ON HAUSDORFF-GROMOV CONVERGENCE AND A THEOREM OF PAULIN

Autor(en): **Bridson, M.R. / Swarup, G. A.**

Objekttyp: **Article**

Zeitschrift: **L'Enseignement Mathématique**

Band (Jahr): **40 (1994)**

Heft 3-4: **L'ENSEIGNEMENT MATHÉMATIQUE**

PDF erstellt am: **12.07.2024**

Persistenter Link: https://doi.org/10.5169/seals-61114

L'Enseignement Mathématique, t. 40 (1994), p. 267-289

ON HAUSDORFF-GROMOV CONVERGENCE AND A THEOREM OF PAULIN

by M.R. BRIDSON [1]) and G.A. SWARUP

ABSTRACT. We give an elementary account of ideas related to Hausdorff-Gromov convergence and explain how, among other things, these ideas can be used to prove a theorem of F. Paulin: If a group Γ is word hyperbolic and its outer automorphism group is infinite, then Γ acts by isometries on an **R**-tree with virtually cyclic segment stabilizers and no global fixed points.

INTRODUCTION

The purpose of this article is to give an essentially self-contained proof of the following theorem of F. Paulin. (The technical terms appearing in this theorem are explained below.)

THEOREM (Paulin). *If Γ is a word hyperbolic group and $Out(\Gamma)$ is infinite, then Γ acts by isometries on an* **R***-tree with virtually cyclic segment stabilizers and no global fixed points.*

We feel that this theorem and (more especially) the techniques involved in its proof are central to the study of word hyperbolic groups and related topics. This is illustrated, for example, by the variety of ways in which these ideas have entered the work of Rips and Sela. The techniques in question centre on Gromov's generalisation of Hausdorff convergence, as developed in Paulin's thesis and Bestvina's work on degeneration of hyperbolic structures. In light of the continuing importance of these techniques, it seemed to us desirable that an elementary and self-contained account of them should be made available.

Let us recall the definitions of the terms appearing in the statement of Paulin's theorem. Let X be a metric space. A topological arc in X is called

[1]) The first author was supported in part by NSF grant DMS-9203500 and FNRS (Switzerland).

a geodesic segment if, with the induced metric, it is isometric to a compact interval of the real line. X is said to be a geodesic space if every pair of points in X can be joined by a geodesic segment. A geodesic triangle in X consists of three points (vertices) together with a choice of geodesic segment (side) joining each pair of them. A geodesic triangle is said to be δ-slim if each of its sides is contained in the δ-neighbourhood of the other two. A geodesic metric space X is said to be δ-hyperbolic if every geodesic triangle in X is δ-slim. An **R**-tree is a 0-hyperbolic space, in other words, a geodesic metric space in which every geodesic triangle is degenerate, i.e., is a tripod. The most primitive example of an **R**-tree is an ordinary simplicial tree in which each of the edges is metrized so as to have length 1. For many purposes, particularly in this article, it is useful to think of a δ-hyperbolic space as a somewhat thickened version of an **R**-tree and to keep in mind the idea that if one looks at the space from so far away that distances of the order of δ appear negligible, then a δ-hyperbolic space takes on the appearance of an **R**-tree.

Recall that the Cayley graph $X(\Gamma, S)$ of a group Γ with respect to a choice of finite generating set $S \subset \Gamma - \{e\}$ is the metric graph whose vertex set is Γ and which has one edge of unit length joining $\gamma \in \Gamma$ to γs whenever $s \in S$. A group is said to be *word hyperbolic* if its Cayley graph $X(\Gamma, S)$ is δ-hyperbolic for some δ. (The hyperbolicity of $X(\Gamma, S)$, but not the specific value of δ, is independent of the choice of S — see [GH] or [Sho].) The class of word hyperbolic groups was introduced by Gromov in [G2], and has proved to be a fruitful context in which to extend many elegant results of hyperbolic geometry, particularly results about geometrically finite groups of isometries of real hyperbolic space that do not contain any parabolic elements.

This article is organised as follows. In Section 1 we describe some basic facts about various generalisations of Hausdorff convergence. In the proof of such elementary facts one discerns a general pattern of argument that can be applied more generally, and it is this pattern, rather than specific compactness criteria, that seems to be most useful in a wider context. In Section 2 we illustrate this point in proving Paulin's theorem.

Paulin has developed an equivariant version of Hausdorff-Gromov convergence, which he calls Gromov convergence, and has used this to give elegant formulations of compactness theorems of Thurston and Culler-Morgan. The compactness criterion which he originally developed in this context relies upon the existence of convex hulls in the spaces under consideration; in spaces such as the Cayley graph of a word hyperbolic group one cannot in general form a precise convex hull for finite sets. This difficulty is the subject of Section 3. Section 4 contains some concluding remarks and a brief discussion of recent work which draws on ideas similar to those discussed in this article.

SECTION 1: HAUSDORFF-GROMOV CONVERGENCE

Until further notice, we fix a compact metric space X and denote by $\mathscr{C}(X)$ the set of closed subsets of X. We shall always denote the open ε-neighbourhood in X of $A \subset X$ by $V_\varepsilon(A)$.

The starting point for our discussion is the following classical construction.

1.1 DEFINITION. *The Hausdorff metric on* $\mathscr{C}(X)$ *is defined by:*

$$D(A, B) = \inf\{\varepsilon \mid A \subseteq V_\varepsilon(B) \text{ and } B \subseteq V_\varepsilon(A)\} .$$

1.2 PROPOSITION. *D is indeed a metric and $\mathscr{C}(X)$ equipped with this metric is compact.*

Proof. The only nontrivial point to check is that $\mathscr{C}(X)$ is compact.

Consider a sequence C_i in $\mathscr{C}(X)$. We must exhibit a convergent subsequence. First notice that given any $\varepsilon > 0$ there exists an integer $N(\varepsilon)$ such that, in its induced metric from X, every $A \in \mathscr{C}(X)$ can be covered by $N(\varepsilon)$ open balls of radius ε. Indeed, because X is compact one can cover it with $N(\varepsilon)$ balls of radius $\varepsilon/2$, then for each such ball which intersects A one chooses a point in the intersection and takes the ball of radius ε about that point. Thus for every positive integer n and every C_i, by taking duplicates if necessary, we may assume that C_i is covered by precisely $N(1/n)$ balls of radius $1/n$, with centres $x_n(i,j)$ for $j = 1, \ldots, N(1/n)$. Furthermore, it is clear from our description of how to choose the $x_n(i,j)$ that this can be done so as to ensure that $x_{n+1}(i,j) = x_n(i,j)$ if $j \leqslant N(1/n)$, thus we may drop the subscript n.

At this stage we have constructed sequences of points $\{x(i,j)\}_j \subset C_i$, each of which has the property that for all $n \in \mathbf{N}$ the balls of radius $1/n$ about the first $N(1/n)$ terms in the sequence cover C_i.

$$\begin{array}{l} C_1 \ni x(1,1), x(1,2), \ldots, x(1,j), \ldots \\ C_2 \ni x(2,1), x(2,2), \ldots, x(2,j), \ldots \\ \quad\vdots \\ C_i \ni x(i,1), x(i,2), \ldots, x(i,j), \ldots \\ \quad\vdots \end{array}$$

Now, because X is compact, we may pass to a subsequence of the C_i in order to assume that the sequence $x(i, 1)$ converges in X, to $x(\omega, 1)$ say. Let C_i^1 denote this subsequence. Inductively, we may pass to further subsequences C_i^k in order to assume that for $j = 1, \ldots, k$ each of the sequences $\{x(i, j)\}_i$ converges in X to $x(\omega, j)$. Let C_ω be the closure in X of $\{x(\omega, j) \mid j \in \mathbf{N}\}$. We claim that the *diagonal sequence* C_k^k converges to C_ω in $\mathscr{C}(X)$. To simplify the notation we write C_k in place of C_k^k.

Observe first that because there is a uniform bound of $1/n$ (independent of l and k) on the distance from $x(k, l)$ to $\Sigma(k, n) := \{x(k, j)\}_{j \leqslant N(1/n)}$, for all l and k we have that the D-distance from $\{x_{\omega, l}\}$ to $\Sigma(\omega, n) := \{x(\omega, j)\}_{j \leqslant N(1/n)}$ is at most $1/n$. Hence the D-distance from C_ω to $\Sigma(\omega, n)$ is at most $1/n$.

