RELATIVE HYPERBOLICITY AND BOUNDED COHOMOLOGY (PRELIMINARY)

IGOR MINEYEV AND ASLI YAMAN

Abstract. Let Γ be a finitely generated group and $\Gamma' = \{\Gamma_i \mid i \in I\}$ be a family of its subgroups. We utilize the notion of a *tuple* $\mathcal{T} = (\Gamma, \Gamma', X, \mathcal{V}')$ that makes the statements and arguments for the pair (Γ, Γ') parallel to the non-relative case. For a given tuple, we define *snakes* and the *snake metric*. The language of tuples and snakes seems to be convenient for dealing with relative hyperbolicity.

For a tuple $\mathcal{T}$ we consider properties of being *finitely generated*, *finitely presented*, and of having *fine triangles* (cf. [23, 24]). *Fine triangles* are the ones that are thin with respect to the snake metric. The pair (Γ, Γ') is called *hyperbolic* if there is a finitely generated tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ with fine triangles and with $X^{(1)}$ fine.

We give a definition of relative hyperbolicity of Γ with respect to Γ' which slightly generalizes the definition of Bowditch, and show that this notion coincides with hyperbolicity of the pair (Γ, Γ').

We describe the *relative standard projective resolution*, or the *snake resolution*, $\mathbf{St}^{\varsigma}(\Gamma, \Gamma')$. It is used to define both relative cohomology and relative bounded cohomology.

We generalize the argument in [17, 18] to show that if (Γ, Γ') is hyperbolic then $\mathrm{H}^2_{\mathrm{b}}(\Gamma, \Gamma'; V) \to \mathrm{H}^2(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V. The same holds for bounded $\mathbb{R}\Gamma$-modules, bounded $\mathbb{C}\Gamma$-modules, and Banach modules. Moreover, this statement extends to several characterizations of hyperbolicity of the pair (Γ, Γ').

Contents

1. Introduction. 2
2. Graphs and complexes. 4
2.1. Simplicial graphs. 4
2.2. Complexes 5
2.3. Angles. 5
2.4. Fineness and fineness at some scale. 6
2.5. Thin triangles. 6
2.6. The snake metric d_ς. 8
2.7. Fine triangles. 8
2.8. The edge graph $\mathcal{G}^{\varsigma}_{L}$. 9
2.9. The ideal complex. 10
3. Group actions on graphs and complexes 13
3.1. Finitely generated actions. 13
3.2. Finitely presented actions. 15
4. Tuples 15

4.1. Graph tuples and tuples. 15
4.2. Finiteness conditions. 16
4.3. Tuples with thin and fine triangles 16
4.4. Isoperimetric functions for tuples. 16
5. Relative hyperbolicity. 18
5.1. Hyperbolic tuples. 19
6. Relative straightening. 21
6.1. Notations. 21
6.2. The functions $f(a,b;i)$ and $\bar{f}(a,b)$. 21
6.3. The 1-chain $q[a,b]$. 23
6.4. Homological isoperimetric inequalities. 24
7. Cohomology and bounded cohomology. 25
7.1. Notations. 26
7.2. Bounded modules. 26
7.3. Projectivity. 26
7.4. Functors B and bHom. 27
7.5. Cohomology of a pair (Γ,Γ'). 27
8. The snake resolution. 28
8.1. $\mathbf{St}$: the standard resolution for $I\Gamma$. 28
8.2. $\mathbf{St}'$: the standard resolution for (Γ',Γ). 29
8.3. $\mathbf{St}^{\varsigma}$: the snake resolution. 30
8.4. Relative bounded cohomology. 31
8.5. A resolution for tuples. 31
8.6. Chain maps between resolutions. 32
8.7. Filling 0-cycles. 33
8.8. The cone. 33
8.9. The relative cone. 33
9. The cohomological characterization of relative hyperbolicity. 34
References 37

1. Introduction.

Since the introduction of the notion of relative hyperbolicity by Gromov [12] there has been a lot of progress in the area [10, 11, 4, 24, 8, 26, 6, 9]. When dealing with relative hyperbolicity one is lead to work with the Cayley graph of Γ, and its coned-off graph defined by Farb [10], i.e. the graph obtained by coning-off left cosets of peripheral subgroups. One of the goals of this paper is to set up a convenient language to deal with relative hyperbolicity. This is of importance, since the proofs of theorems about relatively hyperbolic groups tend to be long and technical. It is especially so when one is considering relative hyperbolicity with respect to a family of subgroups, rather just to one subgroup. The following terms that we use in the paper seem to be convenient: *fine triangle*, *snake metric*, *tuple*, *hyperbolic tuple*, *hyperbolic* *pair*, *ideal complex*. In the homological part of the paper, the *snake resolution* and the *relative cone* are used. Those are relative versions of, respectively, a resolution and a cone.

We propose a natural framework to deal with relative hyperbolicity: a *graph tuple* $(\Gamma, \Gamma', \mathcal{G}, \mathcal{V}')$. Here Γ is a group and Γ' is a family of its subgroups, called *peripheral subgroups*. $\mathcal{G}$ is a graph and $\mathcal{V}'$ is a distinguished set of vertices in $\mathcal{G}$. The vertices of $\mathcal{V}'$ correspond to peripheral subgroups, i.e. conjugates of elements of Γ'. *A priori* we do not assume that Γ' is finite; this will be a consequence of other conditions. When the graph $\mathcal{G}$ is replaced with a simplicial (or cell) complex, the notion of a graph tuple generalizes to a *tuple*.

Let $\mathcal{E}$ be the edge set of $\mathcal{G}$. We work with the path metric d on $\mathcal{G}$, and introduce the *snake metric* d_ς on $\mathcal{E}$ in 2.6. $\mathcal{G}^\varsigma_L$ is the graph obtained by taking $\mathcal{E}$ as its *vertex* set and formally connecting e to e' by an edge if their d_ς-distance is at most L. When $\mathcal{G}^\varsigma_L$ is connected and locally finite, d_ς induces a metric on it, and $\mathcal{G}^\varsigma_L$ plays the role of a Cayley graph of Γ.

Take two copies of ideal triangles and identify their corresponding boundaries. The result is a 2-sphere with three punctures, and three cusps with respect to the hyperbolic metric induced from the metric on the triangles. The universal cover of this space is the hyperbolic plane triangulated into ideal triangles, and we forget about the metric and remember only the simplicial structure. An *ideal complex* X associated with a hyperbolic pair (Γ, Γ') (Definition 11) generalizes the above example. The name comes from the fact that the simplices of X are ideal, i.e. their vertices are the ones fixed by the peripheral subgroups. Also, this complex is *ideal* for our purposes.

The notion of *fine triangles* uses both metrics, as follows. Any d-geodesic triangle in $\mathcal{G}$ has a canonical map onto a tripod. We call a geodesic triangle in $\mathcal{G}$ *δ-fine* if, for any two edges e, e' in the triangle sent to the same edge in the tripod, one has $d_\varsigma(e, e') \leq \delta$.

Throughout the paper we consistently use the prime $'$ for peripheral things, and "the snake" ς for the things "in between".

The use of tuples, the snake metric and the fine triangles property streamline the "relative language" and make statements and arguments about relative hyperbolicity parallel to those in the non-relative case. What before were statements about groups and spaces, now can be easily translated to statements about tuples. The notion of a tuple also allows for generalizations to other kinds of spaces.

Even though many statements are parallel to the non-relative case, their proofs are often not. The main reason is that one has to deal with locally infinite graphs, so usual finiteness considerations do not apply. To give a simple example, in a Cayley graph of a finitely generated group, it is the case that there are only finitely many loops of a given length, up to the group action. This property fails in the relative case for the coned-off graph, so finer arguments are often required.

One can naturally talk about *finitely generated tuples* (see 4.1). We say that a pair (Γ, Γ') is *hyperbolic* if there exists a finitely generated graph tuple $(\Gamma, \Gamma', \mathcal{G}, \mathcal{V}')$ with fine triangles and with $\mathcal{G}$ fine (Definition 37). This notion of hyperbolicity allows some peripheral subgroups to be finite. We show that this notion of hyperbolicity for pairs is equivalent to a version of relative hyperbolicity:

Theorem 39. *The following statements are equivalent.*

(a) (Γ, Γ') *is a hyperbolic pair in the sense of Definition 37.*
(b) Γ *is hyperbolic relative to* Γ' *in the sense of Definition 34.*

With this new language at hand, we generalize the cohomological characterization of hyperbolic groups in [17, 18]:

Theorem 40. *Let* Γ *be a group and* Γ' *be a family of its subgroups. The following statements are equivalent.*

(a) (Γ, Γ') *is hyperbolic.*
(b) *There exists a finitely presented tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$ *such that* X *admits a (combinatorial) isoperimetric function (for edge-loops), and the map* $H^2_b(\Gamma, \Gamma'; V) \to H^2(\Gamma, \Gamma'; V)$ *is surjective for all bounded* $\mathbb{Q}\Gamma$*-modules* V*.*
(b′) *There exists a finitely presented tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$ *such that* X *admits a (combinatorial) isoperimetric function (for edge-loops), and the map* $H^n_b(\Gamma, \Gamma'; V) \to H^n(\Gamma, \Gamma'; V)$ *is surjective for all bounded* $\mathbb{Q}\Gamma$*-modules* V *and all* $n \geq 2$*.*
(c) *There exists a finitely presented tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$ *such that* $X^{(1)}$ *is fine and the map* $H^2_b(\Gamma, \Gamma'; V) \to H^2(\Gamma, \Gamma'; V)$ *is surjective for all bounded* $\mathbb{Q}\Gamma$*-modules* V*.*
(c′) *There exists a finitely presented tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$ *such that* $X^{(1)}$ *is fine and the map* $H^n_b(\Gamma, \Gamma'; V) \to H^n(\Gamma, \Gamma'; V)$ *is surjective for all bounded* $\mathbb{Q}\Gamma$*-modules* V *and all* $n \geq 2$*.*

Bounded $\mathbb{Q}\Gamma$*-modules in this statement can be replaced with bounded* $\mathbb{R}\Gamma$*-modules, bounded* $\mathbb{C}\Gamma$*-modules, or Banach modules.*

We wish to express our gratitude to Robert Bieri for providing several references on relative cohomology.

2. Graphs and complexes.

2.1. **Simplicial graphs.** If $\mathcal{G}$ is a graph, we denote its set of vertices by $\mathcal{V}$ and its set of edges by $\mathcal{E}$. We will always assume in this work that the graphs admit no multiple edges and no two edges have the same end points, i.e. that $\mathcal{G}$ is a simplicial graph. Hence there exists, if any, a unique edge connecting any two vertices.

The *valence* of a vertex is the number (in $\mathbb{N} \cup \{\infty\}$) of edges containing this vertex. We denote by $\mathcal{V}_\infty$ the vertices of infinite valence.

The edges in $\mathcal{E}$, otherwise mentioned, are assumed to be non-oriented. For $e \in \mathcal{E}$ with vertices a, b we use the notations (a, b) for non oriented edge. In this notation the order of vertices is not important. For any adjacent vertices a, b, since there are no multiple edges, $e = (a, b)$ is unique.

The edges of $\mathcal{G}$ are assigned length 1. This determines the *length* of any simplicial path α in $\mathcal{G}$, denoted $\ell(\alpha)$, which the number of times it passes over edges. Let d be the path metric on $\mathcal{G}$, defined as the infimum of the lengths of the edge paths connecting two points in $\mathcal{G}$.

For $e_1, e_2 \in \mathcal{E}$ we measure also the d-distance between e and e', considering them as subsets in the graph $\mathcal{G}$, i.e. $d(e_1, e_2) := \inf\{d(x_1, x_2) \mid x_1 \in e_1 \text{ and } x_2 \in e_2\}$. This is the same as the minimal distance between the end points of e_1 and e_2. Similarly $d(e, y) := \inf\{d(x, y) \mid x \in e\}$. for $e \in \mathcal{E}$ and $x \in \mathcal{V}$.

In a graph, a *simple path* is an injective edge path and a *ray* is an infinite simple path. A *loop* is a closed simplicial path. A *circuit* is an injective loop.

Given a path α in $\mathcal{G}$ we denote by $\mathcal{V}(\alpha)$ its set of vertices and by $\mathcal{E}(\alpha)$ its set of edges.

For a path α and two vertices a, b on α, $\alpha_{[a,b]}$ is the subpath of α connecting a and b, and if α is an infinite path starting at x we denote by $\alpha_{[a)}$ the subpath of α remaining after removing $\alpha_{[x,a]}$.

A *geodesic* path in $\mathcal{G}$ is an injective simplicial path which has the shortest length among all the paths connecting its endpoints. We will always parameterize a path α by the interval $[0, \ell(\alpha)]$. For vertices $a, b \in \mathcal{V}$, $\mathrm{Geod}(a, b)$ will denote the set of all geodesic paths in $\mathcal{G}$ going from a to b.

We say that two vertices $u, v \in \mathcal{G}$ are *adjacent* if they lie on a same edge, and that an edges e is *incident* to $x \in \mathcal{V}$ if x is an end point for e.

Let $e_1, e_2 \in \mathcal{E}$. We say that the pair (e_1, e_2) is *admissible* (or the edges e_1, e_2 are admissible) if e_1 and e_2 share a common vertex a. Clearly for $e \in \mathcal{E}$ the pair (e, e) is admissible.

A sequence of edges $e_0, \ldots, e_n$ in $\mathcal{E}$ is admissible if each pair (e_{i-1}, e_i) is admissible. Note that with this definition it is possible in an admissible sequence to have edge repetitions and two admissible edges that are non consecutive in the sequence.

2.2. **Complexes.** We will work in the category of simplicial complexes. If needed, the results of this paper can be stated in the more general category of combinatorial cell complexes, described as follows. A cellular map between cell complexes is *combinatorial* if it maps each open cell homeomorphically onto an open cell. A *combinatorial cell complex* is obtained inductively on dimension using combinatorial attaching maps (see for example [5, 1.8A]).

Given a simplicial complex X, $\mathcal{G}$ will always denote the 1-skeleton of X, so accordingly, $\mathcal{V}$ and $\mathcal{E}$ will mean the sets of vertices and edges in X. In a simplicial complex if $x_1, \ldots, x_n$ are the vertices of an n-simplex, then $(x_1, \ldots, x_n)$ denote this simplex.

Given a simplex σ in X, the *star of* σ, $Star_X(\sigma)$, is the union of the interiors of the simplices of X having σ as a face and the *closed star of* σ, $\overline{Star}_X(\sigma)$, is the union of the (closed) simplices of X having σ as a face. The link $Link_X(\sigma)$ of σ in X is $\overline{Star}_X(\sigma) \setminus Star_X(\sigma)$.

2.3. **Angles.** Let $\mathcal{G}$ be a graph. Given two admissible edges $e_1 = (a, b)$ and $e_2 = (a, c)$, the *angle*, $\mathrm{ang}_a(e_1, e_2)$, between e_1 and e_2 at vertex a is the length of a shortest path from b to c in $\mathcal{G} \setminus \{a\}$ ($+\infty$ if there are none).

We will frequently omit the subscript a since there is no ambiguity at which vertex the angle is defined.

One can similarly define the angle between two paths α_1 and α_2 sharing an endpoint a. Suppose $e \in \alpha_1$, $e' \in \alpha_2$ are the first edges on these paths sharing the vertex a. Then the *angle*, $\mathrm{ang}_a(\alpha_1, \alpha_2)$, between α_1 and α_2 at the vertex a, is $\mathrm{ang}_a(e, e')$. Given a path α connecting b, c in $\mathcal{G}$ and a vertex a in α, the *angle* at the vertex a, $\mathrm{ang}_a(\alpha)$ is $\mathrm{ang}_a(\alpha_{[a,b]}, \alpha_{[a,c]})$. The *maximal angle* of α, $\mathrm{maxang}(\alpha)$, is the maximum of the angles between pairs of consecutive edges of α.

The following remarks will be useful.

Proposition 1. *Given three admissible edges e_1, e_2, e_3 all adjacent to a vertex a in a graph $\mathcal{G}$ one has*

- $\mathrm{ang}_a(e_1, e_2) = \mathrm{ang}_a(e_2, e_1)$,

- $\text{ang}_a(e_1, e_3) \le \text{ang}_a(e_1, e_2) + \text{ang}_a(e_2, e_3)$.

Proposition 2. *Given $l \ge 2$, any circuit of length l has a maximal angle at most than $l - 2$.*

Proof. If (e_1, e_2) is an admissible pair in the circuit, the circuit itself gives a path of length $\mu - 2$ connecting the end points of e_1, e_2. □

2.4. **Fineness and fineness at some scale.** This notion was introduced by Bowditch in [4]. A graph $\mathcal{G}$ is *fine* if for any given integer n there is $K \in [0, \infty)$ such that for any edge e in $\mathcal{G}$ the cardinality of the circuits of length n that contain e is at most K.

We consider a weaker version of the notion. A graph $\mathcal{G}$ is *fine at scale* n if there is $K \in [0, \infty)$ such that for any edge e in $\mathcal{G}$ the number of circuits of length at most n that contain e is at most K.

Let $\mathcal{G}$ be a graph. Consider the graph $\mathcal{G}^n$ defined as follows. $\mathcal{V}(\mathcal{G}^n) := \mathcal{V}(\mathcal{G})$ and distinct x, y in $\mathcal{V}(\mathcal{G}^n)$ are connected by an edge if and only if either $(x, y) \in \mathcal{E}(\mathcal{G})$ or x, y lie in some circuit of length at most n. Clearly $\mathcal{G} \subseteq \mathcal{G}^n$.

2.5. **Thin triangles.** We will say that $\mathcal{G}$ *has thin triangles* if there exists a constant $\delta \ge 0$ such that all the geodesic triangles in $\mathcal{G}$ are δ*-thin* in the following sense: if $[a, b] \in \text{Geod}(a, b)$, $[b, c] \in \text{Geod}(b, c)$, and $[c, a] \in \text{Geod}(a, c)$ for $a, b, c \in \mathcal{V}$, and if points $\bar{a} \in [b, c]$, $v, \bar{c} \in [a, b]$, $w, \bar{b} \in [a, c]$ satisfy

$$d(b, \bar{c}) = d(b, \bar{a}), \quad d(c, \bar{a}) = d(c, \bar{b}), \quad d(a, v) = d(a, w) \le d(a, \bar{c}) = d(a, \bar{b}),$$

then $d(v, w) \le \delta$. Having thin triangles is equivalent to $\mathcal{G}$ being Gromov-hyperbolic.

One can equivalently formulate this in terms of Gromov product. The *Gromov product* in $(\mathcal{G}, d)$ is defined by

$$(a|b)_c := \frac{1}{2}\big(d(a, c) + d(b, c) - d(a, b)\big), \qquad a, b, c \in \mathcal{G}.$$

$\mathcal{G}$ is *hyperbolic* if there exists a constant $\delta \ge 0$ such that for all $a, b, c \in \mathcal{V}$, if $[a, b] \in \text{Geod}(a, b)$, $[a, c] \in \text{Geod}(a, c)$ and if $u \in [a, b]$ and $v \in [a, c]$ satisfy $d(a, u) = d(a, v) \le (b|c)_a$, then $d(u, v) \le \delta$.

Given a geodesic $\gamma \in \mathcal{G}$, $\mathcal{V}(\gamma)$ and $\mathcal{E}(\gamma)$ denote the set of vertices and edges that occur in γ. The following lemma is a collection of known results proved and elaborated in different languages by people working on relative hyperbolicity. Here in order to complete the presentation we give an explicit proof of the statements.

