bar{v}_i^n$ is the farthest from v_i^n point in $[v_{i-1}^n, v_i^n]$ contained in $\mathcal{N}_{\nu\vartheta}([v_i^n, v_{i+1}^n])$. Let $\tilde{v}_i^n$ be the farthest from v_i^n point $t_n \in [v_i^n, v_{i+1}^n]$ such that $\operatorname{dist}(\bar{v}_i^n, t_n) \leq \nu\vartheta$. It follows that $\operatorname{dist}(\bar{v}_i^n, \tilde{v}_i^n) = \nu\vartheta$. We modify P_n by replacing $[\bar{v}_i^n, v_i^n] \cup [v_i^n, \tilde{v}_i^n]$ with a geodesic $[\bar{v}_i^n, \tilde{v}_i^n]$. We repeat the argument for each of the vertices of P_n, and in the end we obtain a sequence of polygons P_n' with at most $2m$ edges each. As the Hausdorff distance between P_n' and P_n is at most $\epsilon_n + 1 + \nu\vartheta$, $\lim^\omega(P_n') = \lim^\omega(P_n)$.

Let us show that for $n \in \mathfrak{N}$, P_n' is $(\vartheta, 2, \nu)$-fat.

Verification of property (F_1) for $n \in \mathfrak{N}$.

There are two types of edges in P_n', the edges of the form $[\tilde{v}_i^n, \bar{v}_{i+1}^n]$, which we shall call *restricted edges*, and the edges of the form $[\bar{v}_i^n, \tilde{v}_i^n]$, which we shall call *added edges*. We denote by RE_n the union of the restricted edges of P_n' and by AE_n the union of the added edges of P_n'.

Let $[\tilde{v}_i^n, \bar{v}_{i+1}^n]$ be a restricted edge. We first show that for $n \in \mathfrak{N}$,

$$\operatorname{dist}\left([\tilde{v}_i^n, \bar{v}_{i+1}^n] \setminus \mathcal{N}_{2\vartheta}(\{\tilde{v}_i^n, \bar{v}_{i+1}^n\}), RE_n \setminus [\tilde{v}_i^n, \bar{v}_{i+1}^n]\right) \geq \vartheta.$$

Suppose there exists y in $[\tilde{v}_i^n, \bar{v}_{i+1}^n] \setminus \mathcal{N}_{2\vartheta}(\{\tilde{v}_i^n, \bar{v}_{i+1}^n\})$ contained in $\mathcal{N}_\vartheta([\tilde{v}_j^n, \bar{v}_{j+1}^n])$ which is inside $\mathcal{N}_\vartheta([v_j^n, v_{j+1}^n])$, with $j \neq i$. Then $y \in \mathcal{N}_{\epsilon_n+1}(\{v_i^n, v_{i+1}^n\})$. The choice of $\bar{v}_{i+1}^n$ implies that $y \in \mathcal{N}_{\epsilon_n+1}(v_i^n)$. Therefore $\operatorname{dist}(v_i^n, [v_j^n, v_{j+1}^n]) \leq \epsilon_n + 1 + \vartheta$. The previous inequality implies that $j = i - 1$ for $n \in \mathfrak{N}$.

Hence there exists $t \in [\tilde{v}_{i-1}^n, \bar{v}_i^n]$ such that $\operatorname{dist}(t, y) < \vartheta$. By the definition of $\bar{v}_i^n$ we have $t = \bar{v}_i^n$. This contradicts the choice of $\tilde{v}_i^n$.

Now let us show that for $n \in \mathfrak{N}$,

$$\operatorname{dist}\left([\tilde{v}_i^n, \bar{v}_{i+1}^n] \setminus \mathcal{N}_{2\vartheta}(\{\tilde{v}_i^n, \bar{v}_{i+1}^n\}), AE_n\right) \geq \vartheta.$$

Suppose there exists z in $[\tilde{v}_i^n, \bar{v}_{i+1}^n] \setminus \mathcal{N}_{2\vartheta}(\{\tilde{v}_i^n, \bar{v}_{i+1}^n\})$ contained in $\mathcal{N}_\vartheta([\bar{v}_j^n, \tilde{v}_j^n])$. It follows that z belongs to the $(\epsilon_n + \nu\vartheta + 1)$-neighborhood of v_j^n and that $\mathrm{dist}(v_j^n, [v_i^n, v_{i+1}^n]) \leq \epsilon_n + \nu\vartheta + 1$. For $n \in \mathfrak{N}$ this implies that $j \in \{i, i+1\}$. Suppose $j = i$ (the other case is similar). Let $t \in [\bar{v}_i^n, \tilde{v}_i^n]$ with $\mathrm{dist}(t, z) \leq \vartheta$. Then $\mathrm{dist}(\tilde{v}_i^n, t) \geq \mathrm{dist}(\tilde{v}_i^n, z) - \mathrm{dist}(t, z) \geq 2\vartheta - \vartheta \geq \mathrm{dist}(t, z)$. It follows that $\mathrm{dist}(\bar{v}_i^n, z) \leq \mathrm{dist}(\bar{v}_i^n, t) + \mathrm{dist}(t, z) \leq \mathrm{dist}(\bar{v}_i^n, t) + \mathrm{dist}(\tilde{v}_i^n, t)\mathrm{dist}(\bar{v}_i^n, \tilde{v}_i^n) = \nu\vartheta$. This contradicts the choice of $\tilde{v}_i^n$.

Now consider an added edge $[\bar{v}_i^n, \tilde{v}_i^n] \subset B(v_i^n, \epsilon_n + 1 + \nu\vartheta)$. Let $n \in \mathfrak{N}$. If there exists $u \in [\bar{v}_i^n, \tilde{v}_i^n] \setminus \mathcal{N}_{2\vartheta}(\{\bar{v}_i^n, \tilde{v}_i^n\})$ contained in $\mathcal{N}_\vartheta([\bar{v}_j^n, \tilde{v}_j^n])$ with $j \neq i$ then $u \in \mathcal{N}_{\epsilon_n + 1 + (\nu+1)\vartheta}(v_j^n)$. It follows that $\mathrm{dist}(v_i^n, v_j^n) \leq \mathrm{dist}(v_i^n, u) + \mathrm{dist}(u, v_j^n) \leq 2\epsilon_n + 2 + (2\nu + 1)\vartheta$. This contradicts the fact that $n \in \mathfrak{N}$.

If there exists $s \in [\bar{v}_i^n, \tilde{v}_i^n] \setminus \mathcal{N}_{2\vartheta}(\{\bar{v}_i^n, \tilde{v}_i^n\})$ contained in the ϑ-tubular neighborhood of $[\tilde{v}_j^n, \bar{v}_{j+1}^n]$ then $v_i^n \in \mathcal{N}_{(\nu+1)\vartheta + \epsilon_n + 1}([v_j^n, v_{j+1}^n])$, which together with the hypothesis $n \in \mathfrak{N}$ implies that $j \in \{i-1, i\}$. The fact that $\mathrm{dist}(s, \tilde{v}_i^n) \geq 2\vartheta$ together with the choice of $\tilde{v}_i^n$ implies that $\mathrm{dist}(s, [\tilde{v}_i^n, \bar{v}_{i+1}^n]) \geq 2\vartheta$. The fact that $\mathrm{dist}(s, \bar{v}_i^n) \geq 2\vartheta$ together with the choice of $\bar{v}_i^n$ implies that $\mathrm{dist}(s, [\tilde{v}_{i-1}^n, \bar{v}_i^n]) \geq 2\vartheta$. Therefore $j \notin \{i-1, i\}$, a contradiction.

Verification of property (F_2) for $n \in \mathfrak{N}$.

Let $\bar{v} = \bar{v}_i^n$ be a vertex of P_n' and let $v = v_i^n$. We have that $\mathcal{O}_{\bar{v}}(P_n') = (RE_n \setminus [\tilde{v}_{i-1}^n, v]) \cup (AE_n \setminus [\bar{v}, \tilde{v}_i^n])$. The set $RE_n \setminus [\tilde{v}_{i-1}^n, \bar{v}]$ is composed of $[\tilde{v}_i^n, \bar{v}_{i+1}^n]$ and of a part RE_n' contained in $\mathcal{O}_v(P_n)$. By construction we have $\mathrm{dist}(\bar{v}, [\tilde{v}_i^n, \bar{v}_{i+1}^n]) \geq \nu\vartheta$. On the other hand $\mathrm{dist}(\bar{v}, RE_n') \geq \mathrm{dist}(v, RE_n') - \mathrm{dist}(\bar{v}, v) \geq \mathrm{dist}(v, \mathcal{O}_v(P_n)) - \epsilon_n - 1 \geq \varrho d_n - \epsilon_n - 1$, which is larger that $\nu\vartheta$ for $n \in \mathfrak{N}$.

Since $AE_n \setminus [\bar{v}, \tilde{v}_i^n] \subset \mathcal{N}_{\epsilon_n + 1 + \nu\vartheta}(\mathcal{V}_n \setminus \{v\})$ it follows that

$$\mathrm{dist}(\bar{v}, AE_n \setminus [\bar{v}, \tilde{v}_i^n]) \geq \mathrm{dist}(v, \mathcal{V}_n \setminus \{v\}) - \epsilon_n - 1 - (\epsilon_n + 1 + \nu\vartheta) \geq \varrho d_n - (2\epsilon_n + 2 + \nu\vartheta) \geq \nu\vartheta$$

for $n \in \mathfrak{N}$.

Now let $\tilde{v} = \tilde{v}_i^n$ be a vertex of P_n'. We have that $\mathcal{O}_{\tilde{v}}(P_n') = (RE_n \setminus [\tilde{v}, \bar{v}_{i+1}^n]) \cup (AE_n \setminus [\bar{v}, \tilde{v}])$. As before, we show that $\mathrm{dist}(\tilde{v}, AE_n \setminus [\bar{v}, \tilde{v}]) \geq \nu\vartheta$ for $n \in \mathfrak{N}$.

The set $RE_n \setminus [\tilde{v}, \bar{v}_{i+1}^n]$ is composed of $[\tilde{v}_{i-1}^n, \bar{v}]$ and of RE_n'. As above, $\mathrm{dist}(\tilde{v}, RE_n') \geq \nu\vartheta$ for $n \in \mathfrak{N}$. The distance $\mathrm{dist}(\tilde{v}, [\tilde{v}_{i-1}^n, \bar{v}])$ is at least $\nu\vartheta$ by the choice of $\bar{v}$.

We conclude that for $n \in \mathfrak{N}$ the polygon P_n' is $(\vartheta, 2, \nu)$-fat. $\square$

4 A characterization of asymptotically tree-graded spaces

In this section, we find metric conditions for a metric space to be asymptotically tree-graded with respect to a family of subsets.

Theorem 4.1 (a characterization of asymptotically tree-graded spaces). *Let (X, dist) be a geodesic metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X. The metric space X is asymptotically tree-graded with respect to $\mathcal{A}$ if and only if the following properties are satisfied:*

(α_1) *For every $\delta > 0$ the diameters of the intersections $\mathcal{N}_\delta(A_i) \cap \mathcal{N}_\delta(A_j)$ are uniformly bounded for all $i \neq j$.*

(α_2) *For every θ from $\left[0, \frac{1}{2}\right)$ there exists a number $M > 0$ such that for every geodesic $\mathfrak{q}$ of length ℓ and every $A \in \mathcal{A}$ with $\mathfrak{q}(0), \mathfrak{q}(\ell) \in \mathcal{N}_{\theta\ell}(A)$ we have $\mathfrak{q}([0, \ell]) \cap \mathcal{N}_M(A) \neq \emptyset$.*

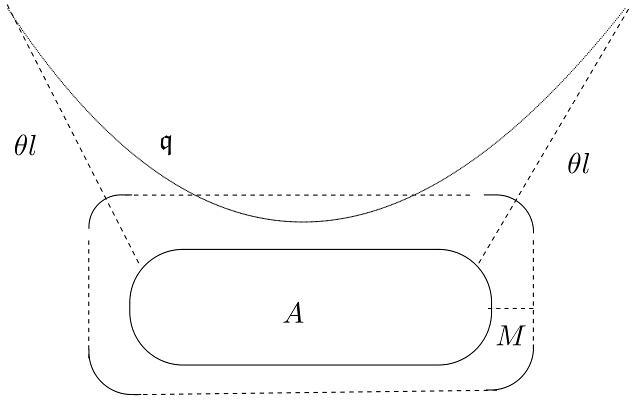

Figure 3: Property (α_2).

(α_3) *For every $k \geq 2$ there exist $\vartheta > 0$, $\nu \geq 8$ and $\chi > 0$ such that every k-gon P in X with geodesic edges which is $(\vartheta, 2, \nu)$-fat satisfies $P \subset \mathcal{N}_\chi(A)$ for some $A \in \mathcal{A}$.*

Remarks 4.2. (1) If the space X is asymptotically uniquely geodesic (for instance asymptotically $CAT(0)$) then in (α_3) it is enough to consider $k = 3$ (only triangles).

(2) From the proof of Theorem 4.1, it will be clear that conditions (α_2), (α_3) can be replaced by the following stronger conditions:

(α_2') For every $L \geq 1$, $C \geq 0$, and $\theta \in \left[0, \frac{1}{2}\right)$ there exists $M > 0$ such that for every (L, C)-quasi-geodesic $\mathfrak{q}$ defined on $[0, \ell]$ and every $A \in \mathcal{A}$ such that $\mathfrak{q}(0), \mathfrak{q}(\ell) \in \mathcal{N}_{\theta\ell/L}(A)$ we have $\mathfrak{q}([0, \ell]) \cap \mathcal{N}_M(A) \neq \emptyset$;

(α_3') For every $L \geq 1$, $C \geq 0$ and $k \geq 2$, and for every $\sigma \geq 1$ and $\nu \geq 4\sigma$, there exist $\vartheta_0 > 0$ such that for every $\vartheta \geq \vartheta_0$ every k-gon P with (L, C)-quasi-geodesic edges which is $(\vartheta\,,\,\sigma\,,\,\nu)$-fat is contained in $\mathcal{N}_\chi(A)$ for some $A \in \mathcal{A}$, where $\chi = \sigma L^2 \vartheta + c$ with c a constant independent of ϑ.

(3) Also from the proof of Theorem 4.1, it will be clear that for every $\epsilon \leq \frac{1}{2}$ the condition (α_2) can be replaced by the following weaker condition :

(α_2^ϵ) For every θ from $[0, \epsilon)$ there exists a number $M > 0$ such that for every geodesic $\mathfrak{q}$ of length ℓ and every $A \in \mathcal{A}$ with $\mathfrak{q}(0), \mathfrak{q}(\ell) \in \mathcal{N}_{\theta\ell}(A)$ we have $\mathfrak{q}([0, \ell]) \cap \mathcal{N}_M(A) \neq \emptyset$.

(Notice that condition (α_2) is the same as the condition $(\alpha_2^{\frac{1}{2}})$.)

(4) If $\mathcal{A} = \{A_i \mid i \in I\}$ satisfies conditions (α_1), (α_2), (α_3), then the family $\mathcal{N}_c(\mathcal{A}) = \{\mathcal{N}_c(A_i) \mid i \in I\}$ also satisfies these conditions, for every $c > 0$.

Proof of Theorem 4.1. First we show that conditions (α_1), (α_2^ϵ) (for an arbitrary $\epsilon \leq \frac{1}{2}$) and (α_3) imply that X is asymptotically tree-graded with respect to $\mathcal{A}$.

Lemma 4.3 ((α_1) and (α_2^ϵ) imply uniform quasi-convexity). *Let (X, d) be a geodesic metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X satisfying properties (α_1) and (α_2^ϵ) for some ϵ. Let $M_0 = M_0(\theta)$ be the number from property (α_2^ϵ) corresponding to $\theta = \frac{2}{3}\epsilon$.*

There exists $t > 0$ such that for every $A \in \mathcal{A}$, $M \geq M_0$ and $x, y \in \mathcal{N}_M(A)$, every geodesic joining x and y in X is contained in $\mathcal{N}_{tM}(A)$.

Proof. Suppose, by contradiction, that for every $n \in \mathbb{N}$ there exist $M_n \geq M_0$, $x_n, y_n \in \mathcal{N}_{M_n}(A_n)$ and a geodesic $[x_n, y_n]$ not contained in $\mathcal{N}_{nM_n}(A_n)$. For every $n \geq 1$ let D_n be the infimum of the distances between points $x, y \in \mathcal{N}_{M_n}(A)$ for some $A \in \mathcal{A}$ such that $[x, y] \not\subset \mathcal{N}_{nM_n}(A)$ for some geodesic $[x, y]$.

We note that $D_n \geq 2(n-1)M_n \geq 2(n-1)M_0$, hence $\lim_{n\to\infty} D_n = \infty$. For every $n \geq 1$, choose $x_n, y_n \in \mathcal{N}_{M_n}(A_n)$ such that $\operatorname{dist}(x_n, y_n) = D_n + 1$. Also choose $a_n, b_n \in [x_n, y_n]$ such that $\operatorname{dist}(x_n, a_n) = \operatorname{dist}(y_n, b_n) = \frac{\theta(D_n+1)}{2}$. Then $\operatorname{dist}(a_n, A_n) \leq \operatorname{dist}(a_n, x_n) + \operatorname{dist}(x_n, A_n) \leq \frac{\theta(D_n+1)}{2} + M_n \leq \frac{\theta(D_n+1)}{2} + \frac{D_n+1}{2(n-1)}$. Likewise $\operatorname{dist}(b_n, A_n) \leq \frac{\theta(D_n+1)}{2} + \frac{D_n+1}{2(n-1)}$. On the other hand $\operatorname{dist}(a_n, b_n) \geq \operatorname{dist}(x_n, y_n) - \operatorname{dist}(x_n, a_n) - \operatorname{dist}(y_n, b_n) \geq (1-\theta)(D_n+1)$. For n large enough we have $\frac{\theta}{2} + \frac{1}{2(n-1)} \leq \frac{2}{3}\theta$. We apply (α_2^ϵ) with $\theta = \frac{2}{3}\epsilon$ to $[a_n, b_n]$ and we deduce that there exists $z_n \in [a_n, b_n] \cap \mathcal{N}_{M_0}(A_n)$. We have that either $[x_n, z_n] \not\subset \mathcal{N}_{nM_n}(A_n)$ or $[z_n, y_n] \not\subset \mathcal{N}_{nM_n}(A_n)$, while $\operatorname{dist}(x_n, z_n), \operatorname{dist}(z_n, y_n) \leq (1 - \frac{\theta}{2})(D_n + 1) < D_n$ for n large enough. This contradicts the choice of D_n. □

Lemma 4.4. *Let (X, d) be a geodesic metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X satisfying properties (α_1) and (α_2^ϵ) for some ϵ. Then in every asymptotic cone $\operatorname{Con}^\omega(X; e, d)$, every set $\lim^\omega(A_n)$ is connected and a geodesic subspace.*

Proof. Indeed, consider any two points $x = \lim^\omega(x_n), y = \lim^\omega(y_n)$ in $\lim^\omega(A_n)$, and geodesics $\mathfrak{q}_n$ connecting x_n, y_n in X. Then by Lemma 4.3, $\mathfrak{q}_n$ is inside $\mathcal{N}_M(A_n)$ for some fixed M. Therefore the geodesic $\lim^\omega(\mathfrak{q}_n)$ is inside $\lim^\omega(\mathcal{N}_M(A_n)) = \lim^\omega(A_n)$. □

Lemma 4.5. *Let (X, d) be a geodesic metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X satisfying properties (α_1) and (α_2^ϵ). Then in every asymptotic cone $Con^\omega(X; e, d)$ the collection of subsets $\mathcal{A}_\omega$ satisfies (T_1).*

Proof. Suppose that, in an asymptotic cone $\operatorname{Con}^\omega(X; e, d)$ of X, the intersection $\lim^\omega(A_{i_n}) \cap \lim^\omega(A_{j_n})$ contains two distinct points $\lim^\omega(x_n), \lim^\omega(y_n)$ but $A_{i_n} \neq A_{j_n}$ ω-almost surely. For every $n \geq 1$ consider a geodesic $[x_n, y_n]$. Its length ℓ_n is $O(d_n)$ while δ_n defined as the maximum of the distances $\operatorname{dist}(x_n, A_{i_n}), \operatorname{dist}(x_n, A_{j_n}), \operatorname{dist}(y_n, A_{i_n}), \operatorname{dist}(y_n, A_{j_n})$, is $o(d_n)$. According to Lemma 4.3, $[x_n, y_n]$ is contained in $\mathcal{N}_{t\delta_n}(A_{i_n}) \cap \mathcal{N}_{t\delta_n}(A_{j_n})$ for some $t > 0$.

Consider $a_n, b_n \in [x_n, y_n]$ at distance $6t\delta_n$ from x_n and y_n, respectively. Property (α_2^ϵ) can be applied twice, to $[x_n, a_n] \subset [x_n, y_n]$ and A_{i_n} (resp. A_{j_n}) for n large enough. It implies that there exist $z_n \in [x_n, a_n] \cap \mathcal{N}_{M_0}(A_{i_n})$ and $z'_n \in [x_n, a_n] \cap \mathcal{N}_{M_0}(A_{j_n})$ (where M_0 is the same as in Lemma 4.3). A similar argument for $[b_n, y_n] \subset [x_n, y_n]$ and A_{i_n} (resp. A_{j_n}) implies that there exist $u_n \in [b_n, y_n] \cap \mathcal{N}_{M_0}(A_{i_n})$ and $u'_n \in [b_n, y_n] \cap \mathcal{N}_{M_0}(A_{j_n})$. Hence $[a_n, b_n] \subset [z_n, u_n] \subset \mathcal{N}_{tM_0}(A_{i_n})$ and $[a_n, b_n] \subset [z'_n, u'_n] \subset \mathcal{N}_{tM_0}(A_{j_n})$. It follows that $[a_n, b_n] \subset \mathcal{N}_{tM_0}(A_{i_n}) \cap \mathcal{N}_{tM_0}(A_{j_n})$, while $\operatorname{dist}(a_n, b_n) = O(d_n)$. This contradicts property (α_1). □

Lemma 4.6. (asymptotic (T_1) and (α_3) implies asymptotic (T_2)) *Let (X, dist) be a geodesic metric space and let $\mathcal{A} = \{A_i \mid i \in I\}$ be a collection of subsets of X. Suppose that property (α_3) holds. Then every simple geodesic triangle in any asymptotic cone $\operatorname{Con}^\omega(X; e, d)$ is contained in one of the sets from $\mathcal{A}_\omega$.*

Proof. Let Δ be a simple geodesic triangle in $\operatorname{Con}^\omega(X; e, d)$. Let $\varepsilon_m = \frac{1}{2^m}$ be fixed, for every large enough integer m. By Proposition 3.34, we can find k_0 and a simple triangle $\Delta_m = \Delta_{\varepsilon_m} = \lim^\omega(P_n^m)$ satisfying properties (a),(b), and (c) for ϑ and $\nu \geq 8$ given by (α_3) for $k_0(\varepsilon_m)$. It follows that ω-almost surely, P_n^m are contained in $\mathcal{N}_\chi(A_n)$ for some $A_n \in \mathcal{A}$. We conclude that $\Delta_m \subset A_\omega = \lim^\omega(A_n)$. By property (b) all triangles Δ_m have at least 3 distinct points in common (e.g. the midpoints of the edges of Δ). This and property (T_1) of the collection $\mathcal{A}_\omega$ imply that the set A_ω is independent of m. Since Δ is a Hausdorff limit of Δ_m and A_ω is closed (see Remark 3.10), we deduce that $\Delta \subset A_\omega$. $\square$

Lemmas 4.4, 4.5 and 4.6 show that (α_1), (α_2^ϵ), (α_3) imply that the space X is asymptotically tree-graded. Now we prove the (stronger version of the) converse statement.

Lemma 4.7 (asymptotic (T_1) implies (α_1)). *Let (X, dist) be a geodesic metric space asymptotically satisfying (T_1) with respect to $\mathcal{A}$. Then X satisfies (α_1) with respect to $\mathcal{A}$.*

Proof. By contradiction, suppose X asymptotically satisfies (T_1) but for some $\delta > 0$ there exists a sequence of pairs of points x_n, y_n in $\mathcal{N}_\delta(A_{i_n}) \cap \mathcal{N}_\delta(A_{j_n})$, where A_{i_n} and A_{j_n} are distinct sets in $\mathcal{A}$, with $\lim_{n\to\infty}\mathrm{dist}(x_n, y_n) = \infty$. Set the observation point e to be $(x_n)^\omega$, and let $d_n = \mathrm{dist}(x_n, y_n)$ for every $n \geq 1$. Then $M_1 = \lim^\omega(A_{i_n})$ and $M_2 = \lim^\omega(A_{j_n})$ are not empty, so these are distinct pieces in $\mathrm{Con}^\omega(X; e, d)$. The limits $x = \lim^\omega(x_n)$ and $y = \lim^\omega(y_n)$ are distinct points in $\mathrm{Con}^\omega(X; e, d)$ that belong to both M_1 and M_2. This contradicts (T_1). $\square$

Definition 4.8 (almost closest points). Let $x \in X$, $A, B \subseteq X$. A point $y \in A$ is called an *almost closest to x point in A* if $\mathrm{dist}(x, y) \leq \mathrm{dist}(x, A) + 1$. Points $a \in A$, $b \in B$ are called *almost closest representatives of A and B* if $\mathrm{dist}(a, b) \leq \mathrm{dist}(A, B) + 1$.

Definition 4.9 (almost projection). Let x be a point in X and $A \subset X$. The *almost projection of x on A* is the set of almost closest to x points in A. For every subset B of X we define the *almost projection* $\mathrm{proj}_A(B)$ of B onto A as $\bigcup_{b\in B}\mathrm{proj}_A(b)$.

Remark 4.10. If all $A \in \mathcal{A}$ were closed sets and the space X was proper (i.e. all balls in X compact) then we could use closest points and usual projections instead of almost closest points and almost projections.

Lemma 4.11. *If the space X is asymptotically tree-graded with respect to $\mathcal{A}$ then for every $x \in X$, $A \in \mathcal{A}$, with $\mathrm{dist}(x, A) = 2d$*

$$\mathrm{diam}(proj_A(\mathcal{N}_d(x)) = o(d).$$

Proof. Suppose there exists $\varepsilon > 0$ and $x_n \in X, A_n \in \mathcal{A}$ with $\mathrm{dist}(x_n, A_n) = 2d_n, \lim_{n\to\infty} d_n = \infty$, and the projection $\mathrm{proj}_{A_n}(\mathcal{N}_{d_n}(x_n))$ is of diameter at least εd_n. Let $e = (x_n)$ and $d = (d_n)$. In the asymptotic cone $\mathrm{Con}^\omega(X; e, d)$, we have the point $x = \lim^\omega(x_n)$ at distance 2 of $A = \lim^\omega(A_n)$, two points $y, z \in \mathcal{N}_1(x)$, and two points y', z' in A such that y', z' are the respective projections of y, z onto A, but $\mathrm{dist}(y', z') \geq \varepsilon$. This contradicts Lemma 2.6. $\square$

Lemma 4.12 (asymptotically tree-graded implies (α_2')). *Let (X, dist) be a geodesic metric space which is asymptotically tree-graded with respect to $\mathcal{A}$. Then X satisfies (α_2').*

Proof. Fix $L \geq 1, C \geq 0$. By contradiction, suppose that for some fixed $\theta \in \left[0, \frac{1}{2}\right)$ there exists a sequence of (L, C)-quasi-geodesics $\mathfrak{q}_n \colon [0, \ell_n] \to X$ and a sequence of sets $A_n \in \mathcal{A}$, such that $\mathfrak{q}_n(0), \mathfrak{q}_n(\ell_n) \in \mathcal{N}_{\theta\ell_n/L}(A_n)$ and $\mathrm{dist}\left(\mathfrak{q}_n([0, \ell_n]), A_n\right) = 2D_n$, $\lim_{n\to\infty} D_n = \infty$. Since $\mathrm{dist}\left(\mathfrak{q}_n([0, \ell_n]), A_n\right) \leq L\ell_n + \frac{\theta\ell_n}{L}$ this implies $\lim_{n\to\infty}\ell_n = \infty$.

Let $t_0 = 0 < t_1 < \cdots < t_{m-1} < t_m = \ell_n$ be such that $\frac{D_n - C}{2L} \leq \mathrm{dist}(t_i, t_{i+1}) \leq \frac{D_n - C}{L}$ for all $i \in \{0, 1, \ldots, m-1\}$. We have $m \leq \frac{3L\ell_n}{D_n}$ for large enough n. Let y_i be an almost projection of $\mathfrak{q}_n(t_i)$ onto A_n. According to Lemma 4.11, $\mathrm{dist}(y_i, y_{i+1}) = o(D_n)$. Consequently $\mathrm{dist}(\mathfrak{q}_n(0), \mathfrak{q}_n(\ell_n)) \leq \mathrm{dist}(\mathfrak{q}_n(0), y_0) + \Sigma_{i=0}^{m-1}\mathrm{dist}(y_i, y_{i+1}) + \mathrm{dist}(y_m, \mathfrak{q}_n(\ell_n)) \leq \frac{2\theta\ell_n}{L} + m \cdot o(D_n) \leq \frac{2\theta\ell_n}{L} + 3Lo(1)\ell_n$. On the other hand $\mathrm{dist}(\mathfrak{q}_n(0), \mathfrak{q}_n(\ell_n)) \geq \frac{\ell_n}{L} - C$. This is a contradiction with $\theta < \frac{1}{2}$. $\square$

It remains to prove that being asymptotically tree-graded implies (α_3).

Definition 4.13 (almost geodesics). If an (L, C)-quasi-geodesic $\mathfrak{q}$ is L-Lipschitz then $\mathfrak{q}$ will be called an (L, C)-*almost geodesic*.

Remark 4.14. Every (L, C)-quasi-geodesic in a geodesic metric space is at bounded (in terms of L, C) distance from an $(L+C, C)$-almost geodesic with the same end points [Bo, Proposition 8.3.4].

Lemma 4.15 ($\mathcal{A}$ is uniformly quasi-convex with respect to quasi-geodesics). *Let X be a geodesic metric space which is asymptotically tree-graded with respect to a collection of subsets $\mathcal{A}$. For every $L \geq 1$ and $C \geq 0$, there exists $t \geq 1$ such that for every $d \geq 1$ and for every $A \in \mathcal{A}$, every (L, C)-quasi-geodesic joining two points in $\mathcal{N}_d(A)$ is contained in $\mathcal{N}_{td}(A)$.*

Proof. Suppose by contradiction that there exists a sequence $\mathfrak{q}_n : [0, \ell_n] \to X$ of (L, C)-quasi-geodesics with endpoints $x_n, y_n \in \mathcal{N}_{d_n}(A_n)$ such that there exists $z_n \in \mathfrak{q}_n([0, \ell_n])$ with $k_n = \mathrm{dist}(z_n, A_n) \geq nd_n \geq n$. By Remark 4.14, we can assume that each $\mathfrak{q}_n$ is an $(L + C, C)$-almost geodesic. This allows us to choose $z_n \in \mathfrak{q}_n([0, \ell_n])$ so that $\mathrm{dist}(z_n, A_n)$ is maximal. In $\mathrm{Con}^\omega(X; (z_n), (k_n))$, the limit set $\mathfrak{q} = \lim^\omega(\mathfrak{q}_n)$ is either a topological arc with endpoints in $\lim^\omega(A_n)$ and not contained in $\lim^\omega(A_n)$, or a bi-Lipschitz ray with origin in $\lim^\omega(A_n)$ or a bi-Lipschitz line (Remark 3.15). Notice also that $\mathfrak{q}$ is contained in $\mathcal{N}_1(\lim^\omega(A_n))$. In all three cases we obtain a contradiction with Corollary 2.7. □

Let (X, dist) be a geodesic space that is asymptotically tree-graded with respect to the collection of subsets $\mathcal{A}$.

Notation: For every $L \geq 1, C \geq 0$, we denote by $M(L, C)$ the constant given by (α_2') for $\theta = \frac{1}{3}$. We also denote by dist the distance function in any of the asymptotic cones of X.

Conventions: To simplify the notations and statements, in the sequel we shall not mention the constants $L \geq 1$ and $C \geq 0$ for each quasi-geodesic anymore. We assume that all constants provided by the following lemmas in the section depend on L and C.

Lemma 4.16. *Let $\mathfrak{q}_n \colon [0, \ell_n] \to X$, $n \geq 1$, be a sequence of (L, C)-quasi-geodesics in X and let A_n, $n \geq 1$, be a sequence of sets in $\mathcal{A}$. Suppose that $\mathrm{dist}(\mathfrak{q}_n(0), A_n) = o(\ell_n), \mathrm{dist}(\mathfrak{q}_n(\ell_n), A_n) = o(\ell_n)$ ω-almost surely. Then there exists $t_n^1 \in [0, \frac{1}{3}\ell_n]$, $t_n^2 \in [\frac{2}{3}\ell_n, \ell_n]$ such that $\mathfrak{q}_n(t_n^i) \in \mathcal{N}_M(A_n)$, $i = 1, 2$, where $M = M(L, C)$, ω-almost surely.*

Proof. By Lemma 4.15, the quasi-geodesic $\mathfrak{q}_n$ is inside $\mathcal{N}_{t_n}(A_n)$ for $t_n = o(\ell_n)$. It remains to apply (α_2') to the quasi-geodesics $\mathfrak{q}_n([0, \frac{1}{3}\ell_n])$ and $\mathfrak{q}_n([\frac{2}{3}\ell_n, \ell_n])$. □

Lemma 4.17 (linear divergence). *For every $\varepsilon > 0$ and every $M \geq M(L, C)$ there exists $t_\varepsilon > 0$ such that if $A \in \mathcal{A}$, $\mathfrak{q}$ is a quasi-geodesic with origin $a \in \mathcal{N}_M(A)$, such that $\mathfrak{q} \cap \mathcal{N}_M(A) = \{a\}$ and $t \geq t_\varepsilon$ then*

$$\mathrm{dist}(\mathfrak{q}(t), A) > (1 - \varepsilon)\mathrm{dist}(\mathfrak{q}(t), a) .$$

Proof. We suppose that for some $\varepsilon > 0$ there exists a sequence $A_n \in \mathcal{A}$, a sequence $\mathfrak{q}_n$ of quasi-geodesics with origin $a_n \in \mathcal{N}_M(A_n)$ such that $\mathfrak{q}_n \cap \mathcal{N}_M(A_n) = \{a_n\}$, and a sequence of numbers $t_n \to \infty$ with the property

$$\mathrm{dist}(\mathfrak{q}_n(t_n), A_n) \leq (1 - \varepsilon)\mathrm{dist}(\mathfrak{q}_n(t_n), a_n) .$$

In $\mathrm{Con}^\omega(X;(a_n),(t_n))$, we obtain the points $a=\lim^\omega(a_n)\in\lim^\omega(A_n)$ and $b=\lim^\omega(\mathfrak{q}_n(t_n))$, joined by the bi-Lipschitz arc $\mathfrak{q}([0,1])=\lim^\omega(\mathfrak{q}_n([0,t_n]))$, such that

$$\mathrm{dist}(b,\lim^\omega(A_n))\le(1-\varepsilon)\mathrm{dist}(b,a).$$

It follows that the projection of b on $\lim^\omega(A_n)$ is a point $c\ne a$. Corollary 2.9 implies that $\mathfrak{q}([0,1])$ contains c and Corollary 2.8 implies that a sub-arc $\mathfrak{q}([0,2\beta])$ of $\mathfrak{q}([0,1])$ is contained in $\lim^\omega(A_n)$. We apply Lemma 4.16 to the sub-quasi-geodesic $\mathfrak{q}_n([0,\beta t_n])$ and obtain that this sub-quasi-geodesics intersects $\mathcal{N}_M(A_n)$ in a point different from a_n, a contradiction. □

Lemma 4.18. *For every $\varepsilon>0$, $\delta>0$ and $M\ge M(L,C)$ there exists $D>0$ such that for every $A\in\mathcal{A}$ and every two quasi-geodesics $\mathfrak{q}_i\colon[0,\ell_i]\to X$, $i=1,2$, that connect $a\in\mathcal{N}_M(A)$ with two points b_1 and b_2 respectively, if the diameter of $\mathfrak{q}_1\cap\mathcal{N}_M(A)$ does not exceed δ, $b_2\in\mathcal{N}_M(A)$, and* $\mathrm{dist}(a,b_2)\ge D$ *then*

$$\mathrm{dist}(b_1,b_2)\ge\frac{1}{L+\varepsilon}(\ell_1+\ell_2).$$

Proof. Suppose there exist sequences $\mathfrak{q}_i^{(n)}:[0,\ell_i^{(n)}]\to X$, $i=1,2$, $n\ge1$, of pairs of quasi-geodesics joining $a^{(n)}\in\mathcal{N}_M(A_n)$ to $b_i^{(n)}$ such that $\mathfrak{q}_1^{(n)}\cap\mathcal{N}_M(A_n)$ has diameter at most δ, $b_2^{(n)}\in\mathcal{N}_M(A_n)$, $\lim_{n\to\infty}\mathrm{dist}(a^{(n)},b_2^{(n)})=\infty$, but

$$\mathrm{dist}(b_1^{(n)},b_2^{(n)})\le\frac{1}{L+\varepsilon}(\ell_1^{(n)}+\ell_2^{(n)}).\tag{5}$$

Denote $\mathrm{dist}(a^{(n)},b_1^{(n)})$ by f_n and $\mathrm{dist}(a^{(n)},b_2^{(n)})$ by d_n. Since $\ell_1^{(n)}\le L(f_n+C)$, $\ell_2^{(n)}\le L(d_n+C)$, for every large enough n the inequality (5) implies that

$$\mathrm{dist}(b_1^{(n)},b_2^{(n)})\le(1-\gamma)(f_n+d_n).\tag{6}$$

for some $\gamma>0$.

Case I. Suppose that $\lim^\omega(\frac{f_n}{d_n})<\infty$. In the asymptotic cone $\mathrm{Con}^\omega(X;(a_n),(d_n))$, the two points $\lim^\omega(b_i^{(n)})$, $i=1,2$, are joined by the Lipschitz arc $\lim^\omega(\mathfrak{q}_1^{(n)})\cup\lim^\omega(\mathfrak{q}_2^{(n)})$ (it is Lipschitz as any union of two Lipschitz arcs). Lemma 4.17 implies that

$$\lim^\omega(\mathfrak{q}_1^{(n)})\cap\lim^\omega(\mathfrak{q}_2^{(n)})=\lim^\omega(a^{(n)})$$

(here we use the fact that the diameters of the intersections $\mathfrak{q}_1^{(n)}$ with $\mathcal{N}_M(A_n)$ are uniformly bounded, so we can cut a comparatively little piece of each $\mathfrak{q}_1^{(n)}$ to make it satisfy the conditions of Lemma 4.17).

Thus the points $\lim^\omega(b_i^{(n)})$ are joined by the simple arc $\lim^\omega(\mathfrak{q}_1^{(n)})\cup\lim^\omega(\mathfrak{q}_2^{(n)})$. This and property (T_2') imply that every geodesic joining $\lim^\omega(b_1^{(n)})$ and $\lim^\omega(b_2^{(n)})$ contains $\lim^\omega(a^{(n)})$. Therefore

$$\mathrm{dist}(\lim^\omega(b_1^{(n)}),\lim^\omega(b_2^{(n)}))=\mathrm{dist}(\lim^\omega(b_1^{(n)}),\lim^\omega(a^{(n)}))+\mathrm{dist}(\lim^\omega(a^{(n)}),\lim^\omega(b_2^{(n)})).$$

This contradicts the inequality (6).