Thus, for any $n > 0$, whenever k is large enough to ensure that $d(x(k, j), x(\omega, j)) \leqslant 1/n$ for all $j \leqslant N(1/n)$, we have:

$$D(C_k, C_\omega) \leqslant D(C_k, \Sigma(k, n)) + D(\Sigma(k, n), \Sigma(\omega, n)) + D(\Sigma(\omega, n), C_\omega) \leqslant 3/n . \qquad \square$$

Remark. Already in the above proof we see two of the central themes which recur at the heart of future proofs. First of all, there is the idea of approximating compact sets by finite ones in a uniform way, and secondly there is the use of a diagonal sequence argument to construct a limit object as (the closure of) an increasing union of finite sets.

A more general form of Proposition 1.2, concerning the Chabauty topology, can be found in [CEG]. A quick development of similar ideas is given in C. Hodgson's (unpublished) notes [H].

The following lemma shows how one can rephrase the convergence of compact subspaces in terms of the more familiar notion of convergence of points.

1.3 LEMMA. *A sequence $\{C_n\}_{n \in \mathbf{N}}$ in $\mathscr{C}(X)$ converges to $C \in \mathscr{C}(X)$ if and only if the following two conditions hold:*

(1) *for all $x \in C$ there exists a sequence $x_n \in C_n$ such that $x_n \to x$ in X;*

(2) *every sequence $y_{n(i)} \in C_{n(i)}$ with $n(i) \to \infty$ has a convergent subsequence whose limit point is an element of C.*

Proof. The necessity of conditions (1) and (2) is clear. Conversely, if C_n does not converge to C in $\mathscr{C}(X)$ then, by passing to a subsequence if necessary, we may assume that there exists $\varepsilon > 0$ such that $D(C_n, C) > \varepsilon$ for all n.

There are two cases to consider. First, if for infinitely many values of n it is the case that C_n is not contained in the ε-neighbourhood of C, then by passing to a further subsequence we obtain $x_n \in C_n - V_\varepsilon(C)$. Since X is compact, one can abstract a convergent subsequence of the x_n which converges to some $x_\omega \notin V_\varepsilon(C)$, thus (2) fails.

The other possibility which we must consider is that for infinitely many values of n there exists $z_n \in C - V_\varepsilon(C_n)$. But in this case one can take a convergent subsequence, say $z_m \to z_\omega \in C$, and then $D(z_\omega, C_n) \geqslant \varepsilon$ for arbitrarily large values of n, thus (1) fails. □

We wish to consider what it means for a sequence of compact metric spaces to converge to a limit space when there is no obvious ambient space containing the sequence. For this we need the following definition.

1.4 DEFINITION. *An ε-approximation between two metric spaces A_1 and A_2 is a subset $R \subseteq A_1 \times A_2$ such that:*

(1) *the projection of R to A_i is onto for $i = 1, 2$;*

(2) *if $(x, y), (x', y') \in R$ then $|d_{A_1}(x, x') - d_{A_2}(y, y')| < \varepsilon$.*

If there exists an ε-approximation between A_1 and A_2 then we write $A_1 \sim_\varepsilon A_2$. The Hausdorff-Gromov distance between A_1 and A_2 is:

$$D_H(A_1, A_2) := \inf\{\varepsilon \mid A_1 \sim_\varepsilon A_2\}.$$

If there exists no ε such that $A_1 \sim_\varepsilon A_2$, then $D_H(A_1, A_2)$ is infinite.

Remark. Sometimes, in the course of an argument, 'approximations' R arise which are similar to those in the above definition, but which do not (quite) project onto A_1 and A_2. For example, it may happen that one has a naturally defined ε-approximation between dense subsets of A_1 and A_2; in this case, given any $\varepsilon' > \varepsilon$, one can extend the given relation to obtain an ε′-approximation between A_1 and A_2. Similarly, if one has a relation $R \subset A_1 \times A_2$ whose projections to A_1 and A_2 are ε-dense (in the sense that every point is within a distance ε of these projections) then R can be extended to a 3ε-approximation between A_1 and A_2. In what follows, when necessary, we shall implicitly assume that approximations which arise are adjusted so as to make their projections surjective, in accordance with Definition 1.4.

Terminology. We say that a sequence of metric spaces C_n converges to C in the Hausdorff-Gromov topology, and write $C_n \to C$, if and only if $D_H(C_n, C) \to 0$ as $n \to \infty$. Given a relation $R \subseteq A_1 \times A_2$, the phrase

$(x, y) \in R$ will often be written xRy, or 'x is related to y', or 'x corresponds to y'.

In closer analogy to the Hausdorff distance between compact subsets of a fixed metric space, one has the so-called Hausdorff distance between compact metric spaces A_1 and A_2. For this one considers all metric spaces X which contain isometric copies of A_1 and A_2. As in (1.1) one can consider the distance between A_1 and A_2 in $\mathscr{C}(X)$, and the Hausdorff distance between A_1 and A_2 is defined to be $D_h(A_1, A_2) := \inf_X \{D_{\mathscr{C}(X)}(A_1, A_2)\}$. It is not hard to show that for compact spaces $D_H = 2D_h$.

It is clear that the Hausdorff-Gromov distance between a metric space and any dense subset of it is zero, and hence limits of sequences of spaces are not unique in general. However:

1.5 PROPOSITION. *Two compact metric spaces A and B are isometric if and only if the Hausdorff-Gromov distance between them is zero.*

Proof. We shall show that if $D_H(A, B) = 0$ then A and B are isometric, the other implication is trivial. Let $\{a_n\}$ be a countable dense subset of A and let R_m be a $(1/m)$-approximation between A and B. We choose $b_{m,n} \in B$ so that $a_n R_m b_{m,n}$. We can pass to a subsequence of $\{b_{m,1}\}_m$ and assume that $b_{m,1} \to b_1$ in B. By passing to a further subsequence we may assume that $b_{m,2} \to b_2$, and so on. For all n, n', m we have that $| d_A(a_n, a_{n'}) - d_B(b_{m,n}, b_{m,n'}) | < 1/m$, and hence $d_A(a_n, a_{n'}) = d_B(b_n, b_{n'})$. Thus, we obtain the desired isometry $A \to B$ by taking the unique continuous extension of $a_n \mapsto b_n$. □

Thus, if we confine ourselves to compact metric spaces then limits are unique whenever they exist. We saw in (1.2) that if a sequence of compact metric spaces is contained in an ambient compact space then it has a convergent subsequence. Recall that closed subspaces of a fixed compact metric space are uniformly compact in the following sense.

1.6 DEFINITION. *We say that a family $\{C_i\}_{i \in I}$ of compact metric spaces is* uniformly compact *if there is a uniform bound on their diameters, and for every $\varepsilon > 0$ there exists an integer $N(\varepsilon)$ such that each of the $\{C_i\}$ can be covered by $N(\varepsilon)$ balls of radius ε.*

A set of points in C_i which has the property that the ε-balls around these points cover C_i is called an *ε-net* for C_i. The corresponding cover is called an *ε-cover.*

Notice that the integer $N(\varepsilon)$, which is sometimes called the ε-count, plays a significant role in our proof of (1.2). Gromov has shown [G1]:

1.7 THEOREM. *If a sequence* $\{C_i\}_{i\in\mathbf{N}}$ *of compact metric spaces is uniformly compact then there is a subsequence which converges in the Hausdorff-Gromov metric.*

If one insists that the limit be complete, then it will be compact. It is possible to establish Gromov's criterion by a direct adaptation of the proof of (1.2) presented above. (The major difficulty in doing so is that one can no longer use the presence of the ambient compact space to deduce the existence of the points $x(\omega, j)$, and instead one must pass to suitable subsequences to ensure that for all j, j' the sequence of numbers $d(x(i,j), x(i,j'))_{i\in\mathbf{N}}$ converges; $x(\omega, j)$ should then be defined to be a certain sequence $\{x(i,j)\}_{i\in\mathbf{N}}$; the limits of the above sequences of numbers give a (pseudo-)metric on the set of the $x(\omega, j)$, and after identifying points which are a distance zero apart and taking the completion, one obtains the desired compact limit space.)