Lemma 3. *Let $\mathcal{G}$ be a graph with the path metric d having δ-thin triangles. There exists a constant κ depending only on δ such that given vertices a, b, c, and geodesics $\alpha \in \text{Geod}(b, c)$, $\beta \in \text{Geod}(a, c)$ and $\gamma \in \text{Geod}(a, b)$, we have the following:*

(1) *If $\text{ang}_z(\alpha) > \kappa$ for some $z \in \alpha$ distinct from b and c, then $z \in \beta$ or $z \in \gamma$.*
(2) *If $z \in \alpha$, $d(c, z) < (a|b)_c$ and $\text{ang}_z(\alpha) > \kappa$, then $z \in \beta$.*
(3) *If $\text{ang}_c(\alpha, \beta) > \kappa$, then $c \in \gamma$.*
(4) *If $b = a$ i.e γ is a null geodesic, then $\text{ang}_c(\alpha, \beta) \le \kappa$.*

Proof. We set $\kappa = 100\delta + 100$

(1) We will show that if z is in α and not in β or γ then $\text{ang}_z(\alpha) \leq \kappa$. Without loss of generality we suppose that $d(c, z) \leq (a|b)_c$.

Consider $x', x'' \in \mathcal{V}(\alpha)$ with $d(z, x') = d(z, x'') = 2\delta + 1$ and $d(c, x'') < d(c, x')$. If there are no such vertices then set $x' = b$ and $x'' = c$. Similarly consider y', y'' on β with $d(c, y') = d(c, x')$ and $d(c, y'') = d(c, x'')$, if there are no such vertices then set $y' = a$ and $y'' = c$. Note that $d(x'', c) = d(y'', c) \leq (a|b)_c$.

Now by hyperbolicity we have $d(x'', y'') \leq \delta$, and if $\gamma'' \in \text{Geod}(x'', y'')$ then $z \notin \gamma''$, since otherwise one has either $d(z, x'') \leq \delta$, or $x'' = c = z$, depending on whether γ'' is an empty path or not. Either case gives a contradiction to the choice of x'', since $d(z, x'') = 2\delta + 1$ and $z \neq c$ $(\notin \beta)$. Moreover, either $d(x', y') \leq \delta$ or there exist x_1, y_1 on γ such that $d(x', x_1) \leq \delta$ and $d(y', y_1) \leq \delta$.

In the first case pick $\gamma' \in \text{Geod}(x', y')$. Note that $z \notin \gamma'$, since otherwise either $d(z, x') \leq \delta$ or $z = x' = b$, depending on whether γ' is an empty path or not; either case gives a contradiction with $d(z, x') = 2\delta + 1$ and $c \neq b$ $(\notin \gamma)$. Consider the loop $\alpha_{[x'',x']} \cdot \gamma' \cdot \beta_{[y',y'']} \cdot \gamma''$ that has length at most $10\delta + 4$, containing z. We know that z does not belong twice to these loop. Hence we can find a circuit of length at most $10\delta + 4 \leq \kappa$ containing z, that controls the angle at z of α.

In the second case we consider the geodesic segments $\gamma_1 \in \text{Geod}(x', x_1)$ and $\gamma_2 \in \text{Geod}(y', y_1)$. Again we see that γ_1 and γ_2 do not contain z since otherwise we would have $d(z, x') \leq \delta$ or $z = x' = b$ and $d(z, y') \leq \delta$ or $z = y' = a$, which, by the argument similar to the above, contradicts to the choice of x' and y'. Now consider the loop $\gamma'' \cdot \alpha_{[x'',x']} \cdot \gamma_1 \cdot \gamma_{[x_1,y_1]} \cdot \gamma_2 . \beta_{[y',y'']}$, which has length at most $22\delta + 8$. It does not pass through z twice and thus there is a circuit of length at most $22\delta + 8 \leq \kappa$ controlling the angle at z of α.

(2) If $\text{ang}_z(\alpha) > K$ then by the first part of the lemma we know that $z \in \beta \cup \gamma$. Clearly if $d(c, z) < (a|b)_c$ then $z \notin \gamma$.

(3) In the above arguments we set $c = z$, and consider the vertices x', y' as in the proof of (1). We see that either the loop $\alpha_{[c,x']} \cdot \gamma' \cdot \beta_{[y',c]}$ that has length at most $5\delta + 2$ or the loop $\alpha_{[c,x']} \cdot \gamma_1 \cdot \gamma_{[x_1,y_1]} \cdot \gamma_2 \cdot \beta_{[y',c]}$ that has length at most $12\delta + 4$ does not pass through c twice. Hence $\text{ang}_c(\alpha, \beta) \leq 12\delta + 4 \leq \kappa$.

(4) is a direct corollary of (3). ☐

As it will be formulated later in the snake-metric language, the following result can be interpreted by stating that all the edges in the thin part of geodesics remain uniformly bounded with respect to the snake metric.

Lemma 4. *Let $\mathcal{G}$ be a graph with the path metric d having δ-thin triangles and κ the constant depending only on δ given by Lemma 3. Let $\alpha \in \text{Geod}(c, b)$, $\beta \in \text{Geod}(c, a)$, $\gamma \in \text{Geod}(a, b)$ and $x \in \mathcal{V}(\alpha)$, $e \in \mathcal{E}(\alpha)$, $x' \in \mathcal{V}(\beta)$, $e' \in \mathcal{E}(\beta)$ with $d(c, x) = d(c, e) = d(c, x') = d(c, e') < (a|b)_c$.*

If ω is in $\text{Geod}(x, x')$ then $\text{maxang}(\omega) \leq \kappa$. Moreover if $x \neq x'$ then $\text{ang}_x(\omega, e) \leq \kappa$ and $\text{ang}_x(\omega, e') \leq \kappa$; and if $x = x'$ then $\text{ang}_x(e, e') \leq \kappa$.

Proof. First suppose that $z \neq z'$. We prove that $\text{ang}_x(\omega, e) \leq \kappa$. Now if $\text{ang}_x(\omega, e) > \kappa$ then by lemma 3 (2) x belongs to a geodesic connecting x' to b. In particular the path $\omega.[x, b]_\alpha$ is actually a geodesic. Now again since $\text{ang}_x(\omega, e) > \kappa$ by lemma 3 (2) $x \in \beta$or γ. Now either cases we have a contradiction since if $x \in \beta$ then $x = x'$ and if $x \in \gamma$ then $d(c, x) \geq (a|b)_c$. Similarly one shows that $\mathrm{ang}_x(\omega, e') \leq \kappa$.

Now we suppose $x = x'$. If $\mathrm{ang}_x(e, e') > \kappa$ then again by lemma 3 (2) $x \in \gamma$ which gives the contradiction with $d(c, x) < (a|b)_c$.

It remains to prove that $\mathrm{maxang}(\omega) \leq \kappa$. If $x = x'$ there is nothing to prove. So let suppose $x \neq x'$. Note that if $\mathrm{ang}_u(\omega) > \kappa$ then $u \in [c, x]_\alpha$ or $[c, x']_\beta$. Let u, u' be the vertices in $\alpha \cap \omega$ and $\beta \cap \omega$ such that $d(x, u)$ and $d(x', u')$ are maximal with $\mathrm{ang}_u(\omega) > \kappa$ and $\mathrm{ang}_{u'}(\omega) > \kappa$.

Denote by α' the path $[c, u]_\alpha.[u, x]_\omega.[x, b]_\alpha$ and by β' $[c, u']_\beta.[u', x']_\omega.[x', a]_\beta$. Clearly these are geodesics in $\mathcal{G}$. Let $c' \in \alpha' \cap \beta'$ with $d(c, c') \leq d(c, x)$ maximal.

If $u \neq u'$ then by choice $\mathrm{maxang}([u, u']_\omega) \leq \kappa$. Moreover, the argument as in the first and second paragraphs shows that if $u \neq u'$ then $\mathrm{ang}([u, u']_\omega, [u, x]_\omega$ and if $u = u'$ then $\mathrm{ang}([u, x]_\omega, [u, x']_\omega$ are all at most κ.

It remains to show that $\mathrm{maxang}([u, x]_\omega$ and $maxang([u, x']_\omega$ are at most κ. If for $z \in [u, x]_\omega$ distinct from u and x we have $\mathrm{maxang}([u, x]_\omega > \kappa$ then $z \in \beta'$ by lemma 3 (1). Since $d(x', z) > d(u, x') > d(u', x')$, $z \notin [u', x']_\omega$. Note also $z \notin [x', a]_\beta$, since if not $d(c, z) < d(c, x) = d(c, x') \leq d(c, z)$. Finally we see that $x \notin [c, u']_\beta$ obtaining a contradiction since if not $x = u$. □

2.6. **The snake metric d_ς.** Given a graph $\mathcal{G}$ with its metric d, let $e_0, \ldots, e_n$ be an admissible sequence in $\mathcal{E}$. The *angle length* of an admissible sequence is $\sum_{i=1}^n \mathrm{ang}(e_{i-1}, e_i)$. For an arbitrary $e, e' \in \mathcal{E}$ let $d_\varsigma(e, e')$ be the minimal angle length of an admissible sequence with $e_0 = e$ and $e_n = e'$. d_ς is a metric on the set $\mathcal{E}$. (d_ς takes infinite values on pairs that cannot be connected by an admissible sequence, but we will not bother calling d_ς a "generalized metric"). The metric d_ς should be called a *snake metric*, and an admissible sequence $(e_0, \ldots, e_n)$ with $d_\varsigma(e_i, e_{i+1}) = 1$ realizing the snake distance between e_0 and e_n should be called a *snake geodesic*.

Any edge path in $\mathcal{G}$ can be viewed as an admissible sequence of edges, so the notion of angle length, or the total angle, makes sense for paths in $\mathcal{G}$.

Lemma 5. *For any $e, e' \in \mathcal{E}$, $d(e, e') \leq d_\varsigma(e, e')$.*

Proof. Given e, e' edges we consider an admissible sequence $e_1, \ldots, e_n$ that realizes $d_\varsigma(e, e')$. If $n = 1$, then $d(e, e') = d_\varsigma(e, e') = 0$, so we assume $n \geq 2$. Since a subset of this sequence gives a simple path from an end point of e to an end point of e' we clearly have $d(e, e') \leq n - 2$. Moreover since for all i, $\mathrm{ang}(e_{i-1}, e_i) \geq 1$ we also have $d_\varsigma(e, e') = \sum_{i=1}^n \mathrm{ang}(e_{i-1}, e_i) \geq n - 1$. Thus $d(e, e') \leq d_\varsigma(e, e') - 1 < d_\varsigma(e, e')$. □

2.7. **Fine triangles.** Given a geodesic γ, $\mathcal{E}(\gamma)$ will denote the set of edges that occur in γ.

Definition 6. *A graph $\mathcal{G}$ is said to have fine triangles if there exists $\delta \in [0, \infty)$ depending only on $\mathcal{G}$ with the following property. If $\alpha \in \mathrm{Geod}(c, b)$, $\beta \in \mathrm{Geod}(c, a)$, $e_1 \in \mathcal{E}(\alpha)$, $e_2 \in \mathcal{E}(\beta)$, and $d(c, e_1) = d(c, e_2) < (a|b)_c$, then $d_\varsigma(e_1, e_2) \leq \delta$.*

While working on this paper we learned that Osin, in a language different from the one we use here, shows that the BCP property together with hyperbolicity ensures the fine triangles condition; he considers the fine triangles property up to an additional constant σ [23], so our notion of fine triangles is slightly stronger. It is well known in the theory of relatively hyperbolicity that the BCP property is equivalent to the fineness property of $\mathcal{G}$ [8, 24].

Proposition 7. *$\mathcal{G}$ is a graph with thin triangles if and only if it is a graph with fine triangles.*

Proof. We first show that having thin triangles implies having fine triangles. Thus we need to show that there exits a constant δ' depending only on the constant of hyperbolicity δ such that if $\alpha \in \mathrm{Geod}(c,b)$, $\beta \in \mathrm{Geod}(c,a)$, $e \in \mathcal{E}(\alpha)$, $e' \in \mathcal{E}(\beta)$, and $d(c,e) = d(c,e') < (a|b)_c$, then $d_\varsigma(e,e') \leq \delta'$.

Let $x \in \mathcal{V}(\alpha)$ and $y \in \mathcal{V}(\beta)$ respectively the end points of e and e' such that $d(c,x) = d(c,y) = d(c,e)$. Clearly $d(c,x) = d(c,y) < (a|b)_c$. Let ω be a geodesic realizing $d(x,y)$ which is at most δ by thiness. By Lemma 4 we see that $\mathrm{maxang}(\omega)$, $\mathrm{ang}_x(e,\omega)$, $\mathrm{ang}_y(e',\omega)$ are all at most κ. In particular, $d_\varsigma(e,e') \leq \mathrm{ang}_x(e,\gamma) + \sum_{i=1}^n \mathrm{ang}(e_{i-1},e_i) + \mathrm{ang}_y(e',\gamma)$, where $e_1,\ldots,e_n$ are the consecutive edges of γ viewed as an admissible sequence in $\mathcal{G}$. Thus we have $d_\varsigma(e,e') \leq \delta' = (\delta+1)(\kappa)$, which is a constant depending only on δ.

The other direction follows from Lemma 5. Since $d(e,e) \leq d_\varsigma(e,e)$, the thin triangles condition follows from the fine triangles condition. ☐

2.8. **The edge graph $\mathcal{G}^\varsigma_L$.** Let $\mathcal{G}$ be a graph. For $L \in [0,\infty)$, the *L-edge graph*, or the *L-graph* of $\mathcal{G}$, is the graph $\mathcal{G}^\varsigma_L$ defined as follows. The vertex set of $\mathcal{G}^\varsigma_L$ is $\mathcal{E}$. Two vertices $e,e' \in \mathcal{E}$ of $\mathcal{G}^\varsigma_L$ are connected by an edge in $\mathcal{G}^\varsigma_L$ if $\{e,e'\}$ is an admissible pair of edges in $\mathcal{G}$ with $d_\varsigma(e,e') \leq L$.

The length of an edge (e_1,e_2) in $\mathcal{G}^\varsigma_L$ is defined to be $d_\varsigma(e_1,e_2)$. $\mathcal{G}^\varsigma_L$ is given the path metric corresponding to this assignment of lengths. The restriction of this path metric to the vertices of $\mathcal{G}^\varsigma_L$ coincides with d_ς, so we denote this metric d_ς as well.

The language of edge graphs and snake metrics makes it easier to deal with relative hyperbolicity. An edge path γ in $\mathcal{G}$ can be thought of as a sequence of edges, therefore a sequence of vertices in $\mathcal{G}^\varsigma$; this allows working with the two metrics simultaneously. It might also happen that all the vertices of $\mathcal{G}$ correspond to *peripheral* subgroups (see Section 5). As we see later, such graphs $\mathcal{G}$ are useful for showing a cohomological characterization of relative hyperbolicity. The language of edge graphs can be conveniently used in this case. All the known results about relative hyperbolicity can be equivalently restated in this language.

An L-graph is generally neither connected nor locally finite. However the following lemmas give us sufficient hypothesis for this to hold.

Lemma 8. *If $\mathcal{G}$ is a graph fine at scale $L+2$, then the balls in the L-graph G^ς_L are uniformly finite.*

Proof. It suffices to prove that the valence of vertices in connected components of $\mathcal{G}^\varsigma_L$ are uniformly bounded. Let e be a vertex in $\mathcal{G}^\varsigma_L$. If e' is a vertex adjacent to e in $\mathcal{G}^\varsigma_L$, then by definition there are admissible edges in $\mathcal{G}$ such that $\mathrm{ang}(e,e') \leq L$. Since $\mathcal{G}$ is fine at scale L, there is a constant K independent of the choice of e such that there are at most K circuits in $\mathcal{G}$ of length $L+2$ containing e, hence at most K such e'. ☐

Lemma 9. *If $\mathcal{G}$ is a graph with δ-fine triangles and if it is fine at scale $\delta+2$, then* $\mathrm{Geod}(u,v)$ *is finite for all $u,v \in \mathcal{V}(\mathcal{G})$.*

Proof. Suppose $\mathcal{G}$ has δ-fine triangle. For edges e,e' on a geodesics $\alpha,\beta \in \mathrm{Geod}(u,v)$ with $d(u,e') = d(u,e)$ we have $d_\varsigma(e,e') \leq \delta$. Consider the δ-graph $\mathcal{G}^\varsigma_\delta$. Clearly any such two edges e,e' are on a same connected component of $\mathcal{G}^\varsigma_\delta$. Since $\mathcal{G}$ is fine at scale $\delta+2$, by Lemma 8 we obtain that $\mathcal{G}_\delta^\varsigma$ uniformly locally finite, hence there are only finitely many such pairs of edges $\{e, e'\}$. In particular, the cardinality of such edges depend only on δ since the bounds on the cardinality of the balls are uniform. □

Given a vertex v in $\mathcal{G}$ and an L-edge graph $\mathcal{G}_L^\varsigma$, we denote by $Link_L^\varsigma(v)$ the full subgraph of $\mathcal{G}_L^\varsigma$ whose vertices are all the edges in $\mathcal{G}$ containing v.

Lemma 10. *Let X be a simplicial complex with the following properties.*

- *X is simply connected.*
- *$X \setminus \{v\}$ is connected for each $v \in \mathcal{V} = X^{(0)}$.*

Then for all $L \geq 1$, $Link_L^\varsigma(v)$ is connected for all $v \in \mathcal{V}$. In particular the L-graph $\mathcal{G}_L^\varsigma$ is connected.

Proof. Given two admissible edges $e = (v, x), e' = (v, y)$ in the graph $\mathcal{G}$, there exists a path α connecting x to y in $X \backslash \{v\}$. Since X is simply connected, we can assume that α lies in $Link_{\mathcal{G}}(v)$, i.e. there exists an admissible sequence $e = e_1, \ldots, e_n = e'$ such that $e_i = (v, x_i)$ where x_i are vertices of α in $Link_{\mathcal{G}}(v)$. Clearly for all i we have $ang(e_i, e_{i+1}) \leq 1$ since $(e_i, e_{i+1}, (x_i, x_{i+1}))$ is a 2-simplex in X. Thus e_i and e_{i+1} are connected by an edge in $\mathcal{G}_L^\varsigma$, hence in $Link_L^\varsigma(v)$, for $L \geq 1$.

To see that $\mathcal{G}_L^\varsigma$ is connected it suffices to remark that $\mathcal{G}$ is connected, and there is a path in $\mathcal{G}$ connecting any two given edges. Thus we consider this path as an admissible sequence, and for each pair of edges that are consecutive in the sequence and incident to the vertex v we connect them in the link $Link_L^\varsigma(v)$ to obtain the result. □

2.9. **The ideal complex.** In the case when a group Γ is relatively hyperbolic with respect to a subgroup and the subgroup is admits a finite dimensional classifying space, Dahmani [7] showed the existence of a locally finite finite-dimensional contractible complex for Γ when the subgroup admit a classifying space. He also states in [7, Theorem 6.2], quoting an observation of Bowditch, that a similar argument gives such a complex without the assumption on the subgroup, and for a family of subgroups.

Below we exhibit a finite-dimensional contractible complex, called the *ideal complex*, associated to a group relatively hyperbolic with respect to a family of subgroups. Our construction uses the fineness property rather than the BCP property. The existence of such a complex is independent of the group structure; it can be constructed using only the fineness property and hyperbolicity for a given graph.

Definition 11. *Let $\mathcal{G}$ be a fine graph with δ-thin triangles and μ be a constant. To each subset $S \subseteq \mathcal{V}(\mathcal{G})$ of cardinality $n+1$ satisfying that any pair of its points can be joined by a geodesic path in $\mathcal{G}$ with maximal angle and length at most μ, associate an n-simplex $\sigma(S)$.*

The ideal complex X is the one obtained from the set of simplices by gluing along the face maps $\sigma(S) \hookrightarrow \sigma(T)$ that correspond to inclusions $S \subseteq T$.

Note that there is a natural bijection between the 0-skeleta of $\mathcal{G}$ and of X, so we will always identify them. Also there is natural injection of $\mathcal{G}$ to the 1-skeleton of X. Recall that we write $(a_1, \ldots, a_n)$ to refer to a simplex $\sigma(\{a_1, \ldots, a_n\})$ in X.