Case II. Suppose that $\lim_\omega\frac{f_n}{d_n}=\infty$. In the asymptotic cone $\mathrm{Con}^\omega(X;(a^{(n)}),(f_n))$, we denote $a=\lim^\omega(a^{(n)})=\lim^\omega(b_2^{(n)})\in\lim^\omega(A_n)$ and $b=\lim^\omega(b_1^{(n)})$. Then inequality (6) implies that $\mathrm{dist}(a,b)\le(1-\gamma)\mathrm{dist}(a,b)$, a contradiction. □

Lemma 4.19. *For every $M \geq M(L,C)$, $\varepsilon > 0$ and $\delta > 0$ there exists $D' > 0$ such that for every $A \in \mathcal{A}$, and every two quasi-geodesics $\mathfrak{q}_i \colon [0,\ell_i] \to X$, $i = 1,2$, joining a in $\mathcal{N}_M(A)$ with b_i, if the diameter of $\mathfrak{q}_1 \cap \mathcal{N}_M(A)$ does not exceed δ, $b_2 \in \mathcal{N}_M(A)$, $\operatorname{dist}(a,b_2) \geq D'$, then the union $\mathfrak{q}_1 \sqcup \mathfrak{q}_2$ of these two quasi-geodesics is an $(L+\varepsilon, K)$-quasi-geodesic, where $K = 2D'$.*

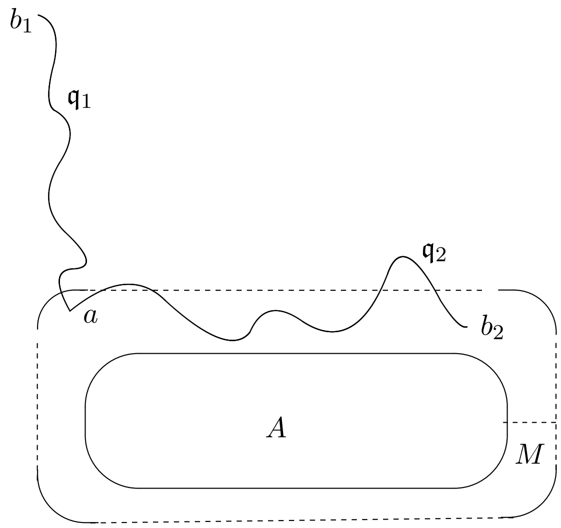

Figure 4: Lemma 4.19

Proof. Let $\mathfrak{q} = \mathfrak{q}_1 \sqcup \mathfrak{q}_2 \colon [0, \ell_1 + \ell_2] \to X$. For every $[t_1,t_2] \subset [0, \ell_1+\ell_2]$ we have

$$\operatorname{dist}(\mathfrak{q}(t_1), \mathfrak{q}(t_2)) \leq L(t_2 - t_1) + 2C$$

by the triangular inequality. This implies $\operatorname{dist}(\mathfrak{q}(t_1), \mathfrak{q}(t_2)) \leq (L+\varepsilon)(t_2-t_1) + K$, for $K \geq 2C$. We need to prove that for some well chosen K we have

$$\frac{1}{L+\varepsilon}(t_2 - t_1) - K \leq \operatorname{dist}(\mathfrak{q}(t_1), \mathfrak{q}(t_2)) . \tag{7}$$

We consider the constant D given by Lemma 4.18 and set $D' = 2L^2(D+C)+C$ and $K = 2D'$. The hypothesis $\operatorname{dist}(a,b_2) \geq D'$ implies that $\ell_2 \geq 2L(D+C)$.

Let $[t_1,t_2] \subset [0,\ell_1+\ell_2]$. If $t_2 - t_1$ is smaller than $2L(D+C)$ then (7) obviously holds. Suppose that $t_2 - t_1 \geq 2L(D+C)$. If $[t_1,t_2] \cap [\ell_1, \ell_1+\ell_2]$ is an interval of length at least $L(D+C)$ then the distance between $\mathfrak{q}(\ell_1)$ and $\mathfrak{q}(t_2)$ is bigger than D. Lemma 4.18 implies (7).

The same inequality is true if (t_1,t_2) does not contain ℓ_1. Suppose that $[t_1,t_2] \cap [\ell_1, \ell_1+\ell_2]$ is a nontrivial interval of length at most $L(D+C)$. Then

$$\operatorname{dist}(\mathfrak{q}(t_1), \mathfrak{q}(t_2)) \geq \operatorname{dist}(\mathfrak{q}(t_1), \mathfrak{q}(\ell_1)) - \operatorname{dist}(\mathfrak{q}(t_2), \mathfrak{q}(\ell_1)) \geq \frac{1}{L}(\ell_1 - t_1) - D' \geq \frac{1}{L}(t_2 - t_1) - 2D'$$

and (7) holds. $\square$

Definition 4.20 (saturations). For every (L,C)-quasi-geodesic $\mathfrak{q}$ in X we define the *saturation* $\operatorname{Sat}(\mathfrak{q})$ as the union of $\mathfrak{q}$ and all $A \in \mathcal{A}$ with $\mathcal{N}_M(A) \cap \mathfrak{q} \neq \emptyset$.

Lemma 4.21. *Let $\mathfrak{q}_n$ be a sequence of (L, C)-quasi-geodesics in X. In every asymptotic cone $\mathrm{Con}^\omega(X; e, d)$ if the limit $\lim^\omega(\mathrm{Sat}\,(\mathfrak{q}_n))$ is not empty then it is either a piece $\lim^\omega(A_n)$ from $\mathcal{A}_\omega$, or the union of $\mathfrak{p} = \lim^\omega(\mathfrak{q}_n)$ and a collection of pieces from $\mathcal{A}_\omega$ such that each piece intersects $\lim^\omega(\mathfrak{q}_n)$ in at least one point and all pieces from $\mathcal{A}_\omega$ that intersect $\lim^\omega(\mathfrak{q}_n)$ in a non-trivial sub-arc are in the collection (recall that by Corollary 2.8 if a piece in a tree-graded space intersects an arc in more than two points then it intersects the arc by a sub-arc).*

Proof. **Case I.** Suppose that $\lim_\omega \frac{\mathrm{dist}(e_n, \mathfrak{q}_n)}{d_n} < \infty$. Let $u_n \in \mathfrak{q}_n$ be an almost closest point to e_n in $\mathfrak{q}_n$.

Suppose that a piece $A = \lim^\omega(A_n)$ intersects $\mathfrak{q} = \lim^\omega(\mathfrak{q}_n)$ in an arc $\mathfrak{q}([t_1, t_2])$, $t_1 < t_2$. This arc is a limit of sub-quasi-geodesics $\mathfrak{q}'_n$ of $\mathfrak{q}_n$ defined on intervals of length $(t_2 - t_1)d_n$. The ends of $\mathfrak{q}'_n$ are at distance $o(d_n)$ from A_n ω-almost surely. Lemma 4.16 implies that ω-almost surely $A_n \subseteq \mathrm{Sat}(\mathfrak{q}_n)$ since $\mathrm{diam}(\mathcal{N}_M(A_n) \cap \mathfrak{q}_n) = O(d_n)$.

Suppose A is such that $A_n \subseteq \mathrm{Sat}(\mathfrak{q}_n)$ and $\lim_\omega \frac{\mathrm{dist}(e_n, A_n)}{d_n} < \infty$. Let a_n be an almost nearest point to u_n in $\mathfrak{q}_n \cap \mathcal{N}_M(A_n)$. Lemma 4.15 implies that the sub-arc $\mathfrak{q}'_n$ of $\mathfrak{q}_n$ with endpoints u_n and a_n is contained ω-almost surely in $\mathcal{N}_{t_n}(A_n)$ for some number $t_n = O(d_n)$. If $\lim_\omega \frac{\mathrm{dist}(u_n, a_n)}{d_n} = \infty$ then by applying Lemma 4.16 we obtain (ω-almost surely) a point in $\mathfrak{q}_n \cap \mathcal{N}_M(A_n)$ nearer to u_n than a_n by a distance $O(d_n)$, a contradiction. Hence $\lim_\omega \mathrm{dist}(u_n, a_n)/d_n < \infty$. Then $a = \lim^\omega(a_n)$ exists and is an intersection point of A with $\mathfrak{q}$.

Case II. Suppose that $\lim_\omega \frac{\mathrm{dist}(e_n, \mathfrak{q}_n)}{d_n} = \infty$. Let $A_n \subset \mathrm{Sat}(\mathfrak{q}_n)$ be such that $\lim_\omega \frac{\mathrm{dist}(e_n, A_n)}{d_n} < \infty$. We have $A = \lim^\omega(A_n) \subseteq \lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$. Suppose there exists $B = \lim^\omega(B_n) \subset \lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ with $B \neq A$ whence $B_n \neq A_n$ ω-almost surely.

For every $n \geq 1$, let y_n be an almost closest to e_n point in A_n. Also pick $b_n = \mathfrak{q}_n(t_n) \in \mathcal{N}_M(B_n)$. If $\mathrm{dist}(t_n, \mathfrak{q}_n^{-1}(\mathcal{N}_M(A_n))) = 0$ then we set $s_n = t_n$. Otherwise let s_n be the almost closest to t_n number in $\mathfrak{q}_n^{-1}(\mathcal{N}_M(A_n))$. We assume that $s_n \leq t_n$ otherwise we can reverse the orientation of $\mathfrak{q}_n$. Then the diameter of the intersection of $\mathfrak{q}_n([s_n, t_n])$ with $\mathcal{N}_M(A_n)$ is bounded in terms of L, C. By Lemma 4.19, $\mathfrak{r}_n = [y_n, \mathfrak{q}_n(s_n)] \cup \mathfrak{q}_n([s_n, t_n])$ is an $(L + \varepsilon, K)$-quasi-geodesic where $[y_n, \mathfrak{q}_n(s_n)]$ is any geodesic connecting y_n and $\mathfrak{q}_n(s_n)$ in X.

Notice that $\mathrm{dist}(y_n, B_n) \leq O(d_n)$, $\mathfrak{q}_n(t_n) \in B_n$. Then by Lemma 4.15, $\mathfrak{r}_n \subseteq \mathcal{N}_{O(d_n)}(B_n)$ ω-almost surely. Applying Lemma 4.16 we find y'_n, a'_n in $[y_n, \mathfrak{q}_n(s_n)]$ with $\mathrm{dist}(y'_n, a'_n) = O(d_n)$ which belong to both $\mathcal{N}_M(A_n)$ and $\mathcal{N}_M(B_n)$. This contradicts property (α_1).

Thus we can conclude that there is no sequence $B_n \subset \mathrm{Sat}(\mathfrak{q}_n)$ with $B_n \neq A_n$ ω-almost surely, such that $\lim_\omega \frac{\mathrm{dist}(e_n, B_n)}{d_n} < \infty$. Hence in this case $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_n)) = A$. $\square$

Lemma 4.22. *For every $d > 0$, every (L, C)-quasi-geodesic $\mathfrak{q}$ and every $A \in \mathcal{A}$, $\mathcal{N}_M(A) \cap \mathfrak{q} = \emptyset$, the diameter of $\mathcal{N}_d(A) \cap \mathcal{N}_d(\mathrm{Sat}(\mathfrak{q}))$ is bounded in terms of d, L, C.*

Proof. Suppose that for some $d > 0$ and some (L, C) there exist sequences of (L, C)-quasi-geodesics $\mathfrak{q}_n$, of sets $A_n \in \mathcal{A}$, $A_n \not\subset \mathrm{Sat}(\mathfrak{q}_n)$, and of points $x_n, y_n \in \mathcal{N}_d(A_n) \cap \mathcal{N}_d(\mathrm{Sat}(\mathfrak{q}_n))$ such that the sequence $\mathrm{dist}(x_n, y_n) = p_n$ is unbounded. Consider the corresponding asymptotic cone $\mathrm{Con}^\omega(X; (x_n), (p_n))$. The limit sets $\lim^\omega(A_n)$ and $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ contain points $x = \lim^\omega(x_n)$ and $y = \lim^\omega(y_n)$ in common, $\mathrm{dist}(x, y) = 1$. By Lemma 4.21, either $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ is $\lim^\omega(A'_n)$ with $A'_n \in \mathcal{A}$, $A'_n \neq A_n$ ω-almost surely, or $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ is equal to $Y(\mathfrak{q})$ where $\mathfrak{q}$ is the arc $\lim^\omega(\mathfrak{q}_n)$, and $\lim^\omega(A_n) \not\subset \lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$. In the first case we get a contradiction with property (T_1) for $\mathcal{A}$. In the second case we get a contradiction with Lemma 2.21, part (2). $\square$

Lemma 4.23 (uniform variant of Lemma 4.11 for saturations). *For every $x \in X$ and every (L, C)-quasi-geodesic $\mathfrak{q}$ in X with* $\mathrm{dist}(x, \mathrm{Sat}(\mathfrak{q})) = 2d$,

$$\mathrm{diam}(proj_{\mathrm{Sat}(\mathfrak{q})}(\mathcal{N}_d(x)) = o(d).$$

Proof. By contradiction, suppose that there exists a sequence of quasi-geodesics $\mathfrak{q}_n$ and points x_n with $\lim_\omega \mathrm{dist}(x_n, \mathrm{Sat}(\mathfrak{q}_n)) = 2d_n$ such that $\lim_\omega d_n = \infty$, and the almost projection of $\mathcal{N}_{d_n}(x_n)$ on $\mathrm{Sat}(\mathfrak{q}_n)$ has diameter at least td_n for some fixed t. In the asymptotic cone $\mathrm{Con}^\omega(X, (x_n), (d_n))$ we have, according to Lemma 4.21, that $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ is either one piece or a set of type Y. We apply Lemma 2.21, part (2), and get a contradiction. □

Lemma 4.24 (uniform property (α_2') for saturations). *For every $\lambda \geq 1$, $\kappa \geq 0$ and $\theta \in \left[0, \frac{1}{2}\right)$ there exists R such that for every (λ, κ)-quasi-geodesic $\mathfrak{c} : [0, \ell] \to X$ joining two points in $\mathcal{N}_{\theta\ell/L}(\mathrm{Sat}(\mathfrak{q}))$, where $\mathfrak{q}$ is a quasi-geodesic, we have $\mathfrak{c}([0, \ell]) \cap \mathcal{N}_R(\mathrm{Sat}(\mathfrak{q})) \neq \emptyset$ (in particular, the constant R does not depend on $\mathfrak{q}$).*

Proof. One can simply repeat the argument of Lemma 4.12 but use Lemma 4.23 instead of Lemma 4.11. □

Lemma 4.25 (uniform quasi-convexity of saturations). *For every $\lambda \geq 1$, $\kappa \geq 0$, there exists τ such that for every $R \geq 1$, for every quasi-geodesic $\mathfrak{q}$ the saturation $\mathrm{Sat}(\mathfrak{q})$ has the property that every (λ, κ)-quasi-geodesic $\mathfrak{c}$ joining two points in its R-tubular neighborhood is entirely contained in its τR-tubular neighborhood.*

Proof. By Remark 4.14, it is enough to prove the statement for (λ, κ)-almost geodesics $\mathfrak{c}$. Suppose there exists a sequence of quasi-geodesics $\mathfrak{q}_n$, a sequence of numbers $R_n \geq 1$, a sequence $\mathfrak{c}_n$ of (λ, κ)-almost geodesics joining the points x_n, y_n in the R_n-tubular neighborhood of $\mathrm{Sat}(\mathfrak{q}_n)$ such that $\mathfrak{c}_n$ is not contained in the nR_n-tubular neighborhood of $\mathrm{Sat}(\mathfrak{q}_n)$.

Let $z_n \in \mathfrak{c}_n$ be such that $d_n = \mathrm{dist}(z_n, \mathrm{Sat}(\mathfrak{q}_n))$ is maximal. By Lemma 4.21, in the asymptotic cone $\mathrm{Con}^\omega(X; (z_n), (d_n))$, we have that $S = \lim^\omega(\mathrm{Sat}(\mathfrak{q}_n))$ is either one piece or a set $Y(\mathfrak{q})$ of type Y. On the other hand by Remark 3.15 $\lim^\omega(\mathfrak{c}_n)$ is either a topological arc with endpoints in S and not contained in it, or a bi-Lipschitz ray with origin in S or a bi-Lipschitz line. In addition, $\lim^\omega(c_n)$ is contained in $\mathcal{N}_1(S)$. In all three cases Lemma 2.21, part (2), and Corollary 2.7 give a contradiction. □

Lemma 4.26 (saturations of polygonal lines). *Let X be a geodesic metric space. Then the following is true for every $k \geq 1$.*

1. *For every $n \geq 1$, let $\bigcup_{i=1}^k \mathfrak{q}_i^{(n)}$ be a polygonal line composed of (L, C)-quasi-geodesics $\mathfrak{q}_i^{(n)}$. Then in every asymptotic cone the limit set $\lim^\omega(\bigcup_{i=1}^k \mathrm{Sat}(\mathfrak{q}_i^{(n)})) = \bigcup_{i=1}^k \lim^\omega(\mathrm{Sat}(\mathfrak{q}_i^{(n)}))$ is either a piece or a connected union of sets of type Y (as in Lemma 2.21, part (3)).*
2. *The results in Lemmas 4.23, 4.24, 4.25 are true if we replace $\mathrm{Sat}(\mathfrak{q})$ with $\bigcup_{i=1}^k \mathrm{Sat}(\mathfrak{q}_i)$, where $\bigcup_{i=1}^k \mathfrak{q}_i$ is a polygonal line composed of (L, C)-quasi-geodesics.*
3. *For every $\delta > 0$, for every polygonal line $\bigcup_{i=1}^k \mathfrak{q}_i$ composed of (L, C)-quasi-geodesics, and every $A \in \mathcal{A}$ such that $A \not\subset \bigcup_{i=1}^k \mathrm{Sat}(\mathfrak{q}_i)$, the intersection $\mathcal{N}_\delta(A) \cap \mathcal{N}_\delta\left(\bigcup_{i=1}^k \mathrm{Sat}(\mathfrak{q}_i)\right)$ has a uniformly bounded diameter in terms of $A, \mathfrak{q}_1, \ldots, \mathfrak{q}_k$.*

Proof. We prove simultaneously (1), (2) and (3) by induction on k. For $k = 1$ all three statements are true. Suppose they are true for $i \leq k$. We prove them for $k+1$. We note that (1) implies (2) in the same way as Lemma 4.21 implies the cited Lemmas, and the implication (1) $\Rightarrow$ (3) follows from Lemma 2.21, part (3) (the argument is essentially the same as in Lemma 4.22). Thus it is enough to prove part (1).

Let $\mathrm{Con}^\omega(X; e, d)$ be an asymptotic cone. We suppose that

$$\lim_\omega \frac{\mathrm{dist}\left(e_n, \bigcup_{i=1}^{k+1} \mathrm{Sat}\left(q_i^{(n)}\right)\right)}{d_n} < \infty$$

(otherwise the ω-limit is empty). There are two possible situations.

Case I. Suppose that there exists an integer i between 2 and k such that

$$\lim_\omega \frac{\mathrm{dist}\left(e_n, \mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)\right)}{d_n} < \infty.$$

By the inductive hypothesis $\lim^\omega(\bigcup_{j=1}^{i} \mathrm{Sat}\left(\mathfrak{q}_j^{(n)}\right))$ is a set of type Y, and so is the set

$$\lim^\omega(\bigcup_{j=i}^{k+1} \mathrm{Sat}\left(\mathfrak{q}_j^{(n)}\right)).$$

These two sets have a common non-empty subset $\lim^\omega(\mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right))$. Since a connected union of two sets of type Y is again a set of type Y, statement (1) follows.

Case II. Suppose that for every i between 2 and k, we have

$$\lim_\omega \frac{\mathrm{dist}\left(e_n, \mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)\right)}{d_n} = \infty.$$

If the same is true either for $i = 1$ or for $i = k+1$ one can apply Lemma 4.21. Thus suppose that for $i = 1, k+1$, we have

$$\lim_\omega \frac{\mathrm{dist}\left(e_n, \mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)\right)}{d_n} < \infty.$$

By Lemma 4.21, for $i = 1, k+1$, for the limit set $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_i^{(n)}))$ one of the following two possibilities occurs:

($\mathbf{A}_i$) it is equal to $\lim^\omega(A_n)$, where $A_n \in \mathcal{A}, A_n \subseteq \mathrm{Sat}(\mathfrak{q}_i^{(n)})$;

($\mathbf{B}_i$) it is equal to $Y(\mathfrak{q}_i)$ as in Lemma 2.21, part (2), where $\mathfrak{q}_i = \lim^\omega(\mathfrak{q}_i^{(n)})$.

It remains to show that the union $\lim^\omega(\mathrm{Sat}(\mathfrak{q}_1^{(n)})) \cup \lim^\omega(\mathrm{Sat}(\mathfrak{q}_{k+1}^{(n)}))$ is connected.

Suppose that we are in the situation ($\mathbf{B}_1$). Let $u_n \in \mathfrak{q}_1^{(n)}$ be an almost nearest point from e_n. Then $\mathrm{dist}(u_n, e_n) = O(d_n)$. Let $v_n \in \bigcup_{j=2}^{k+1} \mathrm{Sat}\left(\mathfrak{q}_j^{(n)}\right)$ be an almost nearest point to e_n. By our assumption, ω-almost surely $v_n \in \mathrm{Sat}\left(\mathfrak{q}_{k+1}^{(n)}\right)$ and $\mathrm{dist}(v_n, e_n) = O(d_n)$. Hence $\mathrm{dist}(u_n, v_n) = O(d_n)$. Let R_k be the constant given by the variant of Lemma 4.24 for polygonal lines composed of k (L,C)-quasi-geodesics with $(\lambda,\kappa)=(L,C)$, $\theta=\frac{1}{3}$ (that R_k exists by the induction hypothesis). Let a_n be an almost nearest point from u_n in $\mathfrak{q}_1^{(n)}\cap\mathcal{N}_{R_k}\left(\bigcup_{j=2}^{k+1}\mathrm{Sat}\left(\mathfrak{q}_j^{(n)}\right)\right)$. Let $\mathfrak{p}^{(n)}$ be the sub-quasi-geodesic of $\mathfrak{q}_1^{(n)}$ with endpoints u_n and a_n. According to the part (2) of the proposition (which by the induction assumption is true for k), $\mathfrak{p}^{(n)}\subset\mathcal{N}_{td_n}\left(\bigcup_{j=2}^{k+1}\mathrm{Sat}(\mathfrak{q}_j^{(n)})\right)$ for some t independent on n. If $\mathrm{dist}(u_n,a_n)\gg d_n$ then according to Lemma 4.24 there exists another point on $\mathfrak{p}^{(n)}\cap\mathcal{N}_{R_k}\left(\bigcup_{j=2}^{k+1}\mathrm{Sat}\left(\mathfrak{q}_j^{(n)}\right)\right)$ whose distance from u_n is smaller than $\mathrm{dist}(a_n,u_n)$ by $O(d_n)$, a contradiction. Therefore $\mathrm{dist}(u_n,a_n)\le O(d_n)$ and the limit point $\lim^\omega(a_n)$ is a common point of $\mathfrak{q}_1$ and $\lim^\omega(\bigcup_{i=2}^{k+1}\mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right))=\lim^\omega(\mathrm{Sat}\left(\mathfrak{q}_{k+1}^{(n)}\right))$.

The same argument works if we are in the situation $(\mathbf{B}_{k+1})$. Therefore we suppose that we are in the situations $(\mathbf{A}_1)$ and $(\mathbf{A}_{k+1})$. We have that $\lim^\omega(\mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)), i=1,k+1$, is equal to $\lim^\omega(A_i^{(n)})$, where $A_i^{(n)}\in\mathcal{A}, A_i^{(n)}\subseteq\mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)$. Suppose that $A_1^{(n)}\neq A_{k+1}^{(n)}$ ω-almost surely. Let $v_i^{(n)}\in\mathrm{Sat}\left(\mathfrak{q}_i^{(n)}\right)$ be an almost nearest point from e_n. By hypothesis $v_i^{(n)}\in A_i^{(n)}$.

The two assumptions:

$$\lim_\omega\frac{\mathrm{dist}(e_n,\mathrm{Sat}(\mathfrak{q}_i^{(n)}))}{d_n}=\infty,$$

$i\in\{2,\dots,k\}$, and

$$\lim^\omega(\mathrm{Sat}(\mathfrak{q}_{k+1}^{(n)}))=\lim^\omega(A_{k+1}^{(n)})$$

imply that $A_1^{(n)}\not\subset\bigcup_{i=2}^{k+1}\mathrm{Sat}(\mathfrak{q}_i^{(n)})$ ω-almost surely.

Suppose that $[0,\ell_1^{(n)}]$ is the domain of $\mathfrak{q}_1^{(n)}$. The following two cases may occur.

Case I. If the distance from $\ell_1^{(n)}$ to the pre-image $(\mathfrak{q}_1^{(n)})^{-1}(A_1^{(n)})$ is at most $LC+1$ then we denote $\mathfrak{q}_1^{(n)}(\ell_1^{(n)})$ by a_n. We have that $\mathrm{dist}(a_n,\mathfrak{q}_1^{(n)}\cap A_1^{(n)})\le L^2C+L+C$, which implies by Lemma 4.15 that a geodesic $\mathfrak{p}_n=[v_1^{(n)},a_n]$ is contained in the $t(L^2C+L+C)$-tubular neighborhood of $A_1^{(n)}$.

Case II. If the distance from $\ell_1^{(n)}$ to $(\mathfrak{q}_1^{(n)})^{-1}(A_1^{(n)})$ is larger than $LC+1$, then we consider $t_n\in[0,\ell_1^{(n)}]$ at distance $LC+1$ of $(\mathfrak{q}_1^{(n)})^{-1}(A_1^{(n)})$ such that all points in $[t_n,\ell_1^{(n)}]$ are at distance at least $LC+1$ of $(\mathfrak{q}_1^{(n)})^{-1}(A_1^{(n)})$. We denote by a_n the point $\mathfrak{q}_1^{(n)}(t_n)$. According to Lemma 4.15 we have that a geodesic $[v_1^{(n)},a_n]$ is contained in the $t(L^2C+L+C)$-tubular neighborhood of $A_1^{(n)}$.

By our assumption, $\lim_\omega\frac{\mathrm{dist}(v_1^{(n)},a_n)}{d_n}=\infty$. Lemma 4.19 implies that $[v_1^{(n)},a_n]$ and the restriction of $\mathfrak{q}_1^{(n)}$ to $[t_n,\ell_1^{(n)}]$ form an $(L+\varepsilon,K)$-quasi-geodesic ω-almost surely. We denote it by $\mathfrak{p}_n$.

Both in Case I and in Case II we have obtained an $(L+\varepsilon,K)$-quasi-geodesic $\mathfrak{p}_n$ with one of the endpoints $v_1^{(n)}$ and the other one contained in $\mathfrak{q}_2^{(n)}$. The distance from $v_1^{(n)}$ to $\bigcup_{i=2}^{k+1}\mathrm{Sat}(\mathfrak{q}_i^{(n)})$ does not exceed $\mathrm{dist}(v_1^{(n)},v_{k+1}^{(n)})$, hence it is at most $O(d_n)$. It follows that $\mathfrak{p}_n\subset\mathcal{N}_{O(d_n)}\left(\bigcup_{i=2}^{k+1}\mathrm{Sat}(\mathfrak{q}_i^{(n)})\right)$. In particular $[v_1^{(n)},a_n]$ is contained in the same tubular neighborhood. Since the length λ_n of $[v_1^{(n)},a_n]$ satisfies $\lim_\omega\frac{\lambda_n}{d_n}=\infty$, by applying Lemmas 4.24 and 4.25 we obtain that a sub-segment $[\alpha_n,\beta_n]$ of $[v_n^1,a_n]$ of length $\frac{\lambda_n}{2}$ is contained in $\mathcal{N}_{\tau R}\left(\bigcup_{i=2}^{k+1}\mathrm{Sat}\left(\mathfrak{q}_n^i\right)\right)$, where R is an universal constant. On the other hand we have $[\alpha_n,\beta_n]\subset\mathcal{N}_{t(L^2C+L+C)}(A_1^{(n)})$.

This contradicts the inductive hypothesis (3). We conclude that if we are in situation $(\mathbf{A}_1)$ then $\lim_\omega \frac{\mathrm{dist}\left(e_n\,,\mathrm{Sat}\left(\mathfrak{q}_n^{k+1}\right)\right)}{d_n} = \infty$. $\square$

Corollary 4.27. *Let Δ be a quasi-geodesic triangle. Then every edge $\mathfrak{a}$ of Δ is contained in an M-tubular neighborhood of $Sat(\mathfrak{b}) \cup Sat(\mathfrak{c})$, where $\mathfrak{b}$ and $\mathfrak{c}$ are the two other edges of Δ and M is an universal constant.*

Lemma 4.28. *For every $R > 0, k \in \mathbb{N}$ and $\delta > 0$ there exists $\varkappa > 0$ such that if $\bigcup_{i=1}^k \mathfrak{q}_i$ is a polygonal line composed of quasi-geodesics and $A, B \in \mathcal{A}, A \cup B \subset \bigcup_{i=1}^k \mathrm{Sat}(\mathfrak{q}_i), A \neq B$, the following holds. Let $a \in \mathcal{N}_R(A)$ and $b \in \mathcal{N}_R(B)$ be two points that can be joined by a quasi-geodesic $\mathfrak{p}$ such that $\mathfrak{p} \cap \mathcal{N}_R(A)$ and $\mathfrak{p} \cap \mathcal{N}_R(B)$ has diameter at most δ. Then $\{a, b\} \subset \mathcal{N}_\varkappa\left(\bigcup_{i=1}^k \mathfrak{q}_i\right)$.*

Proof. Suppose $\mathfrak{q}_i$ is defined on the interval $[0, \ell_i]$. Let $\mathfrak{r} : [0, \sum_{i=1}^k \ell_i] \to X$ be the map defined by $\mathfrak{r}(\sum_{i=1}^{j-1} \ell_i + t) = \mathfrak{q}_j(t)$, for all $t \in [0, \ell_j]$ and all $j \in \{2, \dots, k\}$. It satisfies

$$\mathrm{dist}(\mathfrak{r}(t), \mathfrak{r}(s)) \leq L|t - s| + kC\,. \tag{8}$$

Let x be a point in $\mathfrak{r} \cap \mathcal{N}_M(B)$ and $t_x \in [0, \sum_{i=1}^k \ell_i]$ such that $\mathfrak{r}(t_x) = x$. We have two cases.

(a) If the distance from t_x to the pre-image $\mathfrak{r}^{-1}(\mathcal{N}_M(A))$ does not exceed $LC + 1$ then $x \in \mathcal{N}_{M+L^2C+L+kC}(A)$ by (8). By Lemma 4.19, if $\mathrm{dist}(a, x)$ is larger than D' then the union of $\mathfrak{p}$ and a geodesic $[a, x]$ form an $(L + \varepsilon, K)$-quasi-geodesic, with endpoints in $\mathcal{N}_{R+M}(B)$. It follows that this quasi-geodesic and in particular $[a, x]$ are contained in $\mathcal{N}_{t(M+R)}(B)$. On the other hand $[a, x]$ is contained in $\mathcal{N}_{t(M+R+L^2C+L+kC)}(A)$. If $\mathrm{dist}(a, x)$ is larger than the diameter given by (α_1) for $\delta = t(M + R + L^2C + L + kC)$ then we obtain a contradiction with (α_1).

(a) Suppose that the distance from t_x to $\mathfrak{r}^{-1}(\mathcal{N}_M(A))$ is larger than $LC + 1$. Consider s_0 at distance $LC + 1$ from $\mathfrak{r}^{-1}(\mathcal{N}_M(A))$ such that every s between s_0 and t_x is at distance at least $LC + 1$ from $\mathfrak{r}^{-1}(\mathcal{N}_M(A))$. It follows that $\mathfrak{r}([s_0, t_x])$ or $\mathfrak{r}([t_x, s_0])$ is disjoint of $\mathcal{N}_M(A)$. Let $y = \mathfrak{r}(s_0)$. The restriction $\mathfrak{r}'$ of $\mathfrak{r}$ to $[s_0, t_x]$ or $[t_x, s_0]$ can be written as $\bigcup_{j=1}^m \mathfrak{q}'_j$, where $m \leq k$ and each $\mathfrak{q}'_j$ coincides with one of the $\mathfrak{q}_i$'s or a restriction of it. We note that $A \not\subset \mathrm{Sat}(\mathfrak{r}')$.

If the distance from a to y is larger than the constant D' given by Lemma 4.19 then $\mathfrak{p}$ and a geodesic $[a, y]$ form an $(L + \varepsilon, K)$-quasi-geodesic. Lemma 4.26, part (2), implies that this quasi-geodesic, and in particular $[a, y]$, is contained in the τR-tubular neighborhood of $\mathrm{Sat}(\mathfrak{r}')$. On the other hand, $[a, y]$ is contained in the $t(R + M + L^2C + L + kC)$-tubular neighborhood of A. For $\mathrm{dist}(a, y)$ larger than the diameter given by Lemma 4.26, (3), for $\delta = \max\left(t(R + M + L^2C + L + kC), \tau R\right)$ we obtain a contradiction. $\square$

Lemma 4.29. *Let X be an asymptotically tree-graded metric space with respect to $\mathcal{A}$. Then X satisfies (α'_3).*

Proof. Let $k \geq 2$, $\sigma \geq 1$ and $\nu \geq 4\sigma$. Fix a sufficiently large number ϑ (it will be clear later in the proof how large ϑ should be). Let P be a k-gon with quasi-geodesic edges that is $(\vartheta\,, \sigma\,, \nu)$-fat. Changing if necessary the polygon by a finite Hausdorff distance, we may suppose that its edges are $(L + C, C)$-almost geodesics.

Let $\mathfrak{q} : [0, \ell] \to X$ be an edge with endpoints $\mathfrak{q}(0) = x, \mathfrak{q}(\ell) = y$. We denote $\mathfrak{q}_1, \mathfrak{q}_2, \dots, \mathfrak{q}_{k-1}$ the other edges in the clockwise order. By Lemma 4.26, part (2),

$$\mathfrak{q} \subset \mathcal{N}_{\tau R}\left(\bigcup_{i=1}^{k-1} \mathrm{Sat}\,(\mathfrak{q}_i)\right).$$

We take $\vartheta > \tau R$. Then for every point $z \in \mathfrak{q} \setminus \mathcal{N}_{\sigma\vartheta}(\{x, y\})$ there exists $A \subset \mathrm{Sat}\,(\mathfrak{q}_i)\,, i \in \{1, 2, \ldots, k-1\}$ such that $z \in \mathcal{N}_{\tau R}(A)$. If such a point z is contained in $\mathcal{N}_{\tau R}(A) \cap \mathcal{N}_{\tau R}(B)$, $A \neq B$, then Lemma 4.28 implies that $z \in \mathcal{N}_{\varkappa}(\bigcup_{i=1}^{k-1} \mathfrak{q}_i)$, where $\varkappa$ depends on τR and k. If we choose $\vartheta > \varkappa$ then this gives a contradiction.

Let $t_{\mathfrak{q}}$ be the supremum of the numbers $t \in [0, \ell]$ contained in $\mathfrak{q}^{-1}(\mathcal{N}_{\sigma\vartheta}(x))$. Let $s_{\mathfrak{q}}$ be the infimum of the numbers in $[t_{\mathfrak{q}}, \ell]$ contained in $\mathfrak{q}^{-1}(\mathcal{N}_{\sigma\vartheta}(y))$. Let $a_{\mathfrak{q}} = \mathfrak{q}(t_{\mathfrak{q}})$ and $b_{\mathfrak{q}} = \mathfrak{q}(s_{\mathfrak{q}})$. We note that $\mathrm{dist}(a_{\mathfrak{q}}, x) = \sigma\vartheta$ and $\mathrm{dist}(b_{\mathfrak{q}}, y) = \sigma\vartheta$. According to the argument in the paragraph above, $\mathfrak{q}([t_{\mathfrak{q}}, s_{\mathfrak{q}}])$ is covered by the family of open sets $\mathcal{N}_{\tau R}(A)$, with $A \subset \mathrm{Sat}\,(\mathfrak{q}_i)\,, i \in \{1, 2, \ldots, k-1\}$, and the traces of these sets on $\mathfrak{q}([t_{\mathfrak{q}}, s_{\mathfrak{q}}])$ are pairwise disjoint. The connectedness of $\mathfrak{q}([t_{\mathfrak{q}}, s_{\mathfrak{q}}])$ implies that there exists A as above such that $\mathfrak{q}([t_{\mathfrak{q}}, s_{\mathfrak{q}}]) \subset \mathcal{N}_{\tau R}(A)$.

Thus, for every edge $\mathfrak{q}$ a sub-arc $\mathfrak{q}' : [t_{\mathfrak{q}}, s_{\mathfrak{q}}] \to X$ with endpoints $a_{\mathfrak{q}}, b_{\mathfrak{q}}$ is contained in $\mathcal{N}_{\tau R}(A)$ for some $A \subset \mathrm{Sat}\,(\mathfrak{q}_i)\,, i \in \{1, 2, \ldots, k-1\}$ (A may depend on $\mathfrak{q}$). We note that $t_{\mathfrak{q}}$ and $\ell - s_{\mathfrak{q}}$ are less than $\sigma\vartheta L + C$, hence $\mathfrak{q}|_{[0,t_{\mathfrak{q}}]} \in \mathcal{N}_{\sigma\vartheta L^2+LC+C}(a_{\mathfrak{q}})$ and $\mathfrak{q}|_{[s_{\mathfrak{q}},\ell]} \in \mathcal{N}_{\sigma\vartheta L^2+LC+C}(b_{\mathfrak{q}})$.

Suppose that we have two consecutive edges $\mathfrak{q}_1, \mathfrak{q}_2$ with endpoints x, y and y, z respectively, such that $\mathfrak{q}_1' \subset \mathcal{N}_{\tau R}(A)$ and $\mathfrak{q}_2' \subset \mathcal{N}_{\tau R}(B)$, $A \neq B$. We denote $\mathfrak{q}_3, \mathfrak{q}_4, \ldots, \mathfrak{q}_k$ the other edges in the clockwise order. We have $\mathfrak{q}_i' : [t_{\mathfrak{q}_i}, s_{\mathfrak{q}_i}] \to X$ with endpoints $a_{\mathfrak{q}_i}, b_{\mathfrak{q}_i}$. Suppose $b_{\mathfrak{q}_1} = \mathfrak{q}_1' \cap \mathcal{N}_{\sigma\vartheta}(y)$ and $a_{\mathfrak{q}_2} = \mathfrak{q}_2' \cap \mathcal{N}_{\sigma\vartheta}(y)$.

Let $\bar{\mathfrak{q}}_1$ be the restriction of $\mathfrak{q}_1'$ to $[t_{\mathfrak{q}_1}, t_{\mathfrak{q}_1} + 3L\tau R]$ and $\tilde{\mathfrak{q}}_1 = [x, a_{\mathfrak{q}_1}] \cup \bar{\mathfrak{q}}_1$. We note that since $\mathrm{dist}(a_{\mathfrak{q}_1}, b_{\mathfrak{q}_1}) \geq \mathrm{dist}(x, y) - 2\sigma\vartheta \geq \nu\vartheta - 2\sigma\vartheta \geq 2\sigma\vartheta$, we have $s_{\mathfrak{q}_1} - t_{\mathfrak{q}_1} \geq \frac{2\sigma\vartheta}{L} - C$, so for ϑ large enough we have $s_{\mathfrak{q}_1} - t_{\mathfrak{q}_1} \geq 10L\tau R$ and the restriction $\bar{\mathfrak{q}}_1$ makes sense.

Likewise we construct $\tilde{\mathfrak{q}}_2 = \bar{\mathfrak{q}}_2 \cup [b_{\mathfrak{q}_2}, z]$, where $\bar{\mathfrak{q}}_2$ is the restriction of $\mathfrak{q}_2'$ to the last sub-interval of length $3L\tau R$.