In [G1] Gromov established his compactness criterion by a different argument, embedding the sequence $\{C_i\}_{i\in\mathbf{N}}$ as compact subspaces of a fixed compact space. We emphasized the alternative proof sketched above for two reasons. First of all, the strategy of proof is much the same as that which we shall employ in Section 2 in order to construct the **R**-tree referred to in the statement of Paulin's theorem. Secondly, the argument suggested above highlights the degree of flexibility which one has in constructing limit spaces. In particular, if one has a sequence of well-understood spaces, then it is possible to make points in the limit correspond to specific points in the limiting spaces, and hence one can then use the geometry of the limiting spaces to elucidate the structure of the limit.

Convention. Given a convergent sequence of spaces $C_i \to C$ and ε_i-relations $R_i \subseteq C_i \times C$ with $\varepsilon_i \to 0$, one says that the sequence $\{x_i\}_{i\in\mathbf{N}}, x_i \in C_i$ converges to $x_\infty \in C$ if $x_i R_i x_\infty$. Under these circumstances, we write $x_i \to x_\infty$ and say that x_i approximates x_∞ in C_i.

It is clear from the preceding discussions that Hausdorff-Gromov convergence is very natural in the context of compact metric spaces, however it is a less satisfactory concept of convergence for non-compact spaces. One obvious disadvantage is that the distance between a compact space and an unbounded space is always infinite. Thus, for example, Hausdorff-Gromov convergence is insufficient to capture the intuitive notion that as the radius of a sphere of constant curvature tends to infinity, the sphere looks increasing like Euclidean space. There are at least two useful ways of extending the notion of Hausdorff-Gromov convergence so that it is better adapted to the study of non-compact spaces. The first, which was introduced by Gromov [G1], gives a notion of convergence for *proper* metric spaces (i.e., metric spaces in which closed and bounded subsets are compact) with a choice of basepoint.

1.8 DEFINITION. *Let* $\{X_i\}_{i \in \mathbf{N}}$ *be a sequence of proper metric spaces with basepoints* $x_i \in X_i$. *The sequence of pointed spaces* $\{(X_i; x_i)\}_{i \in \mathbf{N}}$ *is said to converge to* $(X; x)$ *if for every* $r > 0$ *the sequence of compact metric balls* $\{B(x_i, r)\}_{i \in \mathbf{N}}$ *converges to* $B(x, r) \subseteq X$ *in the Hausdorff-Gromov metric.*

We call this notion of convergence 'pointed Hausdorff-Gromov convergence'.

Remark. If we fix a basepoint x_n on the m-dimensional sphere S_n of radius n, then $(S_n; x_n)$ converges to the flat space $(\mathbf{E}^m; 0)$. In particular, this example shows that pointed Hausdorff-Gromov convergence does *not* imply that the corresponding (unpointed) metric spaces converge in the Hausdorff-Gromov metric. For instance, in this example $D_H(S_n, \mathbf{E}^m)$ is infinite for all n.

Gromov's compactness criterion for sequences of compact spaces implies that if for every $r > 0$ the balls $\{B(x_i, r)\}_{i \in \mathbf{N}}$ are uniformly compact, then $\{(X_i; x_i)\}_{i \in \mathbf{N}}$ has a convergent subsequence. But if one insists that the limit space $(X; x)$ be complete, then it is necessarily proper (and unique, by an easy extension of (1.5)). Thus one needs an alternative notion of convergence in situations where the spaces which appear as a limit of proper spaces are not locally compact. Such a situation arises in the study of degenerations of hyperbolic structures [Sha]. A suitable notion of convergence in such cases was introduced by Paulin in his thesis [P1] (see also, Bestvina [B]). Paulin calls this notion *Equivariant Gromov Convergence.* The idea is that finite subsets of the limit should be equivariantly approximated by finite subsets of the limiting sequence (see Section 4 below). It is important to emphasize that even when the group in question is the trivial group, equivariant Gromov convergence does not imply Hausdorff-Gromov convergence. Indeed, in the cases of most interest one typically obtains limit spaces which are not locally compact ($\mathbf{R}$-trees).

With respect to equivariant Gromov convergence, limits are not unique in general, but Paulin has shown that under suitably strong convexity hypotheses one can establish the existence of limits by means of an extension of the compactness criterion of Gromov referred to above.

In this article our main interest lies with constructing group actions on **R**-trees by producing these trees as a limit of δ-hyperbolic spaces. We are not interested in the type of convergence which occurs so much as we are in the properties of the limit. In fact, in our situation, one can deduce these properties simply by looking at Hausdorff-Gromov convergence on compact subsets. The proof of the following proposition gives an illustration of the techniques involved. For this proof we shall need the following terminology.

Terminology. If R is a relation in $A \times B$ and $C \subseteq A$, then we define *the R-image of C in B* (or, more briefly, the image of C in B) to be $\text{proj}_B(\text{proj}_A^{-1}(C) \cap R)$. Given $D \subseteq B$, the image of D in A is defined similarly. Note that if $R \subset A \times B$ is an ε-approximation, then for every subset $C \subseteq A$ with R-image $D \subseteq B$, the restricted relation $R \cap (C \times D)$ is an ε-approximation between C and D.

We recall some basic definitions. A metric space is said to be a *geodesic* space if every pair of points $x, y \in X$ can be joined by a topological arc which, with the induced metric, is isometric to $[0, d(x, y)] \subseteq \mathbf{R}$. Such a topological arc is called a *geodesic segment.* In general, one does not require such geodesic segments to be unique, but despite this it is often convenient to use the notation $[x, y]$ for a definite choice of geodesic segment from x to y.

Given a graph X (i.e., a 1-dimensional CW complex) one can turn it into a geodesic metric space by fixing a homeomorphism from each 1-cell to [0, 1] and pulling back the metric; one can use these local metrics to measure the length of paths, and one obtains a geodesic metric space by defining the distance between two points to be the infimum of the lengths of paths joining them.

Given a connected subgraph or a connected compact $Y \subseteq X$ one defines the *induced path metric* on Y by setting the distance between two points equal to the length of paths *in* Y which connect them. It is easy to see that this endows Y with the structure of a geodesic metric space, and if X is a locally finite graph then Y, thus metrized, is a proper geodesic metric space. It also follows easily from the definition that the distance between two points in the induced path metric on Y is at least as great as the distance between these points in X.

The definition of a δ-hyperbolic space was given in the introduction. The definition which we gave is called the δ-*slim* condition in [Sho]. It is not difficult to show that this is equivalent to requiring that there exists a constant δ' so that every *non-degenerate* geodesic triangle in the given geodesic metric space admits a map to a tripod (a metric graph with 3 edges, 3 vertices of valence 1, and one vertex of valence 3) so that this map restricts to an isometric embedding on each side of the triangle, and the fibres of the map have diameter at most δ' (see [Sho], p. 16). The following proposition is from [P4].

1.9 PROPOSITION. *Let* $\{C_i\}_{i \in \mathbf{N}}$ *and* C *be compact metric spaces such that* C_i *converges to* C *in the Hausdorff-Gromov topology.*

(1) *If* C_i *are geodesic metric spaces, then* C *is a geodesic metric space;*

(2) *if* C_i *are, in addition,* δ_i*-hyperbolic with* $\delta_i \to 0$, *then* C *is an* **R**-*tree.*

Proof. Let R_i be an ε_i-approximation between C_i and C with $\varepsilon_i \to 0$. Given $x, y \in C$, we choose $x_i, y_i \in C_i$ with $x_i R_i x, y_i R_i y$. Thus $|d(x, y) - d(x_i, y_i)| < \varepsilon_i$. Note that the numbers $d(x_i, y_i)$ are bounded. Let $w_i : I_i \to [x_i, y_i]$ be an isometry of $I_i = [0, d(x_i, y_i)]$ to a choice of geodesic $[x_i, y_i]$ joining x_i and y_i. We have $d(x_i, y_i) \to d(x, y)$ and $I_i \to I_\infty = [0, d(x, y)]$. Let L_i be the R_i-image of $[x_i, y_i]$ and let K_i be the closure of L_i in C. Then, K_i is ε_i-close to $[x_i, y_i]$ and hence to I_i (in the Hausdorff-Gromov metric). By 1.1, a subsequence of $\{K_i\}_{i \in \mathbf{N}}$, which we still denote by K_i, converges in the Hausdorff metric, to $K \subset C$ say. But $d_H(I_\infty, K_i) \leqslant d_H(I_\infty, I_i) + d_H(I_i, K_i)$, which goes to 0 as $i \to \infty$. Thus K is isometric to I. Since $x, y \in K$ (in fact they belong to all L_i) and since $d(x, y) = l(I_\infty)$, the isometry $I_\infty \to K$ gives a geodesic joining x and y. This proves assertion (1).