For the rest of the section d will denote the path metric on $\mathcal{G}$. Let $\kappa = 100\delta + 100$ be the constant given by Lemma 3. For each edge $e = (a, b)$ in the 1-skeleton $X^{(1)}$ of X, a geodesic path in $\mathcal{G}$ with maximal angle and d-length at most μ connecting a to b will be referred to as a *geodesic representing* e in $\mathcal{G}$.

Lemma 12. *Given $a, b \in \mathcal{V}(X)$ connected by an edge in X, there are only finitely many $c \in \mathcal{V}(X)$ connected both to a and to b by edges in X.*

Proof. Let $c \in \mathcal{V}(X)$ be connected to a and b by edges in $X^{(1)}$ and let α, β and γ be geodesics representing in $\mathcal{G}$ the edges $(b, c), (a, c)$ and (a, b), respectively.

We first show that α, β, γ can be chosen so that $\text{ang}_c(\alpha, \beta)$, $\text{ang}_b(\alpha, \gamma)$ and $\text{ang}_a(\beta, \gamma)$ are all at most $\lambda = max\{\mu, \kappa\}$. Indeed, if $\text{ang}_c(\alpha, \beta) > \kappa$ then by Lemma 3(3) we have $c \in \gamma$. Thus we can let $\alpha = [c, a]_\gamma$ and $\beta = [c, b]_\gamma$ be geodesics representing in $\mathcal{G}$ the edges (a, c) and (b, c) of $X^{(1)}$.

Now since the maximal angle of γ is at most μ, we have the required result as follows. If e, e' are two edges lying on any of the geodesics representing the edges $(b, c), (a, c)$ and (a, b) and chosen as above then they satisfy $d_\varsigma(e, e') \leq 2\lambda^2$, since the representing geodesics provide admissible sequences between them. Hence they all lies in a ball of the L-edge graph $\mathcal{G}_L^\varsigma$ where $L > 2\lambda^2$. This completes the proof since the balls in the L-edge graph are finite by Lemma 8. ☐

In particular we obtain the following results.

Corollary 13. *If K is a subgraph of $X^{(1)}$ and K is complete, then $\mathcal{V}(K)$ is finite.*

Corollary 14. *Given an edge $e \in X$, there are only finitely many simplices in X that contain e. In particular X is finite dimensional.*

We prove that X is contractible using a modification of the argument of Rips for the non-relative case.

Lemma 15. *Let K be a finite subcomplex of X. Given an edge $e = (x, y)$ in K, let α be a geodesic representing in $\mathcal{G}$ the edge e and suppose that there exists $z \in \mathcal{V}(\alpha)$ such that $\text{ang}_z(\alpha) > \kappa$. Choose z so that $d(z, x)$ is minimal among all $z \in \mathcal{V}(\alpha)$ satisfying $\text{ang}_z(\alpha) > \kappa$.*

Then K is homotopic in X to a subcomplex K' of X with $\mathcal{V}(K') = \mathcal{V}(K) \cup \{z\}$ and $\mathcal{E}(\overline{Star}_{K'}(x)) = \big(\mathcal{E}(\overline{Star}_K(x)) \setminus \{e\}\big) \cup \{(x, z)\}$.

This lemma says that we can homotop K to another subcomplex where the edge e of K is replaced by two edges (x, z) and (z, y).

Proof. Note first that (x, y, z) is a 2-simplex of X; we denote it s. For all $w \neq z$ in $X^{(0)}$, if (x, y, w) is a 2-simplex in X, then (x, y, w, z) is a 3-simplex in X. Indeed, if β and γ are geodesics in $\mathcal{G}$ representing the edges (x, w) and (w, y), respectively, then since $\text{ang}_z(\alpha) > \kappa$, by Lemma 3(1) $z \in \beta$ or $z \in \gamma$. Thus there is a geodesic, namely a subsegment of β or γ, connecting z to w with length and maximal angle at most μ. In other words, $\overline{Star}_X(e) = \overline{Star}_X(s)$.

When $z \notin K$, let $N(s)$ be the complex whose simplices are $(z, x = a_1, y = a_2, a_3, \ldots, a_n)$ whenever $(x = a_1, y = a_2, a_3, \ldots, a_n)$ is a simplex in $\overline{Star}_K(e)$. The remark above shows that $N(s)$ is indeed a subcomplex of X. When $z \in K$ we set $N(s) := \overline{Star}_K(e)$. Denote

$K' := (K \cup N(s)) \setminus Star_X(e)$ and $M := N(s) \setminus Star_X(e)$. Thus $K' = (K \setminus Star_K(e)) \cup M$. We claim that M and $\overline{Star_K}(e)$ are homotopic in X, which would imply that K and K' are homotopic. By definition $\overline{Star_K}(e) = N(s) \setminus Star_X(z)$. So we will show that the inclusions of $\overline{Star_K}(e)$ and M in $N(s)$ are homotopic in X by proving that they are all null-homotopic in X.

By definition when $z \notin K$, $N(s)$ can be seen as the complex obtained by coning off all simplices of the subcomplex $\overline{Star_K}(e)$ to z, hence both $N(s)$ and $\overline{Star_K}(e)$ are null-homotopic since they can be contracted to z in X. When $z \in K$ this statement is true by the definition.

To see that M is null-homotopic we first note that each simplex in M contains either the edge (x, z) or (y, z). Now we do induction on the dimension of M. Let σ be a simplex of maximal dimension in M. If σ has dimension 1 then M is the union of edges (x, z) and (z, y) hence contractible on z. For higher dimension n we consider $\overline{M \setminus \sigma}$, whose simplices contain either (x, z) or (y, z), and we contract σ onto $\sigma \cap \overline{(M \setminus \sigma)}$. This intersection is contained in the boundary of σ and has dimension $n - 1$. We perform this for each maximal simplex of highest dimension in M, and apply the induction hypothesis to prove the claim. □

Theorem 16. *Let $\kappa = 100\delta + 100$ be the constant given by Lemma 3. If $\mu \geq 3\kappa$ then the ideal complex X is contractible.*

Proof. It suffices to take any finite subcomplex K of X and to show that it is contractible in X. Fix a base point $v \in \mathcal{V}(K)$. Let x be a vertex in K maximizing the distance in $\mathcal{G}$ of v to vertices of K, i.e $l := d(v, x)$ is maximal. We argue by induction on l. In each induction step we apply the same argument to each x that maximize $d(x, v)$, in order to decrease l. Since K is finite, at each step there are only finitely many such x.

For $l = 1$, let $x \in \mathcal{V}(K)$ be such that $d(x, v) = 1$. Now for all $y \in \mathcal{V}(K)$ distinct from v, $d(y, v) = 1$ since $d(x, v)$ is maximal. Since K is finite it is contractible to v in finitely many steps.

Suppose $l \geq 1$. Consider a geodesic γ in $\mathcal{G}$ connecting x to v and the vertex $u \in \gamma$ such that either

- $d(u, x) = \mu/2$ with $\text{maxang}([x, u]_\gamma) \leq \kappa$ and $\text{ang}_u(\gamma) \leq \kappa$, or
- $d(u, x) \leq \mu/2$ and $d(u, x)$ is minimal among all u satisfying $\text{ang}_u(\gamma) > \kappa$.

Note that by definition there exists an edge in X connecting x to u since $[x, u]_\gamma$ has length at most μ and maximal angle at most μ. We want to contract x to u. We must check that for each each edge (x, y) in K, there is a simplex (x, y, u) in X, i.e that y and u are connected by an edge in X.

Let α be a geodesic path representing (x, y) in $\mathcal{G}$ with maximal angle and length at most μ. Suppose there is $z \in \mathcal{V}(\alpha)$ with $\text{ang}_z(\alpha) > \kappa$, and so that $d(z, x)$ minimal among all such z. Then by Lemma 15, K can be homotoped to another complex K' with $\mathcal{V}(K') = \mathcal{V}(K) \cup \{z\}$ and $\mathcal{E}(\overline{Star_{K'}}(x)) = \mathcal{E}(\overline{Star_K}(x)) \setminus \{e\} \cup \{(x, z)\}$. Since by the choice $\text{maxang}([x, z]_\alpha) \leq \kappa$, the edge (x, z) is represented by a geodesic with angles at most κ. Moreover if β is a geodesic connecting y to v, since $\text{ang}_z(\alpha) > \kappa$ we see by Lemma 3(1) that either $z \in \beta$, hence $d(z, v) < d(v, y) \leq d(v, x) = l$, or $z \in \gamma$, hence $d(z, v) < d(v, x) = l$.

Repeating this argument finitely many times (there are only finitely many edges adjacent to x in K) for each new complex, i.e replacing an edge (x, y), whose geodesic representative has angles at most κ by the procedure described above we obtain a complex, that we continue to denote by K, for which all the edges adjacent to x have geodesic representative in $\mathcal{G}$ with maximal angle at most κ.

Clearly the final complex K might have more vertices than the initial one, but it remains finite. Moreover the maximal distance l of vertices of K to v and the the set of vertices that realize the maximal distance to v in both complexes remains the same. We therefore suppose that in K all the edges adjacent to x can be represented by geodesics paths with maximal angles at most κ and length at most μ. We first note that $\text{ang}_x(\alpha, \gamma) \leq \kappa$. Indeed, if not, by Lemma 3(3), $x \in \beta$. Since $d(y, v) \leq d(x, v)$, we must have $x = y$, which contradicts the choice of y. We treat the two possible cases for the choice of u separately.

Case I $d(u, x) = \mu/2$ *with* $\text{maxang}([x, u]_\gamma) \leq \kappa$ *and* $\text{ang}_u(\gamma) \leq \kappa$.

By thin triangles condition on $\mathcal{G}$ we have either $d(u, q) \leq \delta$ for some $q \in \alpha$ with $d(x, q) = d(x, u)$ or $d(u, p) \leq \delta$ for some $p \in \alpha$ with $d(v, q) = d(v, u)$. In the first case $d(y, u) \leq \mu - \mu/2 + \delta \leq \mu$. In the second case $d(v, x) \leq d(v, y) - d(y, p) + \delta + \mu/2 \leq d(v, x) - d(y, p) + \delta + \mu/2$. Thus $d(y, u) \leq d(y, p) + \delta \leq \mu/2 + 2\delta \leq \mu$.

We want to find a geodesic in $\mathcal{G}$ connecting y to u with maximal angle at most μ. Let β' a geodesic path connecting u, y in $\mathcal{G}$. If $\text{maxang}(\beta') \leq \kappa (\leq \mu)$ then we have the required result. So suppose there is a point p on β' with $\text{ang}_p(\beta') > \kappa$, hence by Lemma 3(1) we would have $p \in [y, x]_\alpha$ or $p \in [u, x]_\gamma$. Denote p_1 the furthest point from y on $\beta' \cap \alpha$ with $\text{ang}_{p_1}(\beta') > \kappa$, and p_2 be the furthest point from u on $\beta' \cap \gamma$ with $\text{ang}_{p_2}(\beta') > \kappa$. We prove that $[y, p_1]_\alpha \cdot [p_1, p_2]_{\beta'} \cdot [p_2, u]_\gamma$ is a geodesic with required properties. First, $\text{maxang}([y, p_1]_\alpha) \leq \kappa$ and $\text{maxang}([p_2, u]_\gamma) \leq \kappa$ by hypothesis. Assume $p_1 \neq p_2$. By the choice of p_1 and p_2 we have $\text{ang}_p([p_1, p_2]_{\beta'}) \leq \kappa$. Moreover, $\text{ang}_{p_1}([y, p_1]_\alpha, [p_1, p_2]_{\beta'}) \leq \text{ang}_{p_1}(\alpha) + \text{ang}_{p_1}([x, p_1]_\alpha, [p_1, p_2]_{\beta'})$, but $ang_{p_1}([x, p_1]_\alpha, [p_1, p_2]_{\beta'}) \leq \kappa$, since if not, by Lemma 3(2), $p_1 \in [x, p_2]_\gamma$, and hence $p_1 = p_2$ by the definition of p_2, which contradicts the assumption $p_1 \neq p_2$. Thus we have $\text{ang}_{p_1}([y, p_1]_\alpha, [p_1, p_2]_{\beta'}) \leq 2\kappa$. Similarly $\text{ang}_{p_2}([u, p_2]_\gamma, [p_1, p_2]_{\beta'}) \leq \text{ang}_{p_2}(\gamma) + \text{ang}_{p_2}([x, p_2]_\gamma, [p_1, p_2]_{\beta'}) \leq 2\kappa$. If $p_1 = p_2$ then $[p_1, p_2]_{\beta'}$ is a null path and we have $\text{ang}_{p_1}([x, p_1]_\alpha, [p_1, x]_\gamma) \leq \kappa$ by Lemma 3(4), hence $\text{ang}_{p_1}([y, p_1]_\alpha, [p_1, u]_\gamma) \leq \text{ang}_{p_1}(\alpha) + \text{ang}_{p_1}([x, p_1]_\alpha, [p_1, x]_\gamma) + \text{ang}_{p_1}(\gamma) \leq 3\kappa$.

Case II $d(u, x) \leq \mu/2$ *and* $d(u, x)$ *is minimal among all* u *satisfying* $\text{ang}_u(\gamma) > \kappa$.

Since $\text{ang}_u(\gamma) > \kappa$ we have either $u \in \alpha$, in which case $d(y, u) \leq \mu$ and the geodesic $[y, u]_\alpha$ satisfies the properties required by hypothesis, or $u \in [y, v]_\beta$ and $d(y, u) \leq d(x, u) \leq \mu$ since $d(x, v) \geq d(y, v)$. Finally, the same argument as in Case I works to find a geodesic with angles at most μ.

In either case, moving x to u defines a homotopy of K onto another finite complex that does not contain x. If $\{x_i\}_{i=1,\ldots,r}$ are the vertices of K that realize the maximal distance l, the same argument can be applied to each x_i consecutively to decrease l. □

3. Group actions on graphs and complexes

3.1. **Finitely generated actions.** Given a group Γ, let $\mathcal{G}$ be a graph with a simplicial Γ-action. This is equivalent to saying that Γ acts on $\mathcal{G}$ by isometries with respect to the word metric d.

Definition 17. *The action of Γ on $\mathcal{G}$ is finitely generated if the following properties hold.*

- *$\mathcal{G}$ is connected.*
- *$\mathcal{G} \setminus \{v\}$ is connected for each $v \in \mathcal{V}$.*
- *There are only finitely many Γ-orbits in $\mathcal{V}$ and in $\mathcal{E}$.*
- *The stabilizers of the edges in $\mathcal{E}$ are finite.*

Lemma 18. *If the action of a group Γ on a graph $\mathcal{G}$ is finitely generated, then $\mathcal{G}$ is fine at scale n, as in 2.4, if and only if there are only finitely many orbits of circuits of length at most n in $\mathcal{G}$.*

Proof. Suppose $\mathcal{G}$ is fine at scale n and there are infinitely many circuits l_i of length $m \leq n$ in different orbits in $\mathcal{G}$. Since there are only finitely many Γ-orbits in $\mathcal{E}$, without loss of generality we can assume that the loops are all distinct and that they all contain a fixed edge e. This contradicts fineness at scale n.

For the other direction, suppose there are only finitely many orbits of circuits of length at most n in $\mathcal{G}$. If there are distinct circuits l_i of length $m \leq n$ all containing an edge e, we can suppose after passing to a subsequence that $l_i = \gamma_i l$ for $\gamma_i \in \Gamma$ all distinct, and hence $e = \gamma_i e'$ where e' is an edge in l. This contradicts the fact that the stabilizers of edges are finite in $\mathcal{G}$. $\square$

The following is proved in [4, Lemma 4.3]. It shows that in a fine graph, the finite edge stabilizers condition can be replaced by finite pair stabilizers. Here by a *pair stabilizer* we mean the intersection of the stabilizers of two distinct vertices.

Lemma 19. *If the action of a group Γ on a fine graph $\mathcal{G}$ is finitely generated then the pair stabilizers are finite.*

Recall that for $v \in \mathcal{V}$, $Link^{\varsigma}_L(v)$ is the full subgraph of $\mathcal{G}^{\varsigma}_L$ whose vertices are all the edges in $\mathcal{G}$ incident to v.

Lemma 20. *If the action of a group Γ on the graph $\mathcal{G}$ is finitely generated, then for any L, Γ acts on the L-graph $\mathcal{G}^{\varsigma}_L$ by isometries with the following properties.*

- *There are only finitely many orbits of vertices in $\mathcal{G}^{\varsigma}_L$.*
- *The stabilizers of the vertices and the edges in $\mathcal{G}^{\varsigma}_L$ are finite.*
- *For each $v \in \mathcal{V}$, its stabilizer $Stab(v)$ acts on $Link^{\varsigma}_L(v)$ with finite quotient.*

Proof. Since Γ acts on $\mathcal{G}$ by isometries, it clearly also acts on an L-graph $\mathcal{G}^{\varsigma}_L$ by isometries. Moreover, since there are only finitely many orbits of edges in $\mathcal{G}$, there are only finitely many orbits of vertices in $\mathcal{G}^{\varsigma}_L$. By definition the stabilizer of vertices in $\mathcal{G}^{\varsigma}_L$ are finite. Now if the stabilizer of an edge (e, e') in $\mathcal{G}^{\varsigma}_L$ is infinite, then the stabilizers of the edges $e, e' \in \mathcal{G}$ are infinite, which would give a contradiction.

For all $v \in \mathcal{V}$, $Link^{\varsigma}_L(v)$ is invariant under $Stab(v)$. Suppose this action has not a finite quotient. Then there exists an infinite sequence of vertices e_i in $Link^{\varsigma}_L(v)$ in distinct $Stab(v)$-orbits. Since there are only finitely many Γ-orbits of vertices in $\mathcal{G}^{\varsigma}_L$, we can suppose that $e_i = \gamma_i e$ for distinct $\gamma_i \in \Gamma$ and some vertex e in $Link^{\varsigma}_L(v)$. Moreover, after passing to a subsequence and translating by an element of Γ we can also assume that $\gamma_i \in Stab(v)$, which gives a contradiction. □

3.2. **Finitely presented actions.**

Definition 21. *The action of* Γ *on a complex* X *is finitely presented if the following properties hold.*

- *The action of* Γ *on the 1-skeleton of* X *is finitely generated.*
- X *is simply connected.*
- *There are only finitely many orbits of* 2*-simplices.*

Note that when X is simplicial complex the last condition is equivalent to saying that there are only finitely many orbits of 3-circuits in X, hence equivalent by Lemma 18 to saying that $X^{(1)}$ is fine at scale 3.

We denote by $\mathcal{G}$ the 1-skeleton and by $\mathcal{V}(X)$ the 0-skeleton of X.

Lemma 22. *If the* Γ*-action on a simplicial complex* X *is finitely presented then for all* $L \geq 1$ *we have the following.*

- *for all* $v \in \mathcal{V}(X)$*,* $Link^{\varsigma}_L(v)$ *is connected,*
- $\mathcal{G}^{\varsigma}_L$ *is connected.*

and moreover

- $\mathcal{G}^{\varsigma}_1$ *is uniformly locally finite.*

Proof. Lemma 10 implies the first and second statements since X is simply connected and $\mathcal{G} \setminus \{v\}$ is connected for each $v \in \mathcal{V}(X)$. Moreover $\mathcal{G}$ is fine at scale 3 thus Lemma 8 says that $\mathcal{G}^{\varsigma}_1$ has uniformly finite balls. □

Lemma 23. *If the* Γ*-action on a simplicial complex* X *is finitely presented, then for each* $v \in \mathcal{V}$ *its stabilizer* $Stab(v)$ *is finitely generated.*

Proof. Consider the edge graph $\mathcal{G}^{\varsigma}_1$ and the subgraph $Link^{\varsigma}_1(v)$ for $v \in \mathcal{V}(X)$. Now $Link^{\varsigma}_1(v)$ is connected locally finite by Lemma 22. Moreover $Stab(v)$ acts on $Link^{\varsigma}_1(v)$ with finite quotient and with finite edge stabilizers (Lemma 20), which completes the proof. □

4. Tuples

4.1. **Graph tuples and tuples.**

Definition 24. *A graph tuple is a list* $(\Gamma, \Gamma', \mathcal{G}, \mathcal{V}')$ *with the following properties.*

- Γ *is a group.*
- $\Gamma' = \{\Gamma_i \mid i \in \mathrm{I}\}$ *is a family of subgroups of* Γ*, possibly with repetitions.*
- $\mathcal{G}$ *is a graph with a* Γ*-action.*

- *$\mathcal{V}'$ is a Γ-invariant subset of $\mathcal{V}$ containing all the vertices of infinite valence in $\mathcal{V}$, i.e. $\mathcal{V}_\infty \subseteq \mathcal{V}'$,*
- *Each Γ_i is the stabilizer of some vertex $v_i \in \mathcal{V}'$, and $\Gamma_i \mapsto \Gamma v_i$ is a one-to-one correspondence between Γ' and the set of Γ-orbits in $\mathcal{V}'$.*

Definition 25. *A tuple is a list $(\Gamma, \Gamma', X, \mathcal{V}')$ such that X is a simplicial complex and $(\Gamma, \Gamma', X^{(1)}, \mathcal{V}')$ is a graph tuple.*

Any graph tuple is obviously a tuple. *A priori* we do not impose any finiteness conditions on Γ, Γ', or Γ_i. We will work in the category of simplicial complexes, but these notions allow using other categories as well. If needed, similar definitions can be given for cell complexes, combinatorial cell complexes, metric spaces, etc.