Let $[a, b]$ be a geodesic joining the points $a = a_{\mathfrak{q}_2}$ and $b = b_{\mathfrak{q}_1}$. It has length at most $2\sigma\vartheta$. Let $[a', b'] \subset [a, b]$ be a sub-geodesic which intersects $\mathcal{N}_{\tau R}(A)$ in a' and $\mathcal{N}_{\tau R}(B)$ in b' (eventually reduced to a point). Notice that $A \subseteq \mathrm{Sat}(\tilde{\mathfrak{q}}_1)$, $B \subseteq \mathrm{Sat}(\tilde{\mathfrak{q}}_2)$. Lemma 4.28 applied to the polygonal line $\tilde{\mathfrak{q}}_2 \cup \bigcup_{i=3}^k \mathfrak{q}_i \cup \tilde{\mathfrak{q}}_1$ and to the points a', b' implies that $\{a', b'\} \subset \mathcal{N}_{\varkappa}\left(\tilde{\mathfrak{q}}_2 \cup \bigcup_{i=3}^k \mathfrak{q}_i \cup \tilde{\mathfrak{q}}_1\right)$, where $\varkappa$ depends on τR. Since $\mathrm{dist}(y, \{a', b'\})$ is at most $2\sigma\vartheta$, it follows that $y \in \mathcal{N}_{\varkappa+2\sigma\vartheta}\left(\tilde{\mathfrak{q}}_2 \cup \bigcup_{i=3}^k \mathfrak{q}_i \cup \tilde{\mathfrak{q}}_1\right) \subset \mathcal{N}_{\varkappa+3\sigma\vartheta+3L^2\tau R+C}\left(\bigcup_{i=3}^k \mathfrak{q}_i\right)$. On the other hand property (F_2) implies that $\mathrm{dist}(y, \bigcup_{i=3}^k \mathfrak{q}_i) \geq \nu\vartheta \geq 4\sigma\vartheta$. For ϑ large enough this gives a contradiction.

We conclude that there exists $A \in \mathcal{A}$ such that $\bigcup_{i=1}^k \mathfrak{q}_i' \subset \mathcal{N}_{\tau R}(A)$. Hence P is inside the $(\tau R + \sigma\vartheta L^2 + LC + C)$-tubular neighborhood of A. $\square$

The following corollary immediately follows from the proof of Theorem 4.1

Corollary 4.30 (there is no need to vary the ultrafilter in Definition 3.19). *Let X be a metric space, let $\mathcal{A}$ be a collection of subsets in X. Let ω be any ultrafilter over $\mathbb{N}$. Suppose that every asymptotic cone $\mathrm{Con}^{\omega}(X; e, d)$ is tree-graded with respect to the collection of sets $\lim^{\omega}(A_n)$, $A_n \in \mathcal{A}$. Then X is asymptotically tree-graded with respect to $\mathcal{A}$.*

5 Quasi-isometric behavior

One of the main interests in the property of being asymptotically tree-graded with respect to a collection of subsets resides in the rigid behavior with respect to quasi-isometry.

5.1 Asymptotically tree-graded spaces

Theorem 5.1 (being asymptotically tree-graded is a geometric property.). *Let X be a metric space and let $\mathcal{A}$ be a collection of subsets of X. Let $\mathfrak{q}$ be a quasi-isometry $X \to X'$.*

Then:

(1) If X satisfies properties (α_1) and (α_2') with respect to $\mathcal{A}$ then X' satisfies properties (α_1) and (α_2^ϵ), for a sufficiently small ϵ, with respect to $\mathfrak{q}(\mathcal{A}) = \{\mathfrak{q}(A) \mid A \in \mathcal{A}\}$.

(2) If X satisfies (α_3') with respect to $\mathcal{A}$ then X' satisfies (α_3) with respect to $\mathfrak{q}(\mathcal{A})$.

(3) If X is asymptotically tree-graded with respect to $\mathcal{A}$ then X' is asymptotically tree-graded with respect to $\mathfrak{q}(\mathcal{A})$.

Proof. (1) follows from Theorem 4.1 and Remark 4.2.

(2) Assume that $\mathfrak{q}$ is an (L, C)-quasi-isometry and that $\bar{\mathfrak{q}} : X' \to X$ is an (L, C)-quasi-isometry so that $\bar{\mathfrak{q}} \circ \mathfrak{q}$ and $\mathfrak{q} \circ \bar{\mathfrak{q}}$ are at distance at most C from the respective identity maps.

We fix an arbitrary integer $k \geq 2$. Let $\sigma = 2L^2 + 1$ and $\nu = 4\sigma$. Property (α_3') in X implies that for the constants L, C of the quasi-isometries, for the given k, σ and ν there exists ϑ_0 such that for every $\vartheta \geq \vartheta_0$ a k-gon with (L, C)-quasi-geodesic edges in X which is $(\vartheta\,, \sigma\,, \nu)$-fat is contained in $\mathcal{N}_\chi(A)$, where $A \in \mathcal{A}$ and $\chi = \chi(L, C, k, \sigma, \nu, \vartheta)$.

Let $\vartheta_1 = \max(\vartheta_0, 2L^2C + C)$ and let $\theta = L(\vartheta_1 + C)$. Let P be a geodesic k-gon in X' which is $(\theta, 2, \nu)$-fat. Then $\bar{\mathfrak{q}}(P)$ is a k-gon in X with (L, C)-quasi-geodesic edges which is $(\vartheta_1\,, \sigma\,, \nu)$-fat. Consequently, $\bar{\mathfrak{q}}(P) \subset \mathcal{N}_\chi(A)$, where $A \in \mathcal{A}$ and $\chi = \chi(L, C, k, \sigma, \nu, \vartheta_1)$. It follows that $P \subset \mathcal{N}_C(\mathfrak{q} \circ \bar{\mathfrak{q}}(P)) \subset \mathcal{N}_{L\chi+2C}(\mathfrak{q}(A))$.

(3) The statement follows from (1) and (2). It also follows immediately by means of the definition of asymptotically tree graded spaces using asymptotic cones. Indeed, it is easy to see that ω-limits of sequences of subsets commute with quasi-isometries. Since quasi-isometric spaces have bi-Lipschitz equivalent asymptotic cones (Remark 3.16) it remains to note that a metric space that is bi-Lipschitz equivalent to a space that is tree-graded with respect to $\mathcal{P}$, is itself tree-graded with respect to the images of the sets in $\mathcal{P}$ under the bi-Lipschitz map. □

Definition 5.2. Let $\mathcal{B}$ be a family of geodesic metric spaces. We say that $\mathcal{B}$ is *uniformly asymptotically without cut-points* if for every sequence B_n of metric spaces in $\mathcal{B}$ with metrics dist_n and basepoints $b_n \in B_n$, for every ultrafilter ω and for every sequence of scaling constants (d_n) with $\lim_\omega d_n = \infty$, the ultralimit $\lim^\omega(B_n, \frac{1}{d_n}\mathrm{dist}_n)_b$ is without cut-points.

Remarks 5.3. (a) All metric spaces in a family that is uniformly asymptotically without cut-points are asymptotically without cut-points.

(b) If $\mathcal{B}$ is a family of metric spaces composed of finitely many isometry classes of spaces asymptotically without cut-points then $\mathcal{B}$ is uniformly asymptotically without cut-points.

(c) For examples of groups that are asymptotically without cut-points and of families of groups that are uniformly asymptotically without cut-points, see Section 6.

Proposition 5.4. *Let X be asymptotically tree-graded with respect to a collection of subsets $\mathcal{A}$. Let $\mathcal{B}$ be a family of metric spaces which is uniformly asymptotically without cut-points. For every (L, C) there exists $M = M(L, C)$ such that for every $B \in \mathcal{B}$ and every (L, C)-quasi-isometric embedding $\mathfrak{q} \colon B \to X$ there exists $A \in \mathcal{A}$ such that $\mathfrak{q}(B) \subset \mathcal{N}_M(A)$.*

Proof. We argue by contradiction and assume that there is a sequence of metric spaces $B_n \in \mathcal{B}$ and a sequence of (L, C)-quasi-isometric embeddings $\mathfrak{q}_n : B_n \to X$ such that $\mathfrak{q}_n(B_n) \not\subset \mathcal{N}_n(A)$ for all $A \in \mathcal{A}$. Fix a point $b_n \in B_n$. Let $e = (\mathfrak{q}_n(b_n))$, $d = (n)$. In $\mathrm{Con}^\omega(X; e, d)$, the limit set

$\lim^\omega(\mathfrak{q}_n(B_n))$ is a bi-Lipschitz embedding of $\lim^\omega(B_n, \frac{1}{n}\mathrm{dist}_n)_b$, therefore it is without cut-points. Lemma 2.13 implies that

$$\lim^\omega(\mathfrak{q}_n(B_n)) \subset \lim^\omega(A_n)\,, \text{ where } A_n \in \mathcal{A}\,. \tag{9}$$

Let $x_n \in B_n$ be such that $\mathfrak{q}_n(x_n) \in \mathfrak{q}_n(B_n) \setminus \mathcal{N}_n(A_n)$. Let $[b_n, x_n]$ be a geodesic in B_n. The inclusion in (9) implies that $\lim_\omega \frac{\mathrm{dist}(\mathfrak{q}_n(b_n), A_n)}{n} = 0$ and that for every t the point $b_n(t)$ on $[b_n, x_n]$ at distance tn of b_n satisfies $\lim_\omega \frac{\mathrm{dist}(\mathfrak{q}_n(b_n(t)), A_n)}{n} = 0$. Therefore for every $s > t > 0$ the image by $\mathfrak{q}_n$ of every segment $[b_n(s), b_n(t)]$ contains a point in $\mathcal{N}_{M_0}(A_n)$, where M_0 is the constant given by (α_2'), for L and C. Let y_n be the farthest point from b_n in the closure of $[b_n, x_n] \cap \mathfrak{q}_n^{-1}(\mathcal{N}_{M_0}(A_n))$. The previous argument implies that $\lim_\omega \frac{\mathrm{dist}(b_n, y_n)}{n} = +\infty$. Also, $y_n \in \overline{[b_n, x_n] \cap \mathfrak{q}_n^{-1}(\mathcal{N}_{M_0}(A_n))}$ implies that for every $\varepsilon > 0$ the distance from $\mathfrak{q}_n(y_n)$ to A_n is at most $M_0 + L\varepsilon + C$. Hence $\mathfrak{q}_n(y_n) \in \mathcal{N}_{M_0+C+1}(A_n)$. On the other hand, there is a point $z_n \in [b_n, b_n(1)] \cap \mathcal{N}_{M_0}(A_n)$. According to Lemma 4.15, $\mathfrak{q}_n([z_n, y_n]) \subset \mathcal{N}_{\tau(M_0+C+1)}(A_n)$. In $\mathrm{Con}^\omega(X; (\mathfrak{q}_n(y_n)), d)$, $\mathfrak{q} = \lim^\omega(\mathfrak{q}_n([z_n, y_n]))$ is a bi-Lipschitz ray contained in $A = \lim^\omega(A_n)$ and in $\lim^\omega(\mathfrak{q}_n(B_n))$. Since $\lim^\omega(\mathfrak{q}_n(B_n))$ is the image of a bi-Lipschitz embedding of the ultralimit $\lim^\omega(B_n, \frac{1}{n}\mathrm{dist}_n)_y$, it is without cut-points, therefore it is contained in a piece $A' = \lim^\omega(A_n')$. Property (T_1) implies that $A = A'$. In particular $\lim^\omega(\mathfrak{q}_n([y_n, x_n])) \subset A$. We have $\mathrm{dist}\,(\lim^\omega(\mathfrak{q}_n(y_n)), \lim^\omega(\mathfrak{q}_n(x_n))) \geq 1$. Let $u_n \in [y_n, x_n]$ be such that $\mathrm{dist}\,(\lim^\omega(\mathfrak{q}_n(y_n)), \lim^\omega(\mathfrak{q}_n(u_n)))\, \frac{1}{2}$. By (α_2') we find a point $t_n \in \mathfrak{q}_n(y_n, u_n) \cap \mathcal{N}_{M_0}(A_n)$, which contradicts the choice of y_n. $\square$

Corollary 5.5. *Let X be asymptotically tree-graded with respect to a collection of subsets $\mathcal{A}$. Let B be a metric space asymptotically without cut-points. Then every (L, C)-quasi-isometric embedding $\mathfrak{q}\colon B \to X$ maps B into an M-neighborhood of a piece $A \in \mathcal{A}$, where M depends only on L and C.*

Notation: We shall denote the Hausdorff distance between two sets A, B in a metric space by $\mathrm{hdist}(A, B)$.

5.2 Asymptotically tree-graded groups

Definition 5.6. We say that a finitely generated group G is *asymptotically tree-graded with respect to the family of subgroups $\{H_1, H_2, \ldots, H_k\}$* if the Cayley graph $\mathrm{Cayley}(G)$ with respect to some (and hence every) finite set of generators is asymptotically tree-graded with respect to the collection of left cosets $\{gH_i \mid g \in G\,,\, i = 1, 2, \ldots, k\}$.

Remark 5.7. If $\{H_1, H_2, \ldots, H_k\} \neq \{G\}$ and if every H_i is infinite then every H_i has infinite index in G.

Proof. Indeed, a finite index subgroup is at bounded distance of the whole group, which would contradict (α_1). $\square$

Proposition 5.8. *Let $G = \langle S \rangle$ be a group that is asymptotically tree-graded with respect to subgroups $H_1, \ldots, H_n$. Then each of the subgroups H_i is finitely generated.*

Proof. Take $h \in H_i$ and consider a geodesic $\mathfrak{g}$ in $\mathrm{Cayley}(G, S)$ connecting 1 and h. By Lemma 4.15 there exists a constant $M > 0$ such that $\mathfrak{g}$ is in the M-tubular neighborhood of H_i. Let $v_1, \ldots, v_k$ be the consecutive vertices of $\mathfrak{g}$. For each $j = 1, \ldots, k$ consider a vertex w_j in H_i at distance $\leq M$ from v_j. Then the distance between w_j and w_{j+1} is at most $2M+1$, $j = 1, \ldots, k-1$. Hence each element $w_j^{-1}w_{j+1}$ belongs to H_i and is of length at most $2M+1$. Since h is a product of these elements, we can conclude that H_j is generated by all its elements of length at most $2M+1$. $\square$

Remark 5.9. Corollary 5.5 gives certain restrictions on the groups that can be quasi-isometrically embedded into asymptotically tree-graded groups. For instance, if G is a group asymptotically tree-graded with respect to a finite family of free Abelian groups of rank at most r, no free Abelian group of rank at least $r+1$ can be quasi-isometrically embedded into G.

Theorem 5.10. *Let X be a space that is asymptotically tree-graded with respect to a collection of subsets $\mathcal{A}$. Assume that*

(1) $\mathcal{A}$ is uniformly asymptotically without cut-points, where all sets $A \in \mathcal{A}$ are endowed with the metric induced from X;

(2) every $A \in \mathcal{A}$ has infinite diameter;

(3) For a fixed $x_0 \in X$ and every $R > 0$ the ball $B(x_0, R)$ intersects finitely many $A \in \mathcal{A}$.

Let G be a finitely generated group which is quasi-isometric to X. Then there exist subsets $A_1, \dots, A_m \in \mathcal{A}$ and subgroups $H_1, \dots, H_m$ of G such that

(I) every $A \in \mathcal{A}$ is quasi-isometric to A_i for some $i \in \{1, 2, \dots, m\}$;

(II) H_i is quasi-isometric to A_i for every $i \in \{1, 2, \dots, m\}$;

(III) G is asymptotically tree-graded with respect to the family of subgroups $\{H_1, H_2, \dots, H_m\}$.

Proof. First we note that we have a quasi-action by quasi-isometries of G on X, which is moreover quasi-transitive.
Notation: Let $g \in G$. We denote by $\mathbf{g}$ the multiplication on the left by g in G.

Lemma 5.11. *Let $\mathfrak{q}\colon G \to X$ and $\bar{\mathfrak{q}}\colon X \to G$ be (L_0, C_0)-quasi-isometries such that $\mathfrak{q} \circ \bar{\mathfrak{q}}$ is at distance C_0 from the identity map on X and the same is true for $\bar{\mathfrak{q}} \circ \mathfrak{q}$ with respect to the identity map on G.*

(1) For every $g \in G$ the map $\mathfrak{q}_g = \mathfrak{q} \circ \mathbf{g} \circ \bar{\mathfrak{q}}$ is an (L, C)-quasi-isometry, where $L = L_0^2$ and $C = L_0 C_0 + C_0$.

(2) For $g, h \in G$ we have that $\mathfrak{q}_g \circ \mathfrak{q}_h$ is at distance C from $\mathfrak{q}_{gh}$.

(3) For every $g \in G$ the map $\mathfrak{q}_g \circ \mathfrak{q}_{g^{-1}}$ is at distance $C + C_0$ from the identity.

(4) For every $x, y \in X$ there exists $g \in G$ such that $\operatorname{dist}(x, \mathfrak{q}_g(y)) \le C_0$.

Proof. Statement (1) follows from the fact that $\mathbf{g}$ acts as an isometry on G. Statement (2) is a consequence of the fact that $\bar{\mathfrak{q}} \circ \mathfrak{q}$ is at distance C_0 from the identity map on G. For (3) we use (2) and the fact that $\mathfrak{q} \circ \bar{\mathfrak{q}}$ is at distance C_0 from the identity map on X.

(4) Let $g = \bar{\mathfrak{q}}(x)$ and $h = \bar{\mathfrak{q}}(y)$. Then $\mathfrak{q}_{hg^{-1}}(x) = \mathfrak{q}(h) = \mathfrak{q}(\bar{\mathfrak{q}}(y))$, which is at distance at most C_0 from y. $\square$

Notation: Let H be a subgroup in G and let $x \in X$. We denote by Hx the set $\{\mathfrak{q}_h(x) \mid h \in H\}$.

Proposition 5.4 and hypothesis (1) imply that there exists $M = M(L, C)$ such that for every $A \in \mathcal{A}$ and every (L, C)-quasi-isometric embedding $\mathfrak{q} : A \to X$ there exists $A' \in \mathcal{A}$ such that $\mathfrak{q}(A) \subset \mathcal{N}_M(A')$.

Lemma 5.12. *(1) If $A, A' \in \mathcal{A}$ satisfy $A \subset \mathcal{N}_r(A')$ for some $r > 0$ then $A = A'$.*

(2) Let $\mathfrak{q} : X \to X$ and $\bar{\mathfrak{q}}$ be (L, C)-quasi-isometries such that $\mathfrak{q} \circ \bar{\mathfrak{q}}$ and $\bar{\mathfrak{q}} \circ \mathfrak{q}$ are at distance at most K from the identity map on X. If $A, A' \in \mathcal{A}$ are such that $\mathfrak{q}(A) \subset \mathcal{N}_r(A')$ or $A' \subset \mathcal{N}_r(\mathfrak{q}(A))$ for some $r > 0$ then $\mathfrak{q}(A) \subset \mathcal{N}_M(A')$, $\bar{\mathfrak{q}}(A') \subset \mathcal{N}_M(A)$ and $\mathrm{hdist}(\mathfrak{q}(A), A')\, , \mathrm{hdist}(\bar{\mathfrak{q}}(A'), A) \leq LM + C + K$.

Proof. (1) follows from property (α_1) and hypothesis (2) of Theorem 5.10.

(2) Suppose $A' \subset \mathcal{N}_r(\mathfrak{q}(A))$. By Proposition 5.4, there exists $\bar{A}$ such that $\mathfrak{q}(A) \subset \mathcal{N}_M(\bar{A})$. Then $A' \subset \mathcal{N}_{r+M}(\bar{A})$, which implies that $A' = \bar{A}$. We may therefore reduce the problem to the case when $\mathfrak{q}(A) \subset \mathcal{N}_r(A')$.

The set $\bar{\mathfrak{q}}(A')$ is contained in $\mathcal{N}_M(A'')$ for some $A'' \in \mathcal{A}$. Also $\bar{\mathfrak{q}} \circ \mathfrak{q}(A) \subset \mathcal{N}_{Lr+C}(\bar{\mathfrak{q}}(A'))$, which implies that $A \subset \mathcal{N}_{Lr+C+M+K}(A'')$. This and (1) imply that $A = A''$. It follows that $\bar{\mathfrak{q}}(A') \subset \mathcal{N}_M(A)$, which implies that $A' \subset \mathcal{N}_{LM+C+K}(\mathfrak{q}(A))$.

Proposition 5.4 implies that there exists $\tilde{A} \in A$ such that $\mathfrak{q}(A) \subset \mathcal{N}_M(\tilde{A})$. Hence $A' \subset \mathcal{N}_{(L+1)M+C+K}(\tilde{A})$, so $A' = \tilde{A}$. We conclude that $\mathfrak{q}(A) \subset \mathcal{N}_M(A')$ and

$$\mathrm{hdist}(\mathfrak{q}(A), A')\, , \mathrm{hdist}(\bar{\mathfrak{q}}(A'), A) \leq LM + C + K.$$

□

Notation: We denote the constant $LM + 2C + C_0$ by D.

Definition 5.13. For every $r > 0$ and every $A \in \mathcal{A}$ we define the *r-stabilizer of A* as

$$\mathrm{St}_r(A) = \{g \in G \mid \mathrm{hdist}(\mathfrak{q}_g(A), A) \leq r\}\, .$$

Corollary 5.14. *(a) For every $g \in G$ and $A, A' \in \mathcal{A}$ such that $\mathfrak{q}_g(A) \subset \mathcal{N}_r(A')$ or $A' \subset \mathcal{N}_r(\mathfrak{q}_g(A))$, where $r > 0$, we have $\mathrm{hdist}(\mathfrak{q}_g(A), A') \leq D$.*

(b) For every $A \in \mathcal{A}$ and for every $r > D$, $\mathrm{St}_r(A) = \mathrm{St}_D(A)$. Consequently $\mathrm{St}_D(A)$ is a subgroup of G.

(c) Let $A, B \in \mathcal{A}$ and $g \in G$ be such that $\mathrm{hdist}(\mathfrak{q}_g(A), B)$ is finite. Then

$$\mathrm{St}_D(A) = g^{-1}\mathrm{St}_D(B)g.$$

Proof. Statement (a) is a reformulation in this particular case of part 2 of Lemma 5.12, and (b) is a consequence of (a).

(c) For every $r > 0$ there exists R large enough so that we have $\mathrm{St}_r(B) \subset g\mathrm{St}_R(A)g^{-1}$.

Applying the previous result again for g^{-1}, B, A, together with (b), we obtain the desired equality. □

Let $\mathcal{F} = \{A_1, \dots, A_k\}$ be the collection of all the sets in $\mathcal{A}$ that intersect $B(x_0, M + C_0)$. We show that this set satisfies (I). Let A be an arbitrary set in $\mathcal{A}$ and let $a \in A$. There exists $g \in G$ such that $\mathfrak{q}_g(a) \in B(x_0, C_0)$, by Lemma 5.11, (4). On the other hand, there exists $A' \in \mathcal{A}$ such that $\mathfrak{q}_g(A) \subset \mathcal{N}_M(A')$. It follows that A' intersects $B(x_0, C_0 + M)$, hence it is in $\mathcal{F}$. Corollary 5.14, (a), implies that $\mathrm{hdist}(\mathfrak{q}_g(A), A') \leq D$, consequently A is quasi-isometric to A'.

For every $i \in \{1, 2, \dots, k\}$ define

$$I(A_i) = \{j \in \{1, 2, \dots, k\} \mid \text{there exists } g \in G \text{ such that } \mathrm{hdist}(\mathfrak{q}_g(A_i), A_j) \leq D\}.$$

For every $j \in I(A_i)$ we fix $g_j \in G$ such that $\mathrm{hdist}(\mathfrak{q}_{g_j} A_i, A_j) \leq D$. Let $\Gamma(A_i) = \{g_j\}_{j \in I(A_i)}$ and let $K(A_i) = \max_{j \in I(A_i)} \mathrm{dist}(g_j \bar{\mathfrak{q}}(x_0), \bar{\mathfrak{q}}(x_0))$.

We define the constant $K = L_0 \max_{i \in \{1,2,\ldots,k\}} K(A_i) + (2L_0+1)\delta_0$, where $\delta_0 = L_0 C_0 + 2C_0$.

The following argument uses an idea from [KaL$_2$, §5.1].

Lemma 5.15. *For every $A \in \mathcal{A}$ the D-stabilizer of A acts K-transitively on A, that is A is contained in the K-tubular neighborhood of every orbit $\mathrm{St}_D(A)a$, where $a \in A$.*

Proof. Let a and b be two arbitrary points in A. Lemma 5.11, (4), implies that there exist $g, \gamma \in G$ such that $\mathfrak{q}_g(a)$, $\mathfrak{q}_\gamma(b) \in B(x_0, C_0)$. This implies that

$$\mathrm{dist}(\mathbf{g} \circ \bar{\mathfrak{q}}(a), \bar{\mathfrak{q}}(x_0)) \leq \delta_0\,,\ \mathrm{dist}(\gamma \circ \bar{\mathfrak{q}}(b), \bar{\mathfrak{q}}(x_0)) \leq \delta_0\,. \tag{10}$$

There exist $i, j \in \{1,2,\ldots,k\}$ such that $\mathrm{hdist}(\mathfrak{q}_g(A), A_i)$, $\mathrm{hdist}(\mathfrak{q}_\gamma(A), A_j) \leq D$. Then $\mathfrak{q}_{\gamma g^{-1}}(A_i)$ is at finite Hausdorff distance from A_j, which implies that $\mathrm{hdist}(\mathfrak{q}_{\gamma g^{-1}}(A_i), A_j) \leq D$ and that $j \in I(A_i)$. Let g_j be such that $\mathrm{hdist}(\mathfrak{q}_{g_j}(A_i), A_j) \leq D$. It follows that $g\gamma^{-1} g_j \in \mathrm{St}_D(A_i)$. The relation $\mathrm{hdist}(\mathfrak{q}_g(A), A_i) \leq D$ and Corollary 5.14, (c), imply that $\gamma^{-1} g_j g \in \mathrm{St}_D(A)$. We have that

$$\mathrm{dist}(\mathfrak{q}_{\gamma^{-1} g_j g}(a), b) \leq L_0 \mathrm{dist}(\gamma^{-1} g_j g \bar{\mathfrak{q}}(a), \bar{\mathfrak{q}}(b)) + C_0 + L_0 C_0 \leq L_0 \mathrm{dist}(g_j g \bar{\mathfrak{q}}(a), \gamma \bar{\mathfrak{q}}(b)) + \delta_0\,.$$

This and inequalities (10) imply that

$$\mathrm{dist}(\mathfrak{q}_{\gamma^{-1} g_j g}(a), b) \leq L_0 \mathrm{dist}(g_j \bar{\mathfrak{q}}(x_0), \bar{\mathfrak{q}}(x_0)) + (2L_0+1)\delta_0 \leq K\,.$$

□

Corollary 5.16. *For every $A \in \mathcal{A}$ the normalizer of $\mathrm{St}_D(A)$ in G is $\mathrm{St}_D(A)$.*

Proof. Let $g \in G$ be such that $\mathrm{St}_D(A) = g^{-1}\mathrm{St}_D(A) g$. Let $B \in \mathcal{A}$ be such that $\mathrm{hdist}(\mathfrak{q}_g(A), B) \leq D$. Corollary 5.14, (c), implies that $\mathrm{St}_D(A) = \mathrm{St}_D(B) = S$. Let $a \in A$ and $b \in B$. We have $\mathrm{hdist}(Sa, Sb) \leq L\mathrm{dist}(a,b) + C$ and also $\mathrm{hdist}(A, Sa) \leq K$ and $\mathrm{hdist}(B, Sb) \leq K$, therefore $\mathrm{hdist}(A, B) \leq 2K + L\mathrm{dist}(a,b) + C$. Lemma 5.12, (1), implies that $B = A$ and $g \in \mathrm{St}_D(A)$. □

Lemma 5.17. *For every $i \in \{1,2,\ldots,m\}$ we have*

$$\mathrm{hdist}(\bar{\mathfrak{q}}(A_i), \mathrm{St}_D(A_i)) \leq \kappa\,,$$

where κ is a constant depending on L_0, C_0, M and $\mathrm{dist}(\mathfrak{q}(1), x_0)$.

Proof. Let $x_i \in A_i \cap B(x_0, M + C_0)$. For every $g \in \mathrm{St}_D(A_i)$ we have $\mathrm{dist}(\mathfrak{q}_g(x_i), A_i) \leq D$, hence $\mathrm{dist}(\mathbf{g} \circ \bar{\mathfrak{q}}(x_i), \bar{\mathfrak{q}}(A_i)) \leq L_0 D + 2C_0$. It follows that $\mathrm{dist}(g, \bar{\mathfrak{q}}(A_i)) \leq L_0 D + 2C_0 + \mathrm{dist}(1, \bar{\mathfrak{q}}(x_i))$. Or $\mathrm{dist}(1, \bar{\mathfrak{q}}(x_i)) \leq L_0 \mathrm{dist}(\mathfrak{q}(1), x_i) + (L_0+1)C_0 \leq L_0 M + (2L_0+1)C_0 + L_0 \mathrm{dist}(\mathfrak{q}(1), x_0)$.

Let $\bar{\mathfrak{q}}(b) \in \bar{\mathfrak{q}}(A_i)$. According to Lemma 5.15, there exists $g \in \mathrm{St}_M(A_i)$ such that

$$\mathrm{dist}(b, \mathfrak{q}_g(x_i)) \leq K.$$

Hence $\mathrm{dist}(\bar{\mathfrak{q}}(b), \mathbf{g} \circ \bar{\mathfrak{q}}(x_i)) \leq L_0 K + 2C_0$ and $\mathrm{dist}(\bar{\mathfrak{q}}(b), g) \leq L_0 K + 2C_0 + \mathrm{dist}(1, \bar{\mathfrak{q}}(x_i))$. □

Corollary 5.18. *Let $A \in \mathcal{A}$. There exists $g \in G$ and $i \in \{1,2,\ldots,m\}$ such that*

$$\mathrm{hdist}(\bar{\mathfrak{q}}(A), g\mathrm{St}_D(A_i)) \leq \kappa + L_0 D + 2C_0.$$

We continue the *proof of Theorem 5.10.* Consider the minimal subset $\{B_1, ..., B_m\}$ of $\{A_1, ..., A_k\}$ such that for each A_i there exists B_{j_i} and γ_i such that $\mathrm{hdist}(A_i, \mathfrak{q}_{\gamma_i}(B_{j_i})) \le D$. Let $\mathcal{B} = \{B_1, \dots, B_m\}$. We denote $S_i = \mathrm{St}_D(B_i)$, $i \in \{1, 2, \dots, m\}$. Let us show that G is asymptotically tree-graded with respect to $S_1, ..., S_m$.

Indeed, by Theorem 5.1, Cayley(G) is asymptotically tree-graded with respect to $\{\bar{\mathfrak{q}}(A), A \in \mathcal{A}\}$. Corollary 5.18 implies that each $\bar{\mathfrak{q}}(A)$ is at uniformly bounded Hausdorff distance from $g\mathrm{St}_D(A_i)$ for some $i \in \{1, 2, \dots, k\}$ and $g \in G$. Corollary 5.14, (c), implies that $\mathrm{St}_D(A_i) = \gamma_i S_{j_i} \gamma_i^{-1}$, with the notations introduced previously. It follows that $\mathrm{hdist}(g\mathrm{St}_D(A_i), g\gamma_i S_{j_i}) \le \max_{i \in \{1,\dots,k\}} \mathrm{dist}(1, \gamma_i^{-1})$. We conclude that $\bar{\mathfrak{q}}(A)$ is at uniformly bounded Hausdorff distance from $g\gamma_i S_{j_i}$. Thus G is asymptotically tree-graded with respect to $S_1, ..., S_m$. □

Corollary 5.19. *Let G be a group that is asymptotically tree-graded with respect to the family of subgroups $\{H_1, H_2, \dots, H_k\}$, where H_i is an infinite group asymptotically without cut-points for every $i \in \{1, 2, \dots, k\}$. Let G' be a finitely generated group which is quasi-isometric to G. Then G' is asymptotically tree-graded with respect to a finite collection of subgroups $\{S_1, \dots, S_m\}$ such that each S_i is quasi-isometric to one of the H_j.*

Remarks 5.20. If the groups H_i in Corollary 5.19 are contained in classes of groups stable with respect to quasi-isometries (for instance the class of virtually nilpotent groups of a fixed degree, some classes of virtually solvable groups) then S_i are in the same classes.

Corollary 5.21. *If a group is tree-graded with respect to a family of subsets $\mathcal{A}$ satisfying conditions (1), (2), (3) in Theorem 5.10, then it is relatively hyperbolic with respect to subgroups $H_1, ..., H_m$ such that every H_i is quasi-isometric to some $A \in \mathcal{A}$.*

Remark 5.22. (a) If in Theorem 5.10 we have that the cardinal of $\mathcal{A}$ is at least two then for every $i \in \{1, 2, \dots, m\}$, H_i has infinite index in G.

(b) If in Corollary 5.19 we have $\{H_1, \dots, H_k\} \ne \{G\}$ then for every $j \in \{1, 2, \dots, m\}$, S_j has infinite index in G'.

Proof. (a) Suppose that $\{H_1, \dots, H_k\} = \{G\}$. According to the proof of Theorem 5.10, it follows that $G = \mathrm{St}_D(B)$ for some $B \in \mathcal{A}$. Lemma 5.17 then implies that $\mathrm{hdist}(\bar{\mathfrak{q}}(B), G) \le \kappa$, whence $\mathrm{hdist}(B, X) \le 3C_0 + L_0\kappa$. This contradicts the property (α_1) satisfied by $\mathcal{A}$.

Therefore $\{H_1, \dots, H_k\} \ne \{G\}$. Now the statement follows from Remark 5.7.

Statement (b) follows from (a). □

6 Cut-points in asymptotic cones of groups

Theorem 5.10 shows that we need to study groups which are asymptotically without cut-points. The next two theorems provide examples of such groups.

6.1 Groups with central infinite cyclic subgroups

Throughout this section we consider G to be a finitely generated group containing a central infinite cyclic subgroup $H = \langle a \rangle$. We fix a finite set of generators X of G and the corresponding word metric on G.

Lemma 6.1. *For every asymptotic cone* $\mathrm{Con}^\omega(G;e,d)$ *of* G *and every* $\epsilon > 0$*, there exists an element* $h = (h_n)^\omega$ *in* $G_e^\omega \cap H^\omega$ *which acts isometrically on* $\mathrm{Con}^\omega(G;e,d)$*, such that for every* $x \in \mathrm{Con}^\omega(G;e,d)$*,* $\mathrm{dist}(hx,x) = \epsilon$*.*

Proof. Let w be a word in X representing a in G. It is obvious that for every $r > 0$ there exists $h = a^n \in H$ such that $|h|$ is in the interval $[r - |w|, r + |w|]$. For every $n \geq 1$ we consider $h_n \in H$ such that $|h_n| \in [\epsilon d_n - |w| \, , \, \epsilon d_n + |w|]$. According to Remark 3.17, the element $h = (h_n)^\omega$ in G_e^ω acts as an isometry on $\mathrm{Con}^\omega(G;e,d)$. Moreover, for every $g = \lim^\omega(g_n) \in \mathrm{Con}^\omega(G;e,d)$ we have that $\mathrm{dist}(hg,g) = \lim_\omega \frac{\mathrm{dist}(h_n g_n \, , g_n)}{d_n} = \lim_\omega \frac{\mathrm{dist}(g_n h_n \, , g_n)}{d_n} \lim_\omega \frac{|h_n|}{d_n} = \epsilon$. □

Lemma 6.2. *If an asymptotic cone* C *of* G *has a cut-point then* C *is isometric to a point or a (real) line.*

Proof. Let $C = \mathrm{Con}^\omega(G;e,d)$ be an asymptotic cone that has a cut-point, where $e = (1), d = (d_n)$. Let h in $G_e^\omega \cap H^\omega$ be as in Lemma 6.1 for $\epsilon = 1$. Lemma 2.28 implies that the asymptotic cone C is tree-graded with respect to a collection $\mathcal{P}$ of pieces that are either points or geodesic sets without cut-points.

If all sets in $\mathcal{P}$ are points then C is an $\mathbb{R}$-tree. If this tree contains a vertex of degree > 2, then it does not admit an isometry h such that $\mathrm{dist}(h(x),x) = 1$ for every x. Thus in this case C is isometric to $\mathbb{R}$ or to a point.

So we may suppose that $\mathcal{P}$ contains pieces that are not points. Let M be such a piece.

Case I. Suppose $h(M) = M$. Let x be an arbitrary point in M. By Lemma 2.28, part (c), there exists $y \in C \setminus M$ such that x is the projection of y on M. Let $\delta = \mathrm{dist}(x,y)$. Since h acts as an isometry, it follows that $y' = h(y)$ projects on M in $x' = h(x)$ and that $\delta \mathrm{dist}(x',y')$. We have $\mathrm{dist}(x,x') = \mathrm{dist}(y,y') = 1$. On the other hand Lemma 2.26 implies that $[y,x] \cup [x,x'] \cup [x',y']$ is a geodesic. Consequently $\mathrm{dist}(y,y') = 1 + 2\delta$, a contradiction.

Case II. Suppose $h(M) \neq M$. Then $h(M)$ is another piece of the tree-graded space C, by Proposition 2.14. Let x be the projection of $h(M)$ on M and let y be the projection of M on $h(M)$. Let $z \in M \setminus \{x\}$ and $z' = h(z)$. By moving z a little, for instance along the geodesic $[z,x]$, we can ensure that $z' \neq y$. Every geodesic joining z and z' contains x and y, by Lemma 2.4. Let t be a point in $C \setminus M$ that projects on M in z (it exists by Lemma 2.28, part (c)). The projection of $t' = h(t)$ onto $h(M)$ is then z'. Lemma 2.26 implies that $[t,z] \cup [z,x] \cup [x,y] \cup [y,z'] \cup [z',t']$ is a geodesic, whence $\mathrm{dist}(t,t') = 1 + 2\mathrm{dist}(t,z)$. This contradicts the fact that $\mathrm{dist}(t,t') = 1$. □

Theorem 6.3. *Let* G *be a non-virtually cyclic finitely generated group that has a central infinite cyclic subgroup* H*. Then* G *is asymptotically without cut-point.*

Proof. By contradiction suppose that G is not asymptotically without cut-point. By Lemma 6.2, one of the asymptotic cones $C = \mathrm{Con}^\omega(G;e,d)$ of G is a line or a point. If C is a point then G is finite and we are done. So assume that C is isometric to $\mathbb{R}$. We identify C with $\mathbb{R}$. The subset $C_H = \lim^\omega(H)$ of C is closed, by Remark 3.10. Therefore, if C_H is dense in C then $C_H = C$. Suppose that there exists an interval (a,b) in C which does not intersect C_H. Lemma 6.1 implies that for every $\epsilon > 0$, C_H contains either ϵ or $-\epsilon$. Consequently, C_H contains the interval $(-b,-a)$. This and the fact that the group $G_e^\omega \cap H^\omega$ of isometries of C acts transitively on C_H implies that C_H is an open set. Since it is also a closed set and C is connected, it follows that $C = C_H$.