To prove the second part of the proposition, we first show that if $\delta_i \to 0$, then there is a unique geodesic joining x to y in C, and hence every geodesic in C arises as in the first part of the proof. We fix a geodesic $[x, y]$ which arises as in the first part of the proof, and consider an arbitrary geodesic joining x, y; let z be the midpoint of this second geodesic. We must show that $z \in [x, y]$. By the above construction, we obtain geodesics $[x_i, z_i]$, $[z_i, y_i]$ in C_i converging to geodesics $[x, z]$, $[z, y]$ joining x, z and z, y respectively. Consider in C_i the geodesic triangles with sides $[x_i, y_i]$, $[y_i, z_i]$, $[z_i, x_i]$. Choose z_i', y_i', x_i' on $[x_i, y_i]$, $[z_i, x_i]$, $[y_i, z_i]$ respectively so that $d(x_i, z_i') = d(x_i, y_i')$, $d(z_i, y_i') = d(z_i, x_i')$ and $d(y_i, z_i') = d(y_i, x_i')$. It is not difficult to see that $d(y_i', z_i')$, $d(z_i', x_i')$, $d(x_i', y_i')$ are all less than $4\delta_i$ (cf. [Sho], p. 17, proof of slim implies thin).

Thus, as $i \to \infty$, we see that x_i', y_i', z_i' converge to the same point, say z', on $[x, y]$. Thus $d(x_i, y_i') + d(y_i, x_i') - d(x_i, y_i)$ converges to zero. Since $d(x_i, z_i) + d(y_i, z_i) - d(x_i, y_i)$ also converges to zero, we have that $d(y_i', z_i) + d(z_i, x_i')$ converges to zero. Since $d(z_i, z_i') \leqslant d(z_i', y_i') + d(y_i', z_i) \leqslant 4\delta_i + d(y_i', z_i)$ we see that the z_i converge to the point z' on our original geodesic segment $[x, y]$. Thus z, the midpoint of our arbitrary geodesic from x to y, coincides with the midpoint of our fixed geodesic. Repeating the argument we see that these geodesics must agree at a dense set of points, and hence everywhere. Since geodesic triangles in C_i are δ_i-slim, and geodesics in C all arise as limits of geodesics in C_i, we see that geodesic triangles in C must be 0-slim, and hence C is an **R**-tree. ☐

Remark. If one has a sequence of δ_i-hyperbolic spaces C_i, with $C_i \to C$ and $\delta_i \to \delta > 0$, then one can extend the preceding argument to show that C is δ'-hyperbolic (with $\delta' = 19\delta$, for example).

SECTION 2: THE PROOF OF PAULIN'S THEOREM

In this section we shall prove the following theorem of F. Paulin [P4].

2.1 THEOREM (Paulin). *If Γ is a word hyperbolic group and $Out(\Gamma)$ is infinite, then Γ acts by isometries on an* **R***-tree with virtually cyclic segment stabilizers and no global fixed points.*

In its outline, the proof given below is very similar to Paulin's original proof, except that we use Hausdorff-Gromov convergence instead of the equivariant Gromov convergence used by Paulin. In particular, this allows us to avoid the difficulties discussed in the next section.

Let S be a finite set of generators for Γ and let $X = X(\Gamma, S)$ denote the Cayley graph of Γ with respect to S, as defined in the introduction. Γ is the vertex set of X and receives the induced metric. The hypothesis that Γ is word hyperbolic means precisely that there exists $\delta > 0$ such that X is a δ-hyperbolic geodesic metric space. Note that with our definition of a Cayley graph, the endpoints of each edge are distinct, and there is at most one edge joining each pair of vertices; hence the action of Γ on itself be left multiplication can be extended linearly across edges in a unique way to give an isometric action of Γ on X.

The proof of Theorem 2.1 will be broken into a number of smaller results. We begin by noting that, because $Out(\Gamma)$ is infinite, we can choose a sequence of automorphisms $\{\phi_i\}_{i \in \mathbf{N}}$ such that none of the ϕ_i is an inner automorphism and no two of the ϕ_i have the same image in $Out(\Gamma)$. For each $i \in \mathbf{N}$ we consider the function $f_i : X \to [0, \infty)$ defined by:

$$f_i(x) = \max_{s \in S} d(x, \phi_i(s)x) . \tag{2.2}$$

This function has been used by Bestvina in his study of degeneration of real hyperbolic structures [B], and our use of this function is similar to his. (A similar idea was used earlier in a different context by Thurston [T, Prop. 1.1].)

Note that f_i takes on integer values at vertices and midpoints of edges in X, and its restriction to half-edges is linear. It follows that f_i attains its infimum (which is an integer) at some point, $x_i \in X$ say. (In the case where Γ is not virtually cyclic one can also see this by showing that f_i is a proper map, i.e., a map with the property that the inverse image of a compact set is compact.)

Let

$$\begin{aligned} \lambda_i &= \max_{s \in S} d(x_i, \phi_i(s)x_i) \\ &= \inf_{x \in X} \max_{s \in S} d(x, \phi_i(s)x) . \end{aligned} \tag{2.3}$$

We fix a definite choice of points x_i with the above property.

For future reference, we note that by passing to a subsequence of the ϕ_i we may assume there is a single element $s_0 \in S$ such that $\lambda_i = d(x_i, \phi_i(s_0)x_i)$ for all $i \in \mathbf{N}$. We also note that with the above choice of x_i, the triangle inequality yields:

$$d(x_i, \phi_i(\gamma)x_i) \leqslant \lambda_i d(e, \gamma) . \tag{2.4}$$

Following Paulin, we next note that because $Out(\Gamma)$ is infinite, the sequence λ_i must be unbounded. For suppose that there were a uniform bound, ρ say, on the value of λ_i. Then for any vertex $y_i \in X$ closest to x_i, we would have $d(e, y_i^{-1}\phi_i(s)y_i) = d(y_i, \phi_i(s)y_i) \leqslant \rho + 2$ for all $s \in S$, $i \in \mathbf{N}$. But there are only finitely many vertices in the ball of radius $\rho + 2$ about e, so this bound would imply the existence of integers $n \neq m$ such that $y_n^{-1}\phi_n(s)y_n = y_m^{-1}\phi_m(s)y_m$ for all $s \in S$. Whence ϕ_n and ϕ_m would be equal in $Out(\Gamma)$, contrary to hypothesis. Thus we have shown that the sequence of numbers $\{\lambda_i\}_{i \in \mathbf{N}}$ is unbounded, so we may pass to a subsequence $\{\lambda_n\}_{n \in \mathbf{N}}$ which is *strictly increasing* and assume that $\lambda_n \to \infty$ as $n \to \infty$.

Consider the sequence of metric spaces $X_k = (X, d_k)$, where $d_k := d/\lambda_k$ is the original metric on X scaled down by λ_k. In what follows we shall intermittently use both the original metric d and the scaled metric d_k, specifying which on each occasion and, where appropriate, using the formal notation (Y, d) for a metric space which consists of the set Y together with a distance function d. But for the moment, the most important distinction between the X_k will be that we shall regard Γ as acting on X_k via ϕ_k, and think of our chosen point x_k, at which the minimax λ_k is attained, as a *basepoint* in X_k. More precisely, we consider the sequence of pointed Γ-spaces (X_k, x_k), where the action of $\gamma \in \Gamma$ on X_k is $x \rightarrow \phi_k(\gamma)x$.