4.2. **Finiteness conditions.**

Definition 26. *A tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is finitely generated, or of type $\mathcal{F}_1$, if the action of Γ on $X^{(1)}$ is finitely generated in the sense of Definition 17.*

A tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is finitely presented, or of type $\mathcal{F}_2$, if the action of Γ on $X^{(1)}$ is finitely generated in the sense of Definition 21.

More generally, a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is of type $\mathcal{F}_n$, $n \geq 2$, if

- *the action of Γ on $\mathcal{G}$ is finitely generated,*
- *$\pi_k(X) = 0$ for all $k \leq n-1$, and*
- *there are only finitely many orbits of k-cells for each $k \leq n$.*

A tuple is of type $\mathcal{F}_\infty$, if it is of type $\mathcal{F}_n$ for any n. A tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is of finite type, or of type $\mathcal{F}$, if it is of type $\mathcal{F}_\infty$ and X is finite dimensional.

These notions descend to pairs. A pair (Γ, Γ') is called *finitely generated*, *finitely presented*, of type $\mathcal{F}_n$, $\mathcal{F}_\infty$, $\mathcal{F}$, if there exists a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ which has the respective property.

The above notion of finite presentation is an equivalent restatement of Osin's definitions [24]. (Γ, Γ') is finitely presented in the sense of Definition 26 iff Γ is relatively finitely presented with respect to Γ' in the sense of [24].

The following result, which is parallel to [24, Theorem 1.1], follows from the definition of a finitely presented tuple together with Lemma 23.

Theorem 27. *If a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is finitely presented, then*

(a) *Γ' is a finite family, and*
(b) *each $\Gamma_i \in \Gamma'$ is finitely generated.*

4.3. **Tuples with thin and fine triangles.** A tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is said to *have thin* or *fine triangles*, if its 1-skeleton $(\Gamma, \Gamma', X^{(1)}, \mathcal{V}')$ has, respectively, thin or fine triangles. A pair (Γ, Γ') *has thin* or *fine triangles* if there exists a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ which has, respectively, thin or fine triangles.

4.4. **Isoperimetric functions for tuples.** Given a graph $\mathcal{G}$ we can think a loop in $\mathcal{G}$ as a cellulation of the circle S^1, together with a cellular map from S^1 into $\mathcal{G}$ which send each 1-cell homeomorphically onto an edge of $\mathcal{G}$. We can also define a *cellular disk* (D, h) in $\mathcal{G}$ of *coarseness*

n as a cellulation of the unit disk D together with a map h from the 1-skeleton of D into $\mathcal{G}$ such that the boundary of each 2-cell has length n (i.e has at most n 1-cells) and get mapped to a loop of length n in $\mathcal{G}$. We measure the area of (D, h) as the numbers of 2-cells in D. We speak of (D, h) as the *spanning disk* for the loop $h|\partial D$, and of the loop $h|\partial D$ as *bounding* (D, h).

Definition 28. *For a graph $\mathcal{G}$ a function $f : \mathbb{N} \to \mathbb{N}$ is an isoperimetric function of coarseness n for $\mathcal{G}$ if each loop of length l in $\mathcal{G}$ has a spanning disk of coarseness n, whose area is bounded by $f(l)$.*

Proposition 29. *If a graph $\mathcal{G}$ has a well defined isoperimetric function f of coarseness n and it is fine at scale n then it is fine.*

Proof. Suppose that there is an isoperimetric function f of coarseness n for $\mathcal{G}$ and $\mathcal{G}$ is fine at scale n but it is not fine. Then there exist $m \in \mathbb{N}$, $e \in \mathcal{E}$ and an infinite sequence of distinct circuits l_i^1 of length m all containing e_1. In particular $\text{area}(l_i^1) \leq f(m) < \infty$ for all i. After passing to a subsequence we can assume that for all i, $\text{area}(l_i^1) = m_1 \leq f(m)$. We argue by induction on the areas of l_i^1. For each i we consider a spanning disc (D_i, h) of area m_1 bounded by l_i^1 and a 2-cell σ_i such that $h|\partial\sigma_i$ contains e only once. Note that such 2-cell exists since l_i^1 is a circuit in $\mathcal{G}$, and $h|\partial\sigma_i$ is a loop of length at most n in $\mathcal{G}$ since h has coarseness n. We denoted $h|\partial\sigma_i$ by c_i. Now we consider the loops α_i in $\mathcal{G}$ obtained from l_i^1 replacing e_1 by $c_i \setminus \{e_1\}$. By construction $\text{area}(\alpha_i) = m_1 - 1 < m_1$ and α_i contains all the edges of c_i distinct from e_1. The loops α_i do not have to be circuits, however since l_i^1 are all distinct we can find, by reducing α_i, an infinite sequence of distinct circuits l_i^2, each containing at least one edge of $c_i \setminus \{e_1\}$. Moreover since $c_i \setminus \{e_1\}$ has only finitely many edges we can suppose after passing to a subsequence that all these circuits l_i^2 contain the same edge e_2 and$\text{area}\, l_i^2 = m_2 < m_1$. Now we repeat the argument this time for l_i^2 to obtain distinct circuits l_i^3 all containing an edge e_3 and with $\text{area}\, l_i^1 = m_3 < m_2$. This induction will give eventually an edge contained in infinitely many disctinct circuits of length n, which is a contradiction by the fineness at scale n. □

The following follows as a corollary of the above proposition.

Corollary 30. *There exists a constant n such that If $\mathcal{G}$ is a graph with δ-fine triangles, and it is fine at scale n then $\mathcal{G}$ is fine.*

Proof. Using Proposition 7 we see that $\mathcal{G}$ has δ'-thin triangles where δ' is a constant depending only on δ, and by a well known result (see for example [5] Chapter III.H Proposition 2.7), $\mathcal{G}$ satisfies a linear isoperimetric inequality of coarseness n where n depends only on δ'. Thus Proposition 29 gives the result required.

□

For simply connected complexes we can reformulate the notion of area and isoperimetric function as follow. In a simply connected complex X, the *area* of a loop is the minimum number times it passes over two-cells during a null-homotopy, minimum taken over all null-homotopies. For a loop we denote its area by $\text{area}(l)$.

Definition 31. *Let* $(\Gamma, \Gamma', X, \mathcal{V}')$ *be a finitely presented tuple. A function* $f : \mathbb{N} \to \mathbb{N}$ *is an isoperimetric function for* $(\Gamma, \Gamma', X, \mathcal{V}')$ *if it is an isoperimetric function for* X *in the non-relative sense, i.e. it is a function such that the area of any loop in* X *of length at most* l *is at most* $f(l)$*.*

It is important that f take finite values. A finitely presented tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ has a linear isoperimetric function in the above sense if and only if Γ has a linear relative isoperimetric function with respect to Γ' in the sense of [24]. The same applies to quadratic, cubic, polynomial, exponential, etc isoperimetric functions.

Proposition 32. *If the action of a group* Γ *on a complex* X *is finitely presented, then* $X^{(1)}$ *is fine if and only if there exists an isoperimetric function for* X*.*

Proof. Denote $\mathcal{G} := X^{(1)}$ and suppose $\mathcal{G}$ is fine. We set $f(l)$ to be the maximal area of circuits of length l, which is finite, since by Lemma 18 for any given l there are only finitely many orbits of l-circuits in X. If α is a loop of length l in $\mathcal{G}$, it can be split α inductively into at most l circuits of length at most l. Since each circuit has area at most $f(l)$, the area of α is bounded by $lf(l)$, which implies that there is a well defined isoperimetric function.

For the other direction, first observe that there are only finitely many 2-cells containing any given edge in $\mathcal{G}$. Indeed, if there are distinct 2-cells c_i all containing a given edge e, since there are only finitely many orbits of two cells in X we can suppose after passing to a subsequence that $c_i = \gamma_i c$ for some γ_i, where γ_i are pairwise distinct. In particular, $\gamma_i e' = e$ for some $e' \in \mathcal{E}(c)$, which would give a contradiction with the fact that the stabilizers of edges are finite.

Now suppose that there is an isoperimetric function for X and that $\mathcal{G}$ is not fine. Then the same argument used in the proof of Proposition 29 will give an edge contained in infinitely many distinct 2-cell, which is a contradiction by the first observation.

□

Proposition 33. *If* $(\Gamma, \Gamma', X, \mathcal{V}')$ *is a finitely presented tuple such that* $X^{(1)}$ *has fine triangles, then* X *satisfies a (combinatorial) linear isoperimetric inequality.*

Proof. The combinatorial linear isoperimetric inequality for relative presentations was shown in [24]. We give here an alternative in our language. $X^{(1)}$ has fine triangles, so by Proposition 7 it has thin triangles. Thus by a known result (see for example [5] Chapter III.H Proposition 2.7) $X^{(1)}$ has linear isoperimetric inequality, and so does X. □

5. Relative hyperbolicity.

Definition 34. *Let* Γ *be a group and* $\Gamma' = \{\Gamma_i \mid i \in I\}$ *be a family of its subgroups.* Γ *is called relatively hyperbolic with respect to* Γ' *if there exists a graph* $\mathcal{K}$ *on which* Γ *acts such that the following conditions are satisfied.*

- Γ *is finitely generated.*
- Γ' *is finite and each* Γ_i *is finitely generated.*
- $\mathcal{K}$ *is fine and has thin triangles.*
- *There are finitely many orbits of edges and each edge stabilizer is finite.*

- *There exists a Γ-invariant subset $\mathcal{V}'$ such that $\mathcal{V}_\infty \subseteq \mathcal{V}' \subseteq \mathcal{V}$ and the stabilizers of vertices in $\mathcal{V}'$ are precisely Γ_i and their conjugates.*

This definition is a slight generalization of relative hyperbolicity first introduced by Gromov in [12], later reformulated by Bowditch. In [4] Bowditch gave a combinatorial formulation of relative hyperbolicity for a group Γ and showed that it is equivalent to the original Gromov's definition. In the original definition it is assumed that the elements of Γ' are infinite subgroups, i.e. $\mathcal{V}' = \mathcal{V}_\infty$. In the above definition we allow the elements of Γ' to be finite as well as infinite. The elements of Γ' and their conjugates will be called *peripheral subgroups*, similarly to the original definition.

In his paper Bowditch also shows the following result (Proposition 4.9 in [4]).

Theorem 35. *A group Γ is hyperbolic relative to Γ' if and only if there exists a simplicial complex X with a Γ-action such that*

- *the vertex stabilizers are exactly the elements of Γ' and their conjugates,*
- *Γ acts on X with finite quotient and finite edge stabilizers,*
- *X is hyperbolic, has no cut vertex, and is locally finite away from its vertex set,*
- *$X^{(1)}$ is fine and every circuit of length 3 in $X^{(1)}$ bounds a 2-simplex.*

This theorem holds for the generalized Definition 34 of relative hyperbolicity with no change in the argument. We give a sketch of the proof together with some references. In Bowditch's paper relative hyperbolicity is formulated in terms of actions on sets. A set on which a group Γ acts is a *hyperbolic Γ-set* if it is a vertex set of a graph $\mathcal{K}$ as in Definition 34. Moreover Proposition 4.11 of his work shows that this property passes to subsets containing all the vertices with infinite stabilizers. Thus we start with a graph $\mathcal{K}$ given by Definition 34 with vertex set $\mathcal{V}$ whose stabilizers are Γ'. Generally the graph $\mathcal{K}$ will have cut points, but it satisfies all the other properties of the 1-skeleton of the complex X in the above theorem. Now, since the vertex stabilizers in $\mathcal{K}$ are finitely generated, for a cut vertex v, its stabilizer acts with finite quotient and finite edge stabilizers on a locally compact graph L_v such that $\mathcal{V}(L_v) = \mathcal{V}(\overline{Star}_{\mathcal{K}}(v)) \backslash Star_{\mathcal{K}}(v))$. We attach the graph L_v around v in $\mathcal{K}$ using the identification of vertices. and repeat this Γ-equivariantly in all translate of v. The new graph, which we still denote $\mathcal{K}$, remains hyperbolic and fine, and for all $v \in \mathcal{K}$ we have $\mathcal{K} \setminus \{v\}$ is connected. We consider the graph K^n (see section 2.4) and glue a 2-cell to each circuits of length 3. This gives a complex as required.

If $\mathcal{G}$ is the 1-skeleton of the complex X given by Theorem 35, then we obtain the following corollary.

Corollary 36. *A group Γ is hyperbolic relative to Γ' if and only if there exists a finitely generated graph tuple $(\Gamma, \Gamma', \mathcal{G}, \mathcal{V}')$ with thin triangles, with $\mathcal{G}$ fine and with $\mathcal{V} = \mathcal{V}'$.*

5.1. **Hyperbolic tuples.**

Definition 37. *A tuple $\mathcal{T} = (\Gamma, \Gamma', X, \mathcal{V}')$ is hyperbolic if*

- *it is finitely generated as in Definition 26,*
- *has fine triangles as in Definition 6, and*
- *$X^{(1)}$ is fine as in 2.4.*

A pair (Γ, Γ') *is hyperbolic if there exists a hyperbolic tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$.

Note that Proposition 30 allows us to replace in this definition fineness by fineness at some scale, which is a priori a weaker condition.

Theorem 38. *For every hyperbolic pair* (Γ, Γ') *there exists a hyperbolic tuple* $(\Gamma, \Gamma', X, \mathcal{V}')$ *of type* $\mathcal{F}$, *where* $\mathcal{V}(X) = \mathcal{V}'$.

Proof. Take a graph tuple $(\Gamma, \Gamma', \mathcal{G}, \mathcal{V}')$ guaranteed by Corollary 36. $\mathcal{G}$ is fine and δ-hyperbolic for some δ. We consider the simplicial complex X constructed in 2.9. Recall that X is constructed by gluing a simplex to each complete subgraph of $\mathcal{G}_0$, which is the graph with the same vertex set as $\mathcal{G}$ and in which two vertices are connected if there is a geodesic path in $\mathcal{G}$ connecting them with angles and length at most μ, where μ is a constant depending only on δ. Theorem 16 says that for μ large enough X is contractible.

We will show that the tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is a hyperbolic tuple of type $\mathcal{F}$. In order to prove that we have to check that the action of Γ on $X^{(1)} = \mathcal{G}_0$ is finitely generated, there are only finitely many orbits of i-cells for $i \leq n$, X is hyperbolic and $X^{(1)}$ is fine.

Since the action of Γ on $\mathcal{G}$ is finitely generated, we see by definition that $\mathcal{G}_0$ is connected, and $\mathcal{G}_0 \setminus \{v\}$ is connected for all $v \in \mathcal{V}(\mathcal{G}_0)$. Moreover there are only finitely many orbits in $\mathcal{V}(\mathcal{G}) = \mathcal{V}(\mathcal{G}_0)$. Let $\{e_1, \ldots, e_r\}$ be a set of orbit representatives of $\mathcal{E}(\mathcal{G})$. For each e_i consider the set S_i of admissible sequences with all elements in the ball $B(e_i)$ of radius 1 of the edge graph $\mathcal{G}^{\varsigma}_L$ where $L \geq \mu^2$. By Lemma 8 this set is finite. Note that for each edge e in $\mathcal{G}_0$, all the edges of a geodesic representative of e in $\mathcal{G}$ lies in one of the translates of S_i, that ensures that there are only finitely many orbits of edges in $\mathcal{G}_0$. Moreover since $\mathcal{G}$ is fine Lemma 19 implies that the the stabilizers of pairs in $\mathcal{V}(\mathcal{G})$ are finite, hence stabilizers of edges in $\mathcal{G}_0$ are finite. This proves that the action of Γ on $X^{(1)} = \mathcal{G}_0$ is finitely generated.

Lemma 14 says that for each edge in $\mathcal{G}_0$ there are only finitely many simplices containing it. This together with finitely many orbits of edges implies finiteness of orbits of i-cells for each i.

We note that $\mathcal{G}_0$ is hyperbolic. Indeed, it is quasiisometric to $\mathcal{G}$ since $d_{\mathcal{G}_0}(x, y) \leq d_{\mathcal{G}}(x, y) \leq \mu d_{\mathcal{G}_0}(x, y)$ for all $x, y \in \mathcal{V}$. Thus we have Y hyperbolic. It remains to show that $X^{(1)} = \mathcal{G}_0$ is fine. This follows from lemmas 33 and 32. □

Theorem 39. *The following statements are equivalent.*

(a) (Γ, Γ') *is a hyperbolic pair in the sense of Definition 37.*
(b) Γ *is hyperbolic relative to* Γ' *in the sense of Definition 34.*

Proof. By Proposition 7 thin triangle property is equivalent to fine triangle property, hence both definitions can be reformulated using either property. Then we see that (**b**) ⇒ (**a**) follows from Corollary 36.
(**a**) ⇒ (**b**) For a hyperbolic pair (Γ, Γ') by Theorem 38 there is a hyperbolic tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ of finite type, in particular finitely generated. Clearly the 1-skeleton of $X^{(1)}$ satisfies all the requirements of Definition 34. □

Lemma 40. *If* $(\Gamma, \Gamma', X, \mathcal{V}')$ *is a finitely presented hyperbolic tuple, then for all* $L \geq 1$ *the* L*-graph* $\mathcal{G}^{\varsigma}_L$ *satisfies the following.*

- *For all* $v \in \mathcal{V}$, $Link^{\varsigma}_L(v)$ *is connected.*

- *$\mathcal{G}^\varsigma_L$ is connected.*
- *$\mathcal{G}^\varsigma_L$ is uniformly locally finite.*

Moreover Γ acts on $\mathcal{G}^\varsigma_L$ with finite quotient and finite edge and vertex stabilizers.

$\mathcal{G}^\varsigma_L$ plays the role of a Cayley graph for Γ.

Proof. The connectedness follows from Lemma 22 and locally finiteness from Lemma 8 since $\mathcal{G}$ is fine, hence fine at any scale. The rest is obtained from Lemma 20 and the following remark. For any group action on a uniformly locally finite graph with only finitely many orbits of vertices, the action has finite quotient. □

6. Relative straightening.

6.1. **Notations.** We work with a hyperbolic tuple $\mathcal{T} = (\Gamma, \Gamma', X, \mathcal{V}')$ and generalize the constructions of [17]. In our relative setting it is convenient to work both with vertices and edges, so definitions will modify accordingly. Since the action of Γ on X in general is not free, we will need to average over certain sets of vertices and edges; the hyperbolicity of the tuple will guarantee that we always average over finite sets.