Thus, we conclude that $C = \lim^\omega(H)$. Let A be the set of minimal length representatives of the cosets over H. For every a from A let w_a be a shortest word in the alphabet X of generators of G representing a in G. Since G/H is infinite, and X is finite, $\{w_a \mid a \in A\}$ is an infinite set of words in a finite alphabet. Therefore there exists an infinite (in one direction) word β in the alphabet X such that every finite prefix of β is a prefix of one of the words w_a (this is the standard compactness principle). For every $n > 1$ let u_n be the prefix of β of length d_n. Since u_n is a prefix of some w_a, it cannot be equal to a shorter word in G. Let u be the ω-limit of the sequence (u_n) in C. There exists a sequence (h_n) of elements of H such that $\lim^\omega(h_n) = u$. Therefore $|h_n^{-1}u_n| = o(d_n)$. Hence ω-almost surely $u_n = v_n h_n$, where $|v_n| < |u_n|$, $n \in \mathbb{N}$. Then $w_a = v_n h_n s$ for some $a \in A$ and some suffix s of w_a. Since h_n is in the center of G, we have $w_a = v_n s h_n$, so $v_n s$ is a representative of the same coset over H as a. Since $v_n s$ is shorter than w_a (ω-almost surely), we get a contradiction with the choice of A. □

Corollary 6.4. *Let G be a group that is asymptotically tree-graded with respect to certain proper subgroups. Then every finitely generated subgroup in the center $Z(G)$ is finite.*

6.2 Groups satisfying a law

Proposition 6.5. *Let $\mathbb{F}$ be a tree-graded space with respect to a collection $\mathcal{P}$ of proper subsets, $\mathbb{F}$ not isometric to $\mathbb{R}$, and let G be a group acting transitively on $\mathbb{F}$. Then G contains a free subgroup.*

Proof. By Lemma 2.28 we can assume that every piece in $\mathcal{P}$ is either a point or does not have a cut-point. If all pieces are points then $\mathbb{F}$ is an $\mathbb{R}$-tree which has a point of degree ≥ 3. Then G contains a non-Abelian free subgroup by [Chis, Proposition 3.7 on page 111].

Hence we can assume that $\mathcal{P}$ contains a non-singleton piece M.

Lemma 6.6. *Let a and b be two distinct points in M. There exists an isometry $g \in G$ such that the following property holds:*

- *$a \neq g(b)$, the projection of $g(M)$ onto M is a and the projection of M onto $g(M)$ is $g(b)$.*

We shall denote this property of g by $P(a, b, M)$.

Proof. There are two cases:

(A) There exist two distinct pieces in $\mathcal{P}$ that intersect.

(B) Any two distinct pieces in $\mathcal{P}$ are disjoint.

By homogeneity, in case (A), every point is contained in two distinct pieces. In case (B) let x, y be two distinct points in M. There exists an isometry $g \in G$ such that $g(x) = y$. Since $g(M)$ intersects M in y it follows that $g(M) = M$. We conclude that in this case the stabilizer of M in G acts transitively on M.

Suppose we are in case (A). Then we can construct a geodesic $\mathfrak{g}\colon [0, s] \to \mathbb{F}$ such that $s = \Sigma_{i=1}^{\infty} s_n$ with $0 < s_n < \frac{1}{n^2}$ and $\mathfrak{g}\left[\Sigma_{i=0}^{n} s_i, \Sigma_{i=0}^{n+1} s_i\right] \subset M_n$ for some pieces M_n, where $M_n \neq M_{n+1}$ for all $n \in \mathbb{N} \cup \{0\}$. Here $s_0 = 0$. Such a geodesic exists by Lemma 2.26. We call such a geodesic *fractal at the arrival point*. By gluing together two geodesics fractal at their respective arrival points, $\mathfrak{g} \cup \mathfrak{g}'$, and making sure that the two respective initial pieces, M_0 and M_0', are distinct, we obtain a geodesic *fractal at the departure and arrival points* or *bifractal*. By homogeneity, every point in $\mathbb{F}$ is the endpoint of a bifractal geodesic.

Let $[a, c]$ be a bifractal geodesic. Corollary 2.8, (b), implies that $[a, c]$ can intersect M in a or in a non-trivial sub-geodesic $[a, c']$. Since $[a, c]$ is fractal at the departure point the latter case cannot occur. It follows that the intersection of $[a, c]$ and M is $\{a\}$. There exists an isometry $g \in G$ such that $g(b) = c$. Since $[a, c]$ is fractal at the arrival point also, it follows that $[a, c] \cap g(M) = \{c\}$. For every $x \in g(M)$ we have that $[a, c] \cup [c, x]$ is a geodesic, by Lemma 2.26.

In particular a is the projection of $g(M)$ on M. A symmetric argument gives that $c = g(b)$ is the projection of M on $g(M)$.

Now suppose that case (B) holds. Lemma 2.28, part (c), implies that a is the projection of a point $x \in \mathbb{F} \setminus M$. Let g be an isometry in G such that $g(b) = x$. If $[a, x]$ intersects $g(M)$ in x then we repeat the previous argument. Assume $[a, x] \cap g(M) = [x', x]$. By the hypothesis in case (B), $x' \neq a$. We have $x' = g(b')$ for some $b' \in M$. Since the stabilizer of M in G acts transitively on M, there exists g' in it such that $g'(b) = b'$. We have that $gg'(M) = g(M)$ projects onto M in a and M projects onto $gg'(M)$ in $x' = gg'(b)$. □

Notation: For every $t \in M$ let $\Pi_t(M)$ be the set of points x in $\mathbb{F} \setminus M$ that project onto M in t.

Lemma 6.7. *Let g satisfy property $P(a, b, M)$. Then:*

(a) the isometry g^{-1} satisfies property $P(b, a, M)$;

(b) for every $t \neq b$ we have $g(\Pi_t(M)) \subset \Pi_a(M)$.

Proof. (a) We apply the isometry g^{-1} to the situation in $P(a, b, M)$.

(b) The set $g(\Pi_t(M))$ projects on $g(M)$ in $g(t) \neq g(b)$. This, property $P(a, b, M)$ and Corollary 2.27 imply that $g(\Pi_t(M))$ projects onto M in a and that $\mathrm{dist}(g(\Pi_t(M)), M) \geq \mathrm{dist}(g(M), M) > 0$. □

We now finish the *proof of Proposition 6.5.* Let a, b, c, d be four pairwise distinct elements in M. Lemma 6.6 implies that there exist $g \in G$ satisfying $P(a, b, M)$ and h satisfying $P(c, d, M)$.

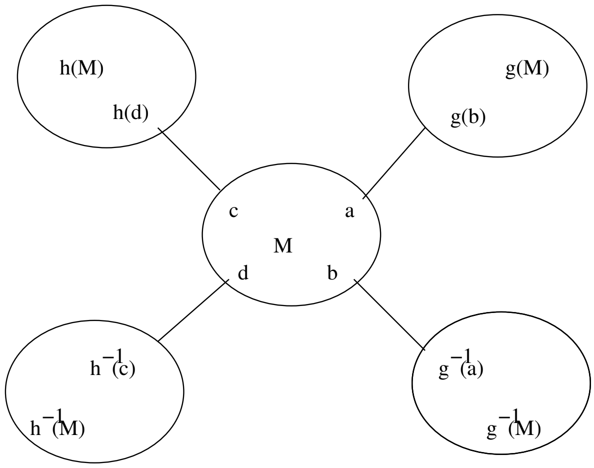

Figure 5: Action of the elements g, g^{-1}, h, h^{-1}.

Then g^{-1} is satisfying $P(b, a, M)$ and h^{-1} is satisfying $P(d, c, M)$ by Lemma 6.7. In particular $g(M) \subset \Pi_a(M)$, $g^{-1}(M) \subset \Pi_b(M)$, $h(M) \subset \Pi_c(M)$, $h^{-1}(M) \subset \Pi_d(M)$.

Since $b \notin \{a, c, d\}$, Lemma 6.7, part (b), implies that $g(\Pi_a(M) \cup \Pi_c(M) \cup \Pi_d(M)) \subset \Pi_a(M)$. The isometries g^{-1}, h, h^{-1} satisfy similar properties. The Tits ping-pong argument allows to conclude that g and h generate a free group. □

Theorem 6.8. *Let $\mathcal{G}$ be a family of finitely generated groups satisfying a law. Then $\mathcal{G}$ is uniformly asymptotically without cut-points.*

Proof. Suppose that an ultralimit $\lim^{\omega}(G_n, \frac{1}{d_n}\mathrm{dist}_n)_e$ has a cut-point, where $\lim_{\omega} d_n = \infty$. The ultrapower $\Pi_e(G_n, \frac{1}{d_n}\mathrm{dist}_n)/\omega$ satisfies the same law as all G_n. On the other hand, the group $\Pi_e(G_n, \frac{1}{d_n}\mathrm{dist}_n)/\omega$ acts transitively on $\lim^{\omega}(G_n, \frac{1}{d_n}\mathrm{dist}_n)_e$. Lemma 2.28 and Proposition 6.5 imply that the group $\Pi_e(G_n, \frac{1}{d_n}\mathrm{dist}_n)/\omega$ contains a free subgroup, a contradiction. □

Examples: Solvable groups of a given degree, Burnside groups of a fixed order are classes of groups satisfying a law.

Corollary 6.9. *Let G be a finitely generated group satisfying a law. Then no asymptotic cone of G has a cut-point.*

Corollary 6.10. *Let X be an asymptotically tree-graded space with respect to a collection of subsets $\mathcal{A}$. For every non-trivial group law and every (L, C) there exists a constant M depending on (L, C) and on the law such that every (L,C)-quasi-isometric embedding of a finitely generated group satisfying the law into X has the image in $\mathcal{N}_M(A)$ for some $A \in \mathcal{A}$.*

7 Fundamental groups of asymptotic cones

In [EO], A. Erschler and D. Osin constructed (modifying an idea from [Ols$_2$]), for every "sufficiently good" metric space M, a two-generated group G with the property that M π_1-embeds isometrically into an asymptotic cone $\mathrm{Con}^{\omega}(G)$. Thus any countable group is a subgroup of the fundamental group of some asymptotic cone of a finitely generated group. In this section we modify, in turn, the construction from [EO] to show that the fundamental group of an asymptotic cone can be isomorphic to the uncountable free power of any countable group. Moreover, that asymptotic cone can be completely described as a tree-graded space. In particular, if, say, M is compact and locally contractible then there exists a 2-generated group one of whose asymptotic cones is tree-graded with respect to pieces isometric to M. We also construct a 2-generated recursively presented group with the maximal possible (under the continuum hypothesis) number of non-homeomorphic asymptotic cones.

7.1 Preliminaries on nets

Let (X, dist) be a metric space. We recall some notions and results from [GLP].

Definition 7.1. A *δ-separated set* A in X is a set such that for every $x_1, x_2 \in A$, $\mathrm{dist}(x_1, x_2) \geq \delta$. A *$\delta$-net* in X is a set B such that $X \in \mathcal{N}_{\delta}(B)$.

Remark 7.2. A maximal δ-separated set in X is a δ-net in X.

Proof. Let N be a maximal δ-separated set in X. For every $x \in X \setminus N$, the set $N \cup \{x\}$ is no longer δ-separated, by maximality of N. Hence there exists $y \in N$ such that $\mathrm{dist}(x, y) < \delta$. □

Definition 7.3. We call a maximal δ-separated set in X a *δ-snet*.

We note that if X is compact then every snet is finite, hence every separated set is finite.

Remark 7.4. Let (X, dist) be a metric space and let $(M_n)_{n\in\mathbb{N}}$ be an increasing sequence of subsets of X. Let $(\delta_n)_{n\in\mathbb{N}}$ be a decreasing sequence of positive numbers converging to zero. There exists an increasing sequence

$$N_1 \subset N_2 \subset \cdots \subset N_n \subset \cdots ,$$

such that N_n is a δ_n-snet in (M_n, dist).

Proof. There exists a δ_1-snet in M_1, which we denote N_1. It is a δ_1-separated set in M_2. Let N_2 be a δ_2-snet in M_2 containing N_1. Then N_2 is a δ_2-separated set in M_3. Inductively we construct an increasing sequence $(N_n)_{n\in\mathbb{N}}$. $\square$

Notation: Let A be a subset in a metric space. We denote by $\Gamma_\kappa(A)$ the metric graph with set of vertices A and set of edges

$$\{(a_1, a_2) \mid a_1, a_2 \in A,\ 0 < \mathrm{dist}(a_1, a_2) \leq \kappa\}\,,$$

such that the edge (a_1, a_2) is of length $\mathrm{dist}(a_1, a_2)$. We shall denote the length of every edge e by $|e|$. We endow $\Gamma_\kappa(A)$ with its length metric.

Notation: Let (X, dist) be a proper geodesic metric space, let O be a fixed point in it and let $\zeta \in (0,1)$. We denote by $B_n = \overline{B(O,n)}$ the closed ball of radius n around O. We consider an increasing sequence of subsets in X,

$$\{O\} \subset N_1 \subset N_2 \subset \cdots \subset N_n \subset \cdots\,,$$

such that N_n is an ζ^n-snet in B_n. Let Γ_n be the finite graph $\Gamma_{\zeta^{[n/2]}}(N_n)$, endowed with its length metric dist_n (here $[\frac{n}{2}]$ is the integer part of $\frac{n}{2}$).

We recall that two metric spaces with fixed basepoints (X, dist_X, x) and (Y, dist_Y, y) are said to be *isometric* if there exists an isometry $\phi : X \to Y$ such that $\phi(x) = y$.

Lemma 7.5. *In the notation as above:*

1. *for every $n \geq 2$, for every $x, y \in N_n$ we have*

$$\mathrm{dist}(x,y) \leq \mathrm{dist}_n(x,y) \leq \left(1 + 6\zeta^k\right)\left(\mathrm{dist}(x,y) + 2\zeta^k\right) + 2\zeta^k\,, \tag{11}$$

   where $k = [\frac{n}{2}]$;

2. *for every observation point $e \in \Pi N_n/\omega$, the spaces $\lim^\omega(N_n, \mathrm{dist}_n)_e$, $\lim^\omega(\Gamma_n, \mathrm{dist}_n)_e$ and $\lim^\omega(B_n, \mathrm{dist})_e$ with the basepoints $\lim^\omega(e)$ are isometric.*

3. *the spaces $\lim^\omega(N_n, \mathrm{dist}_n)$, $\lim^\omega(\Gamma_n, \mathrm{dist}_n)$ with the basepoints $\lim^\omega(O)$ and (X, dist) with the basepoint O are isometric.*

Proof. (1) Let x, y be two fixed points in N_n. If $\mathrm{dist}(x,y) \leq \zeta^{[n/2]}$ then by construction $\mathrm{dist}(x,y) = \mathrm{dist}_n(x,y)$ and both inequalities in (11) are true. Let us suppose that $\mathrm{dist}(x,y) > \zeta^{[n/2]}$.

The distance $\mathrm{dist}_n(x,y)$ in Γ_n is the length of some path composed of the edges $e_1e_2...e_s$, where $x = (e_1)_-$ and $y = (e_s)_+$. It follows that

$$\mathrm{dist}_n(x,y) = \sum_{i=1}^{s} |e_i| \geq \mathrm{dist}(x,y).$$

We conclude that

$$\mathrm{dist}_n(x,y) \geq \mathrm{dist}(x,y)\,.$$

We also note that

$$\mathrm{dist}_n(x,y) \geq \mathrm{dist}_m(x,y) \text{ for every } m \geq n\,, \tag{12}$$

since $N_n \subseteq N_m$.

The distance $\mathrm{dist}(x,y)$ is the length of a geodesic $\mathfrak{c}\colon [0, \mathrm{dist}(x,y)] \to X$. Since $x, y \in N_n \subset \overline{B(O,n)}$, the image of this geodesic is entirely contained in $\overline{B(O,2n)}$. Let $t_0 = 0, t_1, t_2, \dots, t_m = \mathrm{dist}(x,y)$ be a sequence of numbers in $[0, \mathrm{dist}(x,y)]$ such that $0 < t_{i+1} - t_i \leq \frac{\zeta^n}{2}$, for every $i \in \{1, 2, \dots, m-1\}$ and $m \leq \frac{2\mathrm{dist}(x,y)}{\zeta^n} + 1$. Since $\mathrm{dist}(x,y) > \zeta^{[n/2]} > \zeta^n$, we can write $m \leq \frac{3\mathrm{dist}(x,y)}{\zeta^n}$. Let $x_i = \mathfrak{c}(t_i)$, $i \in \{0, 1, 2, \dots, m\}$. For every $i \in \{0, 1, 2, \dots, m\}$ there exists $w_i \in N_{2n}$ such that $\mathrm{dist}(x_i, w_i) \leq \zeta^{2n}$. We note that $w_0 = x, w_m = y$. We can write

$$\mathrm{dist}(x,y) = \sum_{i=0}^{m-1} \mathrm{dist}(x_i, x_{i+1}) \geq \sum_{i=0}^{m-1} [\mathrm{dist}(w_i, w_{i+1}) - 2\zeta^{2n}] \sum_{i=0}^{m-1} \mathrm{dist}(w_i, w_{i+1}) - 2m\zeta^{2n} . \quad (13)$$

We have $\mathrm{dist}(w_i, w_{i+1}) \leq \mathrm{dist}(x_i, w_i) + \mathrm{dist}(x_i, x_{i+1}) + \mathrm{dist}(x_{i+1}, w_{i+1}) \leq 2\zeta^{2n} + \frac{\zeta^n}{2} \leq \zeta^n$ for n large enough. Therefore w_i, w_{i+1} are connected in Γ_{2n} by an edge of length $\mathrm{dist}(w_i, w_{i+1})$. We conclude that

$$\sum_{i=0}^{m-1} \mathrm{dist}(w_i, w_{i+1}) = \sum_{i=0}^{m-1} \mathrm{dist}_{2n}(w_i, w_{i+1}) \geq \mathrm{dist}_{2n}(w_0, w_m) = \mathrm{dist}_{2n}(x,y).$$

This and (13) implies that

$$\mathrm{dist}(x,y) \geq \mathrm{dist}_{2n}(x,y) - 6\mathrm{dist}(x,y)\zeta^n .$$

We have obtained that

$$\frac{1}{1+6\zeta^n}\mathrm{dist}_{2n}(x,y) \leq \mathrm{dist}(x,y) \leq \mathrm{dist}_n(x,y), \text{ for all } x, y \in N_n. \quad (14)$$

Let again x, y be two points in N_n, $k = [n/2]$. There exist $x', y' \in N_k \subset N_n$ such that $\mathrm{dist}(x,x'), \mathrm{dist}(y,y') \leq \zeta^k$. This implies that $\mathrm{dist}(x,x') = \mathrm{dist}_n(x,x') \leq \zeta^k$ and likewise $\mathrm{dist}(y,y') = \mathrm{dist}_n(y,y') \leq \zeta^k$. Hence $\mathrm{dist}_n(x,y) \leq \mathrm{dist}_n(x',y') + 2\zeta^k$.

Inequalities (12) and (14) imply

$$\mathrm{dist}_n(x',y') \leq \mathrm{dist}_{2k}(x',y') \leq (1+6\zeta^k)\mathrm{dist}(x',y') \leq (1+6\zeta^k)(\mathrm{dist}(x,y) + 2\zeta^k).$$

This gives (11).

(2) We have $N_n \subset \Gamma_n \subset \mathcal{N}_{\zeta^{[n/2]}}(N_n)$. Therefore $\lim^\omega (\Gamma_n, \mathrm{dist}_n)_e = \lim^\omega (N_n, \mathrm{dist}_n)_e$. Thus it is enough to prove that $\lim^\omega (N_n, \mathrm{dist}_n)_e$ and $\lim^\omega (B_n, \mathrm{dist})_e$ with the basepoints $\lim^\omega(e)$ are isometric.

We define the map

$$\Psi\colon \lim{}^\omega(x_n) \mapsto \lim{}^\omega(x_n) \quad (15)$$

from $\lim^\omega (N_n, \mathrm{dist}_n)_e$ to $\lim^\omega (B_n, \mathrm{dist})_e$. Inequalities (11) imply that the map Ψ is well defined and that it is an isometric embedding.

We prove that Ψ is surjective. Let $(y_n)^\omega \in \Pi_e B_n/\omega$. For every y_n there exists $x_n \in N_n$ such that $\mathrm{dist}(x_n, y_n) \leq \zeta^n$. Since the sequence $(\mathrm{dist}(y_n, e_n))$ is bounded, the sequence $(\mathrm{dist}(x_n, e_n))$ is also bounded by the second inequality in (11), and so is the sequence $(\mathrm{dist}_n(x_n, e_n))$. We have $\lim^\omega(x_n) \in \lim^\omega (N_n, \mathrm{dist}_n)_e$ and $\Psi(\lim^\omega(x_n)) = \lim^\omega(x_n)$. As $\lim_\omega \mathrm{dist}(x_n, y_n) = 0$ we conclude that $\lim^\omega(x_n) = \lim^\omega(y_n)$.

(3) According to (2) it suffices to prove that $\lim^\omega(B_n, \text{dist})_O$ with the basepoint $\lim^\omega(O)$ and X with the basepoint O are isometric. Let $x \in X$. For n large enough, $x \in \overline{B(O,n)}$. We define the map

$$\Phi \colon x \mapsto \lim^\omega(x) \tag{16}$$

from X to $\lim^\omega(B_n)_O$.

The map Φ is clearly an isometric embedding. Let us show that Φ is surjective. Let $(x_n)_{n\in\mathbb{N}}$ be such that $x_n \in B_n$ and such that $\text{dist}(O, x_n)$ is uniformly bounded by a constant C. It follows that $x_n \in \overline{B(O,C)}$ for all $n \in \mathbb{N}$. Since the space X is proper, $\overline{B(O,C)}$ is compact and there exists an ω-limit x of (x_n). It follows that $\lim_\omega \text{dist}(x_n, x) = 0$, which implies that $\lim^\omega(x_n) = \lim^\omega(x) = \Phi(x)$. □

Notation: We shall denote the point $\lim^\omega(O)$ also by O. This should not cause any confusion.

Remark 7.6. The hypothesis that X is proper is essential for the surjectivity of Φ in the proof of part (3) of Lemma 7.5.

Definition 7.7. For every proper geodesic metric space (X, dist) with a fixed basepoint O, and every sequence of points $e = (e_n)^\omega$, $e_n \in B_n = \overline{B(O,n)}$, we shall call the limit $\lim^\omega(B_n)_e$ an *ultraball* of X with center O and observation point e.

Remark 7.8. Notice that the ultraballs $\lim^\omega(B_n)_e$ and $\lim^\omega(B_n)_{e'}$ with observation points $e = (e_n)^\omega$ and $e' = (e'_n)^\omega$, such that $\text{dist}(e_n, e'_n)$ is uniformly bounded ω-almost surely, are the same spaces with different basepoints (see Remark 3.7).

Remark 7.9. It is easy to prove, using results from [BGS, §I.3] and [KaL$_1$], that an ultraball of a complete homogeneous locally compact CAT(0)-space is either the whole space or a horoball in it (for a definition see [BrH]). In particular the ultraballs of the Euclidean space $\mathbb{R}^n$ are $\mathbb{R}^n$ itself and all its half-spaces.

We are now going to construct a proper geodesic metric space with basepoint (Y_C, dist, O) whose fundamental group is any prescribed countable group C, and such that every ultraball with center O of Y_C either is isometric to the space Y_C itself or is simply connected.

Let $C = \langle S \mid R \rangle$ be a countable group. We assume that $S = \{s_n \mid n \in \mathbb{N}\} = C$, and that R is just the multiplication table of C, i.e. that all relations in R are triangular. For every $n \in \mathbb{N}$, consider X_n the part of the cone $z^2 = x^2 + y^2$ in $\mathbb{R}^3$ which is above the plane $z = n-1$. The intersection of this (truncated) cone with the plane $z = n-1$ will be called its *base*. Cut an *auxiliary* circular hole of radius n (which is the radius of the base of X_{n+1}) at the height n of X_n, that is at the height $2n$ of the initial cone. The resulting space is denoted by Y_n. The vertex of Y_1 is denoted by O.

Now consider the following construction. We start with the space Y_1, glue in the space Y_2 so that the base hole of Y_2 is isometrically identified with the auxiliary hole in Y_1, glue in Y_3 so that the base hole in Y_3 is identified with the auxiliary hole in Y_2, etc. The resulting space with the natural gluing metric is denoted by Y. Now enumerate all relations in $R = \{r_1, r_2, ...\}$. For every $m = 1, 2, ..., r_m$ has the form $x_i x_j x_k^{-1}$. Choose a natural number $k = k(m)$ such that the base holes of Y_i, Y_j, Y_k are at the distance $\leq k$ in Y and such that $k(m) > k(m-1)$. Consider the circles y_i, y_j, y_k obtained by cutting Y_i, Y_j, Y_k by planes parallel to the base hole at distance k from O, connect these circles with O by geodesics. Glue in an Euclidean disc D_n to the circles y_i, y_j, y_k and connecting geodesics such that the boundary ∂D_n is glued according to the relation r_m. We supply the resulting space Y_C with the natural geodesic metric dist.

We keep the above notation for balls $B_n = \overline{B(O,n)}$, and metric spaces N_n and Γ_n for this space Y_C.

The following properties of the space (Y_C, dist) are obvious.

Lemma 7.10. *(1) The space Y_C is geodesic and proper.*

(2) For every $d > 0$ there exists a number $r > 0$ such that every ball of radius d in Y_C, whose center is outside $B(O,r)$, is contractible.

(3) The fundamental group of Y_C is isomorphic to C.

Lemma 7.11. *The ultraball $\lim^\omega(B_n)_e$ of Y_C with center O is simply connected if $\mathrm{dist}(e_n, O)$ is unbounded ω-almost surely, otherwise it is isometric to Y_C.*

Proof. Indeed, if a point $e = (e_n)$ from X^ω is such that $\mathrm{dist}(e_n, O)$ is bounded ω-almost surely then the corresponding ultraball is isometric to Y_C by Remark 7.8. Suppose that

$$\lim_\omega \mathrm{dist}(e_n, O) = \infty.$$

Let U be the corresponding ultraball. Then every closed ball $\overline{B_U(e,r)}$ in U is the ω-limit of $\overline{B_{Y_C}(e_n,r)} \cap B_n$. By Lemma 7.10, the balls $B_{Y_C}(e_n, r)$ are contractible ω-almost surely. Therefore $\overline{B_U(e,r)}$ is contractible. Since every loop in U is contained in one of the balls $\overline{B_U(e,r)}$, U is simply connected. $\square$

7.2 Construction of the group

Let A be an alphabet and $\mathbb{F}_A$ a free group generated by A. For every $w \in \mathbb{F}_A$ we denote by $|w|$ the length of the word w.

Definition 7.12 (property $C^*(\lambda)$). A set $\mathcal{W}$ of reduced words in $\mathbb{F}_A$, that is closed under cyclic permutations and taking inverses, is said to satisfy *property $C^*(\lambda)$* if the following hold.

(1) If u is a subword in a word $w \in \mathcal{W}$ so that $|u| \geq \lambda |w|$ then u occurs only once in w;

(2) If u is a subword in two distinct words $w_1, w_2 \in \mathcal{W}$ then $|u| \leq \lambda \min(|w_1|, |w_2|)$.

We need the following result from [EO].

Proposition 7.13. *[EO] Let $A = \{a, b\}$. For every $\lambda > 0$ there exists a set $\mathcal{W}$ of reduced words in $\mathbb{F}_A$, closed with respect to cyclic permutations and taking inverses, satisfying the following properties:*

(1) $\mathcal{W}$ satisfies $C^*(\lambda)$;

(2) for every $n \in \mathbb{N}$, the set $\{w \in \mathcal{W} \mid |w| \geq n\}$ satisfies $C^*(\lambda_n)$ with $\lim_{n\to\infty} \lambda_n = 0$;

(3) $\lim_{n\to\infty} \mathrm{card}\{w \in \mathcal{W} \mid |w| = n\} = \infty$.

Notation: Let us fix $\lambda = \frac{1}{500}$, and a set of words $\mathcal{W}$ provided by Proposition 7.13.

Let $\kappa(n) = \mathrm{card}\{w \in \mathcal{W} \mid |w| = n\}$. We have that $\lim_{n\to\infty} \kappa(n) = \infty$.

Fix a number $\zeta \in (0,1)$. For every $n \in \mathbb{N}$, let Γ_n be a finite metric graph with edges of length at least ζ^n and at most $\zeta^{[n/2]}$ and diameter at most $10n$ for n large enough. We endow Γ_n with the length metric dist_n. Let N_n be the set of vertices of Γ_n and let O_n be a fixed vertex in N_n. Let E_n be the number of edges of Γ_n.

Definition 7.14 (fast increasing sequences). An increasing sequence (d_n) of positive numbers is called *fast increasing with respect to the sequence of graphs* (Γ_n) if it satisfies the following:

(1) for every $i \geq [\zeta^n d_n]$, $\kappa(i) \geq E_n$;

(2) $\lim_{n\to\infty} \frac{\zeta^n d_n}{d_{n-1}} = \infty$;

(3) $\lim_{n\to\infty} \frac{E_n}{\zeta^n d_n} = 0$.

Fast increasing sequences of numbers clearly exist.

Let us fix a fast increasing sequence $d = (d_n)$ with respect to the sequence of graphs (Γ_n). To every edge $e = (x, y)$ in Γ_n we attach a word $w_n(e)$ in $\mathcal{W}$ of length $[d_n|e|]$ such that

(1) $w_n(e^{-1}) = w_n(e)^{-1}$;

(2) $w_n(e) \neq w_n(e')$ if $e \neq e'$.

We can choose these words because for every edge $e = (x, y)$ in Γ_n, we have $[d_n \mathrm{dist}(x, y)] \geq [\zeta^n d_n]$ and because we have enough words in $\mathcal{W}$ of any given length (part (1) of Definition 7.14).

Definition 7.15 (the presentation of the group G). We define the set of relations R_n as follows: for every loop $p = e_1 e_2 ... e_s$ in Γ_n we include in R_n the free reduction of the word

$$w_n(p) = w_n(e_1)w_n(e_2)\dots w_n(e_s).$$

Let $R = \bigcup_{n\in\mathbb{N}} R_n$ and let $G = \langle a, b \mid R\rangle$.

Notation: We denote by Cayley(G) the left invariant Cayley graph of G with respect to the presentation $G = \langle a, b \mid R\rangle$, that is the vertices are elements of G and the (oriented) edges are (g, gx) for every $x \in \{a, b, a^{-1}, b^{-1}\}$. The edge (g, gx) in Cayley(G) is usually labeled by x, so Cayley(G) can be viewed as a labeled graph. Every path in $Cay(G)$ is labeled by a word in $\mathbb{F}_{\{a,b\}}$. The length of a path p in Cayley(G) is denoted by $|p|$. The distance function in Cayley(G) is denoted by dist, it coincides with the word metric on G.

Notation: For every word w in the free group $F_{\{a,b\}}$ we denote by g_w the element in G represented by w.

As in [EO] and [Ols$_1$], we introduce the following types of words.

Definitions 7.16 (words of rank n). Every freely reduced product

$$w = w_n(e_1)w_n(e_2)\dots w_n(e_m), \tag{17}$$

where $e_1, ..., e_m$ are edges in Γ_n is called a *word of rank* n. The words $w_n(e_i)$ will be called the *blocks* of w.

Every freely reduced product

$$w_n(p) = w_n(e_1)w_n(e_2)\dots w_n(e_m),$$

where $p = e_1e_2...e_m$ is a path in Γ_n, is called a *net word of rank* n.

Remark 7.17. The words $w_n(e)$ have length at least $[\zeta^n d_n] \geq [d_{n-1}] \geq \frac{d_1}{\zeta^{n-1}} - 1 \geq n$ for n large enough. This and the small cancellation assumptions from Proposition 7.13 imply that at most $2\lambda_n$ of the length of the block $w_n(e)$ can cancel in the product (17) provided none of its neighbor factors is $w_n(e^{-1})$. In particular, if a path p in Γ_n has no backtracking, at most $2\lambda_n$ of the length of any factor $w_n(e)$ cancels in the word $w_n(p)$.

Notation: For every path p in Γ_n starting at O_n let $\bar{p}$ be the path in Cayley(G) labeled by $w_n(p)$ starting at 1. We denote by $\Re_n$ the union of all these paths $\bar{p}$. It is easy to see that $\Re_n$ consists of all prefixes of all net words $w_n(p)$, where p is a path in Γ_n starting at O_n.

Definition 7.18 (cells of rank n). By definition of the set of relations R, the boundary label of every cell in a van Kampen diagram Δ over R is a net word. Therefore a cell in Δ is called a *cell of rank* n if its boundary label is a net word of rank n.

Definition 7.19 (minimal diagrams). A van Kampen diagram over R is called *minimal* if it contains the minimal number of cells among all van Kampen diagrams over R with the same boundary label, and the sum of perimeters of the cells is minimal among all diagrams with the same number of cells and the same boundary label.

Notation: The boundary of any van Kampen diagram (cell) Δ is denoted by $\partial\Delta$.

Lemma 7.20. *(1) Every minimal van Kampen diagram Δ over R satisfies the small cancellation property $C'(1/10)$ (that is, the length of any path contained in the boundaries of any two distinct cells in Δ cannot be bigger than $1/10$ of the length of the boundary of any of these cells).*

(2) Every cell π in a minimal van Kampen diagram Δ over R satisfies $|\partial\pi| \le 2|\partial\Delta|$.

Proof. (1) is Lemma 4.2 in [EO].

(2) We prove the statement by induction on the number n of cells in Δ. If $n = 1$ then the statement is obviously true. Suppose it is true for some n. We consider a minimal van Kampen diagram Δ with $n+1$ cells. By Greendlinger's lemma [LS] and Part (1) there exists a cell π and a common path p of $\partial\pi$ and $\partial\Delta$ whose length is bigger than $\frac{7}{10}|\partial\pi|$. It follows that $|\partial\pi| \le 2|\partial\Delta|$. Removing p and the interior of π, we obtain a minimal diagram Δ' with boundary length smaller than $|\partial\Delta|$ and with fewer cells than Δ. It remains to apply the induction assumption to Δ'. □

Notation: We shall denote the graphical equality of words by $\equiv$.

Lemma 7.21. *Let $u \equiv u_1u_2u_3$ be a word of rank n and $u' \equiv u_1'u_2u_3'$ be a word of rank m, $n \ge m$. Suppose $|u_2|$ is at least 5λ times the maximal length of a block in u'. Then $m = n$. In addition, if $u = w_n(p)$ and $u' = w_n(q)$ are net words then the paths p and q in Γ_n have a common edge e: $p = p_1ep_2$, $q = q_1eq_2$, and u_1 (resp. u_1') is a prefix of $w_n(p_1e)$ (resp. $w_n(q_1e)$), u_3 (resp. u_3') is a suffix of $w_n(ep_2)$ (resp. $w_n(eq_2)$).*

Proof. Indeed, the conditions of the lemma imply that one of the blocks of u that either contains u_2 or is contained in u_2 has in common with one of the blocks of u' at least λ of its length. The small cancellation condition $C^*(\lambda)$ implies that the blocks coincide, so $m = n$. The rest of the statement follows immediately from the definition of net words and Remark 7.17. □

Lemma 7.22. *Let u and v be two words in $\{a, b\}$ that are equal in G. Suppose that u is a (net) word of rank n and v is a shortest word that is equal to u in G. Then v is also a (net) word of rank n. In addition, if u is a net word, $u = w_n(p)$, then $v = w_n(q)$ for some simple path q in Γ_n having the same initial and terminal vertices as p.*

Proof. Consider a van Kampen diagram Δ over R with boundary $\partial\Delta = st$ where u labels s, v^{-1} labels t.

By Greendlinger lemma, property $C'(1/10)$ implies that there exists a cell π in Δ such that $\partial\pi$ and $\partial\Delta$ have a common subpath r of length $\frac{7}{10}|\partial\pi|$. Since v is a shortest word that is equal to u in G, no more than $\frac{1}{2}$ of $\partial\pi$ is a subpath of t. Therefore $|r \cap s| \ge \frac{1}{5}|\partial\pi|$. Notice that the label of $\partial\pi$ is the reduced form of a product of at least two blocks. Therefore the label of $r \cap s$

contains at least $(1-4\lambda)/5$ of a block in $\partial\pi$. Lemma 7.21 implies that π is a cell of rank n. After we remove the cell π from Δ we obtain a diagram Δ' corresponding to an equality $u'=v$ of the same type as $u=v$, that is u' is a word of rank n representing the same element in G as u and v, and if $u=w_n(p)$ then $u'=w_n(p')$, where p' is a path in Γ_n with $p'_-=p_-, p'_+=p_+$. Since Δ' has fewer cells than Δ, it remains to use induction on the number of cells in Δ. □

7.3 Tree-graded asymptotic cones

Recall that we consider any sequence of metric graphs Γ_n, $n\geq 1$, satisfying the properties listed before Definition 7.14, that the set of vertices of Γ_n is denoted by N_n, and that we fix basepoints O_n in N_n. For every $x\in N_n$ let p_x be a path from O_n to x in Γ_n. We define

$$\Phi_n\colon N_n\to \Re_n\,,\ \Phi_n(x)=w_n(p_x) \text{ in } G$$

(see notation before Definition 7.18).

The value $\Phi_n(x)$ does not depend on the choice of the path p_x, because $w_n(q)$ is equal to 1 in G for every loop q in Γ_n by the definition of the presentation of G. Hence Φ_n is a map.

Remark 7.23. Notice that every point in $\Re_n$ is at distance at most $\zeta^{[n/2]}d_n(1+\lambda_n)$ from $\Phi_n(N_n)$.

The sequence of maps (Φ_n) clearly defines a map

$$(x_n)^\omega\mapsto(\Phi_n(x_n))^\omega\,.$$

from $\Pi N_n/\omega$ to $\Pi\Re_n/\omega$.

Remark 7.24. Let $a=\Phi_n(x)\,,\ x\in N_n$, and let $b\in G$ such that a and b can be joined in Cayley(G) by a path labeled by $w_n(q)$, where q is a path in Γ_n with $q_-=x$ and $q_+=y$. Then $b=\Phi_n(y)\in\Phi_n(N_n)$.

Lemma 7.25. *Let $e=(e_n)^\omega\in\Pi N_n/\omega$, $e'=(\Phi_n(e_n))^\omega$. The map $\Phi_\omega\colon \lim^\omega(N_n,\mathrm{dist}_n)_e\to \lim^\omega(\Re_n,\mathrm{dist}/d_n)_{e'}$ such that*

$$\Phi_\omega\left(\lim{}^\omega(x_n)\right)=\lim{}^\omega(\Phi_n(x_n))$$

is a surjective isometry.

Proof. For every $x,y\in N_n$, let $p=e_1e_2...e_s$ be a shortest path from x to y in Γ_n. Then $\Phi_n(x)$ and $\Phi_n(y)$ are joined in Cayley(G) by a path labeled by $w_n(p)$. It follows that

$$\mathrm{dist}(\Phi_n(x),\Phi_n(y))\leq\sum_{i=1}^s|w_n(e_i)|\leq d_n\sum_{i=1}^s|e_i|=d_n\mathrm{dist}_n(x,y).$$

By Lemma 7.22, for every $x,y\in N_n$ there exists a geodesic joining $\Phi_n(x)$ to $\Phi_n(y)$ labeled by a net word $w_n(q)$ of rank n. If $q=e_1e_2...e_t$ then

$$w_n(q)=w_n(e_1)\dots w_n(e_t).$$

Therefore

$$\begin{aligned}\mathrm{dist}(\Phi_n(x),\Phi_n(y))&=|w_n(q)|\geq\textstyle\sum_{i=1}^t(1-2\lambda_n)\,|w_n(e_i)|\\ &\geq(1-2\lambda_n)\textstyle\sum_{i=1}^t(d_n|e_i|-1)\geq(1-2\lambda_n)\,(d_n\mathrm{dist}_n(x,y)-t)\\ &\geq(1-2\lambda_n)\,(d_n\mathrm{dist}_n(x,y)-E_n).\end{aligned}$$

Thus for every $x, y \in N_n$:

$$(1-2\lambda_n)\left(d_n \mathrm{dist}_n(x,y) - E_n\right) \le \mathrm{dist}(\Phi_n(x), \Phi_n(y)) \le d_n \mathrm{dist}_n(x,y). \tag{18}$$

According to (18), for every $\lim^\omega(x_n), \lim^\omega(y_n) \in \lim^\omega(N_n, \mathrm{dist}_n)_e$ we have that

$$\lim_\omega \mathrm{dist}_n(x_n, y_n) - \lim_\omega \frac{E_n}{d_n} \le \lim_\omega \frac{\mathrm{dist}(\Phi_n(x_n), \Phi_n(y_n))}{d_n} \le \lim_\omega \mathrm{dist}_n(x_n, y_n)\,. \tag{19}$$

Since $(d_n)_{n\in\mathbb{N}}$ is a fast increasing sequence we have that $\lim_\omega \frac{E_n}{d_n} = 0$. This implies that Φ_ω is well defined and that it is an isometry.