We wish to use the hyperbolic nature of X_k to approximate it by a sequence of star-like compact subsets $X_k(P_i)$ centred at x_k. To this end, we fix a sequence of finite subsets $\{e\} = P_0 \subseteq P_1 \subseteq P_2 \cdots \subseteq P_i \subseteq \cdots$ which exhaust Γ. Let $n_i = |P_i|$ denote the cardinality of P_i. The desired subsets of X_k are defined inductively as follows: $X_k(P_0) = \{x_k\}$, and $X_k(P_i)$ is the union of $n_i - 1$ geodesic segments, those in $X_k(P_{i-1})$ together with a choice of geodesic segment from x_k to each element of $\{\phi_k(\gamma)x_k \mid \gamma \in P_i - P_{i-1}\}$.

We next 'fatten-up' each of the sets $X_k(P_i)$ by taking its closed δ-neighbourhood in the metric d. Henceforth we shall denote this neighbourhood V_k^i. Let $d_{i,k}$ be the induced *path metric* on V_k^i. As we discussed in Section 1, $(V_k^i, d_{i,k})$ is a geodesic metric space. It is also important to notice that the induced path metric which V_k^i receives from d_k is $d_{i,k}/\lambda_k$. The following lemma is suggested by an argument of B. Bowditch [Bo].

2.5 LEMMA. *With the above notation, for all* $x, y \in V_k^i$ *we have:*

$$d(x, y) \leqslant d_{i,k}(x, y) \leqslant d(x, y) + 4\delta .$$

Proof. The left-most inequality comes from the general fact that for any subspace of a geodesic metric space the induced metric is dominated by the induced path metric. In order to establish the other inequality, we first note that $X_k(P_i)$ is δ-convex in (X_k, d), in the sense that if a geodesic segment in X_k joins a pair of points $x, y \in X_k(P_i)$, then this geodesic segment lies entirely within the closed δ-neighbourhood V_k^i of $X_k(P_i)$.

Given $x, y \in V_k^i$, we fix points $z, w \in X_k(P_i)$ closest to x and y respectively. (Such points are not unique in general.) Let $[x, z]$, $[z, w]$ and $[w, y]$ be choices of geodesic segments joining x to z, z to w and w to y, respectively. Each is contained in V_k^i, and hence so is the broken geodesic $[x, z, w, y]$ obtained by concatenating them. The length of this broken geodesic is at most $d(z, w) + 2\delta \leqslant d(x, y) + 4\delta$. Hence $d_{i,k}(x, y) \leqslant d(x, y) + 4\delta$. □

The subspace V^i_k forms a good substitute for the notion of a convex hull for $\phi_k(P_i)x_i$ in X_k. According to the above lemma, geodesics in $(V^i_k, d_{i,k})$ are $(1, 4\delta)$-quasigeodesics in (X_k, d), and hence by [GH, p. 82] there exists a constant $\eta = \eta(\delta)$ (independent of k, i) such that geodesic triangles in $(V^i_k, d_{i,k})$ are η-slim. Thus we have proved the first part of:

2.6 Lemma. *There exists a constant* $\eta = \eta(\delta)$ *such that, for all* $k \in \mathbf{N}$, *with respect to the path metric* $d_{i,k}$ *on* V^i_k, *geodesic triangles in* V^i_k *are* η*-slim. Moreover, for fixed* i, *with respect to the (scaled) path metrics* $d_{i,k}/\lambda_k$, *the metric spaces* $\{V^i_k\}_{k \in \mathrm{N}}$ *are uniformly compact.*

Proof. It remains to prove the assertion of the second sentence. We follow an argument of Bestvina [B]. Until further notice we work with the metric d. Let μ_i be the maximum of the integers $\{d(e, \gamma) \mid \gamma \in P_i\}$. Each of the geodesic segments used to define $X_k(P_i)$ has length at most $\mu_i\lambda_k$ (by (2.4)). Therefore, given $\varepsilon > 0$, we can cover $X_k(P_i)$ by $2n_i\mu_i/\varepsilon$ segments of length at most $\lambda_k\varepsilon/2$. (Recall that $n_i = |P_i|$.) Hence, if $\lambda_k\varepsilon > 2\delta$, then in order to cover V^i_k we need at most $2n_i\mu_i/\varepsilon$ balls of radius $\lambda_k\varepsilon$. But we arranged that $\lambda_k \to \infty$ as $k \to \infty$, so this is true for large k.

Now we change viewpoints and work with the scaled metric d_k on X_k, and the induced path metric on V^i_k. In this setting, the preceding argument shows that for large k one needs only $2n_i\mu_i/\varepsilon$ balls of radius ε to cover V^i_k. Since the path metric on V^i_k and the restriction to V^i_k of d_k differ by at most an additive constant of $4\delta/\lambda_k$, we have thus established the existence of a uniform ε-count for the $\{V^i_k\}_{k \in \mathrm{N}}$ both when equipped with the restriction of the metrics d_k and when equipped with the induced path metrics. Because they are *path* metric spaces, a uniform ε-count also yields a bound on the diameter of the V^i_k. □

Continuing with the proof of Paulin's theorem, we fix an integer j and suppose that we are given a positive constant ε. According to the preceding lemma, we can choose ε-nets $N_\varepsilon(k, j)$ for V^j_k on whose cardinalities there is a bound independent of k. We may also assume that the set $N_\varepsilon(k, j)$ includes $\phi_k(P_j)x_k$. Since, for fixed j, the $N_\varepsilon(k, j)$ are finite metric spaces of uniformly bounded cardinality and diameter, we can pass to a subsequence (using a diagonal type argument, as in Section 1) so as assume that, for all $\gamma, \gamma' \in P_j$, the sequence of numbers $d_{j,k}(\phi_k(\gamma)x_k, \phi_k(\gamma')x_k)$ converges as $k \to \infty$. Passing to a further subsequence which is convergent in the Hausdorff-Gromov topology we obtain a limit metric space $L_{\varepsilon,j}$ (whose cardinality will be no greater than that of the $N_\varepsilon(k, j)$). As a basepoint in the limit space we choose the limit of the sequence x_k, and we christen this point x_∞. For each $\gamma \in P_j$, we denote the limit of the sequence $\phi_k(\gamma)x_k$ by γx_∞.

We next take an $\varepsilon/2$-net for V_k^j which is constructed so as to include the previously chosen ε-net. Passing to a subsequence if necessary, we obtain a finite limit metric space $L_{\varepsilon/2,j}$. We proceed in this manner, taking finer ε-nets, and at each stage including the previous (coarser) ones and extracting convergent subsequences to obtain finite limit metric spaces. The natural inclusions of each ε-net into its refinements gives a natural identification of points in the limit, so it is not too abusive a notation to write:

$$L_{\varepsilon,j} \subset L_{\varepsilon/2,j} \cdots \subset L_{\varepsilon/2^n,j} \subset \cdots$$

We define L_j to be the direct limit of this sequence, that is, $L_j = \bigcup\{L_{\varepsilon/2^n,j} \mid n \in \mathbf{N}\}$. We denote by $\hat{L}_j$ the metric completion of L_j. Since the diameters of the V_k^j are uniformly bounded in the scaled metrics, we see that $\hat{L}_j$ is a complete space of finite diameter, and hence is compact.