For now, edges of $\mathcal{G}$ are assumed to be non-oriented. $\mathcal{V}(\gamma)$ and $\mathcal{E}(\gamma)$ are, respectively, the set of vertices and edges in an edge path γ. For $v, w \in \mathcal{V}$, $W \subseteq \mathcal{V}$ and $t \in \mathbb{Z}$, denote

$$\mathrm{Geod}(v, W) = \bigcup_{w\in W} \mathrm{Geod}(v, w),$$

$$\mathcal{V}[v, W] := \bigcup_{\gamma\in \mathrm{Geod}(v,W)} \mathcal{V}(\gamma), \qquad \mathcal{E}[v, W] := \bigcup_{\gamma\in \mathrm{Geod}(v,W)} \mathcal{E}(\gamma).$$

For a vertex w we will write $\mathcal{E}[v, w]$ and $\mathcal{V}[v, w]$ instead of $\mathcal{E}[v, \{w\}]$ and $\mathcal{E}[v, \{w\}]$.

$$\mathcal{V}[v, w; t] := \{x \in \mathcal{V}[v, w] \mid d(v, x) = t\}, \qquad \mathcal{E}[v, w; t] := \{e \in \mathcal{E}[v, w] \mid d(v, e) = t\},$$

$$\mathcal{E}[v, W; t] := \{e \in \mathcal{E}[v, W] \mid d(v, e) = t\} = \bigcup_{w\in W} \mathcal{E}[v, w; t].$$

Recall that $B^\varsigma(e, r)$ is a closed d_ς-ball at e of radius r, where d_ς is the snake metric on edges. For an edge path γ we will denote $N(\gamma, r) \subseteq \mathcal{G}$ and $N^\varsigma(\gamma, r) \subseteq \mathcal{E}$ the r-neighborhoods of γ in the metrics d and d_ς, respectively.

For $v, w \in \mathcal{V}$, $e \in \mathcal{E}[v, w]$ and $r \in [0, \infty)$, the set

$$Fl(v, w, e; r) := \mathcal{E}[v, \mathcal{V}; d(v, e)] \cap B^\varsigma(e, r) \subseteq \mathcal{E}[v, \mathcal{V}]$$

is the *flower* with respect to v, w, e, r. By Lemma 40, each ball $B^\varsigma(e, r)$, and hence each flower, is a finite set of edges, and for a fixed r, the cardinalities of the flowers are uniformly bounded. Moreover, since there are only finitely many Γ-orbits of edges, its cardinality is bounded by some $\omega = \omega(\mathcal{T}, r)$.

For a set S, $\mathbb{Q}S$ is the $\mathbb{Q}$-vector space spanned by S. The *average* of a finite subset $S' \subseteq S$ is the element of $\mathbb{Q}S$, denoted $\mathrm{av}(S')$, which is the characteristic function of S' divided by $\#S'$.

For $e \in \mathcal{E}[v,\mathcal{V}]$ denote $\mathrm{av}_{Fl}(v,w,e;r) := \mathrm{av}\big(Fl(v,w,e;r)\big) \in \mathbb{Q}\mathcal{E}[v,\mathcal{V}]$. The map $\mathrm{av}_{Fl}(v,w,\cdot\,;r)$ extends by linearity to $\mathrm{av}_{Fl}(v,w,\cdot\,;r) : \mathbb{Q}\mathcal{E}[v,\mathcal{V}] \to \mathbb{Q}\mathcal{E}[v,\mathcal{V}]$. The fine triangles property and Lemma 40 imply that $\mathcal{E}[v,w;t]$ is finite, so

$$\mathrm{av}_{Fl}(v,w,\mathrm{av}(\mathcal{E}[v,w;t]);\delta)$$

is a well-defined function on edges. By the fine triangles property, for any $e \in \mathcal{E}[v,w;t]$ its support satisfies

$$supp\,\mathrm{av}_{Fl}(v,w,\mathrm{av}(\mathcal{E}[v,w;t]);\delta) \subseteq B^{\varsigma}(e,2\delta).$$

6.2. **The functions $f(a,b;i)$ and $\bar{f}(a,b)$.** For $e \in \mathcal{E}[v,\mathcal{V}]$, let $\iota_v(e)$ be the *v-initial vertex* of e, i.e. the vertex of e closest to v. By linearity this extends to a $\mathbb{Q}$-linear map $\iota_v : \mathbb{Q}\mathcal{E}[v,\mathcal{V}] \to \mathbb{Q}\mathcal{V}$. Similarly $\tau_v : \mathbb{Q}\mathcal{E}[v,\mathcal{V}] \to \mathbb{Q}\mathcal{V}$ is the *v-terminal vertex* map. Define a $\mathbb{Q}$-linear projection $pr_a : \mathbb{Q}\mathcal{E}[a,\mathcal{V}] \to \mathbb{Q}\mathcal{E}[a,\mathcal{V}]$ toward a as follows. It suffices to describe pr_a only on edges $e \in \mathcal{E}[v,\mathcal{V}]$.

- If $d(a,e) = 0$, let $pr_a(e) := e$.
- If $d(a,e) > 0$, let $pr_a(e) := \mathrm{av}(\mathcal{E}[a,\iota_a e;t])$, where t is the largest integral multiple of 20δ satisfying $t < d(a,e)$.

The definition of pr_a will only be important for the case when $d(a,e)$ is an integral multiple of 20δ. In this case pr_a moves e toward a exactly by distance 20δ.

Now for all $a,b \in X$ we define a 1-chain $f(a,b) = f(a,b;20\delta)$ in X inductively on $d(a,b)$.

- If $d(a,b) \le 20\delta+1$, the definition of $f(a,b)$ is not important, for example one can set $f(a,b) := \mathrm{av}(\mathcal{E}[a,b;d(a,b)-1])$.
- If $d(a,b) > 20\delta+1$ and $d(a,b)$ does not equal 1 modulo 20δ, let
$$f(a,b) := f\big(a,\mathrm{av}(\mathcal{V}[a,b;t])\big),$$
where t is the largest integer multiple of 20δ satisfying $t < d(a,b)$ (hence $t < d(a,b)-1$), and $f\big(a,\mathrm{av}(\mathcal{V}[a,b;t])\big)$ is defined by linearity in the second variable.
- If $d(a,b) > 20\delta+1$ and $d(a,b)$ equals 1 modulo 20δ, let
$$f(a,b) := pr_a\big(\mathrm{av}_{Fl}(a,b,\mathrm{av}(\mathcal{E}[a,b;d(a,b)-1]);\delta)\big).$$

In the above definitions we used integral multiples of 20δ, which are convenient to describe as numbers of the form $20\delta + 20\delta n$, $n \in \mathbb{Z}$. One can deal equally well with the numbers of the form $i + 20\delta n$ for a fixed $i \in \mathbb{Z}$. Replacing $20\delta + 1$ with $i+1$ and integer multiples of 20δ with numbers congruent to i modulo 20δ in the above definitions we obtain a 1-chain $f(a,b;i)$ such that each edge in its support is at distance i from a.

Proposition 41. *The function f defined above satisfies the following properties:*

1. *$f(a,b;i)$ is a convex combination of edges.*
2. *If $d(a,b) > i$, then $supp\, f(a,b;i) \subseteq Fl(a,b;e,\delta)$ for each $e \in \mathcal{E}[a,b;i]$.*
3. *If $d(a,b) \le i$, then $f(a,b;i) = \mathrm{av}(\mathcal{E}[a,b;d(a,b)-1])$.*
4. *f is Γ-equivariant, i.e. $f(ga,gb;i) = g(f(a,b;i))$ for any $a,b \in X^{(0)}$ and $g \in \Gamma$.*
5. *For each fixed i there exist $L \in [0,\infty)$ and $\lambda \in [0,1)$ such that for any $a,b,c \in X^{(0)}$,*
$$\big|f(a,b;i) - f(a,c;i)\big|_1 \le L\lambda^{(b|c)_a}.$$

This is proved along the lines of [17, Proposition 3] inductively on $d(a,b)$ using the fact that flowers are uniformly finite. The proof goes through because of the following property which is implied by the fine triangles condition.

- Suppose $e_1, e_2 \in \mathcal{E}[a, \mathcal{V}, 20\delta n]$ satisfy $d(\tau_a e_1, \tau_a e_2) \leq 3\delta$, then for any $\bar{e}_1$ and $\bar{e}_2$ in the supports of $pr_a(e_1)$ and $pr_a(e_2)$, respectively, we have $d_\varsigma(\bar{e}_1, \bar{e}_2) \leq \delta$.

For $e \in \mathcal{E}$ let $star_{3\delta}(e) := \mathrm{av}\big(B^\varsigma(e, 3\delta)\big)$; this extends by linearity to a map $star_{3\delta} : \mathbb{Q}\mathcal{E} \to \mathbb{Q}\mathcal{E}$. For $a, b \in X^{(0)}$ let

$$\bar{f}(a,b) := star_{3\delta}\left(\frac{1}{11\delta + 1}\sum_{i=5\delta}^{16\delta} f(a,b;i)\right).$$

Proposition 42. *The function $\bar{f}$ defined above satisfies the following properties:*

1. *$\bar{f}(a,b)$ is a convex combination of edges that are oriented towards a.*
2. *Pick any $\gamma \in \mathrm{Geod}(a,b)$ and let I be the subset of edges $e \in \mathcal{E}(\gamma)$ such that $5\delta \leq d(a,e) \leq 16\delta - 1$. If $d(a,b) > 20\delta$, then $\mathrm{supp}\,\bar{f}(a,b) \subseteq N^\varsigma(I, 2\delta)$*
3. *If $d(a,b) \leq 20\delta$, then $\bar{f}(a,b) := \mathrm{av}(\mathcal{E}[a,b; d(a,b) - 1])$.*
4. *$\bar{f}$ is Γ-equivariant, i.e. $\bar{f}(ga, gb) = g(\bar{f}(a,b))$ for any $a, b \in X^{(0)}$ and $g \in \Gamma$.*
5. *There exist $L \in [0, \infty)$ and $\lambda \in [0,1)$ depending only on the tuple $\mathcal{T}$ such that for any $a, b, c \in X^{(0)}$,*
$$\left|\bar{f}(a,b) - \bar{f}(a,c)\right|_1 \leq L\lambda^{(b|c)_a}.$$
6. *There exists a constant $\lambda' \in [0,1)$ depending only on $\mathcal{T}$ such that if $a, b, c \in X^{(0)}$, satisfy $(a|b)_c \leq 20\delta$ and $(a|c)_b \leq 20\delta$, then*
$$\left|\bar{f}(b,a) - \bar{f}(c,a)\right|_1 \leq 2\lambda'.$$
7. *Let $a, b, c \in X^{(0)}$, $\gamma \in \mathrm{Geod}(a,b)$, and $c \in N(\gamma, 4\delta)$, then $\mathrm{supp}\,\bar{f}(c,a) \subseteq N(\gamma, 4\delta)$.*

This is proved using Proposition 41 similarly to [17, Proposition 7]. The fine triangles property and the following facts guarantee that the proof goes through.

- The number of edges in the support of $\bar{f}(a,b)$ is bounded by a constant depending only on $\mathcal{T}$.
- For all $a, b, c \in X^{(0)}$, if $(a|b)_c \leq 20\delta$ and $(a|c)_b \leq 20\delta$, then there exist edges $e_1 \in \mathrm{supp}\,f(b,a)$ and $e_2 \in \mathrm{supp}\,f(c,a)$ such that $d(a, e_1) = d(a, e_2)$ and $d_\varsigma(e_1, e_2) \leq 3\delta$.

For $e \in \mathcal{E}$, let $\partial_+(e) \in \mathbb{Q}\mathcal{V}$ be the sum of the vertices incident to e, each taken with coefficient $1/2$; this extends to a $\mathbb{Q}$-linear map $\mathbb{Q}\mathcal{E} \to \mathbb{Q}\mathcal{V}$. Then $\partial_+(\bar{f}(b,a))$ is a convex combination of vertices which satisfies the same (properly restated) properties as $f(b,a)$ does in the above proposition.

6.3. **The 1-chain $q[a,b]$.** For each fixed a and b, $\bar{f}(b,a)$ is a function on the unoriented edges of $\mathcal{G}$. To talk about 1-chains we will now assume that there are two possible orientations for each edge in $\mathcal{G}$. A *1-chain* in $\mathcal{G}$ is a function on oriented edges which takes opposite values on oppositely oriented edges.

For $a, b \in X^{(0)}$ let

$$p'[a,b] := \frac{1}{\#\mathrm{Geod}(a,b)} \sum_{\gamma \in \mathrm{Geod}(a,b)} \gamma,$$

where γ is viewed as a 1-chain with boundary $b - a$. This makes sense because $\mathrm{Geod}(a, b)$ is finite by Lemma 9.

For $a, b \in X^{(0)}$ we define a 1-chain $q'[a, b]$ inductively on $d(a, b)$ as follows.

- If $d(a, b) \leq 20\delta$, let $q'[a, b] := p'[a, b]$.
- If $d(a, b) > 20\delta$, let

$$q'[a,b] := q'\big[a, \partial_+\big(\bar{f}(b,a)\big)\big] + p'\big[\partial_+\big(\bar{f}(b,a)\big), b\big],$$

where $q'\big[a, \partial_+\big(\bar{f}(b,a)\big)\big]$ and $p'\big[\partial_+\big(\bar{f}(b,a)\big), b\big]$ are defined by linearity in the second and first variables, respectively. The inductive definition indeed works because

$$supp\, \bar{f}(b,a) \subseteq B(a, d(a,b) - \delta).$$

q' satisfies $\partial q'[a, b] = b - a$, so q' is a *homological bicombing*.

Proposition 43. *The $\mathbb{Q}$-bicombing q' constructed above satisfies the following conditions.*

(1) *q' is $\mathcal{G}$-equivariant.*
(2) *q' is quasigeodesic, i.e. there exists $C \in [0, \infty)$ such that $supp\, q'[a, b] \subseteq N(\gamma, C)$ for any $\gamma \in \mathrm{Geod}(a, b)$.*
(3) *There exist constants $M \geq 0$ and $N \geq 0$ such that, for all $a, b, c \in X^{(0)}$,*

$$\big|q'[a,b] - q'[a,c]\big|_1 \leq M\, d(b,c) + N.$$

The proof is similar to [17, Proposition 8]. The following property is used to run induction.

- If $a, b, c \in X^{(0)}$ satisfy $(a|c)_b > 20\delta$, then for any $x \in supp\, \partial_+\big(f(b, a)\big)$ we have $d(x, c) < d(b, c)$.

Now let

$$q[a,b] := \frac{1}{2}\big(q'[a,b] - q'[b,a]\big).$$

Theorem 44. *Let $\mathcal{T} = (\Gamma, \Gamma', X, \mathcal{V}')$ be a hyperbolic tuple. Then the $\mathbb{Q}$-bicombing q in X defined above satisfies the following properties.*

(1) *q is quasigeodesic.*
(2) *q is Γ-equivariant.*
(3) *q is anti-symmetric, i.e. $q[a, b] = -q[b, a]$ for any $a, b \in X^{(0)}$.*
(4) *There exists a constant T such that, for any $a, b, c \in X^{(0)}$,*

$$\Big|q[a,b] + q[b,c] + q[c,a]\Big|_1 \leq T.$$

The proof is similar to [17, Theorem 10] in the non-relative case, using Proposition 43.

6.4. Homological isoperimetric inequalities.

Definition 45. *A simply-connected complex X satisfies a homological linear isoperimetric inequality over $\mathbb{Q}$ if the boundary map $\partial : C_2(X, \mathbb{Q}) \to C_1(X, \mathbb{Q})$ is undistorted with respect to the ℓ^1-norms. More generally, a simply-connected complex X satisfies a homological linear isoperimetric inequality for i-cycles over $\mathbb{Q}$ if the boundary map $\partial : C_{i+1}(X, \mathbb{Q}) \to C_i(X, \mathbb{Q})$ is undistorted. Similar definitions are given for $\mathbb{Z}$, $\mathbb{R}$, $\mathbb{C}$ coefficients.*

Theorem 46. *If X is a combinatorial cell complex with finitely many types of 2-cells, such that $X^{(1)}$ is hyperbolic, then the boundary map $\partial : C_2(X, \mathbb{Q}) \to C_1(X, \mathbb{Q})$ is undistorted, that is, X satisfies a homological linear isoperimetric inequality (for 1-cycles) over $\mathbb{Q}$. The same holds for chains with coefficients in $\mathbb{Z}$, $\mathbb{R}$, and $\mathbb{C}$.*

Proof. This argument is due to Gersten. The combinatorial linear isoperimetric inequality was shown in Proposition 33. Allcock and Gersten proved in [1] that any 1-cycle c over $\mathbb{R}$ in X can be represented as $c = \sum_i \alpha_i c_i$ where c_i is the chain represented by a simple oriented loop and $\alpha_i \in \mathbb{R}$, $\alpha_i \geq 0$, *coherently*, i.e. so that $|c|_1 = \sum_i \alpha_i |c_i|_1$. The argument can be generalized to $\mathbb{Q}$ and $\mathbb{Z}$ coefficients (Theorem 6 in [18]). The combinatorial linear isoperimetric inequality implies that there is a constant $K \geq 0$ such that each c_i can be filled with an integral 2-chain a_i whose ℓ^1-norm is bounded by $K|c_i|_1$. Then $a := \sum_i \alpha_i a_i$ is a filling of c satisfying $|a|_1 \leq K|c|_i$. □

The support of a chain c, $supp(c)$, is the closure of the union of simplices σ such that $c(\sigma) \neq 0$. For a subset $S \subseteq X$, $diam(S)$ is defined as the diameter of the set $X^{(0)} \cap S$ with respect to the word metric d on $X^{(1)}$. For a chain c, $diam(c)$ will stand for $diam(supp(c))$

The following is similar to [16, Lemma 5.8] for the non-relative case, but the proof is different. In the relative case it is *not* true in general that the number orbits of edge loops of a given length is finite. This is true only for circuits, i.e. *simple* loops. The same problem happens in higher dimensions.

Theorem 47. *Given a fine graph $\mathcal{G}$ with thin triangles, consider an ideal complex X as in Theorem 16. Then for any integer i, there exist functions $R_i, M_i : [0, \infty) \to [0, \infty)$ with the following property. For each cycle $z \in Z_i(X, \mathbb{Z})$ there exists a chain $a \in C_{i+1}(X, \mathbb{Z})$ such that $\partial a = z$, $diam_d(a) \leq R_i(diam_d(z))$, and $|a|_1 \leq M_i(diam_d(z))|z|_1$.*

Proof. The proof of this is obtained tracing the proof of Theorem 16 and goes by induction on l. □

Theorem 48. *If X is an ideal complex, then for each $k \geq 1$ the boundary map $\partial : C_{k+1}(X, \mathbb{Q}) \to C_k(X, \mathbb{Q})$ is undistorted, that is, X satisfies a homological linear isoperimetric inequality for k-cycles over $\mathbb{Q}$. The same holds for chains with coefficients in $\mathbb{R}$ and $\mathbb{C}$.*

Proof. The proof is similar to [16, Lemma 5.9]. The idea is that, inductively on dimension, for each k-simplex σ in X one can form a cone from σ to a fixed vertex v. The cone is a $(k+1)$-chain whose support lies close to any geodesic from v to a vertex in σ, and whose ℓ^1-norm is bounded independently of σ. The construction of the cone is inductive on dimension. One uses concentric spheres at v to cut the cone over $\partial\sigma$ into slices, which are i-cycles with bounded support, then use Theorem 47 (rather than Lemma 5.8 in [16]) to fill each slice with an $(i+1)$-chain of bounded norm and bounded diameter. Once the cone is defined for each simplex, by linearity one can cone-off any cycle, with the ℓ^1-norm of the cone bounded by a multiple of the ℓ^1-norm of the cycle. □

7. Cohomology and bounded cohomology.

The cohomology of a group relative to a subgroup were defined by Auslender [2] and further studied by Takasu [27] and Ribes [25]. Trotter [28] defined homology and cohomology of a group G with respect to any system of homomorphisms $G_i \to G$. Bieri and Eckmann [3], among other things, provided several equivalent ways to define the homology and cohomology of a group relative to any system of its subgroups.