Remark 7.23 implies the surjectivity of the map Φ_ω. $\square$

Notation: We denote by e the element $(1)^\omega \in G^\omega$.

Proposition 7.26. *Let $(\Gamma_n)_{n\in\mathbb{N}}$ be a sequence of metric graphs satisfying the properties listed before Definition 7.14, let $(d_n)_{n\in\mathbb{N}}$ be a fast increasing sequence with respect to $(\Gamma_n)_{n\in\mathbb{N}}$ and let $G = \langle a, b \mid R\rangle$ be the group constructed as above. For every ultrafilter ω the asymptotic cone $\mathrm{Con}^\omega(G; e, d)$ is tree-graded with respect to the set of pieces:*

$$\mathcal{P} = \left\{ \lim^\omega(g_n \Re_n) \mid (g_n)^\omega \in G^\omega \textit{ such that } \lim_\omega \frac{\mathrm{dist}(e, g_n \Re_n)}{d_n} < \infty \right\}, \tag{20}$$

in particular different elements $(g_n)^\omega$ correspond to different pieces from $\mathcal{P}$.

Proof. **Property** (T_1)**.** Suppose that $\lim^\omega(g_n\Re_n) \cap \lim^\omega(g'_n\Re_n)$ contains at least two distinct points, where $(g_n)^\omega, (g'_n)^\omega \in G^\omega$. We may suppose that $(g'_n)^\omega = (1)^\omega$. Let

$$\lim^\omega(a_n), \lim^\omega(b_n) \in \lim^\omega(g_n\Re_n) \cap \lim^\omega(\Re_n)\,,\ \lim^\omega(a_n) \neq \lim^\omega(b_n).$$

The inclusion $\lim^\omega(a_n), \lim^\omega(b_n) \in \lim^\omega(\Re_n)$ implies that

$$\lim^\omega(a_n) = \lim^\omega(\Phi_n(x_n))\,,\ \lim^\omega(b_n) = \lim^\omega(\Phi_n(y_n)).$$

where $x_n, y_n \in N_n$, $\lim^\omega(x_n) \neq \lim^\omega(y_n)$. The inclusion $\lim^\omega(a_n), \lim^\omega(b_n) \in \lim^\omega(g_n\Re_n)$ implies that $\lim^\omega(a_n) = \lim^\omega(g_n\Phi_n(x'_n))$, $\lim^\omega(b_n) = \lim^\omega(g_n\Phi_n(y'_n))$, where $x'_n, y'_n \in N_n$, $\lim^\omega(x'_n) \neq \lim^\omega(y'_n)$.

By Lemma 7.22, for every $n \geq 1$, there exists a geodesic $\mathfrak{p}_1^{(n)}$ in Cayley(G) joining $\Phi_n(x_n)$ with $\Phi_n(y_n)$ labeled by a net word $w_n(p_1^{(n)})$, where $p_1^{(n)}$ is a simple path from x_n to y_n in Γ_n. It follows that $\mathfrak{p}_1^{(n)} \subset \Re_n$. Similarly, there exists a geodesic $\mathfrak{p}_2^{(n)}$ joining $g_n\Phi_n(x'_n)$ to $g_n\Phi_n(y'_n)$ contained in $g_n\Re_n$. The label of this geodesic is a net word $w_n(p_2^{(n)})$. Let $\mathfrak{q}_n$ be a geodesic joining $\Phi_n(x_n)$ to $g_n\Phi_n(x'_n)$ and $\mathfrak{q}'_n$ a geodesic joining $\Phi_n(y_n)$ to $g_n\Phi_n(y'_n)$ in Cayley(G). Both $\mathfrak{q}_n$ and $\mathfrak{q}'_n$ have length $o(d_n)$. The geodesics $\mathfrak{p}_1^{(n)}$ and $\mathfrak{p}_2^{(n)}$ on the other hand have length $O(d_n)$. We consider the geodesic quadrilateral composed of $\mathfrak{p}_1^{(n)}, \mathfrak{q}_n, \mathfrak{p}_2^{(n)}, \mathfrak{q}'_n$ and a minimal van Kampen diagram Δ_n whose boundary label coincides with the label of this quadrangle. Then $\partial\Delta_n$ is a product of four segments which we shall denote s_n, t_n, s'_n, t'_n (the labels of these paths coincide with the labels of the paths $\mathfrak{p}_1^{(n)}, \mathfrak{q}_n, \mathfrak{p}_2^{(n)}, \mathfrak{q}'_n$ respectively).

There exists a unique (covering) map γ from Δ to Cayley(G) that maps the initial vertex of s_n to 1 and preserves the labels of the edges. The map γ maps s_n to $\mathfrak{p}_1^{(n)} \subseteq \Re_n$ and s'_n to $\mathfrak{p}_2^{(n)} \subseteq g\Re_n$.

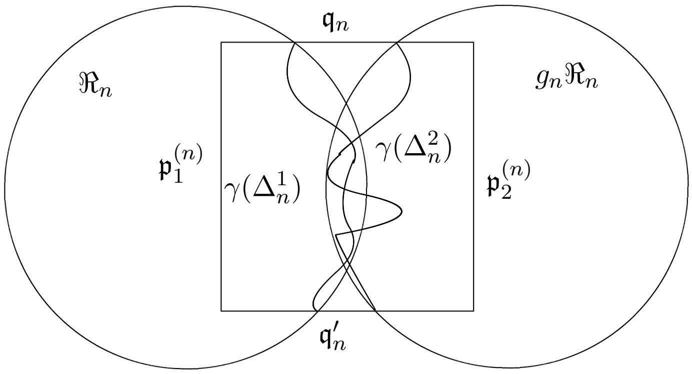

Figure 6: The diagram Δ_n.

Let Δ_n^1 be the maximal (connected) sub-diagram of Δ_n that contains s_n and whose γ-image is contained in $\Re_n$. Likewise, let Δ_n^2 be the maximal sub-diagram of Δ_n that contains s'_n and whose γ-image is contained in $g\Re_n$. The complement $\Delta_n \setminus (\Delta_n^1 \cup \Delta_n^2)$ has several connected components.

Suppose that the complement contains cells, and let Θ_n be one of the non-trivial components of the complement. The boundary of Θ_n is contained in $\partial\Delta_n^1 \cup t_n \cup \partial\Delta_n^2 \cup t'_n$. By Greendlinger's lemma, there exists a cell π in Θ_n such that $\partial\pi \cap \partial\Theta_n$ contains a path u_n of length at least $\frac{7}{10}|\partial\pi|$. Suppose that u_n has more than 15λ of its length in common with $\partial\Delta_n^1$. Then the labels of $\partial\pi$ and $\partial\Delta_n^1$ contain a common subword of length at least 5λ of the length of a block participating in the label of $\partial\pi$. By Lemma 7.21, π has rank n and the γ-image of $\Delta_n^1 \cup \pi$ is in $\Re_n$, a contradiction with the maximality of Δ_n^1. Hence $|u_n \cap \partial\Delta_n^1| \leq 15\lambda|u_n|$. Similar argument applies to Δ_n^2.

Therefore $|u_n \cap \left(\partial\Delta_n^1 \cup \partial\Delta_n^2\right)| \leq 30\lambda|u_n|$. It follows that u_n has more than $\frac{6}{10}|\partial\pi|$ in common with $t_n \cup t'_n$. Since $\gamma(t_n)$ and $\gamma(t'_n)$ are both geodesics, u_n must intersect both of them. We have $|u_n| \leq 30\lambda|u_n| + |t_n| + |t'_n|$, hence $|u_n| = o(d_n)$. Therefore

$$\mathrm{dist}(\Phi_n(x_n), \Phi_n(y_n)) \leq |u_n| + |t_n| + |t'_n| = o(d_n),$$

a contradiction.

Property (T_2)**.** According to Proposition 3.29, it suffices to study sequences of geodesic k-gons P_n in Cayley(G) with all lengths of edges of order d_n, k fixed and $\lim^\omega(P_n)$ a simple geodesic triangle. We need to show that $\lim^\omega(P_n)$ is contained in one piece.

We fix such a sequence $(P_n)_{n\in\mathbb{N}}$ of k-gons in Cayley(G). Let $\mathcal{V}_n$ be the set of vertices of P_n. We consider minimal van Kampen diagrams $\Delta^{(n)}$ and covering maps $\gamma_n\colon \Delta^{(n)} \to \mathrm{Cayley}(G)$ such that $\gamma_n(\partial\Delta^{(n)})$ is P_n. We can consider the boundary of $\Delta^{(n)}$ also as a k-gon whose vertices and sides correspond to the vertices and sides of P_n.

(a) Properties of the diagrams $\Delta^{(n)}$.

By Lemma 7.20, each cell from $\Delta^{(n)}$ has boundary length $\leq O(d_n)$. On the other hand, the cells of rank $k \geq n+1$ have boundary of length at least $[\zeta^{n+1}d_{n+1}]$. Property (2) of the fast increasing sequence (d_n) implies that for n large enough all cells from the diagram $\Delta^{(n)}$ are of rank $k \leq n$.

Suppose that ω-almost surely there exists a cell π of rank $m \leq n-1$ in $\Delta^{(n)}$ the boundary of which intersects two edges $[x,y]$, $[z,t]$ without common endpoint. Recall that the diameter of a cell of rank m is at most $10md_m \leq 10(n-1)d_{n-1}$. Then there exist two points in $\gamma_n[x,y]$

and in $\gamma_n[z,t]$ respectively, which are at distance at most $10(n-1)d_{n-1}$ of each other. In the ω-limit of P_n we obtain that two edges without common endpoint intersect in a point. This contradicts the fact that $\lim^\omega(P_n)$ is a simple loop. We conclude that ω-almost surely all cells whose boundaries intersect two edges without common endpoint are of rank n.

Suppose that the boundary of one of the cells π of rank m in $\Delta^{(n)}$ is not a simple path. Then by applying the Greendlinger lemma to any hole formed by $\partial\pi$, we get a cell π' whose boundary has a common subpath u with $\partial\pi$ such that $|u| \geq \frac{7}{10}|\partial\pi'|$. Then there exists a block w in $\partial\pi'$ such that $|w \cap \partial\pi| \geq \frac{7}{20}|w|$. We apply Lemma 7.21 to $\partial\pi$ and $\partial\pi'$ and we obtain that the ranks of π and π' coincide and that the boundary label of the union $\pi \cup \pi'$ is a net word of rank m corresponding to a loop in Γ_m. Hence the union of the cells π and π' can be replaced by one cell corresponding to a relation from R, a contradiction with the minimality of $\Delta^{(n)}$. Hence the boundary of each cell in $\Delta^{(n)}$ is a simple path.

Suppose that the boundaries of two cells π_1, π_2, in $\Delta^{(n)}$, of rank m_1 and m_2 respectively, intersect in several connected components. We apply the Greendlinger lemma to a hole formed by $\partial\pi_1 \cup \partial\pi_2$ and we get a cell π' whose boundary has a common subpath, of length at least $\frac{7}{10}|\partial\pi'|$, with $\partial\pi_1 \cup \partial\pi_2$. Therefore $\partial\pi'$ has a common subpath with one $\partial\pi_i$, $i \in \{1,2\}$, of length at least $\frac{7}{20}|\partial\pi'|$. Lemma 7.21 implies that the ranks of π_i and π' coincide and that the boundary label of $\pi_i \cup \pi'$ is a net word of rank m_i corresponding to a loop in Γ_{m_i}. Hence $\pi_i \cup \pi'$ can be replaced by one cell, a contradiction with the minimality of $\Delta^{(n)}$. We conclude that the intersection of the boundaries of two cells, if non-empty, is connected.

Suppose that the boundary of a cell π in $\Delta^{(n)}$ of rank m intersects one side $[x,y]$ of $\partial\Delta^{(n)}$ in several connected components. We consider a hole formed by $\partial\pi \cup [x,y]$ and we apply the Greendlinger lemma to it. We obtain a cell π' whose boundary has a common subpath u with $\partial\pi \cup [x,y]$, such that $|u| \geq \frac{7}{10}|\partial\pi'|$. Since $\gamma_n[x,y]$ is a geodesic, u cannot have more than $\frac{5}{7}|u|$ in common with $[x,y]$. Hence $|u \cap \partial\pi| \geq \frac{1}{5}|\partial\pi'|$, which implies that there exists a block w in $\partial\pi'$ such that $|w \cap \partial\pi| \geq \frac{1}{10}|w|$. We apply Lemma 7.21 to π and π' and as previously we obtain a contradiction of the minimality of $\Delta^{(n)}$. Consequently, the intersection of the boundary of a cell in $\Delta^{(n)}$ with a side of $\partial\Delta^{(n)}$, if non-empty, is connected.

(b) Existence of a cell π_n of rank n in $\Delta^{(n)}$ such that $\operatorname{dist}(P_n, \gamma_n(\partial\pi_n)) = o(d_n)$.

Take any vertex $v = v_n$ of the k-gon $\partial\Delta^{(n)}$. Let $[x,v],[v,y]$ be the two consecutive sides of the k-gon $\partial\Delta^{(n)}$. Let $x'_n \in [x,v]$ be such that $\gamma_n(x'_n)$ is the last point on $[\gamma_n(v), \gamma_n(x)]$ (counting from $\gamma_n(v)$) for which there exists a point z on $[\gamma_n(v), \gamma_n(y)]$ with $\operatorname{dist}(\gamma_n(x'_n), z)$ not exceeding $\zeta^{n/2}d_n$. Since $\zeta^{n/2}d_n = o(d_n)$, $\lim^\omega(x'_n) = \lim^\omega(\gamma_n v)$ (recall that the triangle $\lim^\omega(P_n)$ is simple). Therefore $\operatorname{dist}(x'_n, \gamma_n v) = o(d_n)$.

Similarly let $y'_n \in [y,v]$ be such that $\gamma_n(y'_n)$ is the last point on $[\gamma_n(v), \gamma_n(y)]$ for which there exists a point z on $[\gamma_n(v), \gamma_n(x)]$ with $\operatorname{dist}(\gamma_n(y'_n), z) \leq \zeta^{n/2}d_n$. Then $\operatorname{dist}(y'_n, \gamma_n v) = o(d_n)$.

Consider the set Π_v of cells π in $\Delta^{(n)}$ whose boundaries have common points with both $[x,v]$ and $[v,y]$. The boundary of π naturally splits into four parts: a sub-arc of $[x,v]$, a sub-arc of $[v,y]$, and two arcs $c(\pi), c'(\pi)$ which connect points on $[x,v]$ with points on $[v,y]$ and such that $c(\pi)$ and $c'(\pi)$ do not have any common points with $[x,v] \cup [v,y]$ others than their respective endpoints. We assume that $c'(\pi)$ is closer to v than $c(\pi)$.

The cells from Π_v are ordered in a natural way by their distance from v. Take the cell $\pi \in \Pi_v$ which is the farthest from v among all cells in Π_v satisfying

$$\operatorname{dist}(\gamma_n(c(\pi)_-), \gamma_n(c(\pi)_+)) \leq [\zeta^{n/2}d_n].$$

Let us cut off the corner of $\Delta^{(n)}$ bounded by the triangle $\Theta_v = c(\pi) \cup [c(\pi)_-, v] \cup [v, c(\pi)_+]$. Notice that by the definition of x'_n, y'_n, we have $c(\pi)_- \in [x'_n, v]$, $c(\pi)_+ \in [v, y'_n]$. Therefore the lengths of the sides of Θ_v are $o(d_n)$. Also notice that ω-almost surely Θ_v contains all cells of rank $\leq n-1$ from Π_v. That follows from the fact that the diameter of $\Re_k$, $k \leq n-1$, does not exceed $10(n-1)d_{n-1}$, hence for n large enough it does not exceed $[\zeta^{n/2}d_n]$ by property (2) of the definition of a fast increasing sequence.

Let us do this operation for every vertex v of the k-gon $\Delta^{(n)}$. As a result, we get a minimal diagram $\Delta_1^{(n)}$ such that $\gamma_n(\Delta_1^{(n)})$ is a $2k$-gon P'_n with k sides which are sub-arcs of the sides of P_n (we shall call them *long sides*) and k sides which are curves of type $c(\pi)$ whose lengths are $o(d_n)$ (*short sides*). Some of the short sides may have length 0. The ω-limit $\lim^\omega(P'_n)$ coincides with $\lim^\omega(P_n)$. We shall consider $\partial\Delta_1^{(n)}$ as a $2k$-gon with long and short sides corresponding to the sides of P'_n.

Notice that by construction $\Delta_1^{(n)}$ does not have cells of rank $\leq n-1$ which have common points with two long sides of the $2k$-gon $\partial\Delta_1^{(n)}$.

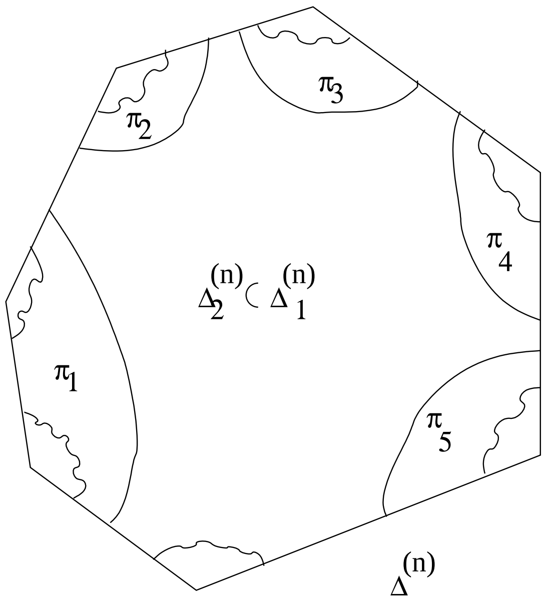

Figure 7: Diagram $\Delta^{(n)}$.

Let π_1, π_2,...,π_m be all Greendlinger $\frac{6}{10}$-cells in $\Delta_1^{(n)}$, i.e. for every $i = 1, ..., m$, the intersection $\partial\pi_i \cap \partial\Delta_1^{(n)}$ contains a subpath u_i of length at least $\frac{6}{10}|\partial\pi_i|$. Let r_i be the rank of the cell π_i, $i = 1, ..., m$. The path u_i cannot have more than $\frac{5}{6}$ of its length in common with a long side of the $2k$-gon $\partial\Delta_1^{(n)}$ because the γ_n-images of these sides are geodesics. By Lemma 7.21, u_i cannot have a subpath of length bigger than 5λ times the length of a block of rank r_i in common with a short side of $\partial\Delta_1^{(n)}$. Since short sides and long sides in $\partial\Delta_1^{(n)}$ alternate ω-almost surely, u_i must have points in common with two long sides of $\partial\Delta_1^{(n)}$. Therefore the number m is at most k and the rank r_i is n for every $i = 1, ..., m$ (ω-almost surely).

Let us cut off all cells $\pi_1, ..., \pi_m$ from the diagram $\Delta_1^{(n)}$. The resulting diagram $\Delta_2^{(n)}$ has a form of a polygon where each side is either a part of a long side of $\Delta_1^{(n)}$ (we call it again *long*) or a part of $\partial\pi_i$ (we call it *special*) or a part of a short side of $\Delta_1^{(n)}$ (we call it *short*). Notice that by the definition of $\Delta_1^{(n)}$, the length of any special side of $\Delta_2^{(n)}$ cannot be smaller than $[\zeta^{n/2}d_n]$ ω-almost surely.

Suppose that the diagram $\Delta_2^{(n)}$ contains cells ω-almost surely. Consider a Greendlinger $\frac{7}{10}$-

cell π of rank m in $\Delta_2^{(n)}$ and the corresponding path $u \subset \partial\pi \cap \partial\Delta_2^{(n)}$. This path cannot have more than $\frac{5}{7}$ of its length in common with a long side of $\Delta_2^{(n)}$, more than 5λ times the length of a block of $\partial\pi$ in common with a special or short side. Therefore u cannot contain a whole special side of $\Delta_2^{(n)}$. Hence u has a subpath u' of length at least $(\frac{7}{10} - 10\lambda)|\partial\pi|$ that intersects only long and short sides of $\Delta_2^{(n)}$. Hence π is a Greendlinger $\frac{6}{10}$-cell in $\Delta_1^{(n)}$. This contradicts the fact that all such cells were removed when we constructed $\Delta_2^{(n)}$.

Thus $\Delta_2^{(n)}$ contains no cells ω-almost surely. In particular, all cells in $\Delta_1^{(n)}$ are of rank n and all of them are Greendlinger $\frac{6}{10}$-cells. For each cell π_i, $i = 1, ..., m$, consider the decomposition $\partial\pi_i = u_i u_i'$. Any two arcs u_i', u_j' $(i \neq j)$, have at most one maximal sub-arc in common. The length of this sub-arc is at most 5λ times the length of a maximal block of rank n (by Lemma 7.21 and the minimality of $\Delta^{(n)}$). Hence (ω-almost surely) the length of any arc u_i' is at most $5k\lambda[\zeta^{n/2} d_n]$. Therefore $\lim_\omega \frac{|u_i'|}{d_n} = 0$. Since $\lim^\omega(P_n')$ is a simple triangle, we can conclude that ω-almost surely for all but one $i \in \{1, ..., m\}$ the length of $\partial\pi_i$ is $o(d_n)$. Indeed otherwise we would have two points on P_n' at distance $O(d_n)$ along the boundary of P_n' but at distance $o(d_n)$ in Cayley(G). The ω-limits of these two points would give us a self-intersection point of $\lim^\omega(P_n')$.

Let us call this exceptional i by i_n. Then $\lim^\omega(P_n')$ coincides with $\lim^\omega(\gamma_n(\partial\pi_{i_n}))$. Since $\gamma_n(\pi_{i_n})$ is contained in $g_n\Re_n$ for some g_n, $\lim^\omega(P_n')$ is contained in one piece $\lim^\omega(g_n\Re_n)$. □

Proposition 7.27 (description of the set of pieces). *Consider the following two collections of metric spaces:*

$$\left\{\lim^\omega(g_n\Re_n)_e \mid (g_n)^\omega \in G^\omega, \lim_\omega \frac{\mathrm{dist}(e, g_n\Re_n)}{d_n} < \infty\right\} \tag{21}$$

and

$$\{\lim^\omega(N_n, \mathrm{dist}_n)_x \mid x \in \Pi N_n/\omega\}\,. \tag{22}$$

We consider each $\lim^\omega(N_n, \mathrm{dist}_n)_x$ *as a space with basepoint* $\lim^\omega(x_n)$ *and each* $\lim^\omega(g_n\Re_n)_e$ *as a space with basepoint* $\lim^\omega(y_n)$, *where* $\lim^\omega(y_n)$ *is the projection of* $\lim^\omega(e)$ *onto* $\lim^\omega(g_n\Re_n)$.

Then every space in one of these collections is isometric, as a metric space with basepoint, to a space in the other collection. Moreover every space in the second collection is isometric to continuously many spaces in the first collection.

Proof. Let $t_n = g_n^{-1}y_n$, $n \geq 1$. Let $y = (y_n)^\omega$ and $t = (t_n)^\omega$. Then $\lim^\omega(g_n\Re_n)_e$ is isometric to $\lim^\omega(g_n\Re_n)_y$ which, in turn, is isometric to $\lim^\omega(\Re_n)_t$. Notice that $t_n \in \Re_n$, ω-almost surely. Remark 7.23 implies that there exists a $u_n \in \Phi_n(N_n)$ such that $\lim_\omega \frac{\mathrm{dist}(u_n, t_n)}{d_n} = 0$. Let $u = (u_n)^\omega$. For every $n \geq 1$, let $x_n \in N_n$ be such that $u_n = \Phi_n(x_n)$, $x = (x_n)^\omega$. Then by Lemma 7.25, $\lim^\omega(g_n\Re_n)_e$ is isometric to $\lim^\omega(N_n)_x$.

The fact that every limit set $\lim^\omega(N_n, \mathrm{dist}_n)_x$ is isometric to a set $\lim^\omega(\Re_n, \mathrm{dist}/d_n)_g$ follows from Lemma 7.25. We write g as $\lim^\omega(g_n^{-1})$ for some $g_n^{-1} \in \Phi_n(N_n)$. The set $\lim^\omega(g_n\Re_n, \mathrm{dist}/d_n)_e$ contains $\lim^\omega(1)$ and with respect to this basepoint it is isometric to $\lim^\omega(N_n, \mathrm{dist}_n)_x$.

We consider an arbitrary element $(\gamma_n)^\omega$ in G_e^ω such that $\lim_\omega \frac{\mathrm{dist}(1,\gamma_n)}{d_n} = 0$. The set $\lim^\omega(\gamma_n g_n\Re_n)_e$ is distinct from the set $\lim^\omega(g_n\Re_n)_e$, as the argument in Proposition 7.26 shows. On the other hand, the metric space $\lim^\omega(\gamma_n g_n\Re_n)_e$ with basepoint $\lim^\omega(\gamma_n) = \lim^\omega(1)$ is isometric to the metric space $\lim^\omega(g_n\Re_n)_e$ with basepoint $\lim^\omega(1)$, hence to $\lim^\omega(N_n, \mathrm{dist}_n)_x$ with basepoint $\lim^\omega(x_n)$. We complete the proof by noting that there are continuously many elements $(\gamma_n)^\omega$ with $\lim_\omega \frac{\mathrm{dist}(1,\gamma_n)}{d_n} = 0$. □

7.4 Free products appearing as fundamental groups of asymptotic cones

The following lemma is obvious.

Lemma 7.28. *The collection of sets* $\{2^k\mathbb{N} + 2^{k-1} \mid k \in \mathbb{N}\}$ *is a partition of* $\mathbb{N}$.

Notation: We denote the set $2^k\mathbb{N}+2^{k-1}$ by $\mathbb{N}_k$, for every $k \in \mathbb{N}$. We denote by $k(n)$ the element $2^k n + 2^{k-1}$ of $\mathbb{N}_k$.

Let $(M_k\, ,\, \mathrm{dist}_k)_{k\in\mathbb{N}}$ be a sequence of proper geodesic locally uniformly contractible spaces, let O_k be a point in M_k and let ζ be a real number in $(0,1)$. Fix $k \in \mathbb{N}$. We apply Remark 7.4 to the sequence of sets $\left(B_n^{(k)}\right)_{n\in\mathbb{N}\cup\{0\}}$, where $B_0^{(k)} = \{O_k\}$ and $B_n^{(k)} = \overline{B(O_k, n)}$, $n \in \mathbb{N}$, and to the sequence of numbers $(\zeta^n)_{n\in\mathbb{N}}$. We obtain an increasing sequence

$$\{O_k\} \subset N_1^{(k)} \subset N_2^{(k)} \subset \cdots \subset N_n^{(k)} \subset \cdots\, , \tag{23}$$

such that $N_n^{(k)}$ is a ζ^n-snet in $(B_n^{(k)}, \mathrm{dist}_k)$. We consider the sequence of graphs $\Gamma_{\zeta^{[n/2]}}\left(N_n^{(k)}\right)$ endowed with the length metric $\mathrm{dist}_n^{(k)}$. We denote $\Gamma_{\zeta^{[n/2]}}\left(N_n^{(k)}\right)$ by $\Gamma_n^{(k)}$.

Remark 7.29. Note that the diameter of $(N_n^{(k)}, \mathrm{dist}_k)$ is at most $2n$, so by (11) the diameter of $(\Gamma_n^{(k)}, \mathrm{dist}_n^{(k)})$ is at most $10n$, for n large enough. Hence the graphs $\Gamma_n^{(k)}$ satisfy the conditions listed before Definition 7.14.

Now consider the sequence $(\Gamma_n\, ,\, \mathrm{dist}_n\, ,\, O_n)$ of finite metric graphs endowed with length metrics and with distinguished basepoints defined as follows: $(\Gamma_n\, ,\, \mathrm{dist}_n\, ,\, O_n) \equiv (\Gamma_n^{(k)}\, ,\, \mathrm{dist}_n^{(k)}\, ,\, O_k)$ when $n \in \mathbb{N}_k$. We consider a sequence (d_n) of positive numbers which is fast increasing with respect to the sequence of graphs (Γ_n). We construct a group $G = \langle a, b \mid R\rangle$ as in Section 7.2, associated to the sequences (Γ_n) and (d_n).

For every $k \in \mathbb{N}$ let μ_k be an ultrafilter with the property that $\mu_k(\mathbb{N}_k) = 1$.

Proposition 7.30. *The asymptotic cone* $\mathrm{Con}^{\mu_k}(G; e, d)$ *is tree-graded with respect to a set of pieces* $\mathcal{P}_k$ *that are isometric to ultraballs of* M_k *with center* O_k*. Ultraballs with different observation points correspond to different pieces from* $\mathcal{P}_k$*.*

Proof. By Proposition 7.26, $\mathrm{Con}^{\mu_k}(G; e, d)$ is tree-graded with respect to

$$\mathcal{P}_k = \left\{\lim^{\mu_k}(g_n\mathfrak{R}_n) \mid (g_n)^{\mu_k} \in G^{\mu_k} \text{ such that } \lim_{\mu_k}\frac{\mathrm{dist}(e, g_n\mathfrak{R}_n)}{d_n} < \infty\right\}. \tag{24}$$

By Proposition 7.27, the collection of representatives up to isometry of the set of pieces (24) coincides with the collection of representatives up to isometry of the set of ultralimits $\lim^{\mu_k}(N_n, \mathrm{dist}_n)_x$, $x \in \Pi N_n/\mu_k$. The hypothesis that $\mu_k(\mathbb{N}_k) = 1$ and the definition of the sequence of graphs (Γ_n) implies that $\lim^{\mu_k}(N_n, \mathrm{dist}_n)_x = \lim^{\mu_k}\left(N_n^{(k)}, \mathrm{dist}_n^{(k)}\right)_{x^{(k)}}$ for some $x^{(k)} \in \Pi N_n^{(k)}/\mu_k$. It remains to apply Lemma 7.5. □

Corollary 7.31. *Suppose that the space* M_k *is compact and locally uniformly contractible. Then the asymptotic cone* $\mathrm{Con}^{\mu_k}(G; e, d)$ *is tree-graded with respect to pieces isometric to* M_k*, and the fundamental group of this asymptotic cone is the free product of continuously many copies of* $\pi_1(M_k)$*.*

Proof. It is a consequence of Proposition 7.30 and Proposition 2.20. □

Corollary 7.32. *There exists a 2-generated group Γ such that for every finitely presented group G, the free product of continuously many copies of G is the fundamental group of an asymptotic cone of Γ.*

Theorem 7.33. *For every countable group C, there exists a finitely generated group G and an asymptotic cone T of G such that $\pi_1(T)$ is isomorphic to an uncountable free power of C. Moreover, T is tree-graded and each piece in it is isometric either to a fixed proper metric space Y_C with $\pi_1(Y_C) = C$ or to a simply connected ultraball of Y_C.*

Proof. Let C be a countable group. By Lemma 7.10, C is the fundamental group of a geodesic, proper, and locally uniformly contractible space Y_C. Moreover, by Lemma 7.11, there exists a point O in Y_C such that every ultraball of Y_C with center O is either isometric to Y_C or it is simply connected. It is easy to see that the cardinality of the set of different ultraballs of Y_C with center O, that are isometric to Y_C, is continuum. Consider the 2-generated group $G = G(Y_C)$ obtained by applying the above construction to $M_k = Y_C$ and $O_k = O$, $k \geq 1$. Then by Proposition 7.30 there exists an asymptotic cone of G that is tree-graded with respect to a set of pieces $\mathcal{P}$ such that the collection of representatives up to isometry of the pieces in $\mathcal{P}$ coincides with the collection of representatives up to isometry of the set of ultraballs of Y_C with center O. By Proposition 2.20, the fundamental group of that asymptotic cone is isomorphic to the free power of C of cardinality continuum. □

7.5 Groups with continuously many non-homeomorphic asymptotic cones

We use the construction in Section 7.2 to obtain a 2-generated recursively presented group which has continuously many non-π_1-equivalent (and thus non-homeomorphic) asymptotic cones. Let us enumerate the set of non-empty finite subsets of $\mathbb{N}$ starting with $\{1\}$ and $\{1, 2\}$, then listing all subsets of $\{1, 2, 3\}$ containing 3, all subsets of $\{1, 2, 3, 4\}$ containing 4, etc. Let F_k, $k \in \mathbb{N}$, be the k-th set in the sequence of subsets.

For every $n \geq 1$ let $\mathcal{T}^n$ be the n-dimensional torus $\mathbb{R}^n/\mathbb{Z}^n$ with its natural geodesic metric and a basepoint $O = (0, 0, ..., 0)$.

For every $k \geq 1$ consider the bouquet of tori $\mathcal{B}_k = \bigvee_{n \in F_k} (\mathcal{T}^n, O)$. This is a compact locally uniformly contractible geodesic metric space with a metric dist_k induced by the canonical metrics on the tori and with the basepoint $O_k = O$.

We repeat the construction of a group $G = \langle a, b \mid R \rangle$ in Section 7.4 for the sequence of proper geodesic spaces with basepoints $(\mathcal{B}_k, \mathrm{dist}_k, O_k)_{k \in \mathbb{N}}$.

Since all $\mathcal{B}_k$ are bouquets of tori, we can choose the snets $N_n^{(k)}$ coming from the same regular tilings of the tori of different dimensions, and from their regular sub-divisions. There is a recursive way to enumerate the snets $N_{k(n)}^{(k)}$. For an appropriate choice of the set of words $\mathcal{W}$ in Proposition 7.13, we obtain a recursively presented group G. The group has the following property.

Proposition 7.34. *The asymptotic cone $\mathrm{Con}^{\mu_k}(G; e, d)$ is tree-graded with respect to a set of pieces $\widetilde{\mathcal{P}}_k$ such that every piece is isometric to one of the tori $\mathcal{T}^n$, $n \in F_k$.*

Proof. Proposition 7.30 implies that the asymptotic cone $\mathrm{Con}^{\mu_k}(G; e, d)$ is tree-graded with respect to a set of pieces $\mathcal{P}_k$ such that all pieces are isometric to $\mathcal{B}_k$. It remains to use Lemma 2.24. □

Notation: We denote $\mathrm{Con}^{\mu_k}(G; e, d)$ by $\mathcal{C}_k$ and $\lim^{\mu_k}(e)$ by e_k.

Let I be an arbitrary infinite subset of $\mathbb{N}$, $I = \{i_1, i_2, \ldots, i_n, \ldots\}$. We consider the increasing sequence of finite sets

$$F_{k_1} \subset F_{k_2} \subset \cdots \subset F_{k_n} \subset \cdots$$

defined by $F_{k_n} = \{i_1, i_2, \ldots, i_n\}$. Correspondingly we consider the sequence of asymptotic cones $(\mathcal{C}_{k_n})_{n\in\mathbb{N}}$. We consider an ultrafilter ω. The ultralimit $\lim^\omega(\mathcal{C}_{k_n})_{(e_{k_n})_{n\in\mathbb{N}}}$ is also an asymptotic cone of G, according to Corollary 3.24. We denote it by $\mathcal{C}_\omega(I)$.

Lemma 7.35. *Let $(\mathcal{T}^{k_i})$ be a sequence of tori $\mathcal{T}^{k_i} = \mathbb{R}^{k_i}/\mathbb{Z}^{k_i}$ with canonical flat metrics. Suppose that $\lim^\omega(k_i) = \infty$ for some ultrafilter ω. Let $\mathcal{T} = \lim^\omega(\mathcal{T}^{k_i})_e$ for some e. Then $\mathcal{T}$ contains isometric π_1-embedded copies of all tori $\mathcal{T}^n$, $n \geq 1$.*

Proof. Since tori are homogeneous spaces, we can assume that e is the sequence of points $(0, 0, \ldots)$. For every $n \geq 1$ the torus $\mathcal{T}^n$ isometrically embeds into $\mathcal{T}^{k_i}$ for ω-almost all i by the map $\phi_i\colon (x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_n, 0, 0, \ldots)$. Consequently $\mathcal{T}^n$ isometrically embeds into $\mathcal{T}$ by $\phi_\omega\colon \bar{x} \mapsto \lim^\omega(\phi_i(\bar{x}))$. Let $\mathfrak{c}$ be a non-0-homotopic loop in $\mathcal{T}^n$. Suppose that $\phi_\omega(\mathfrak{c})$ is 0-homotopic in $\mathcal{T}$. Then there exists a continuous map $\psi\colon \mathbb{D}^2 \to \mathcal{T}$ with $\psi(\partial\mathbb{D}^2) = \phi_\omega(\mathfrak{c})$. For every small positive ε, there exists a triangulation of $\mathbb{D}^2$ such that if e is an edge in the triangulation, the images by ψ of the endpoints of e are at distance at most ε. Let $\mathcal{V}_\varepsilon$ be the set of vertices of such a triangulation. The restricted map $\psi_\varepsilon \psi|_{\mathcal{V}_\varepsilon}$ is an ω-limit of maps $\psi_i\colon \mathcal{V}_\varepsilon \to \mathcal{T}^{k_i}$. For every i and for every edge e in the considered triangulation of $\mathbb{D}^2$ we join with a geodesic in T^{k_i} the images by ψ_i of the endpoints of e. The length of this geodesic is ω-almost surely less than 2ε. To each triangle of the triangulation thus corresponds a geodesic triangle in T^{k_i} of perimeter smaller than 6ε, ω-almost surely. For ε small enough all these geodesic triangles are 0-homotopic in some T^{k_i}. But then $\mathfrak{c}$ is 0-homotopic in T^{k_i}, a contradiction. $\square$

Lemma 7.36. *The asymptotic cone $\mathcal{C}_\omega(I)$ is tree-graded with respect to a set of pieces $\widetilde{\mathcal{P}}_\omega(I)$ such that:*

(1) All pieces are either isometric to one of the tori $\mathcal{T}^i$, $i \in I$, or they have the property that for every $n \in \mathbb{N}$ they contain an isometric π_1-embedded copy of $\mathcal{T}^n$.

(2) The fundamental group of every piece is Abelian.

Proof. Proposition 7.34 implies that for every $n \in \mathbb{N}$, $\mathcal{C}_{k_n}$ is tree-graded with respect to a set of pieces $\widetilde{\mathcal{P}}_{k_n}$ such that every piece is isometric to one of the tori $\{\mathcal{T}^{i_1}, \mathcal{T}^{i_2}, \ldots, \mathcal{T}^{i_n}\}$. Theorem 3.30 implies that $\mathcal{C}_\omega(I) = \lim^\omega(\mathcal{C}_{k_n})_{(e_{k_n})_{n\in\mathbb{N}}}$ is tree-graded with respect to the set of pieces

$$\widetilde{\mathcal{P}}_\omega(I) = \left\{\lim^\omega(M_n) \mid M_n \in \widetilde{\mathcal{P}}_{k_n}\,,\ \mathrm{dist}(e_{k_n}, M_n) \text{ bounded uniformly in } n\right\}. \qquad (25)$$

Let $\lim^\omega(M_n)$ be one of these pieces. Since $M_n \in \widetilde{\mathcal{P}}_{k_n}$, it follows that M_n is isometric to one of the tori $\{\mathcal{T}^{i_1}, \mathcal{T}^{i_2}, \ldots, \mathcal{T}^{i_n}\}$. Let $i(M_n)$ be the dimension of the torus M_n and let dist_n be the geodesic metric on M_n.

(1) We have two possibilities.

I. $\lim^\omega(i(M_n)) = \infty$. In this case we can imply Lemma 7.35 and conclude that $\lim^\omega(M_n)$ contains isometric and π_1-injective copies of tori $\mathcal{T}^N$ for every N.

