

Section 1 : Hausdorff-Gromov Convergence

Objektyp: **Chapter**

Zeitschrift: **L'Enseignement Mathématique**

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

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

PDF erstellt am: **23.05.2024**

Nutzungsbedingungen

Die ETH-Bibliothek ist Anbieterin der digitalisierten Zeitschriften. Sie besitzt keine Urheberrechte an den Inhalten der Zeitschriften. Die Rechte liegen in der Regel bei den Herausgebern.

Die auf der Plattform e-periodica veröffentlichten Dokumente stehen für nicht-kommerzielle Zwecke in Lehre und Forschung sowie für die private Nutzung frei zur Verfügung. Einzelne Dateien oder Ausdrucke aus diesem Angebot können zusammen mit diesen Nutzungsbedingungen und den korrekten Herkunftsbezeichnungen weitergegeben werden.

Das Veröffentlichen von Bildern in Print- und Online-Publikationen ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. Die systematische Speicherung von Teilen des elektronischen Angebots auf anderen Servern bedarf ebenfalls des schriftlichen Einverständnisses der Rechteinhaber.

Haftungsausschluss

Alle Angaben erfolgen ohne Gewähr für Vollständigkeit oder Richtigkeit. Es wird keine Haftung übernommen für Schäden durch die Verwendung von Informationen aus diesem Online-Angebot oder durch das Fehlen von Informationen. Dies gilt auch für Inhalte Dritter, die über dieses Angebot zugänglich sind.

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 $\mathcal{C}(X)$ 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 $\mathcal{C}(X)$ 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 $\mathcal{C}(X)$ equipped with this metric is compact.*

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

Consider a sequence C_i in $\mathcal{C}(X)$. 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 $A \in \mathcal{C}(X)$ 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 $j = 1, \dots, N(1/n)$. 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 $j \leq N(1/n)$, 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 $n \in \mathbb{N}$ the balls of radius $1/n$ about the first $N(1/n)$ terms in the sequence cover C_i .

$$C_1 \ni x(1, 1), x(1, 2), \dots, x(1, j), \dots$$

$$C_2 \ni x(2, 1), x(2, 2), \dots, x(2, j), \dots$$

...

.

$$C_i \ni x(i, 1), x(i, 2), \dots, x(i, j), \dots$$

.

.

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 $j = 1, \dots, k$ 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 $\mathcal{C}(X)$. 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 $\Sigma(k, n) := \{x(k, j)\}_{j \leq N(1/n)}$, for all l and k we have that the D -distance from $\{x_{\omega, l}\}$ to $\Sigma(\omega, n) := \{x(\omega, j)\}_{j \leq N(1/n)}$ 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 $d(x(k, j), x(\omega, j)) \leq 1/n$ for all $j \leq N(1/n)$, we have:

$$\begin{aligned} D(C_k, C_\omega) &\leq D(C_k, \Sigma(k, n)) + D(\Sigma(k, n), \Sigma(\omega, n)) \\ &\quad + D(\Sigma(\omega, n), C_\omega) \leq 3/n. \end{aligned}$$

□

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 $\mathcal{C}(X)$ converges to $C \in \mathcal{C}(X)$ 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 $x_n \rightarrow x$ in X ;*
- (2) *every sequence $y_{n(i)} \in C_{n(i)}$ with $n(i) \rightarrow \infty$ 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 $\mathcal{C}(X)$ 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 $z_m \rightarrow z_\omega \in C$, and then $D(z_\omega, C_n) \geq \varepsilon$ for arbitrarily large values of n , thus (1) fails. \square

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 $C_n \rightarrow C$, if and only if $D_H(C_n, C) \rightarrow 0$ as $n \rightarrow \infty$. 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 $\mathcal{C}(X)$, and the Hausdorff distance between A_1 and A_2 is defined to be $D_h(A_1, A_2) := \inf_X \{D_{\mathcal{C}(X)}(A_1, A_2)\}$. 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 $b_{m,1} \rightarrow b_1$ in B . By passing to a further subsequence we may assume that $b_{m,2} \rightarrow b_2$, 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 $A \rightarrow B$ by taking the unique continuous extension of $a_n \mapsto b_n$. \square

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 $\{C_i\}_{i \in \mathbb{N}}$ 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 $d(x(i, j), x(i, j'))_{i \in \mathbb{N}}$ converges; $x(\omega, j)$ should then be defined to be a certain sequence $\{x(i, j)\}_{i \in \mathbb{N}}$; 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 $\{C_i\}_{i \in \mathbb{N}}$ 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 $C_i \rightarrow C$ and ε_i -relations $R_i \subseteq C_i \times C$ with $\varepsilon_i \rightarrow 0$, one says that the sequence $\{x_i\}_{i \in \mathbb{N}}$, $x_i \in C_i$ converges to $x_\infty \in C$ if $x_i R_i x_\infty$. Under these circumstances, we write $x_i \rightarrow x_\infty$ 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 $\{X_i\}_{i \in \mathbb{N}}$ be a sequence of proper metric spaces with basepoints $x_i \in X_i$. The sequence of pointed spaces $\{(X_i; x_i)\}_{i \in \mathbb{N}}$ is said to converge to $(X; x)$ if for every $r > 0$ the sequence of compact metric balls $\{B(x_i, r)\}_{i \in \mathbb{N}}$ 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 $(\mathbb{E}^m; 0)$. 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 $D_H(S_n, \mathbb{E}^m)$ is infinite for all n .

Gromov’s compactness criterion for sequences of compact spaces implies that if for every $r > 0$ the balls $\{B(x_i, r)\}_{i \in \mathbb{N}}$ are uniformly compact, then $\{(X_i; x_i)\}_{i \in \mathbb{N}}$ 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 (**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 \mathbf{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 $\{C_i\}_{i \in \mathbb{N}}$ 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 $\delta_i \rightarrow 0$, then C is an \mathbf{R} -tree.*

Proof. Let R_i be an ε_i -approximation between C_i and C with $\varepsilon_i \rightarrow 0$. 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 $w_i: I_i \rightarrow [x_i, y_i]$ 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 $d(x_i, y_i) \rightarrow d(x, y)$ and $I_i \rightarrow I_\infty = [0, d(x, y)]$. 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 $\{K_i\}_{i \in \mathbb{N}}$, which we still denote by K_i , converges in the Hausdorff metric, to $K \subset C$ say. But $d_H(I_\infty, K_i) \leq d_H(I_\infty, I_i) + d_H(I_i, K_i)$, which goes to 0 as $i \rightarrow \infty$. 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 $I_\infty \rightarrow K$ gives a geodesic joining x and y . This proves assertion (1).

To prove the second part of the proposition, we first show that if $\delta_i \rightarrow 0$, 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 $i \rightarrow \infty$, 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 $d(z_i, z'_i) \leq d(z'_i, y'_i) + d(y'_i, z_i) \leq 4\delta_i + d(y'_i, z_i)$ 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. \square

Remark. If one has a sequence of δ_i -hyperbolic spaces C_i , with $C_i \rightarrow C$ and $\delta_i \rightarrow \delta > 0$, 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 by 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