Bounded cohomology was first introduced by Johnson in the context of Banach algebras [15]. The reader is referred to works of Ivanov [13, 14], Noskov [20, 21], Monod [19] for definitions and standard results about bounded cohomology. In what follows below, we describe the relative (homogeneous) standard resolution, or the snake resolution, and use it to define both relative cohomology and relative bounded cohomology. Our goal is to present the most geometrically transparent definition possible of these formal notions.

7.1. **Notations.** We will use the rational coefficients $\mathbb{Q}$, but everywhere in the paper $\mathbb{Q}$ can be replaced with $\mathbb{R}$ or $\mathbb{C}$.

A Γ-*set* is a set with Γ-action. For a Γ-set S, $\mathbb{Q}S$ denotes the space of finitely supported functions $S \to \mathbb{Q}$, with the induces left Γ-action. Equivalently, $\mathbb{Q}S$ is the space of finite linear combinations of elements of S with rational coefficients. $\mathbb{Q}S$ is given the ℓ^1-norm

$$\left| \sum_{s\in S} \alpha_s s \right|_1 := \sum_{s\in S} |\alpha_s|.$$

Let Γ be a group and $\Gamma' := \{\Gamma_i \mid i \in \mathrm{I}\}$ be an arbitrary nonempty family of its subgroups, possibly with repetitions. Let $i\Gamma$ be a copy of Γ and denote

$$\mathrm{I}\Gamma := \bigsqcup_{i\in\mathrm{I}} i\Gamma, \qquad \Gamma/\Gamma' := \bigsqcup_{i\in\mathrm{I}} i\Gamma/\Gamma_i.$$

$\mathrm{I}\Gamma$ and Γ/Γ' are Γ-sets by the left Γ-action on each $i\Gamma$. With our convention, $\mathbb{Q}\mathrm{I}\Gamma$ is the space of all finitely supported functions $\mathrm{I}\Gamma \to \mathbb{Q}$ and

$$\mathbb{Q}\Gamma/\Gamma' := \bigoplus_{i\in\mathrm{I}} \mathbb{Q}[i\Gamma/\Gamma_i] = \bigoplus_{i\in\mathrm{I}} \mathbb{Q}\Gamma \otimes_{\Gamma_i} \mathbb{Q}$$

is the space of all finitely supported functions $f : \Gamma/\Gamma' \to \mathbb{Q}$.

Δ will denote the kernel of the augmentation map $\epsilon : \mathbb{Q}\Gamma/\Gamma' \to \mathbb{Q}$, $f \mapsto \sum_{x\in\Gamma/\Gamma'} f(x)$.

7.2. **Bounded modules.** A *bounded* $\mathbb{Q}\Gamma$-*module* is a left $\mathbb{Q}\Gamma$-module V which is a normed vector space over $\mathbb{Q}$, with the norm taking values in $[0,\infty)$, such that the induced Γ action on V is by uniformly bounded $\mathbb{Q}$-linear operators. The category of bounded $\mathbb{Q}\Gamma$-modules is the one whose objects are bounded $\mathbb{Q}\Gamma$-modules and whose morphisms are the $\mathbb{Q}\Gamma$-morphisms that are bounded with respect to the norms in the domain and the range. We will use the name *b-morphism* for morphisms in this category, to distinguish them from the usual morphisms between modules.

7.3. **Projectivity.** A module P is called *projective* if for any morphisms $f : A \to B$ and $\varphi : P \to B$ such that $\operatorname{Im} \varphi \subseteq \operatorname{Im} f$, there exists a morphism $\varphi' : P \to A$ such that $f \circ \varphi' = \varphi$. This is the usual notion of projectivity in the category of modules over a fixed ring.

Following Gersten we call a b-morphism $f : A \to B$ *undistorted* if there is a constant $C \in [0, \infty)$ such that for any $b \in \operatorname{Im} f$ there exists $a \in A$ with $f(a) = b$ and $|a| \le C|b|$. A bounded module P is called *a projective bounded* $\mathbb{Q}\Gamma$*-module*, or a *b-projective* $\mathbb{Q}\Gamma$*-module*, if for any undistorted b-morphism $f : A \to B$ and any b-morphism $\varphi : P \to B$ such that $\operatorname{Im} \varphi \subseteq \operatorname{Im} f$, there exists a b-morphism $\varphi' : P \to A$ such that $f \circ \varphi' = \varphi$. This is projectivity in the category of bounded $\mathbb{Q}\Gamma$-modules.

Lemma 49. *If S is a Γ-set with finite stabilizers, then $\mathbb{Q}S$ is both a projective $\mathbb{Q}\Gamma$-module and a b-projective $\mathbb{Q}\Gamma$-module with respect to the ℓ^1-norm.*

Proof. The statements are easily checked when S is a *free* Γ-set, i.e. when the stabilizers of points are trivial. Now suppose that S is a Γ-set with finite stabilizers. Replacing points in S with finite sets one can produce a free Γ-set S' and a Γ-equivariant surjective map $h : S' \to S$. This induces a $\mathbb{Q}$-linear map $\mathbb{Q}S' \to \mathbb{Q}S$. A $\mathbb{Q}$-linear map $h' : \mathbb{Q}S \to \mathbb{Q}S'$ is defined by taking for each $x \in S$ the uniform distribution of 1 over the finite set $h^{-1}(x)$. h' is a right inverse of h. One uses these maps to deduce projectivity and b-projectivity of $\mathbb{Q}S$ from those of $\mathbb{Q}S'$. ☐

7.4. **Functors** B **and** bHom**.** If S is a Γ-set and V a normed $\mathbb{Q}$-vector space, $\ell^\infty(S, V)$ will denote the space of functions $S \to V$ that are bounded with respect to the norm on V. The norm on $\ell^\infty(S, V)$ is the ℓ^∞-norm

$$|f|_\infty := \sup\{\|f(s)\| \mid s \in S\}.$$

For normed $\mathbb{Q}$-vector spaces U and V, $\operatorname{bHom}(U, V)$ will denote the space of bounded $\mathbb{Q}$-linear maps $U \to V$. $\operatorname{bHom}(U, V)$ is a bounded $\mathbb{Q}\Gamma$-module with respect to the operator norm

$$\|f\| := \sup\{\|f(u)\|/\|u\| \mid u \in U\}$$

and the Γ-action given by

$$(gf)(u) := g(f(g^{-1}u)), \quad g \in \Gamma,\ f \in \operatorname{bHom}(U, V),\ u \in U.$$

Each element of $\ell^\infty(S, V)$ extends by linearity to a linear map $\mathbb{Q}S \to V$. This gives an isomorphism of normed modules

$$\ell^\infty(S, V) \cong \operatorname{bHom}(\mathbb{Q}S, V).$$

If U and V are bounded left $\mathbb{Q}\Gamma$-modules, $\operatorname{bHom}_{\mathbb{Q}\Gamma}(U, V)$ is the subspace of $\operatorname{bHom}(U, V)$ consisting of $\mathbb{Q}\Gamma$-morphisms, i.e. the operators that commute with the $\mathbb{Q}\Gamma$-actions on U and V.

7.5. **Cohomology of a pair** (Γ, Γ'). In [3] Bieri and Eckmann define $H^k(\Gamma, \Gamma'; V)$, the *cohomology of a pair* (Γ, Γ'), or the *relative cohomology* of Γ with respect to Γ', with coefficients in a $\mathbb{Z}\Gamma$-module V, and prove the following.

Proposition 50 ([3, Proposition 1.2]). *Let* (Γ, Γ') *be a pair as above,* $\mathbf{C}$ *a* Γ*-projective resolution of* $\mathbb{Z}$*,* $\mathbf{D}$ *a subcomplex of* $\mathbf{C}$ *which is a* Γ*-projective resolution of* $\mathbb{Z}\Gamma/\Gamma'$ *such that* $\mathbf{D} \hookrightarrow \mathbf{C}$ *induces* $\epsilon : \mathbb{Z}\Gamma/\Gamma' \hookrightarrow \mathbb{Z}$ *and that* $\mathbf{Q} := \mathbf{C}/\mathbf{D}$ *is* Γ*-projective. Then the cohomology sequences of* $\mathbf{C}$ *modulo* $\mathbf{D}$ *and of* Γ *modulo* Γ' *are isomorphic. More precisely, one has, for a* Γ*-module* V*, the following diagram which commutes up to sign.*

$$\begin{array}{ccccccccc} \dots \longrightarrow & H^k(\Gamma, \Gamma'; V) & \longrightarrow & H^k(\Gamma; V) & \longrightarrow & H^k(\Gamma'; V) & \longrightarrow & H^{k+1}(\Gamma, \Gamma'; V) & \longrightarrow \dots \\ & \cong\downarrow & & \cong\downarrow & & \cong\downarrow & & \cong\downarrow & \\ \dots \longrightarrow & H^k(\mathbf{Q}; V) & \longrightarrow & H^k(\mathbf{C}; V) & \longrightarrow & H^k(\mathbf{D}; V) & \longrightarrow & H^{k+1}(\mathbf{Q}; V) & \longrightarrow \dots \end{array}$$

We note that the same holds for $\mathbb{Q}\Gamma$ modules. Given $\mathbf{D}$, $\mathbf{C}$, and $\mathbf{Q}$ as above, we will take $H^k(\mathbf{Q}; V)$ as the *definition* of $H^k(\Gamma, \Gamma'; V)$, the *relative cohomology of* Γ *with respect to* Γ' with coefficients in a $\mathbb{Q}\Gamma$-module V.

8. The snake resolution.

Let Γ be any group and $\Gamma' = \{\Gamma_i \mid i \in \mathrm{I}\}$ be any nonempty family of its subgroups, possibly with repetitions. In this section we will define a short exact sequence of chain complexes $\mathbf{St}' \hookrightarrow \mathbf{St} \twoheadrightarrow \mathbf{St}^{\varsigma}$ which correspond to Γ', Γ, and Γ/Γ', respectively. The snake resolution $\mathbf{St}^{\varsigma}$ is a relative version of the standard (i.e. homogeneous bar) resolution.

8.1. **St: the standard resolution for** $\mathrm{I}\Gamma$. For $i \in \mathrm{I}$ and $k \geq 0$, let $S_k(\mathrm{I}\Gamma)$ be the set of sequences $[x_0, \dots, x_k]$ such that $x_j \in \mathrm{I}\Gamma$ for $j \in \{0, \dots, k\}$. The Γ-action on $S_k(\Gamma)$ is left-diagonal:

$$g[x_0, \dots, x_k] := [gx_0, \dots, gx_k]. \tag{1}$$

Denote

$$\mathrm{St}_k(\mathrm{I}\Gamma) := \mathbb{Q}S_k(\mathrm{I}\Gamma).$$

This gives the exact sequence

$$\mathbf{St} \twoheadrightarrow \mathbb{Q} : \qquad \dots \to \mathrm{St}_2(\mathrm{I}\Gamma) \xrightarrow{\partial_2} \mathrm{St}_1(\mathrm{I}\Gamma) \xrightarrow{\partial_1} \mathrm{St}_0(\mathrm{I}\Gamma) \xrightarrow{\partial_0'} \mathbb{Q} \to 0, \tag{2}$$

where $\mathrm{St}_k(\Gamma', \Gamma) \xrightarrow{\partial_k} \mathrm{St}_{k-1}(\Gamma', \Gamma)$ is the usual boundary map defined on the basis by

$$\partial_k[x_0, \dots, x_k] := \sum_{j=0}^{k} (-1)^j [x_0, \dots, \hat{x_j}, \dots, x_k].$$

In dimension 0 this formally means that ∂_0 is the augmentation map

$$\mathrm{St}_0(\Gamma) \cong \mathbb{Q}\mathrm{I}\Gamma \xrightarrow{\partial_0} \mathbb{Q}, \qquad f \mapsto \sum_{x \in \mathrm{I}\Gamma} f(x).$$

(2) and Lemma 49 show that

$$(3)\qquad \mathbf{St}: \qquad \ldots \to \mathrm{St}_2(\mathrm{I}\Gamma) \xrightarrow{\partial_2} \mathrm{St}_1(\mathrm{I}\Gamma) \xrightarrow{\partial_1} \mathrm{St}_0(\mathrm{I}\Gamma)$$

is a projective resolution of $\mathbb{Q}$.

Let

$$\mathrm{St}^k(\mathrm{I}\Gamma; V) := \ell^\infty(S_k(\mathrm{I}\Gamma), V) \cong \mathrm{bHom}(\mathrm{St}_k(\mathrm{I}\Gamma), V).$$

This is the space of bounded functions $\mathrm{I}\Gamma \to V$; it is a normed $\mathbb{Q}$-vector space with respect to the norm

$$\|f\| := \sup\{\|f(s)\| \mid s \in S_k(\mathrm{I}\Gamma)\}.$$

The elements of $\mathrm{St}^k(\mathrm{I}\Gamma; V)$ are the *(standard) cochains in* $\mathrm{I}\Gamma$ with coefficients in V.

Applying $\mathrm{bHom}(\cdot, V)$ to (2) yields the cochain complex

$$(4)\qquad \mathbf{St}(\mathrm{I}\Gamma; V) \hookleftarrow V: \qquad \ldots \to \mathrm{St}^2(\mathrm{I}\Gamma; V) \xleftarrow{\delta_2} \mathrm{St}^1(\mathrm{I}\Gamma; V) \xleftarrow{\delta_1} \mathrm{St}^0(\mathrm{I}\Gamma; V) \xhookleftarrow{\delta_0} V \leftarrow 0$$

where $\mathrm{St}^k(\mathrm{I}\Gamma; V) \xleftarrow{\delta_k} \mathrm{St}^{k-1}(\mathrm{I}\Gamma; V)$ is the usual coboundary map

$$(\delta_k f)[x_0, \ldots, x_k] := \sum_{j=0}^{k} (-1)^j f([x_0, \ldots, \hat{x_j}, \ldots, x_k]).$$

8.2. **$\mathbf{St}'$: the standard resolution for (Γ', Γ).** This is going to be the standard resolution of Γ' *with respect to* Γ (that is, we will use induction from Γ' to Γ, without mentioning it explicitly). In what follows, it is convenient to think of Γ_i as of a subgroup in $i\Gamma$.

For each $i \in I$ and $k \geq 0$, let $S'_{k,i}(\Gamma', \Gamma)$ be the set of sequences $[x_0, \ldots, x_k]$ such that

- $x_j \in i\Gamma$ for $j \in \{0, \ldots, k\}$, and
- $x_{j-1}^{-1} x_j \in \Gamma_i$ for $j \in \{1, \ldots, k\}$.

The above two conditions equivalently say that all x_j for $j \in \{0, \ldots, k\}$ belong to the same left coset of Γ_i in $i\Gamma$. We have $S'_{k,i}(\Gamma', \Gamma) \subseteq \mathrm{S}_k(\mathrm{I}\Gamma)$. Denote

$$S'_k(\Gamma', \Gamma) := \bigsqcup_i S'_{k,i}(\Gamma', \Gamma) \subseteq S_k(\mathrm{I}\Gamma).$$

$S'_0(\Gamma', \Gamma)$ obviously identifies with $\mathrm{I}\Gamma$. Let

$$\mathrm{St}'_k(\Gamma', \Gamma) := \mathbb{Q}S_k(\Gamma', \Gamma),$$

so in particular $\mathrm{St}'_0(\Gamma', \Gamma) \cong \mathbb{Q}\mathrm{I}\Gamma$. Elements of $\mathrm{St}'_k(\Gamma', \Gamma)$ are called *(standard) chains in* (Γ', Γ). The free Γ-action (1) on $\mathrm{St}_k(\mathrm{I}\Gamma)$ restricts to a free action on $\mathrm{St}'_k(\Gamma', \Gamma)$. This gives the exact sequence

$$(5)\qquad \mathbf{St}' \twoheadrightarrow \mathbb{Q}\Gamma/\Gamma': \qquad \ldots \to \mathrm{St}'_2(\Gamma', \Gamma) \xrightarrow{\partial'_2} \mathrm{St}'_1(\Gamma', \Gamma) \xrightarrow{\partial'_1} \mathrm{St}'_0(\Gamma', \Gamma) \xrightarrow{\partial'_0} \mathbb{Q}\Gamma/\Gamma' \to 0,$$

where $\mathrm{St}'_0(\Gamma', \Gamma) \cong \mathbb{Q}\mathrm{I}\Gamma \xrightarrow{\partial'_0} \mathbb{Q}\Gamma/\Gamma'$ is the augmentation map induced by the surjection

$$\mathrm{I}\Gamma = \bigsqcup_{i \in I} \Gamma \twoheadrightarrow \bigsqcup_{i \in \mathrm{I}} \Gamma/\Gamma_i = \Gamma/\Gamma'$$

and ∂_k is the restriction of the boundary homomorphism in (2). (5) and Lemma 49 show that

(6) $$\mathbf{St}': \qquad \ldots \to \mathrm{St}'_2(\Gamma',\Gamma) \stackrel{\partial'_2}{\to} \mathrm{St}'_1(\Gamma',\Gamma) \stackrel{\partial'_1}{\to} \mathrm{St}'_0(\Gamma',\Gamma)$$

is a projective resolution of $\mathbb{Q}\Gamma/\Gamma'$.

Let

$$\mathrm{St}'^{k}(\Gamma',\Gamma;V) := \mathrm{B}(S'_k(\Gamma',\Gamma),V) \cong \mathrm{Hom}(\mathrm{St}'_k(\Gamma',\Gamma),V).$$

This is the space of bounded functions $S'_k(\Gamma',\Gamma) \to V$; it is a normed $\mathbb{Q}$-vector space with respect to the norm

$$\|f\| := \sup\{\|f(s)\| \mid s \in S_k(\Gamma',\Gamma)\}.$$

The elements of $\mathrm{St}'^{k}(\Gamma',\Gamma;V)$ are the *standard cochains in* (Γ',Γ) with coefficients in V.

$$\mathrm{St}'^{k}(\Gamma',\Gamma;V)^{\Gamma} := \mathrm{B}_{\Gamma}(S'_k(\Gamma',\Gamma),V) \cong \mathrm{Hom}_{\mathbb{Q}\Gamma}(\mathrm{St}'_k(\Gamma',\Gamma),V)$$

is the space of *invariant* standard cochains in (Γ',Γ).