II. $\lim^\omega(i(M_n)) < \infty$. It follows that there exists a finite m such that $i(M_n) \in \{i_1, i_2, \ldots, i_m\}$ ω-almost surely. Remark 3.2 implies that there exists $j \in \{1, 2, \ldots, m\}$ such that $i(M_n) = i_j$ ω-almost surely. Hence ω-almost surely M_n is isometric to $\mathcal{T}^{i_j}$ and $\lim^\omega(M_n)$ is isometric to $\mathcal{T}^{i_j}$.

(2) Every torus $\mathcal{T}^n$ is a topological group, so it admits a continuous binary operation and a continuous unary operation satisfying the standard group axioms. It is not difficult to see that ω-limits of tori also are topological groups. Now the statement follows from the fact that the fundamental group of every topological group is Abelian [Hat]. □

Theorem 7.37. *The two-generated recursively presented group G has continuously many non-π_1-equivalent (and in particular non-homeomorphic) asymptotic cones.*

Proof. Indeed, by Lemma 7.36 and Proposition 2.20 the fundamental group of $\mathcal{C}_\omega(I)$ is a free product of $\mathbb{Z}^i$, $i \in I$, and infinitely dimensional Abelian groups. By Kurosh's theorem [LS], if $j \notin I$ then $\mathbb{Z}^j$ cannot be a free factor of that fundamental group. Hence the asymptotic cones $\mathcal{C}_\omega(I)$ for different subsets I of $\mathbb{N}$ have different fundamental groups. □

Remark 7.38. Each of the continuously many asymptotic cones from Theorem 7.37 is a restrictive asymptotic cone in the sense of Section 3.3. Indeed by Remark 3.25, each of the cones $\mathrm{Con}^{\mu_k}(G; e, d)$ is isometric to a restrictive asymptotic cone $\mathrm{Con}^{\nu_k}(G; e, (n))$. The map ϕ defined in Section 3.3 just before Remark 3.25 is in this case injective. The images of the sets $\mathbb{N}_k$ under this map are pairwise disjoint and $\nu_k(\phi(\mathbb{N}_k)) = 1$. It remains to use Proposition 3.26.

8 Asymptotically tree-graded groups are relatively hyperbolic

Let G be a finitely generated group and let $\{H_1, \dots, H_m\}$ be a collection of subgroups of G. Let S be a finite generating set of G closed with respect to taking inverses.

We denote by $\mathcal{H}$ the set $\bigsqcup_{i=1}^m (H_i \setminus \{e\})$. We note that Cayley$(G, S)$ is a subgraph of Cayley$(G,\, S \cup \mathcal{H})$, with the same set of vertices but a smaller set of edges. We have that $\mathrm{dist}_{S\cup\mathcal{H}}(u, v) \le \mathrm{dist}_S(u, v)$, for every two vertices u, v.

For every continuous path $\mathfrak{p}$ in a metric space X we endow the image of $\mathfrak{p}$ with a pseudo-order $\le_{\mathfrak{p}}$ (possibly not anti-symmetric, but transitive and reflexive relation) induced by the order on the interval of definition of $\mathfrak{p}$. For every two points x, y we denote by $\mathfrak{p}[x, y]$ the subpath of $\mathfrak{p}$ composed of points z such that $x \le_{\mathfrak{p}} z \le_{\mathfrak{p}} y$.

Definition 8.1. Let $\mathfrak{p}$ be a path in Cayley$(G,\, S \cup \mathcal{H})$. An *$\mathcal{H}$-component* of $\mathfrak{p}$ is a maximal sub-path of $\mathfrak{p}$ contained in a left coset gH_i, $i \in \{1, 2, \dots, m\}$, $g \in G$.

The path $\mathfrak{p}$ is said to be *without backtracking* if it does not have two distinct $\mathcal{H}$-components in the same left coset gH_i.

There are two notions of relative hyperbolicity. The weak relative hyperbolicity has been introduced by B. Farb in [Fa]. We use a slightly different but equivalent definition. The proof of the equivalence can be found in [Os].

Definition 8.2. The group G is *weakly hyperbolic relative to* $\{H_1, \dots, H_m\}$ if and only if the graph Cayley$(G,\, S \cup \mathcal{H})$ is hyperbolic.

The strong relative hyperbolicity has several equivalent definitions provided by several authors. The definition that we consider here uses the following property.

Definition 8.3. The pair $(G\,,\, \{H_1, \dots, H_m\})$ satisfies the *Bounded Coset Penetration (BCP) property* if for every $\lambda \ge 1$ there exists $a = a(\lambda)$ such that the following holds. Let $\mathfrak{p}$ and $\mathfrak{q}$ be two λ-bi-Lipschitz paths without backtracking in Cayley$(G,\, S \cup \mathcal{H})$ such that $\mathfrak{p}_- = \mathfrak{q}_-$ and $\mathrm{dist}_S(\mathfrak{p}_+, \mathfrak{q}_+) \le 1$.

(1) Suppose that s is an $\mathcal{H}$-component of $\mathfrak{p}$ such that $\mathrm{dist}_S(s_-, s_+) \geq a$. Then $\mathfrak{q}$ has an $\mathcal{H}$-component contained in the same left coset as s;

(2) Suppose that s and t are two $\mathcal{H}$-components of $\mathfrak{p}$ and $\mathfrak{q}$, respectively, contained in the same left coset. Then $\mathrm{dist}_S(s_-, t_-) \leq a$ and $\mathrm{dist}_S(s_+, t_+) \leq a$.

Definition 8.4. The group G is *(strongly) hyperbolic relative to* $\{H_1, \dots, H_m\}$ if it is weakly hyperbolic relative to $\{H_1, \dots, H_m\}$ and if $(G\,,\,\{H_1, \dots, H_m\})$ satisfies the BCP property.

We are going to prove the following theorem.

Theorem 8.5. *A finitely generated group G is asymptotically tree-graded with respect to subgroups $\{H_1, \dots, H_m\}$ if and only if G is (strongly) hyperbolic relative to $\{H_1, \dots, H_m\}$ and each H_i is finitely generated.*

This section is devoted to the proof of the "only if" statement. Note that the fact that each H_i is finitely generated has been proved before (Proposition 5.8).
The "if" statement is proved in the Appendix.

8.1 Weak relative hyperbolicity

The most difficult part of Theorem 8.5 is the following statement.

Theorem 8.6. *If G is asymptotically tree-graded with respect to $\{H_1, \dots, H_m\}$ then G is weakly hyperbolic relative to $\{H_1, \dots, H_m\}$.*

The main tool is a characterization of hyperbolicity due to Bowditch [Bow_1, Section 3]. For the sake of completeness we recall the results of Bowditch here.

8.1.A A characterization of hyperbolicity

Let $\mathcal{G}$ be a connected graph, with vertex set $\mathcal{V}$ and distance function dist, such that every edge has length 1.

We assume that to every pair $u, v \in \mathcal{V}$ we have associated a subset Λ_{uv} of $\mathcal{V}$. Assume that each Λ_{uv} is endowed with a relation $\leq_{uv}$ such that the following properties are satisfied.

(l_1) $\leq_{uv}$ is reflexive and transitive;

(l_2) for every $x, y \in \Lambda_{uv}$ either $x \leq_{uv} y$ or $y \leq_{uv} x$;

(l_3) for every $u, v \in \mathcal{V}$ we have $\Lambda_{uv} = \Lambda_{vu}$ and $\leq_{uv} = \geq_{vu}$.

We note that the relations $\leq_{uv}$ may not be anti-symmetric.

Notation: For $x, y \in \Lambda_{uv}$ with $x \leq_{uv} y$, we write

$$\Lambda_{uv}[x, y] = \Lambda_{uv}[y, x] = \{z \in \Lambda_{uv} \mid x \leq_{uv} z \leq_{uv} y\}\,.$$

We also assume that we have a function $\phi : \mathcal{V} \times \mathcal{V} \times \mathcal{V} \to \mathcal{V}$ with the following properties.

(c_1) (symmetry) $\phi \circ \sigma = \phi$ for every 3-permutation σ;

(c_2) $\phi(u, u, v) = u$ for all $u, v \in \mathcal{V}$;

(c_3) $\phi(u, v, w) \in \Lambda_{uv} \cap \Lambda_{vw} \cap \Lambda_{uw}$.

Suppose moreover that there exists a constant $K > 0$ such that the following conditions are satisfied.

(I) For every $u, v, w \in \mathcal{V}$, the Hausdorff distance between the sets $\Lambda_{uv}[u, \phi(u, v, w)]$ and $\Lambda_{uw}[u, \phi(u, v, w)]$ is at most K;

(II) If $p, q \in \mathcal{V}$ are such that $\mathrm{dist}(p, q) \leq 1$ then $\mathrm{diam}\, \Lambda_{uv}[\phi(u, v, p)\,,\, \phi(u, v, q)]$ is at most K;

(III) If $w \in \Lambda_{uv}$ then $\mathrm{diam}\, \Lambda_{uv}[w\,,\, \phi(u, v, w)]$ is at most K.

We call $(\Lambda_{uv}\,,\, \leq_{uv})$ the *line from u to v*. We call $\phi(u, v, w)$ the *center of u, v, w*.

Proposition 8.7 ([Bow$_1$], Proposition 3.1). *If the graph $\mathcal{G}$ admits a system of lines and centers satisfying the conditions above then $\mathcal{G}$ is hyperbolic with the hyperbolicity constant depending only on K. Moreover, for every $u, v \in \mathcal{V}$, the line Λ_{uv} is at uniformly bounded Hausdorff distance from any geodesic joining u to v, where the previous bound depends only on K.*

8.1.B Generalizations of already proven results and new results

Lemma 8.8. *Let $\mathfrak{q} : [0, t] \to X$ be an (L, C)-quasi-geodesic. Let x be a point in its image and let a, b be its endpoints. Then*

$$\mathrm{dist}(a, b) \geq \frac{1}{L_1}[\mathrm{dist}(a, x) + \mathrm{dist}(x, b)] - C_1\,, \tag{26}$$

where $L_1 = L^2$ and $C_1 = C\left(\frac{2}{L} + 1\right)$.

Proof. Let $s \in [0, t]$ be such that $\mathfrak{q}(s) = x$. We have that $\mathrm{dist}(a, b) \geq \frac{1}{L}t - C$. On the other hand $s \geq \frac{1}{L}\mathrm{dist}(a, x) - C$ and $t - s \geq \frac{1}{L}\mathrm{dist}(x, b) - C$ imply that $t \geq \frac{1}{L}[\mathrm{dist}(a, x) + \mathrm{dist}(x, b)] - 2C$. $\square$

Let (X, dist) be a metric space asymptotically tree-graded with respect to a collection of subsets $\mathcal{A}$. Given $L \geq 1$ and $C \geq 0$ we denote by $M(L, C)$ the constant given by (α_2') for $\theta = \frac{1}{3}$.

Definition 8.9 (parameterized saturations). Given $\mathfrak{q}$ an (L, C)-quasi-geodesic and $\mu \geq 0$, we define the *μ-saturation* $\mathrm{Sat}^\mu(\mathfrak{q})$ as the union of $\mathfrak{q}$ and all $A \in \mathcal{A}$ with $\mathcal{N}_\mu(A) \cap \mathfrak{q} \neq \emptyset$.

Notice that if a metric space X is asymptotically tree-graded with respect to a collection $\mathcal{A} = \{A_i \mid i \in I\}$ then X is also asymptotically tree-graded with respect to $\mathcal{N}_\mu(\mathcal{A}) = \{\mathcal{N}_\mu(A_i) \mid i \in I\}$ for every number $\mu > 0$. This immediately follows from the definition of asymptotically tree-graded spaces. One can also easily see that properties $(\alpha_1), (\alpha_2), (\alpha_3)$ are preserved. Hence the following two lemmas follow from Lemmas 4.21, 4.26 and 4.28.

Lemma 8.10 (uniform quasi-convexity of parameterized saturations). *For every $L \geq 1$, $C \geq 0$ and $\mu \geq M(L, C)$, and for every $\lambda \geq 1$, $\kappa \geq 0$, there exists $\tau = \tau(L, C, \mu, \lambda, \kappa)$ such that for every $R \geq 1$, for every (L, C)-quasi-geodesic $\mathfrak{q}$, the saturation $\mathrm{Sat}^\mu(\mathfrak{q})$ has the property that every (λ, κ)-quasi-geodesic $\mathfrak{c}$ joining two points in its R-tubular neighborhood is entirely contained in its τR-tubular neighborhood.*

Lemma 8.11 (parameterized saturations of polygonal lines). *The statements in Lemmas 4.26 and 4.28 remain true if we replace the saturations by μ-saturations, for every $\mu > 0$.*

Lemma 8.12. *Let $\mathfrak{q} = \mathfrak{q}_1 \cup \mathfrak{q}_2 \cup \cdots \cup \mathfrak{q}_n$ be such that*

(1) $\mathfrak{q}_i$ is an (L, C)-almost-geodesic in X for $i = 1, 2, \ldots, n$;

(2) $\mathfrak{q}_i \cap \mathfrak{q}_{i+1} = \{x_i\}$ *is an endpoint of* $\mathfrak{q}_i$ *and of* $\mathfrak{q}_{i+1}$ *for* $i = 1, \dots n-1$*;*

(3) x_{i-1} *and* x_i *are the two endpoints of* $\mathfrak{q}_i$ *for* $i = 2, \dots n-1$*;*

(4) each $\mathfrak{q}_i$ *satisfies one of the following two properties:*

  (i) the endpoints of $\mathfrak{q}_i$ *are in a set* $A_i \in \mathcal{A}$ *or*

  (ii) $\mathfrak{q}_i$ *has length at most* ℓ*, where* ℓ *is a fixed constant;*

(5) $A_i \neq A_j$ *if* $i \neq j$*.*

Then there exists $L_n \geq L$ *,* $C_n \geq C$ *depending on* n*,* ℓ *and* (L, C)*, such that* $\mathfrak{q}$ *is an* (L_n, C_n)*-almost-geodesic.*

Proof. Clearly $\mathfrak{q}$ is an L-Lipschitz map. We prove by induction on n that $\mathrm{dist}(\mathfrak{q}(t), \mathfrak{q}(s)) \geq \frac{1}{L_n}|t-s| - C_n$ for some $L_n \geq L$ and $C_n \geq C$.

The statement is true for $n = 1$. Assume it is true for some n. Let $\mathfrak{q} = \mathfrak{q}_1 \cup \mathfrak{q}_2 \cup \cdots \cup \mathfrak{q}_n \cup \mathfrak{q}_{n+1}$ be as in the statement of the Lemma. Let $\mathfrak{q}' = \mathfrak{q}_1 \cup \mathfrak{q}_2 \cup \cdots \cup \mathfrak{q}_n$ which, by the induction hypothesis, is an (L_n, C_n)-almost-geodesic.

Suppose that $\mathfrak{q}_{n+1}$ satisfies (4), (ii). Then $\mathfrak{q}$ is an $(L_n\,,\, 2(\ell + C_n))$-almost-geodesic.

Suppose that $\mathfrak{q}_{n+1}$ satisfies (4), (i). Let $A = A_{n+1}$. We take $M_n = M(L_n, C_n)$. Let y be the farthest point from x_n in the intersection $\overline{\mathcal{N}}_{M_n}(A) \cap \mathfrak{q}'$. Consider $\mathfrak{q}_y$ a sub-almost-geodesic of $\mathfrak{q}'$ of endpoints y and x_n. By Lemma 4.15, $\mathfrak{q}_y$ is contained in the $\tau_n M_n$-tubular neighborhood of A. On the other hand, $\mathfrak{q}_y = \mathfrak{q}_i' \cup \mathfrak{q}_{i+1} \cup \cdots \cup \mathfrak{q}_n$, where $\mathfrak{q}_i'$ is a sub-almost-geodesic of $\mathfrak{q}_i$. Again Lemma 4.15 implies that every $\mathfrak{q}_j$ satisfying (4), (i), is contained in $\mathcal{N}_\tau(A_i)$ for some uniform constant τ. Therefore, every such $\mathfrak{q}_j$ composing $\mathfrak{q}_y$ has endpoints at distance at most the diameter of $\mathcal{N}_\tau(A_i) \cap \mathcal{N}_{\tau_n M_n}(A)$, hence at most D_n, for some $D_n = D_n(\tau_n M_n)$. It follows that the distance $\mathrm{dist}(y, x_n)$ is at most $n(\ell + D_n)$. Lemma 4.19 implies that if the endpoints of $\mathfrak{q}_{n+1}$ are at distance at least $D' = D'(L_n, C_n, D_n)$, then $\mathfrak{q}$ is an $(L_n + 1\,,\, 2D')$-almost-geodesic.

If the endpoints of $\mathfrak{q}_{n+1}$ are at distance at most D' then the length of $\mathfrak{q}_{n+1}$ is at most $LD' + C$ and $\mathfrak{q}$ is an $(L_n\,,\, 2(LD' + C + C_n))$-almost-geodesic. □

Lemma 8.13. *For every* $L \geq 1$*,* $C \geq 0$*,* $M \geq M(L, C)$ *and* $\delta > 0$ *there exists* $D_0 > 0$ *such that the following holds. Let* $A \in \mathcal{A}$ *and let* $\mathfrak{q}_i \colon [0, \ell_i] \to X$*,* $i = 1, 2$*, be two* (L, C)*-quasi-geodesics with one common endpoint* b *and the other two respective endpoints* $a_i \in \mathcal{N}_M(A)$*, such that the diameter of* $\mathfrak{q}_i \cap \overline{\mathcal{N}}_M(A)$ *does not exceed* δ *for* $i = 1, 2$*. Then one of the two situations occurs:*

(a) either $\mathrm{dist}(a_1, a_2) \leq D_0$ *or*

(b) $b \in \mathcal{N}_M(A)$ *and* $\ell_i \leq L\delta + C$*.*

Proof. Let $\mathrm{dist}(a_1, a_2) = D$. We show that if D is large enough then we are in situation (b). Remark 4.14 implies that we may suppose that $\mathfrak{q}_i$ are $(L + C, C)$-almost geodesics.

According to Lemma 4.19, there exists D' such that if $D \geq D'$ then $\mathfrak{q}_1 \sqcup [a_1, a_2]$ is an $(L + C + 1, 2D')$-quasi-geodesic. Suppose that $D \geq D'$.

Suppose that b is not contained in $\mathcal{N}_M(A)$. Let $t \in [0, \ell_2]$ be such that $\mathfrak{q}_2(t) \in \overline{\mathcal{N}}_M(A)$ and $\mathfrak{q}_2|_{[0,t]}$ does not intersect $\mathcal{N}_M(A)$. The sub-arc $\mathfrak{q}_2|_{[t,\ell_2]}$ has endpoints at distance at most δ, hence it has length at most $L\delta + C$. It follows that $\mathfrak{q}_1 \sqcup [a_1, a_2] \sqcup \mathfrak{q}_2|_{[t,\ell_2]}$ is an $(L + C + 1\,,\, C_1)$-quasi-geodesic, where $C_1 = C_1(D', \delta)$. Lemma 4.25 implies that $\mathfrak{q}_1 \sqcup [a_1, a_2] \sqcup \mathfrak{q}_2|_{[t,\ell_2]}$ is contained in the τ-tubular neighborhood of $\mathrm{Sat}(\mathfrak{q}_2|_{[0,t]})$, where $\tau = \tau(L, C, D', \delta)$. This implies that $\mathcal{N}_M(A) \cap \mathcal{N}_\tau\left(\mathrm{Sat}(\mathfrak{q}_2|_{[0,t]})\right)$ has diameter at least D. By Lemma 4.22, for D large enough we must have that $A \subset \mathrm{Sat}(\mathfrak{q}_2|_{[0,t]})$. This contradicts the choice of t.

It follows that b is contained in $\mathcal{N}_M(A)$, which implies that $\ell_i \leq L\mathrm{dist}(a_i, b) + C \leq L\delta + C$. □

Corollary 8.14. *For every $L \geq 1$, $C \geq 0$, $M \geq M(L,C)$ and $\delta > 0$ there exists $D_1 > 0$ such that the following holds. Let $A \in \mathcal{A}$ and let $\mathfrak{q}_i \colon [0, \ell_i] \to X$, $i = 1, 2$, be two (L,C)-quasi-geodesics with one common endpoint b and the other two respective endpoints $a_i \in \mathcal{N}_M(A)$, such that the diameter of $\mathfrak{q}_i \cap \overline{\mathcal{N}}_M(A)$ does not exceed δ. Then* $\mathrm{dist}(a_1, a_2) \leq D_1$.

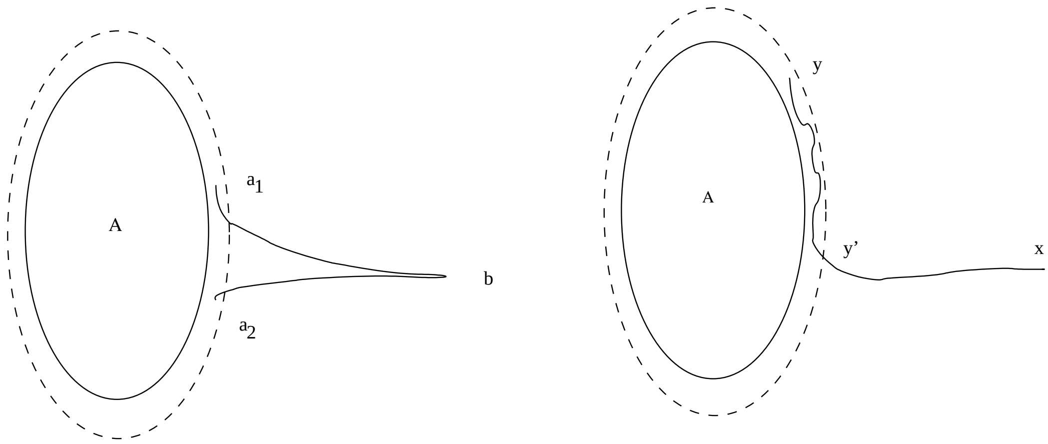

Figure 8: Corollary 8.14 and Lemma 8.15.

Lemma 8.15. *For every $L \geq 1$, $C \geq 0$ and $M \geq M(L,C)$ there exists $\mathfrak{d} = \mathfrak{d}(L,C,M) > 0$ such that the following holds. Let $A \in \mathcal{A}$ and let $\mathfrak{q} \colon [0, \ell] \to X$ be an (L,C)-almost-geodesic with endpoints x and $y \in \mathcal{N}_M(A)$. There exists a sub-arc $\mathfrak{q}'$ of $\mathfrak{q}$ with one endpoint x and the second endpoint in $\mathcal{N}_M(A)$ such that the diameter of $\mathfrak{q}' \cap \overline{\mathcal{N}}_M(A)$ is at most $\mathfrak{d}$.*

Proof. If $x \in \mathcal{N}_M(A)$ then we take $\mathfrak{q}' = \{x\}$. Suppose that $x \notin \mathcal{N}_M(A)$. Let $t = \inf\{t' \in [0, \ell] \mid t' \in \mathfrak{q}^{-1}(\mathcal{N}_M(A))\}$. Then $\mathfrak{q}(t) \in \overline{\mathcal{N}}_M(A)$. Let $s_i \in [0, t]$ be such that $\mathfrak{q}(s_i) \in \overline{\mathcal{N}}_M(A)$, $i = 1, 2$. If $|s_1 - s_2| \geq 3L(M+1)$ then property (α_2') implies that $\mathfrak{q}([s_1, s_2]) \cap \mathcal{N}_M(A) \neq \emptyset$. This contradicts the choice of t. Therefore $|s_1 - s_2| \leq 3L(M+1)$. We deduce that $\mathfrak{q}([0,t]) \cap \overline{\mathcal{N}}_M(A)$ has diameter at most $3L^2(M+1)$.

The definition of t implies that there exists $t_1 > t$ with $t_1 - t \leq \frac{1}{L}$ and $\mathfrak{q}(t_1) \in \mathcal{N}_M(A)$. We take $\mathfrak{q}' = \mathfrak{q}|_{[0,t_1]}$. The diameter of $\mathfrak{q}' \cap \overline{\mathcal{N}}_M(A)$ is at most $\mathfrak{d} = 3L^2(M+1) + 1$. □

8.1.C Hyperbolicity of $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$

Let G be a finitely generated group that is asymptotically tree-graded with respect to the finite collection of subgroups $\{H_1, \dots, H_m\}$. This means that $\mathrm{Cayley}(G, S)$ is asymptotically tree-graded with respect to the collection of subsets $\mathcal{A} = \{gH_i \mid g \in G\,,\, i = 1, 2, \dots, m\}$. We prove that $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$ is hyperbolic, using Proposition 8.7. The following result is central in the argument.

Proposition 8.16. *Let $L \geq 1$, $C \geq 0$, let $\mu \geq M(L,C)$ and let $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$ be three (L,C)-almost-geodesics composing a triangle in* $\mathrm{Cayley}(G,S)$. *We consider the set*

$$\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3) = \mathcal{N}_\kappa(\mathrm{Sat}^\mu(\mathfrak{q}_1)) \cap \mathcal{N}_\kappa(\mathrm{Sat}^\mu(\mathfrak{q}_2)) \cap \mathcal{N}_\kappa(\mathrm{Sat}^\mu(\mathfrak{q}_3))\,.$$

(1) There exists $\kappa_0 = \kappa_0(L, C, \mu)$ such that for every $\kappa \geq \kappa_0$ the set $\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ intersects each of the almost-geodesics $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$. In particular it is non-empty.

(2) For every $\kappa \geq \kappa_0$ there exists D_κ such that the set $\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ has diameter at most D_κ in $\mathrm{Cayley}(G, S \cup \mathcal{H})$.

Proof of (1). Let $\{i, j, k\} = \{1, 2, 3\}$. According to Lemma 8.11, the result in Lemma 4.25 is true if we replace $\mathrm{Sat}(\mathfrak{q})$ by $\mathrm{Sat}^\mu(\mathfrak{q}_i) \cup \mathrm{Sat}^\mu(\mathfrak{q}_j)$. In particular there exists $\tau = \tau(L, C, \mu)$ such that $\mathfrak{q}_k \subset \mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_i)) \cup \mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_j))$. The traces on $\mathfrak{q}_k$ of the two sets $\mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_i))$ and $\mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_j))$ compose a cover of two open sets, none of them empty. Since $\mathfrak{q}_k$ is an almost geodesic, it is connected, hence $\mathfrak{q}_k \cap \mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_i))$ and $\mathfrak{q}_k \cap \mathcal{N}_\tau(\mathrm{Sat}^\mu(\mathfrak{q}_j))$ intersect. The intersection is in $\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ for every $\kappa \geq \tau$. □

We need several intermediate results before proving (2). In the sequel we work with the data given in the statement of the Proposition 8.16, without mentioning it anymore.

Lemma 8.17. *There exist α, β positive constants depending only on L, C, μ and κ such that every point $x \in \mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ is in one of the two situations:*

(i) the ball $B(x, \alpha)$ intersects each of the three almost-geodesics $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$;

(ii) $x \in \mathcal{N}_\kappa(A)$ and $\mathcal{N}_\beta(A)$ intersects each of the three almost-geodesics $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$.

Proof. Let x be an arbitrary point in $\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$. The inclusion $x \in \mathcal{N}_\kappa(\mathrm{Sat}^\mu(\mathfrak{q}_i))$, $i \in \{1, 2, 3\}$, implies that there are two possibilities:

(I_i) $x \in \mathcal{N}_\kappa(\mathfrak{q}_i)$ or

(II_i) $x \in \mathcal{N}_\kappa(A)$, where $A \in \mathcal{A}$, $\mathcal{N}_\mu(A) \cap \mathfrak{q}_i \neq \emptyset$.

If we are in case (I) for the three edges then this means that (i) is satisfied with $\beta = \kappa$.

Suppose that only one edge is in case (II). Suppose it is $\mathfrak{q}_3$. Then $x \in \mathcal{N}_\kappa(\mathfrak{q}_1) \cap \mathcal{N}_\kappa(\mathfrak{q}_2)$ and there exists $A \in \mathcal{A}$ with $\mathcal{N}_\mu(A) \cap \mathfrak{q}_3 \neq \emptyset$ such that $x \in \mathcal{N}_\kappa(A)$. It follows that $\mathcal{N}_\beta(A)$ intersects the three edges for $\beta = \max(\mu, 2\kappa)$, so (ii) is satisfied.

Suppose that two edges are in case (II), for instance $\mathfrak{q}_2$ and $\mathfrak{q}_3$. Consequently, $x \in \mathcal{N}_\kappa(\mathfrak{q}_1)$ and $x \in \mathcal{N}_\kappa(A_2) \cap \mathcal{N}_\kappa(A_3)$, with $\mathcal{N}_\mu(A_i) \cap \mathfrak{q}_i \neq \emptyset$, where $i = 2, 3$. If $A_2 = A_3 = A$ then $\mathcal{N}_\beta(A)$ intersects the three edges for $\beta = \max(\mu, 2\kappa)$, so (ii) is satisfied. If $A_2 \neq A_3$ then, according to Lemma 8.11 (more precisely to Lemma 4.28 which also holds for μ-saturations) we have that $x \in \mathcal{N}_\varkappa(\mathfrak{q}_2 \cup \mathfrak{q}_3)$, where $\varkappa = \varkappa(\mu, \kappa)$. Suppose that $x \in \mathcal{N}_\varkappa(\mathfrak{q}_2)$. Then $\mathcal{N}_\beta(A_3)$ intersects the three edges for $\beta = \max(\mu, 2\kappa, \kappa + \varkappa)$, so (ii) is satisfied.

Suppose that the three edges are in case (II). It follows that $x \in \mathcal{N}_\kappa(A_1) \cap \mathcal{N}_\kappa(A_2) \cap \mathcal{N}_\kappa(A_3)$, with $\mathcal{N}_\mu(A_i) \cap \mathfrak{q}_i \neq \emptyset$, where $i = 1, 2, 3$.

If the cardinal of the set $\{A_1, A_2, A_3\}$ is 1 then we are in situation (ii) with $\beta = \mu$. Suppose the cardinal of the set is 2. Suppose that $A_1 = A_2 \neq A_3$. Lemma 4.28 for μ-saturations implies that $x \in \mathcal{N}_\varkappa(\mathfrak{q}_2 \cup \mathfrak{q}_3) \cap \mathcal{N}_\varkappa(\mathfrak{q}_1 \cup \mathfrak{q}_3)$. If $x \in \mathcal{N}_\varkappa(\mathfrak{q}_3)$ then $A = A_1 = A_2$ has the property that $\mathcal{N}_\beta(A)$ intersects the three edges for $\beta = \max(\mu, \kappa + \varkappa)$, and we are in case (ii). Otherwise $x \in \mathcal{N}_\varkappa(\mathfrak{q}_1) \cap \mathcal{N}_\varkappa(\mathfrak{q}_2)$, hence $\mathcal{N}_\beta(A_3)$ intersects the three edges for $\beta = \max(\mu, \kappa + \varkappa)$.

Assume that the cardinal of the set $\{A_1, A_2, A_3\}$ is 3. Then $x \in \mathcal{N}_\varkappa(\mathfrak{q}_1 \cup \mathfrak{q}_2) \cap \mathcal{N}_\varkappa(\mathfrak{q}_2 \cup \mathfrak{q}_3) \cap \mathcal{N}_\varkappa(\mathfrak{q}_1 \cup \mathfrak{q}_3)$. It follows that x is in the $\varkappa$-tubular neighborhood of at least two edges. Suppose these edges are $\mathfrak{q}_1$ and $\mathfrak{q}_2$. Then $\mathcal{N}_\beta(A_3)$ intersects the three edges for $\beta = \max(\mu, \kappa + \varkappa)$. □

Lemma 8.18. *For every $r > 0$ there exists $\varrho = \varrho(r, L, C)$ such that the following holds. Let $A \neq B$ be such that $A, B \in \mathcal{A}$, and both $\mathcal{N}_r(A)$ and $\mathcal{N}_r(B)$ intersect each of the three almost-geodesic edges of the triangle. Then there exists x such that $B(x, \varrho)$ intersects each of the edges of the triangle.*

Proof. Let $y \in \mathcal{N}_r(A)$ and $z \in \mathcal{N}_r(B)$. Lemma 8.15 implies that up to taking a subsegment of $[y,z]$, we may suppose that the diameters of $[y,z] \cap \mathcal{N}_r(A)$ and of $[y,z] \cap \mathcal{N}_r(B)$ are at most $\mathfrak{d}$, where $\mathfrak{d} = \mathfrak{d}(r)$. We apply Lemma 4.28 for r-saturations and for each $\mathfrak{q}_i$, $i \in \{1,2,3\}$, and we obtain that both $B(y,\varrho)$ and $B(z,\varrho)$ intersect $\mathfrak{q}_i$, where $\varrho = \varrho(r)$. □

Lemma 8.19. *There exists $R = R(L,C)$ such that for every triangle with (L,C)-almost-geodesic edges, one of the following two situations holds.*

(C) There exists x such that $B(x,R)$ intersects each of the three edges;

(P) There exists a unique $A \in \mathcal{A}$ such that $\mathcal{N}_R(A)$ intersects each of the three edges.

Proof. Let $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$ be the three edges. For $\mu = M(L,C)$ and $\kappa_0 = \kappa_0(L,C)$ we have that $\mathcal{C}^\mu_\kappa(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ is nonempty. It remains to apply Lemmas 8.17 and 8.18. □

Notation: We denote the vertices of the triangle by O_1, O_2, O_3, such that $\mathfrak{q}_i$ is opposite to O_i.

Lemma 8.20. *For every $r > 0$ there exists $D = D(r,L,C)$ such that the following holds. Let x be such that $B(x,r)$ intersects the three edges.*

(a) If y is such that $B(y,r)$ intersects the three edges then $\mathrm{dist}_{S\cup\mathcal{H}}(x,y) \le D$.

(b) If $A \in \mathcal{A}$ is such that $\mathcal{N}_r(A)$ intersects the three edges then $\mathrm{dist}_{S\cup\mathcal{H}}(x,A) \le D$.

Proof. Let x_i be nearest points to x in $\mathfrak{q}_i$, $i = 1,2,3$.

(a) We denote $\mathrm{dist}_{S\cup\mathcal{H}}(x,y)$ by D. Let y_i be nearest points to y in $\mathfrak{q}_i$, $i = 1,2,3$. Then $\mathrm{dist}_{S\cup\mathcal{H}}(x_i,y_j) \ge D - 2r$ for every $i,j \in \{1,2,3\}$. Suppose that $D > 2r$. Without loss of generality we may assume that $y_1 \in \mathfrak{q}_1[x_1,O_3]$. We have $\mathrm{dist}_S(x_1,x_2) \le 2r$, hence $\mathfrak{q}_1[x_1,O_3] \subset \mathcal{N}_{2\tau r}\left(\mathrm{Sat}\left(\mathfrak{q}_2[x_2,O_3]\right)\right)$, where $\tau = \tau(L,C)$. In particular y_1 is contained either in $\mathcal{N}_{2\tau r}(\mathfrak{q}_2[x_2,O_3])$ or in $\mathcal{N}_{2\tau r}(B)$ for $B \in \mathcal{A}$ such that $\mathcal{N}_M(B)$ intersects $\mathfrak{q}_2[x_2,O_3]$.

Case (a)I. Suppose that $y_2 \in \mathfrak{q}_2[x_2,O_1]$.

Case (a)I.1. Suppose that $y_1 \in \mathcal{N}_{2\tau r}(\mathfrak{q}_2[x_2,O_3])$. Then there exists $u \in \mathfrak{q}_2[x_2,O_3]$ such that $\mathrm{dist}_S(y_1,u) \le 2\tau r$. It follows that $\mathrm{dist}_S(u,x_2) \ge \mathrm{dist}_{S\cup\mathcal{H}}(u,x_2) \ge D - 2r - 2\tau r$. Inequality (26) implies that

$$\mathrm{dist}_S(u,y_2) \ge \frac{1}{L_1}[\mathrm{dist}(u,x_2) + \mathrm{dist}(x_2,y_2)] - C_1 \ge \frac{1}{L_1}(2D - 4r - 2\tau r) - C_1 .$$

On the other hand $\mathrm{dist}(u,y_2) \le 2\tau r + 2r$. Hence $D \le 2r + \tau r + L_1(r + \tau r + C_1/2)$.

Case (a)I.2. Assume that $y_1 \in \mathcal{N}_{2\tau r}(B)$, where $B \in \mathcal{A}$ is such that $\mathcal{N}_M(B)$ intersects $\mathfrak{q}_2[x_2,O_3]$. Let w_2 be a point in $\mathcal{N}_M(B) \cap \mathfrak{q}_2[x_2,O_3]$.

Suppose that $\mathfrak{q}_2[x_2,y_2] \cap \overline{\mathcal{N}}_{2\tau r}(B) \neq \emptyset$. Let z_2 be a point in the previous intersection. Then $\mathfrak{q}_2[w_2,z_2]$ has its endpoints in $\mathcal{N}_\chi(B)$, with $\chi = \max(M, 2\tau r + 1)$. Consequently $\mathfrak{q}_2[w_2,z_2] \subset \mathcal{N}_{\tau\chi}(B)$. In particular x_2 is contained in $\mathcal{N}_{\tau\chi}(B)$ and $\mathrm{dist}_{S\cup\mathcal{H}}(y_1,x_2) \le \tau(2r+\chi) + 1$, hence $D \le \tau(2r+\chi) + 2r + 1$.

Suppose that $\mathfrak{q}_2[x_2,y_2] \cap \overline{\mathcal{N}}_{2\tau r}(B) = \emptyset$. We have that x_2 is in $\mathfrak{q}_2[w_2,y_2]$. Also, $\mathfrak{q}_2[w_2,y_2]$ has its endpoints in $\mathcal{N}_\chi(B)$, with $\chi = \max(M, 2r(\tau+1))$. Consequently $\mathfrak{q}_2[w_2,y_2] \subset \mathcal{N}_{\tau\chi}(B)$. In particular x_2 is contained in $\mathcal{N}_{\tau\chi}(B)$ and $\mathrm{dist}_{S\cup\mathcal{H}}(y_1,x_2) \le \tau(2r+\chi) + 1$, hence $D \le \tau(2r + \chi) + 2r + 1$.

Case (a)II. Suppose that $y_2 \in \mathfrak{q}_2[x_2, O_3]$. If $y_3 \in \mathfrak{q}_3[x_3, O_1]$ then we repeat the previous argument with y_1 replaced by y_3. If $y_3 \in \mathfrak{q}_3[x_3, O_2]$ then we repeat the previous argument with (y_1, y_2) replaced by (y_3, y_1).

(b) We denote $\mathrm{dist}_{S\cup\mathcal{H}}(x, A)$ by D. We note that for every point y in $\mathcal{N}_r(A) \cap (\mathfrak{q}_1 \cup \mathfrak{q}_2 \cup \mathfrak{q}_3)$ we have that $\mathrm{dist}_S(x_i, y) \geq \mathrm{dist}_{S\cup\mathcal{H}}(x_i, y) \geq D - 2r$ for $i = 1, 2, 3$. We choose $y_i \in \mathcal{N}_r(A) \cap \mathfrak{q}_i$, $i = 1, 2, 3$. Suppose $y_1 \in \mathfrak{q}_1[x_1, O_3]$. Like in case (a), we have that y_1 is contained either in $\mathcal{N}_{2\tau r}(\mathfrak{q}_2[x_2, O_3])$ or in $\mathcal{N}_{2\tau r}(B)$ for some $B \in \mathcal{A}$ such that $\mathcal{N}_M(B)$ intersects $\mathfrak{q}_2[x_2, O_3]$.

Case (b)I. Suppose that $y_2 \in \mathfrak{q}_2[x_2, O_1]$.