By choosing a diagonal type subsequence and renumbering, we obtain the following array of spaces with convergence in both the horizontal and vertical directions:

$$\begin{array}{ccccccccccc}
N_\varepsilon(1,j) & \subseteq & N_{\varepsilon/2}(2,j) & \subseteq & \cdots\cdots & \subseteq & N_{\varepsilon/2^n}(1,j) & \subseteq & \cdots & \subseteq & V_1^j \subseteq X_1 \\
N_\varepsilon(2,j) & \subseteq & N_{\varepsilon/2}(2,j) & \subseteq & \cdots\cdots & \subseteq & N_{\varepsilon/2^n}(2,j) & \subseteq & \cdots & \subseteq & V_2^j \subseteq X_2 \\
\vdots & & \vdots & & & & \vdots & & & & \vdots \\
N_\varepsilon(m,j) & \subseteq & N_{\varepsilon/2}(m,j) & \subseteq & \cdots\cdots & \subseteq & N_{\varepsilon/2^n}(m,j) & \subseteq & \cdots & \subseteq & V_m^j \subseteq X_m \\
\vdots & & \vdots & & & & \vdots & & & & \vdots \\
L_{\varepsilon,j} & \subseteq & L_{\varepsilon/2,j} & \subseteq & \cdots\cdots & \subseteq & L_{\varepsilon/2^n,j} & \subseteq & \cdots & \subseteq & \hat{L}_j
\end{array}$$

Our next goal is to show that as $k \to \infty$ the V_k^j actually converge to $\hat{L}_j$ in the Hausdorff-Gromov topology. We have that $N_{\varepsilon/2^n}(m,j)$ is $\varepsilon/2^{n-1}$ close to V_m^j for all m. After passing to yet another diagonal type subsequence, we may assume that $N_{\varepsilon/2^n}(m,j)$ is $\varepsilon/2^{m-1}$ close to $L_{\varepsilon/2^n,j}$ for all $m \geqslant n$. Thus V_m^j and $L_{\varepsilon/2^n,j}$ are $\varepsilon/2^{n-2}$ close for $m \geqslant n$. On the other hand, $L_{\varepsilon/2^n,j}$ and $L_{\varepsilon/2^{n+1},j}$ are $\varepsilon/2^{n+1}$ close (since any choice of $\varepsilon/2^n$ and $\varepsilon/2^{n+1}$ nets of V_k^j are $\varepsilon/2^{n+1}$ close). Thus $L_{\varepsilon/2^n,j}$ is $\Sigma_{i\geqslant n}\varepsilon/2^i$ close to L_j and $\hat{L}_j$. Hence V_n^j and $\hat{L}_j$ are $\varepsilon/2^{n-3}$ close, so V_n^j converges to $\hat{L}_j$, in the Hausdorff-Gromov topology, as $n \to \infty$.

Notice that, by (1.9) and (2.6), the spaces $\hat{L}_j$ are **R**-trees of finite diameter, because V_k^j is η/λ_k-hyperbolic and $\lambda_k \to \infty$. It is also useful to observe that $\hat{L}_j$ is spanned by γx_∞, with $\gamma \in P_j$. Furthermore, the $X_k(P_j)$ themselves converge to $\hat{L}_j$ because $X_k(P_j)$ and V_k^j are $4\delta/\lambda_k$-close and $\lambda_k \to \infty$. However, in what follows it is most convenient to still work with V_k^j rather than $X_k(P_j)$ when we need to take a choice of geodesic between two points of $X_k(P_j)$. Also, because the scaled path metric on V_k^j and the induced metric d_k/λ_k differ only by $4\delta/\lambda_k$, which tends to 0 as $k \to \infty$, henceforth it is not important to keep track of the difference between these two metrics.

By construction, all of our $\varepsilon/2^n$-nets include the set $\{\phi(\gamma)x_k \mid \gamma \in P_j\}$ and each of the sequences $d_k(\phi(\gamma)x_k, \phi(\gamma')x_k)$ converges. Thus, if we denote by $x_\infty \in \hat{L}_j$ the 'limit' of the x_k, and by γx_∞ the limit of the $\phi(\gamma)x_k$, then we see that $d(\gamma x_\infty, \gamma' x_\infty)$ (distance in $\hat{L}_j$) is independent of j. Since the tree $\hat{L}_j$ is the convex hull of the points γx_∞, we can define an isometric embedding of $\hat{L}_j$ into $\hat{L}_{j+1}$ for all j and hence obtain an **R**-tree by taking the direct limit of the resulting system of inclusions. We denote the direct limit metric space with basepoint (which as the limit of **R**-trees is itself an **R**-tree) by $(X_\infty; x_\infty)$. The final important observation to make is that Γ acts isometrically on X_∞, because it acts isometrically on the subset $\{\gamma x_\infty\}_{\gamma \in \Gamma}$ (by left translation), and the convex hull of this subset is the whole of X_∞.

Let us now examine the nature of the action of Γ on X_∞. We claim that it has the following properties:

(1) There is no point of X_∞ whose stabilizer is the whole of Γ.

(2) The stabilizer of every non-trivial segment in X_∞ is virtually cyclic.

To see that (1) is true, let us see what would happen if it were to fail. Suppose that Γ were to stabilize a point $z_\infty \in X_\infty$. We fix a segment $z_\infty \in [\gamma x_\infty, \gamma' x_\infty] \subseteq \hat{L}_j$. Up to the taking of subsequences, we have that the closures in $\hat{L}_j$ of the images of the geodesic segments $[\gamma x_k, \gamma' x_k] \subseteq V_k^j$ converge (in the Hausdorff metric) to $[\gamma x_\infty, \gamma' x_\infty]$, and we fix points $z_k \in [\gamma x_k, \gamma' x_k]$ which converge to z_∞. We then choose j large enough to ensure that $S \subset P_j$ (recall that S is our fixed finite generating set for Γ), and l large enough to ensure that $P_j P_j \subset P_l$.

We have, for every $s \in S$, geodesics $[s\gamma x_k, s\gamma' x_k] := s \cdot [\gamma x_k, \gamma' x_k]$ in V_k^l, and (by definition of the action on X_∞) the closures of their images in $\hat{L}_l \subseteq X_\infty$ converge to $[s\gamma x_\infty, s\gamma' x_\infty]$. Moreover, $\{s z_k\}_{k \in \mathbf{N}}$ converges to $s \cdot z_\infty = z_\infty$, so for large k we have that $d_k(s \cdot z_k, z_k) < 1/4$ in the *scaled* metric of X_k. Hence $d(s \cdot z_k, z_k) < \lambda_k/4$, for large k, in the original metric on X_k. But this contradicts the definition of λ_k.

Remark. The preceding argument actually shows that for every finite set $P \subseteq \Gamma$ which fixes z_∞, given any $\varepsilon > 0$ one has that for k sufficiently large z_k and γz_k are ε-close, in the scaled metric d_k, for every $\gamma \in P$.

We next need to show that segment stabilizers are virtually cyclic. This seems to be the place where some sort of discreteness assumption on Γ is needed. In the classical real-hyperbolic case, Margulis' Lemma implies the result for discrete actions (see [B] and [P2]). Since we are using Cayley graphs and the group actions are (almost) free there is still some sort of discreteness and Paulin gives a delicate argument to show that segment stabilizers are virtually cyclic. The following algebraic lemma is taken from [P4]:

2.7 LEMMA. *Let* G *be a finitely generated group. If the set of commutators* $\{aba^{-1}b^{-1} \mid a, b \in G\}$ *is finite, then* G *is virtually abelian.*

Proof. The action of G on itself by conjugation determines a map $G \to Aut(\Gamma)$, whose image is $Inn(G)$ and whose kernel is the centre of G; it suffices to prove that $Inn(G)$ is finite. If A is a finite generating set for G, then the action of $g \in G$ by conjugation is determined by its action on the elements $a \in A$. But $g^{-1}ag = (g^{-1}aga^{-1})a$, and by hypothesis there are only finitely many possibilities, M say, for the commutator $g^{-1}aga^{-1}$. Hence the cardinality of $Inn(G)$ is at most $M^{|A|}$. □

We proceed with the proof of assertion (2) on segment stabilizers. We call a subgroup *large* if it contains a non-abelian free subgroup (for hyperbolic groups this is equivalent to not having a cyclic subgroup of finite index). Suppose that a large subgroup G of Γ stabilizes a non-trivial segment $e \subseteq X_\infty$. If e is finite, then a subgroup of index 2 in G fixes e pointwise. If e is infinite, a subgroup of index 2 in G acts as translations on a ray in e and thus a large subgroup of G, obtained by taking commutators, fixes a segment of positive length in e pointwise. Thus, in any case, if a large subgroup of Γ stabilizes a segment, then a (perhaps smaller) large subgroup of Γ fixes a segment e of positive length pointwise. Therefore, in order to complete the proof of Paulin's theorem, it suffices to show that if a subgroup of Γ fixes a segment of X_∞ pointwise, then that subgroup is virtually cyclic. Let D denote the length of such a segment which is fixed pointwise by the subgroup $G \subset \Gamma$, and let z and z' denote the endpoints of the segment.