8.3. **$\mathbf{St}^{\varsigma}$: the snake resolution.** This is the standard resolution for the pair (Γ,Γ'). For resolutions $\mathbf{St}$ in (3) and $\mathbf{St}'$ in (6) we denote $\mathrm{St}^{\varsigma}_{\mathrm{k}}(\Gamma,\Gamma') := \mathrm{St}_k(I\Gamma)/\mathrm{St}'_{\mathrm{k}}(\Gamma',\Gamma)$ for $k \geq 1$; sometimes we will write $\mathrm{St}'_{\mathrm{k}}$, St_k, $\mathrm{St}^{\varsigma}_{\mathrm{k}}$ for simplicity. Then $\mathbf{St}$ and $\mathbf{St}'$ fit in the diagram

(7) $$\begin{array}{llccccccc}
\mathbf{St}': & \ldots & \xrightarrow{\partial'_3} & \mathrm{St}'_2(\Gamma',\Gamma) & \xrightarrow{\partial'_2} & \mathrm{St}'_1(\Gamma',\Gamma) & \xrightarrow{\partial'_1} & \mathrm{St}'_0(\Gamma',\Gamma) \\
 & & & \downarrow & & \downarrow & & \cong\downarrow \\
\mathbf{St}: & \ldots & \xrightarrow{\partial_3} & \mathrm{St}_2(I\Gamma) & \xrightarrow{\partial_2} & \mathrm{St}_1(I\Gamma) & \xrightarrow{\partial_1} & \mathrm{St}_0(I\Gamma) \\
 & & & \downarrow & & \downarrow & & \downarrow \\
 & \ldots & \xrightarrow{\partial^{\varsigma}_3} & \mathrm{St}^{\varsigma}_2(\Gamma,\Gamma') & \xrightarrow{\partial^{\varsigma}_2} & \mathrm{St}^{\varsigma}_1(\Gamma,\Gamma') & \longrightarrow & 0.
\end{array}$$

This extends to the larger diagram

(8) $$\begin{array}{lcccccccc}
\ldots & \xrightarrow{\partial'_3} & \mathrm{St}'_2(\Gamma',\Gamma) & \xrightarrow{\partial'_2} & \mathrm{St}'_1(\Gamma',\Gamma) & \xrightarrow{\partial'_1} & \mathrm{St}'_0(\Gamma',\Gamma) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \cong\downarrow & & \downarrow \\
\ldots & \xrightarrow{\partial_3} & \mathrm{St}_2(I\Gamma) & \xrightarrow{\partial_2} & \mathrm{St}_1(I\Gamma) & \xrightarrow{\partial_1} & \mathrm{St}_0(I\Gamma) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
\ldots & \xrightarrow{\partial^{\varsigma}_3} & \mathrm{St}^{\varsigma}_2(\Gamma,\Gamma') & \xrightarrow{\partial^{\varsigma}_2} & \mathrm{St}^{\varsigma}_1(\Gamma,\Gamma') & \longrightarrow & 0. & &
\end{array}$$

(5) and (2) imply that the first two rows in the above diagram are exact in dimensions $k \geq 1$ and induce the augmentation map $\mathbb{Q}\Gamma/\Gamma' \to \mathbb{Q}$ in dimension 0. Recall that Δ was defined to be the kernel of this map. The $\mathbb{Q}\Gamma$-modules $\mathrm{St}^{\varsigma}_k(\Gamma,\Gamma')$ are free, so the long exact sequence for (8) implies that the last row

(9) $$\mathbf{St}^{\varsigma}: \qquad \ldots \stackrel{\partial^{\varsigma}_3}{\to} \mathrm{St}^{\varsigma}_3(\Gamma,\Gamma') \stackrel{\partial^{\varsigma}_3}{\to} \mathrm{St}^{\varsigma}_2(\Gamma,\Gamma') \stackrel{\partial^{\varsigma}_2}{\to} \mathrm{St}^{\varsigma}_1(\Gamma,\Gamma')$$

is a resolution of Δ with a dimension shift, i.e. it extends to the exact sequence

$$\mathbf{St}^{\varsigma} \twoheadrightarrow \Delta : \qquad \ldots \xrightarrow{\partial_3^{\varsigma}} \mathrm{St}_3^{\varsigma}(\Gamma, \Gamma') \xrightarrow{\partial_3^{\varsigma}} \mathrm{St}_2^{\varsigma}(\Gamma, \Gamma') \xrightarrow{\partial_2^{\varsigma}} \mathrm{St}_1^{\varsigma}(\Gamma, \Gamma') \xrightarrow{\partial_1^{\varsigma}} \Delta \to 0. \tag{10}$$

Then (7) is a short exact sequence of chain complexes that satisfies the assumptions of Proposition 50, therefore resolution (9) can be used to define the relative cohomology $H^*(\Gamma, \Gamma'; V)$. That is, if we apply $\mathrm{Hom}_{\mathbb{Q}\Gamma}(\cdot, V)$ to $\mathbf{St}^{\varsigma}$ and denote

$$\mathrm{St}_{\varsigma}^{k}(\Gamma, \Gamma'; V) := \mathrm{Hom}_{\mathbb{Q}\Gamma}(\mathrm{St}_k^{\varsigma}(\Gamma, \Gamma'), V),$$

then $\mathrm{H^k}(\Gamma, \Gamma'; V)$ is the homology of the resulting cochain complex

$$\mathbf{St}_{\varsigma} \leftarrow 0 : \qquad \ldots \xleftarrow{\delta_{\varsigma}^{3}} \mathrm{St}_{\varsigma}^{2}(\Gamma, \Gamma'; V) \xleftarrow{\delta_{\varsigma}^{2}} \mathrm{St}_{\varsigma}^{1}(\Gamma, \Gamma'; V) \leftarrow 0 \tag{11}$$

at the term $\mathrm{St}_{\varsigma}^{k}(\Gamma, \Gamma'; V)$. Here δ_k is dual to ∂_k.

8.4. **Relative bounded cohomology.** For the *relative bounded cohomology* of Γ with respect to Γ' we use a parallel definition, applying bHom instead of Hom. Denote

$$\mathrm{bSt}_{\varsigma}^{k}(\Gamma, \Gamma'; V) := \mathrm{bHom}_{\mathbb{Q}\Gamma}(\mathrm{St}_k^{\varsigma}(\Gamma, \Gamma'), V)$$

and let $H_b^k(\Gamma, \Gamma'; V)$ be the homology of the cochain complex

$$\mathbf{bSt}_{\varsigma} \leftarrow 0 : \qquad \ldots \xleftarrow{\delta_3^{\varsigma}} \mathrm{bSt}_{\varsigma}^{2}(\Gamma, \Gamma'; V) \xleftarrow{\delta_2^{\varsigma}} \mathrm{bSt}_{\varsigma}^{1}(\Gamma, \Gamma'; V) \leftarrow 0 \tag{12}$$

at the term $\mathrm{St}^k(\Gamma, \Gamma'; V)$. There is a canonical map $H_b^k(\Gamma, \Gamma'; V) \to H^k(\Gamma, \Gamma'; V)$ induced by the inclusion $\mathrm{bHom}(\cdot, V) \subseteq \mathrm{Hom}(\cdot, V)$.

8.5. **A resolution for tuples.** Suppose (Γ, Γ') is a hyperbolic pair. Let X be an ideal complex for the pair (Γ, Γ') as in Definition 11; in particular we have $\mathcal{V}' = X^{(0)}$. One can equally well work with hyperbolic tuples of type $\mathcal{F}_n$; the statements below will hold up to dimension n.

$C_k(X)$ will denote the space of k-chains in X with rational coefficients. It follows from Definition 24 of graph tuples that there is a bijection between $\mathcal{V}'$ and Γ/Γ'. $\mathcal{V}'$ is a subcomplex of X, this gives the obvious chain complex

$$\mathbf{C}' \twoheadrightarrow \mathbb{Q}\Gamma/\Gamma' : \qquad \ldots \to 0 \to 0 \to C_0(\mathcal{V}') \xrightarrow{\cong} \mathbb{Q}\Gamma/\Gamma' \to 0. \tag{13}$$

Let $\partial_0 : C_0(X) \to \mathbb{Q}$ be the augmentation map. Since X is simply connected, the sequence

$$\mathbf{C} \twoheadrightarrow \mathbb{Q} : \qquad C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\partial_0} \mathbb{Q} \to 0 \tag{14}$$

is exact in dimensions ≤ 1. By Lemma 49, $\mathbf{C}$ is a partial projective $\mathbb{Q}\Gamma$-resolution of $\mathbb{Q}$.

The inclusion $\mathcal{V}' \subseteq X$ gives a chain map $\mathbf{C}' \to \mathbf{C}$ which induces the augmentation map $\mathbb{Q}\Gamma/\Gamma' \to \mathbb{Q}$ in dimension -1. The following diagram is $\mathbf{C}' \hookrightarrow \mathbf{C} \twoheadrightarrow \mathbf{C}/\mathbf{C}'$:

$$
\begin{array}{llccccccc}
& \mathbf{C}' : & \ldots \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & C_0(\mathcal{V}') \\
& & & \downarrow & & \downarrow & & \cong\downarrow \\
(15) & \mathbf{C} : & \ldots \xrightarrow{\partial_3} & C_2(X) & \xrightarrow{\partial_2} & C_1(X) & \xrightarrow{\partial_1} & C_0(X) \\
& & & \cong\downarrow & & \cong\downarrow & & \downarrow \\
& & \ldots \xrightarrow{\partial_3^\varsigma} & C_2(X) & \xrightarrow{\partial_2^\varsigma} & C_1(X) & \longrightarrow & 0.
\end{array}
$$

Denote $\mathbf{C}^\varsigma$ the bottom row of (15) in dimensions ≥ 1, i.e. the sequence

$$
(16) \qquad \mathbf{C}^\varsigma : \qquad \ldots \overset{\partial_3^\varsigma}{\to} C_2(X) \overset{\partial_2^\varsigma}{\to} C_1(X).
$$

The diagram (15) extends to

$$
\begin{array}{lcccccccc}
& \ldots \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & C_0(\mathcal{V}') & \longrightarrow & 0 \\
& & \downarrow & & \downarrow & & \cong\downarrow & & \\
(17) & \ldots \xrightarrow{\partial_3} & C_2(X) & \xrightarrow{\partial_2} & C_1(X) & \xrightarrow{\partial_1} & C_0(X) & \longrightarrow & 0 \\
& & \cong\downarrow & & \cong\downarrow & & \downarrow & & \\
& \ldots \xrightarrow{\partial_3^\varsigma} & C_2(X) & \xrightarrow{\partial_2^\varsigma} & C_1(X) & \longrightarrow & 0. & &
\end{array}
$$

The first two rows in (17) induce the augmentation $\varepsilon : \mathbb{Q}\Gamma/\Gamma' \twoheadrightarrow \mathbb{Q}$ in dimension 0. By the long exact sequence, the bottom row is exact in dimensions ≥ 1 and induces Δ in dimension 1, hence there is an exact sequence

$$
(18) \qquad \mathbf{C}^\varsigma \twoheadrightarrow \Delta : \qquad \ldots \overset{\partial_3^\varsigma}{\to} C_2(X) \overset{\partial_2^\varsigma}{\to} C_1(X) \overset{\partial_1^\varsigma}{\to} \Delta \to 0.
$$

This is a partial resolution of Δ with a dimension shift. It can also be thought of simply as a shorter version of (14), because Δ is isomorphic to the kernel of ∂_0. Lemma 49 says that this resolution is both projective and b-projective.

8.6. **Chain maps between resolutions.** We have two chain complexes for the pair (Γ, Γ'): (9) and (16) are resolutions of Δ with a dimension shift. We will define chain maps between them.

Since $\mathbf{C}^\varsigma$ is projective, there exists a chain map ψ_* for the diagram

$$
\begin{array}{llcccccccc}
& \mathbf{St}^\varsigma \twoheadrightarrow \Delta : & \ldots \xrightarrow{\partial_3^\varsigma} & \mathrm{St}_2^\varsigma(\Gamma,\Gamma') & \xrightarrow{\partial_2^\varsigma} & \mathrm{St}_1^\varsigma(\Gamma,\Gamma') & \xrightarrow{\partial_1^\varsigma} & \Delta & \longrightarrow & 0 \\
(19) & & & \psi_2 \uparrow & & \psi_1 \uparrow & & \Vert & & \\
& \mathbf{C}^\varsigma \twoheadrightarrow \Delta : & \ldots \xrightarrow{\partial_3^\varsigma} & C_2(X) & \xrightarrow{\partial_2^\varsigma} & C_1(X) & \xrightarrow{\partial_1^\varsigma} & \Delta & \longrightarrow & 0.
\end{array}
$$

Each ψ_k is a $\mathbb{Q}\Gamma$-morphism and is bounded because $C_k(X)$ is a finitely generated $\mathbb{Q}\Gamma$-module.

A chain map φ_* for the diagram

$$
\begin{array}{llccccccccc}
\mathbf{St}^{\varsigma} \twoheadrightarrow \Delta : & \dots & \xrightarrow{\partial_3^{\varsigma}} & \mathrm{St}_2^{\varsigma}(\Gamma, \Gamma') & \xrightarrow{\partial_2^{\varsigma}} & \mathrm{St}_1^{\varsigma}(\Gamma, \Gamma') & \xrightarrow{\partial_1^{\varsigma}} & \Delta & \longrightarrow & 0 \\
 & & & \Big\downarrow{\scriptstyle \varphi_2} & & \Big\downarrow{\scriptstyle \varphi_1} & & \Big\| & & \\
\mathbf{C}^{\varsigma} \twoheadrightarrow \Delta : & \dots & \xrightarrow{\partial_3^{\varsigma}} & C_2(X) & \xrightarrow{\partial_2^{\varsigma}} & C_1(X) & \xrightarrow{\partial_1^{\varsigma}} & \Delta & \longrightarrow & 0
\end{array} \tag{20}
$$

is defined as follows. Let for each $i \in \mathrm{I}$, let v_i be the vertex stabilized by Γ_i. Define a map $\mathrm{ver} : \mathrm{I}\Gamma \to \mathcal{V}'$ by $\mathrm{ver}(x) := x \cdot v_i$ for $x \in i\Gamma$. This map sends each left coset of Γ_i in $i\Gamma$ to a vertex in $\mathcal{V}'$, and is surjective. Preimages of the vertices in $\mathcal{V}'$ are exactly the left cosets, i.e. the elements of Γ/Γ'. Extending by linearity gives the map $\mathrm{ver} : \mathbb{Q}\mathrm{I}\Gamma \to \mathbb{Q}\mathcal{V}'$.

Define $\varphi_1' : \mathrm{St}_1(\mathrm{I}\Gamma) \to C_1(X)$ by $\varphi_1'([x_0, x_1]) := q[\mathrm{ver}(x_0), \mathrm{ver}(x_1)]$, where q is the bicombing from Theorem 44, and extending by linearity. One checks that φ_1' vanishes on $\mathrm{St}_1'(\Gamma', \Gamma) \subseteq \mathrm{St}_1(\mathrm{I}\Gamma)$, so it induces a morphism $\varphi_1 : \mathrm{St}_1^{\varsigma}(\Gamma, \Gamma') \to C_1(X)$. This defines the map φ_* in dimension 1.

The morphism φ_1 is *not* bounded in general, but the composition $\varphi_1 \circ \partial_2^{\varsigma}$ is, by Theorem 44(4). By Theorem 46, $\partial_2^{\varsigma} = \partial_2 : C_2(X) \to C_1(X)$ is undistorted, and by Lemma 49, $\mathrm{St}_2^{\varsigma}(\Gamma, \Gamma')$ is b-projective, so there exists a *bounded* morphism φ_2 that makes the above diagram commutative. For higher dimensions, $\partial_k^{\varsigma} = \partial_k : C_k(X) \to C_{k-1}(X)$ is undistorted by Theorem 48, so the same argument gives a bounded $\mathbb{Q}\Gamma$-morphism φ_k for $k \geq 2$.

Since $\mathbf{St}^{\varsigma}$ and $\mathbf{C}^{\varsigma}$ are projective, the chain maps

$$\psi_* : (\mathbf{C}^{\varsigma} \twoheadrightarrow \Delta) \to (\mathbf{St}^{\varsigma} \twoheadrightarrow \Delta) \qquad \text{and} \qquad \varphi_* : (\mathbf{St}^{\varsigma} \twoheadrightarrow \Delta) \to (\mathbf{C}^{\varsigma} \twoheadrightarrow \Delta)$$

described above are chain homotopy equivalences. Applying $\mathrm{Hom}_{\mathbb{Q}\Gamma}(\cdot, V)$ to the two chain complexes yields dual cochain complexes C_{ς}^* and $\mathrm{St}_{\varsigma}^*$ and the dual chain maps $\varphi^* : C_{\varsigma}^* \to \mathrm{St}_{\varsigma}^*$ and $\psi^* : \mathrm{St}_{\varsigma}^* \to C_{\varsigma}^*$. Since φ_* and ψ_* are mutually inverse chain homotopy equivalences, the chain map $\varphi^* \circ \psi^*$ induces the identity map on the relative cohomology $H^k(\Gamma, \Gamma'; V)$ for $k \geq 2$.

8.7. **Filling 0-cycles.** A *0-cycle* is a 0-chain whose sum of coefficients is 0. Any $c \in \mathrm{St}_0(\mathrm{I}\Gamma)$ can be uniquely written as $c = c_+ - c_-$, where $c_+, c_- \in \mathrm{St}_0(\mathrm{I}\Gamma)$ have non-negative coefficients and mutually disjoint supports. If c is a 0-cycle, then $|c_+|_1 = |c_-|_1 = \frac{1}{2}|c|_1$. If c_+ and c_- are explicitly written as $c_- = \sum_{x \in \mathrm{I}\Gamma} \alpha_x^- x$, $c_+ = \sum_{y \in \mathrm{I}\Gamma} \alpha_y^+ y$, define

$$\Phi[c] := \frac{1}{|c_+|_1} \sum_{x \in \mathrm{I}\Gamma} \sum_{y \in \mathrm{I}\Gamma} \alpha_x^- \alpha_y^+ [x, y] = \frac{1}{|c_-|_1} \sum_{x \in \mathrm{I}\Gamma} \sum_{y \in \mathrm{I}\Gamma} \alpha_x^- \alpha_y^+ [x, y].$$

One checks that

$$\partial(\Phi[c]) = c \qquad \text{and} \qquad \big|\Phi[c]\big|_1 = \frac{1}{2}|c|_1, \tag{21}$$

so $\Phi(c)$ is a filling of c.

8.8. **The cone.** A *k-cycle* in **St** is an element of $\mathrm{ZSt}_k(\mathrm{I}\Gamma) := Ker(\partial_k : \mathrm{St}_k(\mathrm{I}\Gamma) \to \mathrm{St}_{k-1}(\mathrm{I}\Gamma))$, and a *k-boundary* in **St** is an element of $\mathrm{BSt}_k(\mathrm{I}\Gamma) := Im(\partial_{k+1} : \mathrm{St}_{k+1}(\mathrm{I}\Gamma) \to \mathrm{St}_k(\mathrm{I}\Gamma))$. We have $\mathrm{ZSt}_k(\mathrm{I}\Gamma) = \mathrm{BSt}_k(\mathrm{I}\Gamma)$, i.e. the two notions coincide.

For each 1-chain $b = \sum_{y_0,y_1\in G} \beta_{[y_0,y_1]}[y_0, y_1]$ in **St**, *the cone* over b with vertex y is the 2-chain

$$[y, b] := \sum_{y_0,y_1\in G} \beta_{[y_0,y_1]}[y, y_0, y_1].$$

If b happens to be a cycle, then $\partial[y, b] = b$.

8.9. **The relative cone.** A *k-cycle* in $\mathbf{St}^\varsigma$, or a *relative standard k-cycle*, is an element of

$$\mathrm{ZSt}^\varsigma_\mathrm{k}(\Gamma, \Gamma') := Ker\big(\partial^\varsigma_k : \mathrm{St}^\varsigma_\mathrm{k}(\Gamma, \Gamma') \to \mathrm{St}^\varsigma_{\mathrm{k}-1}(\Gamma, \Gamma')\big),$$

and a *k-boundary* in $\mathbf{St}^\varsigma$, or a *relative standard k-boundary*, is an element of

$$\mathrm{BSt}^\varsigma_\mathrm{k}(\Gamma, \Gamma') := Im\big(\partial^\varsigma_{k+1} : \mathrm{St}^\varsigma_{\mathrm{k}+1}(\Gamma, \Gamma') \to \mathrm{St}^\varsigma_\mathrm{k}(\Gamma, \Gamma')\big).$$

Here we denote $\mathrm{St}^\varsigma_0(\Gamma, \Gamma') = \Delta$. We have $\mathrm{ZSt}^\varsigma_\mathrm{k}(\Gamma, \Gamma') = \mathrm{BSt}^\varsigma_\mathrm{k}(\Gamma, \Gamma')$, i.e. the two notions coincide.

Since St'_k is a direct summand of St_k, there is an $\mathbb{Q}\Gamma$-morphism $j : \mathrm{St}^\varsigma_\mathrm{k} \to \mathrm{St}_\mathrm{k}$ which is a section of the projection morphism $pr : \mathrm{St}_\mathrm{k} \to \mathrm{St}^\varsigma_\mathrm{k}$. For a 1-chain a in **St** and a left coset $s \in \Gamma/\Gamma'$, $\partial^s a$ will denote the restriction of ∂a to $s \subseteq \mathrm{I}\Gamma$.