Case (b)I.1. Assume that $y_1 \in \mathcal{N}_{2\tau r}(\mathfrak{q}_2[x_2, O_3])$. Then there exists $u \in \mathfrak{q}_2[x_2, O_3]$ such that $\mathrm{dist}_S(y_1, u) \leq 2\tau r$. It follows that $u \in \mathcal{N}_{r(1+2\tau)}(A)$ which together with $y_2 \in \mathcal{N}_r(A)$ implies that $\mathfrak{q}_2[u, y_2] \in \mathcal{N}_{\tau r(1+2\tau)}(A)$. In particular $x_2 \in \mathcal{N}_{\tau r(1+2\tau)}(A)$, therefore $D \leq r + \tau r(1 + 2\tau)$.

Case (b)I.2 Suppose $y_1 \in \mathcal{N}_{2\tau r}(B)$, with $B \in \mathcal{A}$ such that $\mathcal{N}_M(B)$ intersects $\mathfrak{q}_2[x_2, O_3]$. Let w_2 be a point in $\mathcal{N}_M(B) \cap \mathfrak{q}_2[x_2, O_3]$.

Suppose that $\mathfrak{q}_2[x_2, y_2] \cap \overline{\mathcal{N}}_{2\tau r}(B) \neq \emptyset$. As in the proof of part (a), Case I.2, we obtain that $\mathrm{dist}_{S\cup\mathcal{H}}(y_1, x_2) \leq \tau(2r + \chi) + 1$, whence $D \leq \tau(2r + \chi) + 2r + 1$.

Suppose that $\mathfrak{q}_2[x_2, y_2] \cap \overline{\mathcal{N}}_{2\tau r}(B) = \emptyset$. Then x_2 is in $\mathfrak{q}_2[w_2, y_2]$. On the other hand, $\mathfrak{q}_2[w_2, y_2]$ has its endpoints in the M-tubular neighborhood of $\mathrm{Sat}^{2\tau r}([y_1, y_2])$. It follows that $\mathfrak{q}_2[w_2, y_2]$, in particular x_2, is in the tM-tubular neighborhood of $\mathrm{Sat}^{2\tau r}([y_1, y_2])$. In $\mathrm{Cayley}(G, S \cup \mathcal{H})$, $\mathrm{Sat}^{2\tau r}([y_1, y_2])$ is contained in the $(2\tau r + 1)$-tubular neighborhood of $[y_1, y_2]$. Since in $\mathrm{Cayley}(G, S)$ we have that $[y_1, y_2] \subset \mathcal{N}_{\tau r}(A)$, we deduce that in $\mathrm{Cayley}(G, S \cup \mathcal{H})$, x_2 is in the $(tM + 3\tau r + 1)$-tubular neighborhood of A. Hence $D \leq tM + (3\tau + 1)r + 1$.

Case (b)II. Suppose that $y_2 \in \mathfrak{q}_2[x_2, O_3]$. Then we can use the same argument as in Case II of part (a). □

Proof of Proposition 8.16, (2). By Lemma 8.19 we are either in case (C) or in case (P).

Case (C). Let $y \in \mathcal{C}_\kappa^\mu(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$. According to Lemma 8.17 we have either (i) or (ii). Suppose that (i) is satisfied. Then, by Lemma 8.20, (a), $\mathrm{dist}_{S\cup\mathcal{H}}(x, y) \leq D$, where $D = D(\alpha, R, L, C)$.

Suppose that (ii) is satisfied, that is $y \in \mathcal{N}_\kappa(B)$ and $\mathcal{N}_\beta(B)$ intersects each of the three almost-geodesics $\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3$. Lemma 8.20, (b), implies that $\mathrm{dist}_{S\cup\mathcal{H}}(x, B) \leq D$, where $D = D(\beta, R, L, C)$. Therefore $\mathrm{dist}_{S\cup\mathcal{H}}(x, y) \leq D + \kappa + 1$.

Case (P). Let $y \in \mathcal{C}_\kappa^\mu(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$. Suppose that y satisfies (i). Lemma 8.20, (b) implies that $\mathrm{dist}_{S\cup\mathcal{H}}(y, A) \leq D$, with $D = D(\alpha, R, L, C)$.

If y satisfies (ii) of Lemma 8.17, then the unicity stated in (P) implies that $y \in \mathcal{N}_\kappa(A)$, hence that $\mathrm{dist}_{S\cup\mathcal{H}}(y, A) \leq \kappa$.

We may conclude that in all cases the diameter of the set $\mathcal{C}_\kappa^\mu(\mathfrak{q}_1, \mathfrak{q}_2, \mathfrak{q}_3)$ in the metric $\mathrm{dist}_{S\cup\mathcal{H}}$ is uniformly bounded. □

We now define a system of lines and centers in $\mathrm{Cayley}(G, S \cup \mathcal{H})$ such that the properties in Section 8.1.A are satisfied.

First of all, for every pair of vertices u, v in $\mathrm{Cayley}(G, S \cup \mathcal{H})$ we choose and fix a geodesic $[u, v]$ in $\mathrm{Cayley}(G, S)$ joining the two points. Let $M_0 = M(1, 0)$ and let κ_0 be the constant given by Proposition 8.16 for $\mu = M_0$. We may suppose that $\kappa_0 \geq M_0$. For every pair of vertices u, v in $\mathrm{Cayley}(G, S \cup \mathcal{H})$, we define Λ_{uv} as $\mathcal{N}_{\kappa_0}(\mathrm{Sat}([u, v]))$. The relation on it is defined as follows: to every $x \in \mathcal{N}_{\kappa_0}(\mathrm{Sat}([u, v]))$ we associate one nearest point (projection) $x' \in [u, v]$ and we put $x \leq_{uv} y$ if x' is between u and y'. Properties (l_1), (l_2), (l_3) are obviously satisfied.

We define the function ϕ by choosing, for every three vertices u, v, w in Cayley(G, S) a point C_{uvw} in $\mathcal{C}_{\kappa_0}^{M_0}([u,v],[u,w],[v,w])$ and defining $\phi(u,v,w) = \phi \circ \sigma(u,v,w) = C_{uvw}$ for every 3-permutation σ. We choose $C_{uuv} = u$.

Properties (c_1), (c_2), (c_3) are satisfied. Before proceeding further, we prove some intermediate results.

Lemma 8.21. *For every $\alpha > 0$ there exists $\lambda = \lambda(\alpha)$ such that the following holds. Let $[u,v]$ be a geodesic and let $A \in \mathcal{A}$ be such that $\mathcal{N}_\alpha(A) \cap [u,v] \neq \emptyset$. Let x be a point in $\mathcal{N}_\alpha(A)$ and let $x' \in [u,v]$ be a projection of x. Then $x' \in \mathcal{N}_\lambda(A)$.*

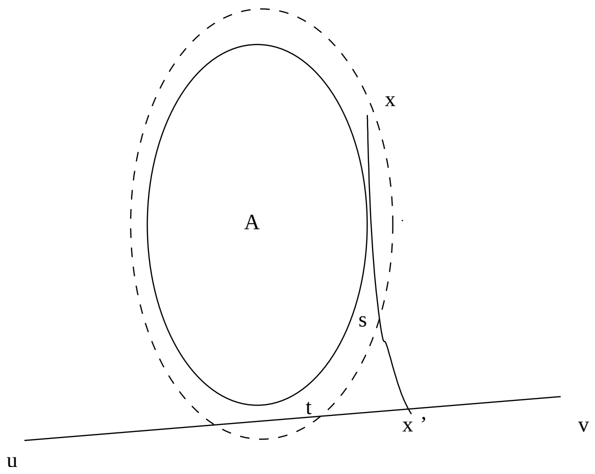

Figure 9: Projection of a point in the saturation.

Proof. Suppose that $x' \notin \mathcal{N}_\alpha(A)$. Lemma 8.15 implies that there exist $t \in [u,v] \cap \mathcal{N}_\alpha(A)$ and $s \in [x', x] \cap \mathcal{N}_\alpha(A)$ such that the sets $[x', t] \cap \overline{\mathcal{N}}_\alpha(A)$ and $[x', s] \cap \overline{\mathcal{N}}_\alpha(A)$ have diameters at most $\mathfrak{d}$, where $\mathfrak{d} = \mathfrak{d}(\alpha)$. Corollary 8.14 implies that $\mathrm{dist}(s,t) \leq D_1$. On the other hand, since $\mathrm{dist}(x,x') \leq \mathrm{dist}(x,t)$, it follows that $\mathrm{dist}(s,x') \leq \mathrm{dist}(s,t) \leq D_1$. We conclude that $\mathrm{dist}_S(x', A) \leq D_1 + \alpha$. $\square$

Corollary 8.22. *Let x be a point in $\mathcal{N}_\kappa(\mathrm{Sat}^\mu([u,v]))$ and let $x' \in [u,v]$ be a projection of x. Then $\mathrm{dist}_{S\cup\mathcal{H}}(x,x') \leq \chi$, where $\chi = \chi(\kappa, \mu)$.*

Proof. Since $x \in \mathcal{N}_\kappa(\mathrm{Sat}^\mu([u,v]))$ it follows that either $x \in \mathcal{N}_\kappa([u,v])$ or $x \in \mathcal{N}_\kappa(A)$, where $\mathcal{N}_\mu(A) \cap [u,v] \neq \emptyset$. In the first case it follows that $\mathrm{dist}_{S\cup\mathcal{H}}(x,x') \leq \kappa$, while in the second case we may apply Lemma 8.21. $\square$

Corollary 8.23. *Let u, v be a pair of vertices in* Cayley$(G,\, S \cup \mathcal{H})$ *and let $x, y \in \Lambda_{uv}$ and x', y' their chosen respective projections on $[u,v]$. Then, in* Cayley$(G,\, S\cup\mathcal{H})$*, $\Lambda_{uv}[x,y] = \Lambda_{uv}[y,x]$ is at Hausdorff distance χ of $[x',y'] \subset [u,v]$, where $\chi = \chi(G)$.*

Before proving properties (I), (II), (III), we make some remarks and introduce some notations.

Remarks 8.24. 1) For every quasi-geodesic $\mathfrak{q}$ in Cayley(G, S), we have that $\mathrm{Sat}^\mu(\mathfrak{q})$ is in the $(\mu+1)$-tubular neighborhood of $\mathfrak{q}$ in Cayley$(G,\, S \cup \mathcal{H})$.

2) Lemma 8.17 implies that there exist two constants η and c such that for every three geodesics $[u,v],[v,w],[u,w]$ in $\mathrm{Cayley}(G,S)$ every point $x \in \mathcal{C}_{\kappa_0}^{M_0}([u,v],[v,w],[u,w])$ satisfies one of the following two properties:

(i) the ball $B(x,\eta)$ intersects each of the three geodesics $[u,v],[v,w],[u,w]$;

(ii) $x \in \mathcal{N}_{\kappa_0}(A)$ and $\mathcal{N}_c(A)$ intersects each of the three geodesics $[u,v],[v,w],[u,w]$.

We note that the constants η and c depend on M_0 and κ_0, so they depend only on G. We may suppose without loss of generality that $c \geq M_0$.

3) Lemma 8.21 implies that there exists ξ such that if $[u,v]$ is a geodesic, $A \in \mathcal{A}$ is such that $\mathcal{N}_c(A)$ intersects $[u,v]$ and x is a point in $\mathcal{N}_{\kappa_0}(A)$, then any projection of x on $[u,v]$ is in $\mathcal{N}_\xi(A)$. The constant ξ depends on $\max(c,\kappa_0)$, so it depends only on G. Without loss of generality we may suppose that $\xi \geq M_0$.

4) In the sequel we denote the constant $\mathfrak{d}(1,0,c)$ provided by Lemma 8.15 simply by $\mathfrak{d}$.

Proof of properties (I), (II), (III).

(I). Let $x=\phi(u,v,w)$ and let x_1 and x_2 be the chosen projections of x on $[u,v]$ and on $[u,w]$, respectively. According to Corollary 8.23, it suffices to prove that $[u,x_1]$ and $[u,x_2]$ are at uniformly bounded Hausdorff distance in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$. The point $x=\phi(u,v,w)$ satisfies either (i) or (ii) from Remark 8.24, part 2.

Suppose x is in case (ii). Then $x \in \mathcal{N}_{\kappa_0}(A)$ such that $\mathcal{N}_c(A)$ intersects the three geodesic edges. Lemma 8.21 implies that $x_1,x_2 \in \mathcal{N}_\xi(A)$. The geodesic $[u,x_1]$ has its endpoints in $\mathcal{N}_\xi(\mathrm{Sat}^\xi[u,x_2])$. Lemma 8.10 implies that $[u,x_1]$ is entirely contained in $\mathcal{N}_{\tau\xi}(\mathrm{Sat}^\xi[u,x_2])$. It follows that $[u,x_1]$ is in the $[(\tau+1)\xi+1]$-tubular neighborhood of $[u,x_2]$ in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$. A similar argument done for $[u,x_2]$ allows to conclude that (I) is satisfied.

Suppose x is in case (i). Then $\mathrm{dist}_S(x,x_i)\leq\eta$ for $i=1,2$. Hence $\mathrm{dist}_S(x_1,x_2)\leq 2\eta$ and $[u,x_i]$ has its endpoints in $\mathcal{N}_{2\eta}(\mathrm{Sat}[u,x_j])$, for $\{i,j\}=\{1,2\}$. We repeat the previous argument.

(II) The fact that $\mathrm{dist}_{S\cup\mathcal{H}}(p,q)\leq 1$ means that either $\mathrm{dist}_S(p,q)\leq 1$ or $p,q\in A_0$, where $A_0\in\mathcal{A}$. Let $x=\phi(u,v,p)$ and $y=\phi(u,v,q)$. We have to show that $\Lambda_{uv}[x,y]$ has uniformly bounded diameter in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$. Let x_0 and y_0 be the respective projections of x and y on $[u,v]$. Corollary 8.23 implies that it suffices to prove that $[x_0,y_0]$ has uniformly bounded diameter in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$, where by $[x_0,y_0]$ we denote the sub-arc of $[u,v]$ of endpoints x_0,y_0.

Suppose that both x and y are in case (i). We have that $x_0\in\mathcal{N}_{2\eta}[u,p]\cap\mathcal{N}_{2\eta}[v,p]$ and that $y_0\in\mathcal{N}_{2\eta}[u,q]\cap\mathcal{N}_{2\eta}[v,q]$. Since $[u,p]\subset\mathcal{N}_\tau(\mathrm{Sat}[u,q])$ and $[v,p]\subset\mathcal{N}_\tau(\mathrm{Sat}[v,q])$, we conclude that $x_0,y_0\in\mathcal{C}_{2\eta+\tau}^{M_0}([u,q],[v,q],[u,v])$, hence that $[x_0,y_0]\subset\mathcal{C}_{\tau(2\eta+\tau)}^{M_0}([u,q],[v,q],[u,v])$. We complete the proof by applying Proposition 8.16.

Suppose x is in case (i) and y is in case (ii). The case when x is in case (ii) and y is in case (i) is discussed similarly. As above we have that $x_0\in\mathcal{C}_{2\eta+\tau}^{M_0}([u,q],[v,q],[u,v])$. We have that $y\in\mathcal{N}_{\kappa_0}(A)$ such that $\mathcal{N}_c(A)$ intersects $[u,q],[v,q],[u,v]$. Lemma 8.21 implies that $y_0\in\mathcal{N}_\xi(A)$. Then $y_0\in\mathcal{C}_\xi^c([u,q],[v,q],[u,v])$. As previously we obtain that $[x_0,y_0]\subset\mathcal{C}_{\tau' r}^s([u,q],[v,q],[u,v])$, where $r=\max(2\eta+\tau\,,\,\xi)$, $s=\max(M_0,c)$ and $\tau'\tau'(s)$. Proposition 8.16 allows to complete the argument.

Suppose that both x and y are in case (ii). Then $x\in\mathcal{N}_{\kappa_0}(A)$ such that $\mathcal{N}_c(A)$ intersects $[p,u],[p,v],[u,v]$. Let $p_1\in[u,p]\cap\mathcal{N}_c(A)$ and $p_2\in[v,p]\cap\mathcal{N}_c(A)$ be such that $[p,p_i]\cap\overline{\mathcal{N}}_c(A)$ has diameter at most $\mathfrak{d}$, $i=1,2$. Likewise we consider $u_1\in[u,v]\cap\mathcal{N}_c(A)$ and $u_2\in[u,p]\cap\mathcal{N}_c(A)$ so that $[u, u_i] \cap \overline{\mathcal{N}}_c(A)$ has diameter at most $\mathfrak{d}$, and $v_1 \in [p, v] \cap \mathcal{N}_c(A)$ and $v_2 \in [u, v] \cap \mathcal{N}_c(A)$ so that $[v, v_i] \cap \overline{\mathcal{N}}_c(A)$ has diameter at most $\mathfrak{d}$. Corollary 8.14 implies that $\operatorname{dist}_S(p_1, p_2), \operatorname{dist}_S(u_1, u_2)$ and $\operatorname{dist}_S(v_1, v_2)$ are at most ζ, where $\zeta = \zeta(G)$.

We have that either $A \subset \operatorname{Sat}[u, q]$ or $\mathcal{N}_c(A) \cap \mathcal{N}_\tau(\operatorname{Sat}[u, q])$ has diameter at most γ, where $\gamma = \gamma(G)$. The latter case implies, together with the inclusion $[u, p] \subset \mathcal{N}_\tau(\operatorname{Sat}[u, q])$, that $\operatorname{dist}(p_1, u_2) \leq \gamma$. Thus, we have that either $A \subset \operatorname{Sat}[u, q]$ or $\operatorname{dist}(p_1, u_2) \leq \gamma$. Likewise, we obtain that either $A \subset \operatorname{Sat}[v, q]$ or $\operatorname{dist}(p_2, v_1) \leq \gamma$.

Suppose that $\operatorname{dist}(p_1, u_2) \leq \gamma$. Then $\operatorname{dist}(p_1, u_1) \leq \gamma + \zeta$, hence $B(p_1, \gamma + \zeta)$ intersects $[p, u], [p, v], [u, v]$. We can argue similarly to the case above when x is in case (i) and y is in case (ii), with x replaced by p_1 and η by $\gamma + \zeta$. We obtain that if p_1' is the chosen projection of p_1 on $[u, v]$ then $[p_1', y_0]$ has the diameter bounded in $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$ by a constant depending on G. Since $[x_0, y_0] \subset [x_0, p_1'] \cup [p_1', y_0]$, it remains to prove that $[x_0, p_1']$ has bounded diameter in $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$. Lemma 8.21 provides for $\alpha = \max(\kappa_0, c)$ a constant $\tilde{\lambda}$. We have that x_0 and p_1' are in $\mathcal{N}_{\tilde{\lambda}}(A)$, hence that $[x_0, p_1'] \subset \mathcal{N}_{\tau\tilde{\lambda}}(A)$. We conclude that the diameter of $[x_0, p_1']$ in $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$ is at most $2\tau\tilde{\lambda} + 1$. A similar argument works if $\operatorname{dist}(p_2, v_1) \leq \gamma$.

Now suppose that $A \subset \operatorname{Sat}[u, q] \cap \operatorname{Sat}[v, q]$. Lemma 8.21 implies that $x_0 \in \mathcal{N}_\xi(A)$. Since y is also in case (ii), we have that $y \in \mathcal{N}_{\kappa_0}(B)$ such that $\mathcal{N}_c(B)$ intersects the three geodesic edges $[q, u], [q, v], [u, v]$ and that $y_0 \in \mathcal{N}_\xi(B)$. We have that $A \cup B \subset \operatorname{Sat}^c[u, q] \cap \operatorname{Sat}^c[v, q] \cap \operatorname{Sat}^c[u, v]$. Lemma 8.10 implies that $[x_0, y_0] \subset \mathcal{C}^c_{\tau\xi}([q, u], [q, v], [u, v])$ and Proposition 8.16 allows to finish the argument.

(III) Let u, v, w be three vertices such that $w \in \mathcal{N}_{\kappa_0}(\operatorname{Sat}[u, v])$. Let $x = \phi(u, v, w)$. Let w_0 and x_0 be the projections of w and x respectively on $[u, v]$. We bound the diameter of $[x_0, w_0]$ in $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$.

We have $x, w \in \mathcal{C}^{M_0}_{\kappa_0}([u, v], [u, w], [v, w])$. Suppose both x and w are in case (i). Then $x_0, w_0 \in \mathcal{C}^{M_0}_{\kappa_0 + \eta}([u, v], [u, w], [v, w])$, consequently $[x_0, w_0] \subset \mathcal{C}^{M_0}_{\tau(\kappa_0 + \eta)}([u, v], [u, w], [v, w])$ and we apply Proposition 8.16 to obtain the conclusion.

Suppose that x is in case (i) and w in case (ii). The case when x is in case (ii) and w in case (i) is similar. The ball $B(x, \eta)$ intersects the three edges and $w \in \mathcal{N}_{\kappa_0}(A)$ such that $\mathcal{N}_c(A)$ intersects the three edges. Lemma 8.21 implies that $w_0 \in \mathcal{N}_\xi(A) \subset \mathcal{C}^c_\xi([u, v], [u, w], [v, w])$. The point x_0 is in $\mathcal{C}^{M_0}_{\eta + \kappa_0}([u, v], [u, w], [v, w])$. It follows that $[x_0, w_0] \subset \mathcal{C}^s_{\tau' r}([u, v], [u, w], [v, w])$, where $r = \max(\eta + \kappa_0\,,\, \xi)$, $s = \max(M_0, c)$ and $\tau' = \tau'(s)$. We apply Proposition 8.16.

Suppose that x and w are both in case (ii). We have that $x \in \mathcal{N}_{\kappa_0}(A)$ and $w \in \mathcal{N}_{\kappa_0}(B)$ such that both $\mathcal{N}_c(A)$ and $\mathcal{N}_c(B)$ intersect the three edges. We also have that $x_0 \in \mathcal{N}_\xi(A)$ and $w_0 \in \mathcal{N}_\xi(B)$, hence $[x_0, w_0] \subset \mathcal{C}^c_{\tau\xi}([u, v], [u, w], [v, w])$. We end the proof by applying Proposition 8.16. $\square$

Proposition 8.7 implies that $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$ is hyperbolic. Moreover we have that Λ_{uv} is at bounded Hausdorff distance from every geodesic connecting u and v in $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$. Since in the previous argument the choice of the geodesics $[u, v]$ in $\operatorname{Cayley}(G, S)$ was arbitrary, we may write the following.

Proposition 8.25. *Every geodesic in* $\operatorname{Cayley}(G, S)$ *joining two points u and v is at bounded Hausdorff distance in* $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$ *from any geodesic joining u and v in* $\operatorname{Cayley}(G,\, S \cup \mathcal{H})$.

8.2 The BCP Property

We now prove the following.

Proposition 8.26. *If G is asymptotically tree-graded with respect to $\{H_1, \dots, H_m\}$ and G is weakly hyperbolic relative to $\{H_1, \dots, H_m\}$ then the pair $(G, \{H_1, \dots, H_m\})$ satisfies the BCP-property.*

Proof. Let $\lambda \geq 1$. Let $\mathfrak{p}$ and $\mathfrak{q}$ be two λ-bi-Lipschitz paths without backtracking in Cayley$(G, S \cup \mathcal{H})$ such that $\mathfrak{p}_- = \mathfrak{q}_-$ and $\mathrm{dist}_S(\mathfrak{p}_+, \mathfrak{q}_+) \leq 1$.

(1) Let s be an $\mathcal{H}$-component of $\mathfrak{p}$ contained in a left coset $A \in \mathcal{A}$, and let $\mathrm{dist}_S(s_-, s_+) = D$. We show that if D is large enough then $\mathfrak{q}$ has an $\mathcal{H}$-component contained in A.

Notations: In this section M denotes $M(\lambda, 0)$, the constant given by (α_2') for $\theta = \frac{1}{3}$ and $(L, C) = (\lambda, 0)$.

The graph Cayley$(G, S \cup \mathcal{H})$ is hyperbolic. Therefore for the given λ there exists $\varkappa = \varkappa(\lambda)$ such that two λ-bi-Lipschitz paths $\mathfrak{p}$ and $\mathfrak{q}$ in Cayley$(G, S \cup \mathcal{H})$ with $\mathrm{dist}_{S \cup \mathcal{H}}(\mathfrak{p}_-, \mathfrak{q}_-) \leq 1$ and $\mathrm{dist}_{S \cup \mathcal{H}}(\mathfrak{p}_+, \mathfrak{q}_+) \leq 1$ are at Hausdorff distance at most $\varkappa$.

Given two vertices u, v in Cayley$(G, S \cup \mathcal{H})$, we denote by $[u, v]$ a geodesic joining them in Cayley(G, S) and by $\mathfrak{g}_{uv}$ a geodesic joining them in Cayley$(G, S \cup \mathcal{H})$.

For a path $\mathfrak{p}$ in Cayley$(G, S \cup \mathcal{H})$, we denote by $\tilde{\mathfrak{p}}$ a path in Cayley(G, S) obtained by replacing every $\mathcal{H}$-component s in $\mathfrak{p}$ by a geodesic in Cayley(G, S) connecting s_- and s_+. We call $\tilde{\mathfrak{p}}$ a *lift* of $\mathfrak{p}$.

Step I. We show that for $D \geq D_0(G)$, some lift $\tilde{\mathfrak{q}}$ of $\mathfrak{q}$ intersects $\mathcal{N}_{M'}(A)$, where $M' = M'(G)$.

We choose u on the arc $\mathfrak{p}[\mathfrak{p}_-, s_-]$ such that either the length of $\mathfrak{p}[u, s_-]$ is $2\lambda(\varkappa + 1)$ or, if the length of $\mathfrak{p}[\mathfrak{p}_-, s_-]$ is less than $2\lambda(\varkappa + 1)$, $u = \mathfrak{p}_-$. Likewise we choose v on the arc $\mathfrak{p}[s_+, \mathfrak{p}_+]$ such that either the length of $\mathfrak{p}[s_+, v]$ is $2\lambda(\varkappa + 1)$ or $v = p_+$. We have that $\mathrm{dist}_{S \cup \mathcal{H}}(u, s_-)$, $\mathrm{dist}_{S \cup \mathcal{H}}(s_+, v) \in [2(\varkappa + 1)\, ,\, 2\lambda^2(\varkappa + 1)]$, in the first cases.

There exist w and z on $\mathfrak{q}$ such that $\mathrm{dist}_{S \cup \mathcal{H}}(u, w) \leq \varkappa$ and $\mathrm{dist}_{S \cup \mathcal{H}}(v, z) \leq \varkappa$. We consider $\mathfrak{g}_{uw}$ and $\mathfrak{g}_{vz}$ geodesics in Cayley$(G, S \cup \mathcal{H})$.

Let u' be the farthest from u point on $\mathfrak{g}_{uw}$ which is contained in the same left coset $B \in \mathcal{A}$ as an $\mathcal{H}$-component σ of $\mathfrak{p}[u, v]$. Suppose that $\sigma \cap \mathfrak{p}[s_-, v] \neq \emptyset$. We have that

$$\mathrm{dist}_{S \cup \mathcal{H}}(u, u') \geq \mathrm{dist}_{S \cup \mathcal{H}}(u, \sigma_+) - 1 \geq \frac{1}{\lambda}\mathrm{length}\,(\mathfrak{p}[u, \sigma_+]) - 1 \geq \frac{1}{\lambda}\mathrm{length}\,(\mathfrak{p}[u, s_-]) - 1 \geq 2\varkappa + 1\,.$$

This contradicts the inequality $\mathrm{dist}_{S \cup \mathcal{H}}(u, u') \leq \varkappa$. Therefore σ is contained in $\mathfrak{p}[u, s_-] \setminus \{s_-\}$. We choose v' the farthest from v point on $\mathfrak{g}_{vz}$ contained in the same left coset as a component σ' of $\mathfrak{p}[u, v]$. In a similar way we prove that σ' is contained in $\mathfrak{p}[s_+, v] \setminus \{s_+\}$. It is possible that $u' = u$, $\sigma = \{u\}$ and/or $v' = v$, $\sigma' = \{v\}$.

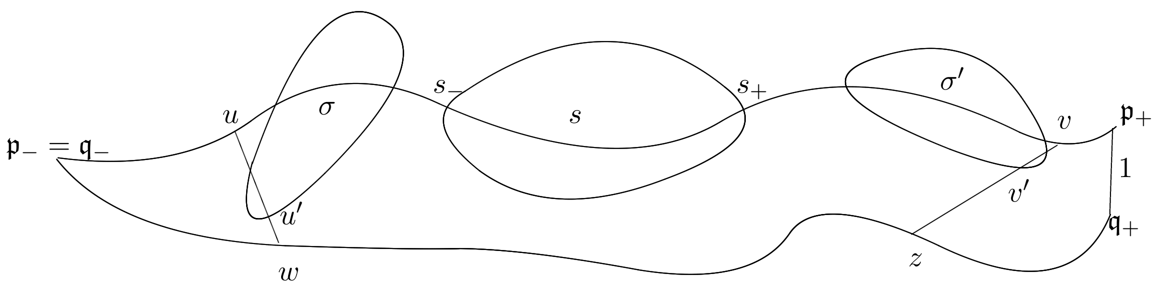

Figure 10: Proof of (1) in BCP Property.

We consider the path in Cayley$(G, S \cup \mathcal{H})$ defined as $\mathfrak{r} = \mathfrak{g}_{wu'} \sqcup \mathfrak{g}_{u'\sigma_+} \sqcup \mathfrak{p}[\sigma_+, \sigma'_-] \sqcup \mathfrak{g}_{\sigma'_- v'} \sqcup \mathfrak{g}_{v'z}$, where $\mathfrak{g}_{wu'}$ and $\mathfrak{g}_{v'z}$ are sub-geodesics of $\mathfrak{g}_{wu}$ and $\mathfrak{g}_{vz}$, respectively, and $\mathfrak{g}_{u'\sigma_+}$ and $\mathfrak{g}_{\sigma'_- v'}$ are composed of one edge. The length of $\mathfrak{r}$ is at most $N = \lambda(4\varkappa+5)+2\varkappa$. It contains the component s. We show that it has no backtracking. By construction and the fact that geodesics do not have backtracking ([Os, Lemma 2.23]), we have that the sub-arcs $\mathfrak{r}[w, v']$ and $\mathfrak{r}[u', z]$ do not have backtracking. Suppose that $\mathfrak{g}_{wu'}$ and $\mathfrak{g}_{v'z}$ have $\mathcal{H}$-components in the same left coset. It follows that there exists $x \in \mathfrak{g}_{wu'}$ and $y \in \mathfrak{g}_{v'z}$ with $\mathrm{dist}_{S\cup\mathcal{H}}(x, y) \leq 1$. Then $\mathrm{dist}_{S\cup\mathcal{H}}(u, v) \leq 2\varkappa+1$. By construction either $\mathrm{length}\, \mathfrak{p}[u, v] \geq 2\lambda(\varkappa + 1) + 1$ or $u = \mathfrak{p}_-$ and $v = \mathfrak{p}_+$. In the latter case, the geodesic $\mathfrak{g}_{wu}$ is trivial, $\mathfrak{g}_{vz}$ is an edge e in $\mathrm{Cayley}(G, S)$, and $\mathfrak{r} = \mathfrak{p} \cup e$ has no backtracking. In the former case we have that $\mathrm{dist}_{S\cup\mathcal{H}}(u, v) > 2\varkappa + 2$, which contradicts the previous inequality.

We conclude that $\mathfrak{r}$ is without backtracking. A lift $\tilde{\mathfrak{r}}$ of it is composed of n consecutive sub-paths,

$$\tilde{\mathfrak{r}} = \tilde{\mathfrak{r}}_1 \cup \cdots \cup \tilde{\mathfrak{r}}_n \,, \tag{27}$$

with $n \leq N$, such that each $\tilde{\mathfrak{r}}_i$ is either

(R_1) a λ-bi-Lipschitz arc both in $\mathrm{Cayley}(G, S)$ and in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$ of length at most N or

(R_2) a geodesic in $\mathrm{Cayley}(G, S)$ with endpoints in some left coset $A_i \in \mathcal{A}$.

Since $\mathfrak{r}$ is without backtracking, we have that $A_i \neq A_j$ when $i \neq j$. Lemma 8.12 implies that $\tilde{\mathfrak{r}}$ is an (L_N, C_N)-almost geodesic.

On the other hand, $\mathrm{dist}_{S\cup\mathcal{H}}(w, z) \leq \mathrm{length}\, \mathfrak{r} \leq N$. Hence the length of $\mathfrak{q}[w, z]$ is at most N_1, where $N_1 = \lambda N$. As above, a lift $\tilde{\mathfrak{q}}[w, z]$ decomposes into m consecutive sub-paths,

$$\tilde{\mathfrak{q}}[w, z] = \tilde{\mathfrak{q}}_1 \cup \cdots \cup \tilde{\mathfrak{q}}_m \,, \tag{28}$$

with $m \leq N_1$, such that each $\tilde{\mathfrak{q}}_i$ is either

(Q_1) a λ-bi-Lipschitz arc both in $\mathrm{Cayley}(G, S)$ and $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$, of length at most N_1 or

(Q_2) a geodesic in $\mathrm{Cayley}(G, S)$ with endpoints in some left coset $B_i \in \mathcal{A}$.

Since $\mathfrak{q}$ is without backtracking, we have that $B_i \neq B_j$ when $i \neq j$. Lemma 8.12 implies that $\tilde{\mathfrak{q}}[w, z]$ is an (L_{N_1}, C_{N_1})-almost geodesic. We denote $L' = \max(L_N, L_{N_1})$ and $C' = \max(C_N, C_{N_1})$. We denote $M' = M(L', C')$. Lemma 4.25 implies that in $\mathrm{Cayley}(G, S)$ the path $\tilde{\mathfrak{r}}$ is contained in the τ'-tubular neighborhood of $\mathrm{Sat}(\tilde{\mathfrak{q}}[w, z]) = \mathrm{Sat}^{M'}(\tilde{\mathfrak{q}}[w, z])$, where $\tau' = \tau'(L', C')$. In particular the component s is contained in $\mathcal{N}_{\tau'}(\mathrm{Sat}(\tilde{\mathfrak{q}}[w, z]))$, hence the set $\mathcal{N}_{\tau'}(\mathrm{Sat}(\tilde{\mathfrak{q}}[w, z])) \cap A$ has diameter at least D. Lemma 4.22 implies that for $D \geq D_0(L', C', \tau')$ we must have that $\mathcal{N}_{M'}(A) \cap \tilde{\mathfrak{q}}[w, z] \neq \emptyset$.

Step II. We show that there exist two points w_1 and z_1 on $\tilde{\mathfrak{q}}[w, z]$ such that $\mathrm{dist}_S(w_1, s_-) \leq D_1$ and $\mathrm{dist}_S(z_1, s_+) \leq D_1$, where $D_1 = D_1(G)$. We do this by means of Corollary 8.14.

Lemma 8.15 implies that there exist $w_1, z_1 \in \tilde{\mathfrak{q}}[w, z] \cap \mathcal{N}_{M'}(A)$ such that $\tilde{\mathfrak{q}}[w, w_1]$ and $\tilde{\mathfrak{q}}[z_1, z]$ intersect $\overline{\mathcal{N}}_{M'}(A)$ in two sets of diameter at most $\mathfrak{d}_1$, where $\mathfrak{d}_1 = \mathfrak{d}_1(L', C', M')$.

We show that $\tilde{\mathfrak{r}}[w, s_-]$ and $\tilde{\mathfrak{r}}[s_+, z]$ intersect $\overline{\mathcal{N}}_{M'}(A)$ in two sets of bounded diameter. We prove it only for $\tilde{\mathfrak{r}}[w, s_-]$, the same argument works for $\tilde{\mathfrak{r}}[s_+, z]$. Let $x \in \tilde{\mathfrak{r}}[w, s_-] \cap \overline{\mathcal{N}}_{M'}(A)$ and let $\mathrm{dist}_S(x, s_-) = \delta$. According to the decomposition (27), we have that $\tilde{\mathfrak{r}}[x, s_-] = \tilde{\mathfrak{r}}_i' \cup \tilde{\mathfrak{r}}_{i+1} \cup \cdots \cup \tilde{\mathfrak{r}}_j$, where $i \leq j$, $i, j \in \{1, 2, \ldots, n\}$ and $\tilde{\mathfrak{r}}_i'$ is eventually a restriction of $\tilde{\mathfrak{r}}_i$ such that x is an endpoint of it. If all the components are of type (R_1), then $\tilde{\mathfrak{r}}[x, s_-]$ has length at most N and $\delta \leq N$. Suppose that at least one component is of type (R_2). We have at most N such components. Then at least one component $\tilde{\mathfrak{r}}_k$ of type (R_2) has the distance between its endpoints at least $\frac{\delta - N}{N}$. On the other hand since $x, s_- \in \mathcal{N}_{M'+1}(A)$ and $\tilde{\mathfrak{r}}[x, s_-]$ is an (L', C')-almost-geodesic, it follows that $\tilde{\mathfrak{r}}[x, s_-] \subset \mathcal{N}_{\tau'(M'+1)}(A)$. In particular $\tilde{\mathfrak{r}}_k$ is contained in the same tubular neighborhood, therefore the diameter of $A_k \cap \mathcal{N}_{\tau'(M'+1)}(A)$ is at least $\frac{\delta - N}{N}$. There exists $\delta_0 = \delta_0(L', C', N)$ such that if $\delta \geq \delta_0$ then $A_k = A$. This contradicts the fact that $\mathfrak{r}$ is without backtracking. We conclude that $\delta \leq \delta_0$.

We apply Corollary 8.14 to $\tilde{\mathfrak{q}}[w, w_1]$ and to $\tilde{\mathfrak{r}}[w, s_-]$ and we obtain that $\mathrm{dist}_S(w_1, s_-) \leq D_1$, where $D_1 = D_1(L', C', \delta_0)$. With a similar argument we obtain that $\mathrm{dist}_S(z_1, s_+) \leq D_1$.

Step III. We show that $\mathfrak{q}$ has a component in A.

We have that $\mathrm{dist}_S(w_1, z_1) \geq D - 2D_1$ and that $\tilde{\mathfrak{q}}[w_1, z_1] \subset \mathcal{N}_{\tau' D_1}(A)$. The decomposition (28) implies that $\tilde{\mathfrak{q}}[w_1, z_1] = \tilde{\mathfrak{q}}_k' \cup \tilde{\mathfrak{q}}_{k+1} \cup \cdots \cup \tilde{\mathfrak{q}}_{l-1} \cup \tilde{\mathfrak{q}}_l'$, where $k \leq l$, $k, l \in \{1, 2, \ldots, N_1\}$ and $\tilde{\mathfrak{q}}_k'$, $\tilde{\mathfrak{q}}_l'$ are eventually restrictions of $\tilde{\mathfrak{q}}_k$, $\tilde{\mathfrak{q}}_l$, respectively, with endpoints w_1 and z_1. If $D - 2D_1 > N_1$ it follows that $\tilde{\mathfrak{q}}[w_1, z_1]$ has at least a component of type (Q_2). Since it has at most N_1 such components, we may moreover say that $\tilde{\mathfrak{q}}[w_1, z_1]$ has at least a component $\tilde{\mathfrak{q}}_i$ with endpoints at distance at least $\frac{D-2D_1-N_1}{N_1}$. Consequently the diameter of $B_i \cap \mathcal{N}_{\tau' D_1}(A)$ is at least $\frac{D-2D_1-N_1}{N_1}$. For D large enough we obtain that $B_i = A$. We conclude that $\mathfrak{q}$ has a component in A.

(2) Suppose that s and t are $\mathcal{H}$-components of $\mathfrak{p}$ and $\mathfrak{q}$, respectively, contained in a left coset $A \in \mathcal{A}$. We show that $\mathrm{dist}_S(s_-, t_-)$ and $\mathrm{dist}_S(s_+, t_+)$ are bounded by a constant depending on G.