We fix $\varepsilon > 0$ small (to be estimated later) compared to D, and k so large that if $z_k, z'_k \in X_k$ correspond to $z, z' \in X$ then $|d(z_k, z'_k) - D| < \varepsilon$. We fix a geodesic segment $[z_k, z'_k]$ from z_k to z'_k *in* X_k. Given any finite subset $P \subseteq G$, we choose a finite subset $Q \subseteq G$ which contains all products of length $\leqslant 4$ in s, t, s^{-1}, t^{-1}, as s and t vary over P. We choose k large enough so that z_k, z'_k are moved by less than ε by each $\gamma \in Q$ with respect to the scaled metric $d_k = d/\lambda_k$. If $D > 3\varepsilon + (24\delta/\lambda_k)$ then if we omit segments of d-length $\lambda_k\varepsilon + 12\delta$ from the ends of $[z_k, z'_k]$, the remaining sub-segment is non-empty; call this segment C_k. We assume that ε is small enough to satisfy the above inequality; we shall place further restrictions on ε later.

Now we use the original metric d on X_k. From the proof 'slim $\Rightarrow$ thin' (see [Sho] p. 17), if $x \in C_k$ then γx is within 12δ of $[z_k, z'_k]$. We denote by $\gamma_* x$ the projection of γx on $[z_k, z'_k]$. Of course, the 'projection' is not uniquely defined, but the preceding sentence is true no matter which closest point on $[z_k, z'_k]$ one chooses — we fix a definite choice for each $x \in C_k$, thus defining a map $\gamma_* : C_k \to [z_k, z'_k]$ for each γ. Next, we omit segments of length $5(\lambda_k\varepsilon + 12\delta)$ from the ends of $[z_k, z'_k]$ and denote the remaining long segment by $E_k \subseteq C_k$. The map $C_k \to [z_k, z'_k]$ just defined restricts to a map $E_k \to [z_k, z'_k]$; we continue to denote this map by γ_*. Notice that this map is a 24δ-isometry, that is to say, it distorts distances by at most an additive constant of 24δ; in fact it is 24δ close to a translation of E_k along $[z_k, z'_k]$. (Here, and in what follows, the terminology η-*close* is used to describe functions f, g with the same domain such that $d(f(x), g(x)) < \eta$ for all points in their common domain.)

Note that on E_k the maps s_*, $s_* t_*$, $s_* t_* (s^{-1})_*$, $s_* t_* (s^{-1})_* (t^{-1})_*$ etc. are well-defined and uniformly close to translations. Choose $M = \mathrm{Max}\{5(\lambda_k\varepsilon + 12\delta), 600\delta\}$. We will denote by e_k the segment obtained from $[z_k, z'_k]$ by omitting segments of length M from the ends. We have $e_k \subset E_k$. To make sure that $e_k \neq \emptyset$ we assume $D - \varepsilon > 5\varepsilon + (60\delta/\lambda_k)$, we also assume $D - \varepsilon > (600\delta/\lambda_k)$. Since $\lambda_k \to \infty$, we can choose large enough k and small enough ε so that the above conditions are satisfied.

We shall consider the restrictions $\gamma_* : e_k \to C_k$ to e_k of the maps γ_* defined above; we retain the notation γ_* for these restricted maps. Our goal is to obtain a bound (independent of $|Q|$) on the number commutators $tst^{-1}s^{-1}$ in Q by estimating how close the action of such a commutator on e_k is to the identity map. We first compare $t_* s_* (t^{-1})_* (s^{-1})_*$ to $tst^{-1}s^{-1}$. Observe that, since the maps s and s_* are 12δ close, ts and $t(s_*)$ are 12δ close (the left-action of Γ on X_k is by isometries in the metric d). Hence, $(ts)_*$ and $t_* s_*$ are 36δ close. Comparing successively $tst^{-1}s^{-1}$, $(tst^{-1}s^{-1})_*$, $(ts)_*(t^{-1}s^{-1})_*$, $t_* s_* (t^{-1})_* (s^{-1})_*$ shows that $tst^{-1}s^{-1}$ and $t_* s_* (t^{-1})_* (s^{-1})_*$ are $(12 + 36 + 108)\delta$ close.

Next, we compare $t_* s_* (t^{-1})_* (s^{-1})_*$ to the identity map on e_k. Since $s_* (t^{-1})_*$ and $(t^{-1})_* s_*$ are 72δ close to the same translation, and translations commute, we have that $t_* s_* (t^{-1})_* (s^{-1})_*$ and $t_* (t^{-1})_* s_* (s^{-1})_*$ are $(144+24)\delta$ close. Moreover, $t_* (t^{-1})_*$ and $s_* (s^{-1})_*$ are 36δ close to the identity. Thus $t_* (t^{-1})_* s_* (s^{-1})_*$ is 108δ close to the identity. Hence $t_* s_* (t^{-1})_* (s^{-1})_*$ is 276δ close to the identity. Combining this with the estimate in the previous paragraph we have that the restriction of $(tst^{-1}s^{-1})$ to e_k is 532δ close to the identity on e_k. Therefore, a vertex close to the midpoint of e_k is moved by less than $532\delta + 2$ by $tst^{-1}s^{-1}$. Thus $tst^{-1}s^{-1}$ lies in the ball of radius $532\delta + 2$ about the identity in Γ, and we have the desired bound on the number of commutators in the arbitrary finite subset $P \subset G$.

Now Lemma 2.7 implies that G is virtually abelian. But every abelian subgroup of a hyperbolic group is virtually cyclic. Hence the segment stabilizers for the action of Γ on X_∞ must be virtually cyclic. This completes the proof of Paulin's theorem. □

SECTION 3: CONVEX HULLS

A subset Σ of a geodesic metric space X is said to be *geodesically convex* if for all $p, q \in \Sigma$ every geodesic segment from p to q is completely contained in Σ. Given a bounded set $Y \subset X$, perhaps the most natural way to define its convex hull is as the intersection of all geodesically convex sets containing Y.

If X is simply connected and non-positively curved then round balls are geodesically convex and hence the convex hull of a bounded set is bounded. However, for more general geodesic metric spaces, even δ-hyperbolic spaces, it may happen that the convex hull of a finite set is the whole of the ambient space X. The following example illustrates how general this problem is.

3.1 PROPOSITION. *Given any finitely generated group Γ there exists a finite generating set S and a finite subset $Y \subset \Gamma$ such that the convex hull of Y in the Cayley graph $X(\Gamma, S)$ is the whole of $X(\Gamma, S)$.*

Proof. Let A be any finite generating set for Γ, and take S to be the set of those elements of Γ which are a distance 1 or 4 from the identity in the Cayley graph of Γ with respect to S. Let Y be the set of elements of Γ which are a distance at most 3 away from the identity in the Cayley graph associated to A.

Notice that the convex hull of Y with respect to S contains the ball of radius 4 as measured in the A metric. Furthermore, a simple induction shows that if this convex hull contains the balls of radius n and $n + 3$ about the identity (as measured in the metric associated to A) then it contains the ball of radius $n + 4$. Thus the convex hull of Y is the whole of $X(\Gamma, S)$. □

SECTION 4: CONCLUDING REMARKS

The type of limit spaces which we considered in Section 2 first arose in work of Morgan and Shalen in which they reinterpreted and generalized Thurston's compactification of Teichmuller space (see [Sha]). The particular topology with respect to which limits are taken in that setting is equivalent to what Paulin has termed "Equivariant Gromov convergence" (see [P1, 2]). It can be shown that the limit tree which we constructed in Section 2 is also a limit in the sense of this topology. We recall Paulin's recent definition:

4.1 DEFINITION. *A sequence of metric spaces* Y_n *which are equipped with actions by isometries of a fixed group* Γ, *converge to a metric space* Y, *which is also equipped with an action of* Γ *by isometries, if and only if, given any finite set* $K \subset Y$, *any* $\varepsilon > 0$, *and any finite subset* $P \subset \Gamma$, *for sufficiently large* n, *one can find subsets* $K_n \subset Y_n$ *and bijections* $x_n \mapsto x$ *from* K_n *to* K, *such that*

$$|d(\gamma x, y) - d_n(\gamma x_n, y_n)| < \varepsilon$$

for all $x, y \in K$ *and all* $\gamma \in P$.