For a relative standard 1-cycle b, the *relative cone* of b with vertex y is the relative 2-chain

$$[y, b]_\varsigma := pr\big[y, b - \sum_{s\in\Gamma/\Gamma'} \Phi[\partial^s(j(b))]\big].$$

One checks that this definition makes sense and, using (21), that

$$\partial^\varsigma[y, b]_\varsigma = b \in \mathrm{St}^\varsigma_1 \qquad \text{and} \qquad \big|[y, b]_\varsigma\big|_1 \le 2|b|_1. \tag{22}$$

Also, if α is a relative 2-cocycle and c is a relative 2-chain, then $c-[y, \partial^\varsigma c]_\varsigma$ is a relative cycle, hence a relative boundary, so $\big\langle\alpha, c - [y, \partial^\varsigma c]_\varsigma\big\rangle = 0$ and

$$\big\langle\alpha, c\big\rangle = \big\langle\alpha, [y, \partial^\varsigma c]_\varsigma\big\rangle. \tag{23}$$

9. The cohomological characterization of relative hyperbolicity.

Denote $C^{(1)}_i(X)$ the space of ℓ^1-summable i-chains in X, with coefficients in either $\mathbb{Z}$, $\mathbb{Q}$,$\mathbb{R}$, or $\mathbb{C}$. Let $B^{(1)}_1(X)$ be the image of the boundary map $\partial_2 : C^{(1)}_2(X) \to C^{(1)}_1(X)$, with the filling norm $|\cdot|_f$ induced from the norm on $C^{(1)}_2(X)$:

$$|b|_f := \inf\{|a|_1 \,|\, a \in C^{(1)}_2(X),\ \partial a = b\}. \tag{24}$$

Theorem 51. *Let X be a simply connected combinatorial complex with a uniform bound on the number of boundary edges in 2-cells. The following statements are equivalent.*

(1) *X has thin triangles.*

(2) *There is $K \ge 0$ such that for any 1-cycle b in X over $\mathbb{Q}$, $|b|_f \le K|b|_1$.*

The same holds for cycles over $\mathbb{Z}$, $\mathbb{R}$, and $\mathbb{C}$.

Proof. (1) $\Rightarrow$ (2) was proved in [18, Theorem 7] for the case when the complex X is simply connected and admits a cocompact action by a finitely presented group. The argument is mostly due to Gersten. The same argument applies under the above assumptions for X.

For (2) $\Rightarrow$ (1), assume that (1) does not hold and consider the following lemma due to Ol'shanskii.

Lemma 52 ([22, Lemma 3]). *Suppose the bisizes of triangles in a geodesic space Y are unbounded. Then for any $t_0 > 0$ there exists a hexagon in Y with thickness $t > t_0$ and perimeter at most $46t$.*

The assumption of having unbounded bisizes in a geodesic metric space is equivalent to non-hyperbolicity of the space, i.e. to not having thin triangles [22], so this assumption is satisfied for our $\mathcal{G} := X^{(1)}$. Then the conclusion of the lemma means that there exist

- a sequence of numbers t tending to ∞,
- a geodesic hexagon $w = w(t)$ in $\mathcal{G}$ for each t,
- a (geodesic) side γ in each w, and
- a vertex $p \in \gamma$,

such that

- $d(e, \gamma') \geq t$, where γ' denotes the union of the sides in w other than γ, and
- the perimeter of w, $l(w)$, is at most $46t$.

This allows running the argument similar to [18, Proposition 8], using hexagons instead of quadrilaterals, to show that (2) does not hold. The idea is to take any filling of w and slice it by concentric spheres at p; then show that the sum of the areas of the slices grows quadratically in t. This method was originally used by Ol'shanskii' to show that groups with subquadratic (combinatorial) isoperimetric functions are hyperbolic. [18, Proposition 8] proves a homological version of that statement, using chains instead of van Kampen diagrams. □

Recall the Definition 26 of finitely presented tuples and Definition 37 of hyperbolic tuples.

Theorem 53. *Let $(\Gamma, \Gamma', X, \mathcal{V}')$ be a finitely presented tuple such that X admits a (combinatorial) isoperimetric function in the sense of Definition 31. Suppose that the map $H^2_b(\Gamma, \Gamma'; V) \to H^2(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V. Then there is $K \geq 0$ such that for any 1-cycle b in X over $\mathbb{Q}$, $|b|_f \leq K|b|_1$. The same is true when V is in the class of bounded $\mathbb{R}\Gamma$-modules, bounded $\mathbb{C}\Gamma$-modules, or Banach modules.*

Proof. We adapt the argument in [18, p.70-72] to the relative case. It suffices to prove the statement for the smallest class, that is the class of Banach modules; those are Banach spaces over $\mathbb{R}$ with a linear Γ-action such that the operator norms of the elements of Γ are uniformly bounded.

Let $V := B_1^{(1)}(X)$. This is a Banach space with respect to the filling norm (24).

The chain maps φ_* and ψ_* defined in 8.6 are in the category of $\mathbb{Q}\Gamma$-modules, so for each k, φ_k and ψ_k are linear maps commuting with the Γ-action. Denote for simplicity $C_k^\varsigma := C_k(X, \mathbb{Q})$ and $\mathrm{St}_k^\varsigma := \mathrm{St}_k^\varsigma(\Gamma, \Gamma'; \mathbb{Q})$. Consider the dual cochain complexes

$$C_\varsigma^k := C_\varsigma^k(X, V) := \mathrm{Hom}_{\mathbb{Q}\Gamma}(C_k^\varsigma, V) \quad \text{and} \quad \mathrm{St}_\varsigma^k := \mathrm{St}_\varsigma^k(\Gamma, \Gamma'; V) := \mathrm{Hom}_{\mathbb{Q}\Gamma}(\mathrm{St}_\varsigma^k, V)$$

with the coboundary maps denoted δ_ς and the dual maps

$$\varphi^* : C^*_\varsigma \leftarrow \mathrm{St}^*_\varsigma \quad \text{and} \quad \psi^* : \mathrm{St}^*_\varsigma \leftarrow C^*_\varsigma.$$

The cochain map $\psi^* \circ \varphi^*$ is homotopic to the identity map, hence $\psi^* \circ \varphi^*$ induces the identity map on cohomology $H^*(G, V)$ in dimensions ≥ 2.

The *universal cocycle* $u \in C^2_\varsigma$ is the 2-cochain $u : C^\varsigma_2 \to V$ which coincides with the composition

$$C_2(X, \mathbb{Q}) \xrightarrow{\partial^\varsigma_2} B_1(X, \mathbb{Q}) \hookrightarrow B^{(1)}_1(X, \mathbb{Q}).$$

One checks that u is indeed a cocycle. By the above observations,

$$u = (\psi^2 \circ \varphi^2)(u) + \delta_\varsigma v \tag{25}$$

for some 1-cochain $v : C^\varsigma_1 \to V$.

Since $\varphi^2(u)$ is a cocycle in St^2_ς and the map $H^2_b(\Gamma, \Gamma; V) \to H^2(\Gamma, \Gamma'; V)$ is surjective by the assumption,

$$\varphi^2(u) = u' + \delta_\varsigma v', \tag{26}$$

for some 1-cochain $v' \in \mathrm{St}^1_\varsigma$ and a *bounded* 2-cocycle $u' \in \mathrm{St}^2_\varsigma$, i.e.

$$|u'|_\infty < \infty. \tag{27}$$

The above information is demonstrated by the diagrams

$$\begin{array}{ccc} \mathrm{St}^\varsigma_* & = & \mathrm{St}^\varsigma_*(\Gamma, \Gamma'; \mathbb{Q}) \\ {\scriptstyle\psi_*}\big\uparrow\big\downarrow{\scriptstyle\varphi_*} & & \\ a, b \in \mathrm{C}^\varsigma_* & = & C_*(X, \mathbb{Q}) \end{array} \quad \text{and} \quad \begin{array}{ccc} u', v' \in \mathrm{St}^*_\varsigma & = & \mathrm{St}^*_\varsigma(\Gamma, \Gamma'; V) \\ {\scriptstyle\psi^*}\big\downarrow\big\uparrow{\scriptstyle\varphi^*} & & \\ u, v \in \mathrm{C}^\varsigma_\varsigma & = & C^*(X, V). \end{array}$$

Let $\langle \cdot, \cdot \rangle : C^k(X, V) \oplus C_k(X, \mathbb{Q}) \to V$ and $\langle \cdot, \cdot \rangle : \mathrm{St}^k_\varsigma(\Gamma, \Gamma'; V) \oplus \mathrm{St}^\varsigma_k(\Gamma, \Gamma'; \mathbb{Q}) \to V$ be the standard pairings.

Pick any 1-boundary $b \in B_1(X, \mathbb{Q})$ and any 2-chain a with $\partial a = b$. The goal is to show that $|b|_f \leq K|b|_1$ for some uniform constant K.

By (25),

$$b = \partial a = \langle u, a \rangle = \langle (\psi^2 \circ \varphi^2)(u) + \delta v, a \rangle = \langle (\psi^2 \circ \varphi^2)(u), a \rangle + \langle v, b \rangle.$$

Pick any $y \in \mathrm{I}\Gamma$. Since $\varphi^2(u)$ is a cocycle, using (22), (23) and (26),

$$\begin{aligned} &\langle (\psi^2 \circ \varphi^2)(u), a \rangle = \langle \varphi^2(u), \psi_2(a) \rangle = \langle \varphi^2(u), [y, \partial^\varsigma(\psi_2(a))]_\varsigma \rangle = \\ &= \langle \varphi^2(u), [y, \psi_1(b)]_\varsigma \rangle = \langle u' + \delta_\varsigma v', [y, \psi_1(b)]_\varsigma \rangle = \langle u', [y, \psi_1(b)]_\varsigma \rangle + \langle v', \partial^\varsigma [y, \psi_1(b)]_\varsigma \rangle = \\ &= \langle u', [y, \psi_1(b)]_\varsigma \rangle + \langle v', \psi_1(b) \rangle = \langle u', [y, \psi_1(b)]_\varsigma \rangle + \langle \psi^1(v'), b \rangle. \end{aligned}$$

Combining the two formulas above with (22),

$$\begin{aligned}
b = {} & \langle u', [y, \psi_1(b)]_\varsigma \rangle + \langle \psi^1(v') + v, b \rangle, \\
|b|_f \leq {} & \Big| \langle u', [y, \psi_1(b)]_\varsigma \rangle \Big|_f + \Big| \langle \psi^1(v') + v, b \rangle \Big|_f \leq \\
& \leq |u'|_\infty \cdot \Big| [y, \psi_1(b)]_\varsigma \Big|_1 + \Big| \psi^1(v') + v \Big|_\infty \cdot |b|_1 = \\
& = 2|u'|_\infty \cdot \big| \psi_1(b) \big|_1 + \Big| \psi^1(v') + v \Big|_\infty \cdot |b|_1 \leq \\
& \leq \Big(2|u'|_\infty \cdot |\psi_1|_\infty + \Big| \psi^1(v') + v \Big|_\infty \Big) \cdot |b|_1.
\end{aligned}$$

This will give the desired inequality once we prove that all the norms in the parentheses are finite. The cochain u' is bounded by definition (by a constant depending only on the choice of Γ and X, see (27)). The maps $\psi_1 : C_1^\varsigma \to \mathrm{St}_1^\varsigma$ and $\psi^1(v') + v : C_1^\varsigma \to V$ are $\mathbb{Q}\Gamma$-morphisms. Their boundedness (by constants depending only on Γ and X) is immediate from the following simple lemma which is proved similarly to [18, Lemma 10].

Lemma 54. *Let S be a Γ-set with finitely many Γ-orbits. Suppose V is a bounded $\mathbb{Q}\Gamma$-module and $f : \mathbb{Q}S \to V$ is a $\mathbb{Q}\Gamma$-morphism. Then f is bounded with respect to the ℓ^1-norm on $\mathbb{Q}S$, i.e. $|f|_\infty < \infty$.*

This finishes the proof of Theorem 53. □

Theorem 55. *Let Γ be a group and Γ' be a family of its subgroups. The following statements are equivalent.*

- (a) (Γ, Γ') *is hyperbolic.*
- (b) *There exists a finitely presented tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ such that X admits a (combinatorial) isoperimetric function (for edge-loops), and the map $H_b^2(\Gamma, \Gamma'; V) \to H^2(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V.*
- (b′) *There exists a finitely presented tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ such that X admits a (combinatorial) isoperimetric function (for edge-loops), and the map $H_b^n(\Gamma, \Gamma'; V) \to H^n(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V and all $n \geq 2$.*
- (c) *There exists a finitely presented tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ such that $X^{(1)}$ is fine and the map $H_b^2(\Gamma, \Gamma'; V) \to H^2(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V.*
- (c′) *There exists a finitely presented tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ such that $X^{(1)}$ is fine and the map $H_b^n(\Gamma, \Gamma'; V) \to H^n(\Gamma, \Gamma'; V)$ is surjective for all bounded $\mathbb{Q}\Gamma$-modules V and all $n \geq 2$.*

Bounded $\mathbb{Q}\Gamma$-modules in this statement can be replaced with bounded $\mathbb{R}\Gamma$-modules, bounded $\mathbb{C}\Gamma$-modules, or Banach modules.

Proof. Implications (**c′**) ⇒ (**c**) and (**b′**) ⇒ (**b**) are obvious. Equivalences (**b**) ⇔ (**c**) and (**b′**) ⇔ (**c′**) follow from Proposition 32.

(**a**) ⇒ (**c′**) Fix any bounded $\mathbb{Q}\Gamma$-module V. Since (Γ, Γ') is hyperbolic, by Theorem 38 there exists a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ of type $\mathcal{F}$ with $\mathcal{V} = \mathcal{V}'$, and therefore as in 8.6 we obtain the chain maps φ_* and ψ_*.

Take any $n \geq 2$ and any relative n-cocycle f in $\mathrm{St}^n_\varsigma(\Gamma, \Gamma'; V)$. The cocycle f is not necessarily bounded, but the composition $f \circ \psi_n : C_n(X, \mathbb{Q}) \to V$ is, by Lemma 54, because there are only finitely many Γ-orbits of n-simplices in X and $f \circ \psi_n$ is a $\mathbb{Q}\Gamma$-morphism.

Since the standard resolution St^ς_* is projective, the composition $\psi_* \circ \varphi_*$ is homotopic to the identity map of St^ς_*, and therefore the dual chain map $\varphi^* \circ \psi^* : \mathrm{St}^*_\varsigma(\Gamma, \Gamma'; V) \to \mathrm{St}^*_\varsigma(\Gamma, \Gamma'; V)$ induces the identity map on the relative cohomology in dimensions $n \geq 2$. This implies that the relative cocycle $(\varphi^n \circ \psi^n)(f)$ is cohomologous to f. But $(\varphi^n \circ \psi^n)(f) = f \circ \psi_n \circ \varphi_n$ is bounded because $f \circ \psi_n$ and φ_n are. This proves the surjectivity of the map $H^n_b(\Gamma, \Gamma'; V) \to H^n(\Gamma, \Gamma'; V)$. **(c)** ⇒ **(a)** Suppose to the contrary that there exists a tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ which is finitely presented, whose 1-skeleton $\mathcal{G} := X^{(1)}$ is fine, but the pair (Γ, Γ') is not hyperbolic. This implies that the tuple $(\Gamma, \Gamma', X, \mathcal{V}')$ is not hyperbolic, i.e. $\mathcal{G}$ does not have fine triangles. Proposition 7 implies that $\mathcal{G}$ does not have thin triangles. Now Theorem 51 and Theorem 53 yield a contradiction. □

References

[1] D. J. Allcock and S. M. Gersten, *A homological characterization of hyperbolic groups*, Invent. Math., 135 (1999), pp. 723–742.

[2] M. Auslander. Thesis, Columbia University, 1954.

[3] R. Bieri and B. Eckmann, *Relative homology and poincaré duality for group pairs*, J. Pure Appl. Algebra, 13 (1978), pp. 277–319.

[4] B. Bowditch, *Relatively hyperbolic groups*. Preprint, 2000.

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

[6] I. Bumagin, *On definitions of relatively hyperbolic groups*, in Geometric methods in group theory, vol. 372 of Contemp. Math., Amer. Math. Soc., Providence, RI, 2005, pp. 189–196.

[7] F. Dahmani, *Classifying spaces and boundaries for relatively hyperbolic groups*, Proc. London Math. Soc. (3), 86 (2003), pp. 666–684.

[8] F. Dahmani, *Les bords des groupes relativement hyperbolique*, PhD thesis, Strasbourg, 2003.

[9] C. Druţu and M. Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology, 44 (2005), pp. 959–1058. With an appendix by Denis Osin and Sapir.

[10] B. Farb, *Relatively hyperbolic automatic groups with applications on negatively curved manifolds*, PhD thesis, Princeton, 1994.

[11] B. Farb, *Relatively hyperboic groups*, Geom. Funct. Anal., 8 (1998), pp. 1–31.

[12] M. Gromov, *Hyperbolic groups*, in Essays in group theory, vol. 8 of Math. Sci. Res. Inst. Publ., Springer, New York, 1987, pp. 75–263.

[13] N. V. Ivanov, *Foundations of the theory of bounded cohomology*, Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI), 143 (1985), pp. 69–109, 177–178. Studies in topology, V.

[14] ——, *Foundations of the theory of bounded cohomology*, Jornal of Soviet Mathematics, 37 (1987), pp. 1090–1115.

[15] B. E. Johnson, *Cohomology in Banach algebras*, no. 127 in Mem. Amer. Math. Soc., American Mathematical Society, Providence, R.I., 1972.

[16] I. Mineyev, *ℓ_∞-cohomology and metabolicity of negatively curved complexes*, Internat. J. Algebra Comput., 9 (1999), pp. 51–77.

[17] ——, *Straightening and bounded cohomology of hyperbolic groups*, Geom. Funct. Anal. (GAFA), 11 (2001), pp. 807–839.

[18] ——, *Bounded cohomology characterizes hyperbolic groups*, Quart. J. Math. Oxford Ser., 53 (2002), pp. 59–73.

[19] N. Monod, *Continuous bounded cohomology of locally compact groups*, vol. 1758 of Lecture Notes in Mathematics, Springer-Verlag, Berlin, 2001.

[20] G. A. Noskov, *Bounded cohomology of discrete groups with coefficients.*, Leningrad Math. J, 5 (1991), pp. 1067–1084.

[21] G. A. Noskov, *The Hochschild-Serre spectral sequence for bounded cohomology*, in Proceedings of the International Conference on Algebra, Part 1 (Novosibirsk, 1989), vol. 131 of Contemp. Math., Providence, RI, 1992, Amer. Math. Soc., pp. 613–629.

[22] A. Y. Ol'shanskii, *Hyperbolicity of groups with subquadratic isoperimetric inequality*, Internat. J. Algebra Comput., 1 (1991), pp. 281–289.

[23] D. Osin, *Asymptotic dimension of relatively hyperbolic groups*, Int. Math. Res. Not., (2005), pp. 2143–2161.

[24] D. V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc., 179 (2006), pp. vi+100.

[25] L. Ribes, *On a cohomology theory for pairs of groups*, Proc. Amer. Math. Soc., 21 (1969), pp. 230–234.

[26] A. Szczepański, *Examples of relatively hyperbolic groups*, Geom. Dedicata, 93 (2002), pp. 139–142.

[27] S. Takasu, *Relative homology and relative cohomology theory of groups*, J. Fac. Sci. Univ. Tokyo. Sect. I, 8 (1959), pp. 75–110.

[28] H. F. Trotter, *Homology of group systems with applications to knot theory*, Ann. of Math. (2), 76 (1962), pp. 464–498.

Igor Mineyev, University of Illinois at Urbana-Champaign
mineyev@math.uiuc.edu
http://www.math.uiuc.edu/~mineyev/math/

AslıYaman, CRM, Apartat 50, E-08193, Bellaterra, Spain
ayaman@crm.es
http://www.crm.es/Visitors/PaginesPersonalsVisitants/Asli_Yaman/