We take $u \in \mathfrak{p}[\mathfrak{p}_-, s_-]$ either such that the length of $\mathfrak{p}[u, s_-]$ is $2\lambda(\varkappa + 1)$ or, if the length of $\mathfrak{p}[\mathfrak{p}_-, s_-]$ is less than $2\lambda(\varkappa + 1)$, $u = \mathfrak{p}_-$. Likewise we take $v \in \mathfrak{p}[s_+, \mathfrak{p}_+]$ either such that the length of $\mathfrak{p}[s_+, v]$ is $2\lambda(\varkappa + 1)$ or, if the length of $\mathfrak{p}[s_+, \mathfrak{p}_+]$ is less than $2\lambda(\varkappa + 1)$, $v = \mathfrak{p}_+$.

Since $\mathrm{dist}_{S \cup \mathcal{H}}(s_-, t_-) \leq 1$ and Cayley$(G, S \cup \mathcal{H})$ is hyperbolic, there exists $w \in \mathfrak{q}[\mathfrak{q}_-, t_-]$ such that $\mathrm{dist}_{S \cup \mathcal{H}}(u, w) \leq \varkappa$. Similarly, $\mathrm{dist}_{S \cup \mathcal{H}}(s_+, t_+) \leq 1$ implies the existence of $z \in \mathfrak{q}[t_+, \mathfrak{q}_+]$ such that $\mathrm{dist}_{S \cup \mathcal{H}}(v, z) \leq \varkappa$. We consider two geodesics $\mathfrak{g}_{uw}$ and $\mathfrak{g}_{vz}$. As in Step 1 of the proof of (1), we show that the path $\mathfrak{g}_{wu} \cup \mathfrak{p}[u, v] \cup \mathfrak{g}_{vz}$ can be modified to give a path $\mathfrak{r}$ with endpoints w and z and of length at most N, without backtracking, containing s, such that any of its lifts, $\tilde{\mathfrak{r}}$, decomposes as in (27) and it is an (L', C')-almost-geodesic. Again as in Step I of the proof of (1), we show that the length of $\mathfrak{q}[w, z]$ is at most N_1 and that any lift $\tilde{\mathfrak{q}}[w, z]$ decomposes as in (28) and it is an (L', C')-almost-geodesic.

With an argument as in Step II of the proof of (1), we show that $\tilde{\mathfrak{r}}[w, s_-]$ and $\tilde{\mathfrak{r}}[s_+, z]$ intersect $\mathcal{N}_{M'}(A)$ in sets of diameter at most δ_0. The same argument can be used to show that $\tilde{\mathfrak{q}}[w, t_-]$ and $\tilde{\mathfrak{q}}[t_+, z]$ intersect $\mathcal{N}_{M'}(A)$ in sets of diameter at most $\delta_0' = \delta_0'(L', C', N_1)$. Corollary 8.14 implies that $\mathrm{dist}_S(s_-, t_-)$ and $\mathrm{dist}_S(s_+, t_+)$ are at most D_1, where $D_1 = D_1(L', C', \delta_0, \delta_0')$. $\square$

8.3 Undistorted subgroups of relatively hyperbolic groups

Theorem 8.27. *Let $G = \langle S \rangle$ be a finitely generated group that is hyperbolic relative to subgroups $H_1, ..., H_n$. Let $G_1 = \langle S_1 \rangle$ be an undistorted finitely generated subgroup of G. Then G_1 is relatively hyperbolic with respect to subgroups $H_1', ..., H_m'$, where each H_i' is one of the intersections $G_1 \cap gH_jg^{-1}$, $g \in G$.*

Proof. Since G_1 is undistorted, there exists a constant $D \geq 1$ such that for every element $g \in G_1$, $|g|_{S_1} \leq D|g|_S$. Here by $|g|_S$ and $|g|_{S_1}$ we denote the length of g in G and G_1 respectively. We can assume that $S_1 \subseteq S$ so that the graph Cayley(G_1, S_1) is inside Cayley(G, S). Then every geodesic in Cayley(G_1, S_1) is a $(D, 0)$-quasi-geodesic of Cayley(G, S).

Step I. Let us prove that for every coset gH_i and every constant $C > 0$ there exists $C' = C'(C, g, i) > 0$ such that $G_1 \cap \mathcal{N}_C(gH_i) \subseteq \mathcal{N}_{C'}(G_1 \cap gH_ig^{-1})$. By contradiction, let $(x_j)_{j \in \mathbb{N}}$ be a sequence of elements in G_1 such that $x_j = gh_jp_j \in G_1$, $h_j \in H_i$, $|p_j|_S < C$, and $\mathrm{dist}(x_j, G_1 \cap gH_ig^{-1}) \geq j$ for every j. Without loss of generality we can assume that $p_j = p$ is constant. Then $x_jx_1^{-1} \in G_1 \cap gH_ig^{-1}$. Hence $\mathrm{dist}(x_j, G_1 \cap gH_ig^{-1}) \leq |x_1|_S$, a contradiction.

Step II. Let $R > 0$ and let gH_i be such that $\mathcal{N}_R(gH_i) \cap G_1 \neq \emptyset$.

We prove that for every $K > 0$ there exists $K' = K'(K, R)$ such that

$$G_1 \cap \mathcal{N}_K(gH_i) \subset \mathcal{N}_{K'}(G_1 \cap g_1\gamma H_i\gamma^{-1})$$

for some $g_1 \in G_1$ and some $\gamma \in G$ with $|\gamma|_S \leq R$.

Fix $K > 0$ and define K' as the maximum of numbers $C'(K, \gamma, i)$ defined in Step I taken over all $i \in \{1, 2, \ldots, n\}$ and all $\gamma \in G$ with $|\gamma|_S \leq R$.

Let $g \in G$ be such that $G_1 \cap \mathcal{N}_R(gH_i) \neq \emptyset$. Let g_1 be an element of the intersection. Then $g_1^{-1}\mathcal{N}_R(gH_i) = \mathcal{N}_R(g_1^{-1}gH_i)$ contains 1, hence $g_1^{-1}gH_i = \gamma H_i$ where $|\gamma|_S \leq R$.

Step I and the choice of K' imply that

$$G_1 \cap \mathcal{N}_K(\gamma H_i) \subset \mathcal{N}_{K'}(G_1 \cap \gamma H_i\gamma^{-1}) .$$

Multiplying this inclusion by g_1 on the left, we obtain

$$G_1 \cap \mathcal{N}_K(gH_i) \subset \mathcal{N}_{K'}(G_1 \cap g_1\gamma H_i\gamma^{-1}).$$

Step III. Let $R = M(D, 0, \frac{1}{3})$ be the constant given by the property (α_2') satisfied by the left cosets $\{gH_i \mid g \in G,\, i = 1, 2, \ldots, n\}$ in Cayley(G, S).

For every $i \in \{1, ..., n\}$ consider the following equivalence relation on the ball $B(1, R)$ in G:

$$\gamma \sim_i \gamma' \text{ iff } G_1\gamma H_i = G_1\gamma' H_i.$$

For each pair (γ, γ') of $\sim_i$-equivalent elements in $B(1, R)$ we choose one $g_1 \in G_1$ such that $\gamma \in g_1\gamma' H_i$. Let $\tilde{C}$ be the maximal length of these elements g_1.

Let $\mathcal{M}$ be the collection of all nontrivial subgroups of G_1 in the set

$$\{G_1 \cap \gamma H_i \gamma^{-1} \mid i \in \{1, 2, \ldots, n\}\, ,\, |\gamma|_S \leq R\}.$$

By Step II, this collection of subgroups has the property that for every $K > 0$ there exists $K' = K'(K, R)$ such that for every $g \in G$ with $\mathcal{N}_R(gH_i) \cap G_1 \neq \emptyset$, we have

$$G_1 \cap \mathcal{N}_K(gH_i) \subset \mathcal{N}_{K'}(g_1 H) \tag{29}$$

for some $g_1 \in G_1$ and $H \in \mathcal{M}$.

We say that two non-trivial subgroups $G_1 \cap \gamma H_i \gamma^{-1}$ and $G_1 \cap \beta H_i \beta^{-1}$ from $\mathcal{M}$ are equivalent if $\gamma \sim_i \beta$.

Let $H_1', ..., H_m'$ be the set of representatives of equivalent classes in $\mathcal{M}$. If $\mathcal{M}$ is empty, we set $m = 1$, $H_1' = \{1\}$.

Notice that for every $H \in \mathcal{M}$ there exists $j \in \{1, ..., m\}$ such that H is at Hausdorff distance at most $\tilde{C}$ from a left coset gH_j' from G_1. Indeed, $H = \gamma H_i \gamma^{-1} \cap G_1$. Let $H_j' = \beta H_i \beta^{-1} \cap G_1$ be equivalent to H. Then $\gamma = g\beta h$ for some $g \in G_1$, $h \in H_i$, where $|g| \leq \tilde{C}$. Then

$$H = g\beta h H_i h^{-1}\beta^{-1}g^{-1} \cap G_1 = gH_j' g^{-1} ,$$

from which we deduce that H is at Hausdorff distance at most $\tilde{C}$ from gH_j'.

Hence (29) remains true if we replace $\mathcal{M}$ by the smaller set $\{H_1', ..., H_m'\}$ and K' by $K' + \tilde{C}$.

We shall prove that G_1 is relatively hyperbolic with respect to $\{H_1', \ldots, H_m'\}$ by checking properties (α_1), $(\alpha_2^{\frac{1}{6D}})$, (α_3) from Theorem 4.1 and Remark 4.2 for the collection of left cosets $\{g_1 H_j' \mid g_1 \in G_1, j = 1, 2, \ldots, m\}$.

Property (α_1). Consider $g_1H'_j \neq g'_1H'_k$. We have that

$$\mathcal{N}_\delta(g_1H'_j)\cap\mathcal{N}_\delta(g'_1H'_k) \subset \mathcal{N}_\delta(g_1\gamma H_{i_j}\gamma^{-1})\cap\mathcal{N}_\delta(g'_1\gamma' H_{i_k}(\gamma')^{-1}) \subset \mathcal{N}_{\delta+R}(g_1\gamma H_{i_j})\cap\mathcal{N}_{\delta+R}(g'_1\gamma' H_{i_k}) .$$

Suppose that $g_1\gamma H_{i_j} = g'_1\gamma' H_{i_k}$. Then $(g_1\gamma)^{-1}g'_1\gamma' \in H_{i_j}$ hence $g_1\gamma H_{i_j} = g'_1\gamma' H_{i_j}$. We deduce that $H_{i_j} = H_{i_k}$. Therefore $g_1\gamma = g'_1\gamma' h$ for some $h \in H_{i_j}$. Hence $\gamma \sim_{i_j} \gamma'$, so $\gamma = \gamma'$. We deduce that $g_1\gamma H_{i_j}\gamma^{-1} = g'_1\gamma H_{i_j}\gamma^{-1}$. So $g_1H'_j = g'_1H'_k$, a contradiction.

Thus, $g_1\gamma H_{i_j} \neq g'_1\gamma' H_{i_k}$. Property (α_1) satisfied by the left cosets $\{gH_i \mid g \in G, i = 1, 2, \dots, n\}$ allows to complete the proof.

Property $(\alpha_2^{\frac{1}{6D}})$. Let $\theta_1 \in \left[0, \frac{1}{6D}\right)$. We may write $\theta_1\frac{\theta}{D}$, with $\theta \in \left[0, \frac{1}{6}\right)$. Let $\mathfrak{g}\colon [0,\ell] \to \mathrm{Cayley}(G_1, S_1)$ be a geodesic of length ℓ in $\mathrm{Cayley}(G_1, S_1)$ with endpoints in $\mathcal{N}_{\theta_1\ell}(g_1H'_j) \subset \mathcal{N}_{\theta_1\ell+R}(g_1\gamma H_i)$, where $|\gamma|_S \leq R$ and $i \in \{1, 2, \dots, n\}$. Then $\mathfrak{g}$ is a $(D, 0)$-quasi-geodesic in $\mathrm{Cayley}(G, S)$.

Suppose that $\ell \leq 6DR$. Then $\mathfrak{g}$ is contained in the $(3DR+R)$-tubular neighborhood of $g_1H'_j$ in $\mathrm{Cayley}(G_1, S_1)$.

Suppose that $\ell > 6DR$. Then the endpoints of $\mathfrak{g}$ are contained in $\mathcal{N}_{(\theta+\frac{1}{6})\frac{\ell}{D}}(g_1\gamma H_i) \subset \mathcal{N}_{\frac{1}{3}\frac{\ell}{D}}(g_1\gamma H_i)$ in $\mathrm{Cayley}(G, S)$. Since the property (α'_2) is satisfied by the cosets of H_i in G, it follows that $\mathfrak{g}$ intersects $\mathcal{N}_R(g_1\gamma H_i)$. Hence $\mathfrak{g}$ intersects $G_1 \cap \mathcal{N}_R(g_1\gamma H_i) = g_1[G_1 \cap \mathcal{N}_R(\gamma H_i)] \subset g_1\mathcal{N}_{R'}(H'_j)$ where $R' = R'(R, R)$ is given by Step II.

We conclude that $\mathfrak{g}$ intersects $\mathcal{N}_{M'}(g_1H'_j)$ in $\mathrm{Cayley}(G_1, S_1)$, for $M' = \sup(DR', 3DR + R)$.

Property (α_3). We use the property (29) of $\{H'_1, \dots, H'_m\}$ and the property (α'_3) satisfied by the cosets of groups H_i.

Fix an integer $k \geq 2$. Let P be a $(\vartheta, 2, 8D)$-fat geodesic k-gon in $\mathrm{Cayley}(G_1, S_1)$ for some ϑ. Then P has $(D, 0)$-quasi-geodesic sides in $\mathrm{Cayley}(G, S)$ and it is $\left(\frac{\vartheta}{D}, 2D, 8D\right)$-fat. Consequently, for ϑ large enough, by property (α'_3) satisfied by the left cosets $\{gH_i \mid g \in G,\, i = 1, \dots, n\}$, the k-gon P is contained in a tubular neighborhood $\mathcal{N}_\varkappa(gH_i)$ in $\mathrm{Cayley}(G, S)$ for some $\varkappa > 0$.

Suppose that all edges of P have lengths at most $3D\varkappa$ in $\mathrm{Cayley}(G_1, S_1)$. Then P has diameter at most $3kD\varkappa$ in the same Cayley graph.

Suppose that one edge $\mathfrak{g}$ of P has length at least $3D\varkappa$. This, the fact that $P \subset \mathcal{N}_\varkappa(gH_i)$ and property (α'_2) satisfied by the left cosets $\{gH_i\}$ implies that $\mathfrak{g}$ intersects $\mathcal{N}_R(gH_i)$, therefore $\mathcal{N}_R(gH_i) \cap G_1 \neq \emptyset$.

Then by (29) there exists $\varkappa' = \varkappa'(\varkappa, R)$ such that

$$G_1 \cap \mathcal{N}_\varkappa(gH_i) \subset \mathcal{N}_{\varkappa'}(g_1H'_j)$$

for some $g_1 \in G_1$ and $j \in \{1, 2, \dots, m\}$. We conclude that in this case $P \subset \mathcal{N}_{\varkappa'}(g_1H'_j)$.

Thus we can take ξ needed in (α_3) to be the maximum of $3kD\varkappa$ and $\varkappa'$. □

Remarks 8.28. (1) If in Theorem 8.27 the subgroup G_1 is infinite and asymptotically without cut-points then G_1 is inside a conjugate of one of the subgroups H_i.

(2) If the subgroup G_1 intersects with all conjugates of the subgroups $H_1, ..., H_n$ by hyperbolic subgroups then G_1 is hyperbolic.

Proof. (1) Indeed, Corollary 5.5 implies that G_1 is contained in the K-tubular neighborhood of a left coset gH_i, where K depends only on the non-distortion constants. For every $g_1 \in G_1$, $G_1 = g_1G_1$ is contained in the K-tubular neighborhoods of g_1gH_i and of gH_i. Since G_1 is infinite, property (α_1) implies that $g_1gH_i = gH_i$. We conclude that G_1 is contained in gH_ig^{-1}.

(2) By Theorem 8.27 G_1 is relatively hyperbolic with respect to hyperbolic subgroups, so every asymptotic cone of G_1 is tree-graded with respect to $\mathbb{R}$-trees, whence it is an $\mathbb{R}$-tree itself. Therefore G_1 is hyperbolic [Gr$_3$]. $\square$

Corollary 8.29. *Let G be a finitely generated group that is relatively hyperbolic with respect to subgroups $H_1, ..., H_m$. Suppose that each H_i is infinite and either asymptotically without cut-points or does not contain a copy of H_1. Let* $\mathrm{Fix}(H_1)$ *be the subgroup of the automorphism group of G consisting of the automorphisms that fix H_1 as a set. Then:*

(1) $\mathrm{Inn}(G)\mathrm{Fix}(H_1)$ *has index at most* $m!$ *in* $\mathrm{Aut}(G)$ *(in particular, if $m = 1$, these two subgroups coincide).*

(2) $\mathrm{Inn}(G) \cap \mathrm{Fix}(H_1) = \mathrm{Inn}_{H_1}(G)$*, where* $\mathrm{Inn}_{H_1}(G)$ *is by definition* $\{i_h \in \mathrm{Inn}(G) \mid h \in H_1\}$.

(3) There exists a natural homomorphism from a subgroup of index at most $m!$ *in* $\mathrm{Out}(G)$ *to* $\mathrm{Out}(H_1)$ *given by $\phi \mapsto i_{g_\phi}\phi|_{H_1}$, where g_ϕ is an element of G such that $i_{g_\phi}\phi \in \mathrm{Fix}(H_1)$, and $|_{H_1}$ denotes the restriction of the automorphism to H_1.*

Proof. (1) Indeed, every automorphism ϕ of G is a quasi-isometry of the Cayley graph of G. Hence $\phi(H_1)$ is an undistorted subgroup of G that is isomorphic to H_1. By Remark 8.28, (1), we have that $\phi(H_1) \subset gH_jg^{-1}$ for some $g \in G$ and $j \in \{1, 2, \ldots, m\}$. In particular $i_g^{-1}\phi(H_1) \subset H_j$. By hypothesis H_j is asymptotically without cut-points. If we denote by ψ the automorphism $i_g^{-1}\phi$, we have that $\psi^{-1}(H_j) \subset \gamma H_k\gamma^{-1}$, for some $\gamma \in G$ and $k \in \{1, 2, \ldots, m\}$. Consequently $H_1 \subset \gamma H_k\gamma^{-1}$. We deduce from the fact that H_1 is infinite and from property (α_1) that $H_1 = \gamma H_k\gamma^{-1}$ and $i_g^{-1}\phi(H_1) = H_j$. In particular every automorphism of G induces a permutation of the set

$$\{H_i \mid H_i \text{ is isomorphic to } H_1\}.$$

Therefore we have an action of $\mathrm{Aut}(G)$ on a subset of $\{H_1, ..., H_m\}$. Let S be the kernel of this action. Then $|\mathrm{Aut}(G) : S| \leq m!$. The composition of any $\phi \in S$ with an inner automorphism i_g^{-1} induced by g^{-1} is in $\mathrm{Fix}(H_1)$. Therefore S is contained in $\mathrm{Inn}(G)\mathrm{Fix}(H_1)$.

(2) Let i_g be an element in $\mathrm{Inn}(G) \cap \mathrm{Fix}(H_1)$. Then g normalizes H_1, hence by [Os], $g \in H_1$.

(3) This immediately follows from (1) and (2). $\square$

9 Appendix. Relatively hyperbolic groups are asymptotically tree-graded. By Denis Osin and Mark Sapir

Here we prove the "if" statement in Theorem 8.5.

Theorem 9.1. *Let G be a group generated by a finite set S, that is relatively hyperbolic with respect to finitely generated subgroups $H_1, ..., H_m$. Then G is asymptotically tree-graded with respect to these subgroups.*

Throughout the rest of this section we assume that G, $H_1, ..., H_m$, an ultrafilter ω, and a sequence of numbers $d = (d_i)$ are fixed, G is generated by a finite set S and is hyperbolic relative to $H_1, ..., H_m$. We denote the asymptotic cone $\mathrm{Con}^\omega(G; e, d)$ by C.

If $(g_i), (h_i)$ are sequences of numbers, we shall write $g_i \leq_\omega h_i$ instead of "$g_i \leq h_i$ ω-almost surely". The signs $=_\omega$, $\in_\omega$ have similar meanings.

As before, $\mathcal{H} = (\bigcup_{i=1}^m H_i) \setminus \{e\}$. For every $i = 1, ..., m$, in every coset of H_i ($i = 1, ..., m$) we choose a smallest length representative. The set of these representatives is denoted by T_i. Let $\mathcal{T}_i$ be the set $\{(g_j)^\omega \mid \lim^\omega(|g_i|_S) < \infty\}$. For each $\gamma = (g_j)^\omega \in \mathcal{T}_i$ we denote by M_γ the ω-limit

$\lim^\omega(g_jH_i)_e$. We need to show that C is tree-graded with respect to all $\mathcal{P} = \{M_\gamma \mid \gamma \in \mathcal{T}_i, i = 1, ..., , m\}$.

We use the notation dist_S and $\mathrm{dist}_{S\cup H}$ for combinatorial metrics on $\mathrm{Cayley}(G, S)$ and $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$. When speaking about geodesics in $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$ we always assume them to be geodesic with respect to $\mathrm{dist}_{S\cup H}$.

The lemma below can be found in [Os, Theorem 3.23].

Lemma 9.2. *There exists a constant $\nu > 0$ such that the following condition holds. Let $\Delta = pqr$ be a geodesic triangle in* $\mathrm{Cayley}(G, S \cup \mathcal{H})$ *whose sides are geodesic (with respect to the metric* $\mathrm{dist}_{S\cup\mathcal{H}}$*). Then for any vertex v on p, there exists a vertex u on the union $q \cup r$ such that*

$$\mathrm{dist}_S(u, v) \le \nu.$$

Lemma 9.3. *Let p and q be paths in* $\mathrm{Cayley}(G,\, S \cup \mathcal{H})$ *such that $p_- = q_-$, $p_+ = q_+$, and q is geodesic. Then for any vertex $v \in q$, there exists a vertex $u \in p$ such that*

$$\mathrm{dist}_S(u, v) \le (1 + \nu) \log_2 |p|.$$

Proof. Let $f : \mathbb{N} \to \mathbb{N}$ be the smallest function such that the following condition holds. Let p and q be paths in $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$ such that $p_- = q_-$, $p_+ = q_+$, q is geodesic, and $|p| \le n$. Then for any vertex $v \in q$, there exists a vertex $u \in p$ such that $\mathrm{dist}_S(u, v) \le f(n)$. Clearly $f(n)$ is finite for each value of the argument. By dividing p into two parts and applying Lemma 9.2, we obtain $f(m + n) \le \max\{f(m),\ f(n)\} + \nu$. In particular, $f(2n) \le f(n) + \nu$ and $f(n + 1) \le f(n) + \nu$.

Suppose that

$$n = \varepsilon_0 + 2\varepsilon_1 + \ldots + 2^k\varepsilon_k,$$

where $\varepsilon_i \in \{0, 1\}$ and $\varepsilon_k = 1$. Then

$$\begin{aligned} f(n) = \quad & f\big(\varepsilon_0 + 2(\varepsilon_1 + \ldots + 2(\varepsilon_{k-1} + 2)\ldots)\big) \le \\ & \underbrace{\nu + \nu + \ldots + \nu}_{2k \text{ times}} + f(1) \le 2\nu \log_2 n. \end{aligned}$$

□

The next lemma can be found in [Os, Lemma 3.1].

Lemma 9.4. *There is a constant α such that for any cycle q in* $\mathrm{Cayley}(G,\, S\cup\mathcal{H})$*, and any set of isolated $\mathcal{H}$-components of $p_1, \ldots, p_k$ of q, we have*

$$\sum_{i=1}^{k} \mathrm{dist}_S((p_i)_-, (p_i)_+) \le \alpha|q|.$$

The following lemma which holds for any (not necessarily relatively hyperbolic) finitely generated group G and any subgroup $H \le G$.

Lemma 9.5. *For any $i = 1, ..., m$, $\theta, \sigma \in \mathcal{T}_i$, if $\theta \ne \sigma$ then the intersection $M_\theta \cap M_\sigma$ consists of at most* 1 *point.*

Proof. Suppose that $x, y \in M_\theta \cap M_\sigma$. Suppose that $\theta = (f_j)^\omega$, $\sigma = (g_j)^\omega$. Then

$$x = \lim^\omega(f_j a_j),\ y = \lim^\omega(f_j s_j)$$

for some $a_j, s_j \in H_i$ and
$$x = \lim^\omega(g_j b_j),\ y = \lim^\omega(g_j t_j)$$
for some $b_j, t_j \in H_i$. Since the sequences $(f_j a_j)^\omega$ and $(g_j b_j)^\omega$ are equivalent, we have $f_j a_j = g_j b_j u_j$, where $|u_j|_S =_\omega o(d_j)$. Similarly $f_j s_j = g_j t_j v_j$, where $|v_j|_S =_\omega o(d_j)$. From these equalities we have
$$a_j^{-1} s_j = u_j^{-1} b_j^{-1} t_j v_i.$$

Let U_j, V_j be shortest words over S representing u_j and v_j respectively. Let also $h_j = a_j^{-1} s_i$ and $k_j = b_j^{-1} t_j$. Then there exists a quadrangle
$$q^j = p_1^j p_2^j p_3^j p_4^j$$
in $\mathrm{Cayley}(G, S)$ such that $\phi(p_1^j) \equiv U_i$, $\phi(p_3^j) \equiv V_j^{-1}$ and p_2^j, p_4^j are edges of $\mathrm{Cayley}(G, S)$ labelled h_j and k_j^{-1} respectively. Note that the cycle q^j contains only two components, namely, p_2^j and p_4^j, as the labels of p_1^j and p_3^j are words over S. Let $A \subseteq \mathbb{N}$ be the set of all j such that the components p_2^j and p_4^j are connected. There are two cases to consider.

Case 1. $\omega(A) = 1$. Note that $\phi(p_1^j)$ represents an element of H_i in G for any $j \in A$, i.e., $s_j \in_\omega H_i$. It follows that $\theta = \sigma$.

Case 2. $\omega(A) = 0$. Note that p_2^j is an isolated component of q^j for any $j \in \mathbb{N} \setminus A$. Since $\omega(\mathbb{N} \setminus A) = 1$, applying Lemma 9.4, we obtain
$$|h_j|_S = \mathrm{dist}_S((p_2)_-, (p_2)_+) \le \alpha |q^j| \le_\omega \alpha(2 + 2o(d_j)) = o(d_j).$$
This yields
$$\mathrm{dist}(x, y) = \lim^\omega(\frac{1}{d_j}\mathrm{dist}_S(f_j a_j, f_j s_j)) = \lim^\omega(\frac{1}{d_j}|h_j|_S) = 0,$$
i.e., $x = y$. □

The following lemma does use the relative hyperbolicity of G.

Lemma 9.6. *For every $i \ne i'$ and every $\theta \in \mathcal{T}_i$, $\sigma \in \mathcal{T}_{i'}$, the intersection $M_\theta \cap M_\sigma$ consists of at most one point.*

Proof. Indeed, repeating the argument from the proof of Lemma 9.5, we immediately get contradiction with the BCP property. □

Lemmas 9.5, 9.6 show that the asymptotic cone C satisfies the property (T_1) with respect to the set $\mathcal{P}$. Now we are going to prove (T_2).

Lemma 9.7. *Let $\mathfrak{g}$ be a simple loop in C. Suppose that $\mathfrak{g} = \lim^\omega(\mathfrak{g}_j)$ for certain loops $\mathfrak{g}_j$ in* $\mathrm{Cayley}(G, S)$, *$|\mathfrak{g}_j| \le C d_j$ for some constant C. Then there exists $i = 1, ..., m$ and $\theta \in \mathcal{T}_i$ such that $\mathfrak{g}$ belongs to M_θ.*

Proof. Let $a \ne b$ be two arbitrary points of $\mathfrak{g}$,
$$a = \lim^\omega(a_j), \qquad b = \lim^\omega(b_j),$$
where a_j, b_j are vertices on $\mathfrak{g}_j$. For every j, we consider a geodesic path $\mathfrak{q}_j$ in $\mathrm{Cayley}(G, S \cup \mathcal{H})$ such that $(\mathfrak{q}_j)_- = a_j$, $(\mathfrak{q}_j)_+ = b_j$.

According to Lemma 9.3, for every vertex $v \in \mathfrak{q}_j$, there exist vertices $x_j = x_j(v) \in \mathfrak{g}_j[a_j, b_j]$ and $y_j = y_j(v) \in \mathfrak{g}_j[b_j, a_j]$ (here $\mathfrak{g}_j[a_j, b_j]$ and $\mathfrak{g}_j[b_j, a_j]$ are segments of $\mathfrak{g}_j = \mathfrak{g}_j[a_j, b_j]\mathfrak{g}_j[b_j, a_j]$) such that

$$\mathrm{dist}_S(v, x_j) \leq 2\nu \log_2 |\mathfrak{g}_j[a_j, b_j]| < 2\nu \log_2(Cd_j) = o(d_j) \tag{30}$$

and similarly

$$\mathrm{dist}_S(v, y_j) \leq 2\nu \log_2 |\mathfrak{g}_j[b_j, a_j]| < 2\nu \log_2(Cd_j) = o(d_j). \tag{31}$$

Summing (30) and (31), we obtain

$$\mathrm{dist}_S(x_j, y_j) \leq \mathrm{dist}_S(x_j, v) + \mathrm{dist}_S(v, y_j) = o(d_j).$$

Thus for any j, there are only two possibilities: either $\lim^\omega(x_j) = \lim^\omega(y_j) = a$ or $\lim^\omega(x_j) = \lim^\omega(y_j) = b$, otherwise the loop $\mathfrak{g}$ is not simple.

For every j, we take two vertices $v_j, w_j \in \mathfrak{q}_j$ such that

$$\lim^\omega(x_j(v_j)) = \lim^\omega(y_j(v_j)) = a,$$

$$\lim^\omega(x_j(w_j)) = \lim^\omega(y_j(w_j)) = b,$$

and $\mathrm{dist}_{S\cup H}(v_j, w_j) = 1$. Since $\lim^\omega(x_j(v_j)) = a$, we have $\mathrm{dist}_S(x_j(v_j), a) =_\omega o(d_j)$. Hence

$$\mathrm{dist}_S(v_j, a_j) \leq \mathrm{dist}_S(v_j, x_j(v_j)) + \mathrm{dist}_S(x_j(v_j), a_j) =_\omega o(d_j).$$

Similarly,

$$\mathrm{dist}_S(w_j, b) =_\omega o(d_j).$$

This means that

$$\lim^\omega(a_j) = \lim^\omega(v_j), \quad \text{and} \quad \lim^\omega(b_j) = \lim^\omega(w_j). \tag{32}$$

For every $i = 1, ..., m$, set $A_i = \{j \in \mathbb{N} \mid v_j^{-1}w_j \in H_i\}$. Let us consider two cases.

Case 1. $\omega(A_i) = 1$ for some i. Set $\theta = (t_i(v_j))^\omega \in \mathcal{T}_i$ where $t_i(v_j)$ is the representative of the coset v_jH_i chosen in the definition of $\mathcal{T}_i$. Then $a, b \in_\omega M_\theta$. Indeed, this is obvious for a since $\lim^\omega(a_j) = \lim^\omega(v_j) \in M_\theta$. Further, since $v_j^{-1}w_j \in_\omega H_i$, we have $t(w_j) =_\omega t(v_j)$. Hence $\lim^\omega(b_j) = \lim^\omega(w_j) \in M_\theta$.

Case 2. $\omega(A_i) = 0$ for every i. Recall that $v_j^{-1}w_j \in S \cup \mathcal{H}$. Thus we have $v_j^{-1}w_j \in_\omega S$. This implies $|v_j^{-1}w_j|_S =_\omega 1$ and $\lim^\omega(v_j) = \lim^\omega(w_j)$. Taking into account (32), we obtain $\lim^\omega(a_j) = \lim^\omega(b_j)$, i.e., $a = b$.

Since a and b were arbitrary points of β, the lemma is proved. □

Now property (T_2) immediately follows from Proposition 3.29 and Lemma 9.7.

References

[AN] S. A. Adeleke, P. M. Neumann. Relations related to betweenness: their structure and automorphisms, Mem. Amer. Math. Soc. 131 (1998), no. 623.

[BGS] W. Ballmann, M. Gromov and V. Schroeder. *Manifolds of Non-positive Curvature*, Springer, 1999.

[Bou] N. Bourbaki. *Topologie générale*, quatrième édition, Hermann, Paris, 1965.

[Bow$_1$] B. Bowditch. Intersection numbers and the hyperbolicity of the curve complex. Preprint. 2003.

[Bow_2] B. Bowditch. *Treelike structures arising from continua and convergence groups*, Memoirs Amer. Math. Soc. Volume 662 (1999).

[Bow_3] B. Bowditch, *private communications.*

[Bri] M. Bridson. Asymptotic cones and polynomial isoperimetric inequalities. Topology 38 (1999), no. 3, 543–554.

[BrH] M. R. Bridson, A. Haefliger. *Metric Spaces of Non-positive Curvature*, Springer, 1999.

[Bo] Dmitri Burago, Yuri Burago, Sergei Ivanov. *A course in metric geometry,* Graduate Studies in Mathematics, 33. American Mathematical Society, Providence, RI, 2001.

[Bu] J. Burillo. *Dimension and fundamental groups of asymptotic cones*, PhD thesis, University of Utah, june 1996.

[CR] I. Chatterji, K. Ruane. Some geometric groups with rapid decay. arXive:math.GR/0310356.

[Chis] Ian Chiswell. *Introduction to Λ-trees.* World Scientific Publishing Co., Inc., River Edge, NJ, 2001.

[CDP] T. Delzant. Sous-groupes distingués et quotients des groupes hyperboliques, Duke Math. J. 83:3 (1996), 661682.

[Dah_1] F. Dahmani. Combination of convergence groups. Geometry and Topology 7 (2003), 933-963.

[Dah_2] F. Dahmani. *Les groupes relativement hyperboliques et leurs bords*, PhD Thesis, University of Strasbourg.

[DP_1] A. Dioubina, I. Polterovich. Explicit constructions of universal R-trees and asymptotic geometry of hyperbolic spaces. preprint, math.DG/9904133.

[DP_2] A Dioubina, I. Polterovich. Structures at infinity of hyperbolic spaces. (Russian) Uspekhi Mat. Nauk 53 (1998), no. 5(323), 239-240; translation in Russian Math. Surveys 53 (1998), no. 5, 1093-1094.

[Dr_1] C. Druţu, *Remplissage dans des réseaux de Q-rang* 1 *et dans des groupes resolubles*, Pacific J. Math 185 (1998), 269-305.

[Dr_2] C. Druţu. Quasi-isometric classification of non-uniform lattices in semisimple groups of higher rank. Geom. Funct. Anal. **10** (2000), no. 2, 327-388.

[Dr_3] C. Druţu. Cônes asymptotiques et invariants de quasi-isométrie pour des espaces métriques hyperboliques. Ann. Inst. Fourier **51** (2001), no. 1, 81-97.

[Dr_4] C. Druţu. Quasi-isometry invariants and asymptotic cones. Int. J. of Algebra and Computation 12 (2002), no. 1 and 2, 99-135.

[EO] A. Erschler, D. Osin. Fundamental groups of asymptotic cones. preprint arXiv:math.GR/0404111 5 april 2004.

[Fa] B. Farb. Relatively hyperbolic groups. Geom. Funct. Analysis **8** (1998), 810-840.

[Gr_1] M. Gromov. Groups of polynomial growth and expanding maps. Publ. Math. IHES **53** (1981), 53-73.

[Gr_2] M. Gromov. Hyperbolic groups. Essays in group theory, 75-263, Math. Sci. Res. Inst. Publ., 8, Springer, New York, 1987.

[Gr_3] M. Gromov. Asymptotic Invariants of Infinite Groups. Geometric Group Theory(vol. 2), G. A. Niblo, M. A. Roller (eds), Proc. of the Symposium held in Sussex, LMS Lecture Notes Series 181, Cambridge University Press 1991.

[GLP] M. Gromov, J. Lafontaine, P. Pansu. *Structures métriques pour les variétés riemanniennes.* Cedic/Fernand Nathan, Paris (1981).

[Hat] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[KaL_1] M. Kapovich, B. Leeb. On asymptotic cones and quasi-isometry classes of fundamental groups of nonpositively curved manifolds. Geom. Funct. Analysis **3** (1995), vol.5, 582-603.

[KaL_2] M. Kapovich, B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. Invent. Math. **128** (1997), no.2, 393-416.

[KaL_3] M. Kapovich, B. Leeb. 3-manifold groups and nonpositive curvature. Geom. Funct. Anal. 8 (1998), no. 5, 841-852.

[KlL] B. Kleiner, B. Leeb. Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings. Publ. Math. IHES **86** (1997), 115-197.

[KSTT] L. Kramer, S. Shelah, K. Tent, S. Thomas. Asymptotic cones of finitely presented groups. preprint, arXive, math.GT/0306420.

[KT] L. Kramer, K. Tent. Asymptotic cones and ultrapowers of Lie groups. preprint, arXive, math.GT/0311101.

[LS] R. Lyndon and P. Schupp. *Combinatorial group theory.* Springer-Verlag, 1977.

[Ols_1] A.Yu. Olshanskii. SQ-universality of hyperbolic groups. (Russian) Mat. Sb. 186 (1995), no. 8, 119–132; translation in Sb. Math. 186 (1995), no. 8, 1199–1211.

[Ols_2] A.Yu.Olshanskii. Distortion functions for subgroups. In: Group Theory Down Under (ed. J.Cossey, C.F. Miller, W.D. Neumann, M. Shapiro), de Gruyter, 1999, 281–291.

[Os] D.V. Osin. Relatively hyperbolic groups. Preprint, 2003.

[Pa] P. Pansu. Croissance des boules et des géodésiques fermées dans les nilvariétés. Ergod. Th. Dynam. Syst. 3 (1983), 415-445.

[Pp] P. Papasoglu. On the asymptotic cone of groups satisfying a quadratic isoperimetric inequality. J. Differential Geom. 44 (1996), no. 4, 789-806.

[PW] P. Papasoglu, K. Whyte. Quasi-isometries between groups with infinitely many ends. Comment. Math. Helv. 77 (2002), no. 1, 133–144.

[Ri] T.R. Riley. Higher connectedness of asymptotic cones. Topology 42 (2003), no. 6, 1289–1352.

[RiSe] E.Rips, Z. Sela. Canonical representatives and equations in hyperbolic groups. Invent. Math. 120 (1995), no. 3, 489–512.

[SBR] M. V. Sapir, J. C. Birget, E. Rips. Isoperimetric and isodiametric functions of groups, Annals of Mathematics, 157, 2 (2002), 345-466.

[Sch] R. E. Schwartz. The quasi-isometry classification of rank one lattices. Inst. Hautes Etudes Sci. Publ. Math. no. 82 (1995), 133-168 (1996).

[Sh] S. Shelah. Classification theory and the number of non-isomorphic models. Studies in Logic and the Foundations of Mathematics, 92. North-Holland Publishing Co., Amsterdam-New York, 1978.

[TV] S. Thomas, B. Velickovic. Asymptotic cones of finitely generated groups. Bull. London Math. Soc. 32 (2000), no. 2, 203-208.

[Tr] V.I. Trofimov. Some asymptotic characteristics of groups. (Russian) Mat. Zametki 46 (1989), no. 6, 85–93, 128; translation in Math. Notes 46 (1989), no. 5-6, 945–951 (1990).

[VDW] L. van den Dries, A. J. Wilkie. On Gromov's theorem concerning groups of polynomial growth and elementary logic. J. of Algebra **89** (1984), 349-374.

[Ya] A. Yaman. A topological characterisation of relatively hyperbolic groups. Journal fur die reine und angewandte Mathematik (Crelle's Journal), to appear.

Cornelia Druţu
Department of Mathematics
University of Lille 1
and UMR CNRS 8524
Cornelia.Drutu@math.univ-lille1.fr

Mark V. Sapir
Department of Mathematics
Vanderbilt University
msapir@math.vanderbilt.edu