Limits are not unique in this topology, even if one allows only limit spaces which are complete (cf. [P2], p. 55).

The technique of Equivariant Gromov convergence has been successfully applied in the following settings:

(1) $Y_n = \mathbf{H}^m$ for every integer n and the action of the (abstract) group Γ is discrete and varies with n;

(2) the spaces Y_n are $\mathbf{R}$-trees with isometric Γ-actions;

(3) each Y_n is equal to the Cayley graph of Γ with respect to a fixed set of generators and the action of Γ is left-multiplication twisted by a sequence of homomorphisms $\varphi_n : \Gamma \to \Gamma$.

The situation which we considered in Section 2 belongs to the third of the above cases. In the first two cases, the spaces under consideration enjoy strong convexity properties that allow one to form compact convex hulls of any finite set of points, and under suitable assumptions on the actions being considered one can show that such convex hulls satisfy uniform compactness conditions. The example of the previous section shows that case (3) is less hospitable to analysis of this type. But the arguments which we presented in Section 2 show that this difficulty can be accommodated by using suitable quasi-convex hulls to imitate the more familiar convex hulls used in cases (1) and (2).

If one is interested purely in elucidating the structure of the group Γ under consideration, then arguments using Hausdorff-Gromov convergence, as in Section 2, provide a direct method for constructing a limit object (X_∞, x_∞) which can be used to study properties of Γ. If, on the other hand, one is interested in some kind of representation space for Γ, then it is more useful to formulate results in terms of the above notion of equivariant Gromov convergence. If one is working in a situation where there is an *a priori* different topology to be considered, then one is left with the task of showing that the above notion of equivariant Gromov convergence agrees with the notion of convergence in this other topology. Such verifications have been successfully carried out for various discrete actions of non-elementary hyperbolic groups Γ by Bestvina [B] and Paulin [P1-3]. In particular Paulin has shown that the above notion of convergence leads to the same topology on small actions of a fixed group on **R**-trees $SLF(\Gamma)$ as the more familiar topology given by length functions.

Somewhat surprisingly, similar Hausdorff-Gromov type arguments do not seem to work so well without some assumption of smallness on the actions considered. For example, the space of all length functions $LF(\Gamma)$ on a fixed group Γ does not seem so amenable to such an analysis. The problem appears to lie with the absence in this generality of any assumption to play the role which the Margulis Lemma plays in the case of discrete actions on $\mathbf{H}^n$. One final remark about length functions: an important subspace of $SLF(\Gamma)$ is $VSLF(\Gamma)$, the space of very small actions introduced by Cohen and Lustig [CL]. Compactness of $VSLF(\Gamma)$ (a result due to Cohen and Lustig) can be proved using equivariant Gromov convergence in the same manner as one shows compactness of $SLF(\Gamma)$; see Paulin [P5].

Rips and Sela ([RS], [S1], [S2]) have made extensive use of variations on the arguments in Section 2 above to study hyperbolic groups. Most strikingly, from Paulin's theorem and deep work of Rips (see [RS], [BF]) one obtains an analogue for hyperbolic groups of the annulus theorem from 3-dimensional topology.

4.2 (Cylindrical Splitting Theorem).

If Γ *is a hyperbolic group with* $Out(\Gamma)$ *infinite, and if* Γ *has one end, then* Γ *is either an HNN-extension or an amalgamated free product over a virtually infinite cyclic group.*

Hyperbolic groups which do not admit a splitting over a virtually cyclic group have been termed *rigid* by Rips and Sela. They show, by a variant of the argument in Section 2 above (termed the Bestvina-Paulin method by Sela) that rigid hyperbolic groups are co-hopf. If Γ is torsion-free and rigid, then they show that there are only finitely many conjugacy classes of embeddings of Γ into any hyperbolic group. Sela [S2] has begun to investigate hopficity for rigid hyperbolic groups. Thus the techniques which we have attempted to exemplify in Section 2 appear to provide an extremely useful tool in the study of hyperbolic groups.

Acknowledgement: This paper arose from discussions at the special semester in Geometric Group Theory at CRM Barcelona in 1992. The writing of it was completed while the first author was visiting the University of Geneva, and the second author was visiting ETH, Zurich. We thank all of these organisations for their support. We would also like to thank Frédéric Paulin for his helpful comments on an earlier version of this article.

REFERENCES

[B] Bestvina, M. Degeneration of hyperbolic space. *Duke Math. Journal 56* (1988), 143-161.

[BF] Bestvina, M. and M. Feighn. Stable actions of groups on real trees. Preprint, 1992.

[Bo] Bowditch, B. Notes on Gromov's hyperbolicity criterion for path-metric spaces. *Group theory from a geometrical viewpoint* (E. Ghys, A. Haefliger, A. Verjovsky, eds.). World Scientific, Singapore, 1991.

[CEG] Canary, R., D.B.A. Epstein and L. Green. Notes on notes of Thurston. *Analytical and Geometric aspects of hyperbolic space.* L.M.S. Lectures Notes 111, CUP, 1992, 3-92.

[CL] Cohen, M. and M. Lustig. Very small actions on **R**-trees and Dehn twist automorphisms. *Topology,* to appear.

[GH] Ghys, E. and P. de la Harpe (editors). *Sur les groupes hyperboliques d'après Mikhael Gromov.* Birkhäuser, Prog. in Mathematics series, 1990.

[G1] Gromov, M. Groups of polynomial growth and expanding maps. *Publ. Math. Inst. Hautes Etud. Sci. 53* (1981), 53-78.

[G2] —— Hyperbolic groups, in *Essays in group theory,* S.M. Gersten (ed.). Springer Verlag, *MSRI series 8* (1987), 75-86.

[GLP] Gromov, M., J. Lafontaine and P. Pansu. *Structures métriques pour les variétés riemanniennes.* Cedic / Fernand Nathan, Paris 1981.

[H] HODGSON, C. Notes on Thurston's orbifold theorem. Princeton Univ. preprint (1987).

[P1] PAULIN, F. *Topologie de Gromov équivariante, structures hyperboliques et arbres réels.* Thesis, Université de Paris Sud, Orsay 1987.

P2] —— Topologie de Gromov équivariante, structures hyperboliques et arbres réels. *Invent. Math. 94* (1988), 53-80.

[P3] —— The Gromov topology on **R**-trees. *Topology and its applications 32* (1989), 197-221.

[P4] —— Outer automorphisms of hyperbolic groups and small actions on **R**-trees. *Arboreal Group Theory* (R. Alperin, ed.). Springer-Verlag MSRI series 19, 1991, 331-343.

[P5] —— Sur les automorphismes extérieurs des groupes hyperboliques. Preprint ENS de Lyon, 1993.

[RS] RIPS, E. and Z. SELA. Structure and rigidity of hyperbolic groups I. *GAFA*, to appear.

[S1] SELA, Z. Structure and rigidity of hyperbolic groups II. Preprint, Princeton 1993.

[S2] —— Acylindrical accessibility for groups. Preprint, Princeton 1993.

[Sha] SHALEN, P. Dendrology and its applications. *Group theory from a geometrical viewpoint* (E. Ghys, A. Haefliger, A. Verjovsky, eds.). World Scientific, Singapore, 1991, 543-616.

[Sho] SHORT, H. (ed.). Notes on word hyperbolic groups. *Group theory from a geometrical viewpoint* (E. Ghys. A. Haefliger, A. Verjovsky, eds.). World Scientific, Singapore, 1991, 3-63.

[T] THURSTON, W.P. Hyperbolic structures on 3-manifolds I: the acylindrical case. *Ann. of math. 124* (1986), 203-246.

(Reçu le 10 mars 1994)

M. R. Bridson

Department of Mathematics
Princeton University
Princeton, NJ 08540, USA
E-mail address: bridson@math.princeton.edu

G. A. Swarup

Department of Mathematics
University of Melbourne
Parkville 3052, Australia
E-mail address: gadde@mundoe.maths.mu.oz.au

vide-leer-empty