

RELATIVELY HYPERBOLIC GROUPS

B. H. BOWDITCH

ABSTRACT. In this paper we develop some of the foundations of the theory of relatively hyperbolic groups as originally formulated by Gromov. We prove the equivalence of two definitions of this notion. One is essentially that of a group admitting a properly discontinuous geometrically finite action on a proper hyperbolic space, that is, such that every limit point is either a conical limit point or a bounded parabolic point. The other is that of a group which admits a cofinite action on a connected fine hyperbolic graph. We define a graph to be “fine” if there are only finitely many circuits a given length containing any given edge, and we develop some of the properties of this notion. We show how a relatively hyperbolic group can be assumed to act on a proper hyperbolic space of a particular geometric form. We define the boundary of a relatively hyperbolic group, and show that the limit set of any geometrically finite action of the group is equivariantly homeomorphic to this boundary. This generalises a result of Tukia for geometrically finite kleinian groups. We also describe when the boundary is connected.

1. INTRODUCTION

The aim of this paper is to give an account of the notion of a “relatively hyperbolic group”, or more precisely, a group which is “hyperbolic relative to” a preferred class of “peripheral subgroups”. This notion was introduced by Gromov in his original paper on hyperbolic groups, [Grom]. Since then, it has been a major feature of geometric group theory, with many naturally occurring groups falling into this class. It simultaneously generalises the idea of a hyperbolic group (the case where there are no peripheral subgroups) and that of a geometrically finite kleinian group (where the peripheral subgroups are the maximal parabolic subgroups). Elaborations on the notion of relative hyperbolicity were given by Farb [Fa] and Szczepański [Sz]. In this paper, we describe some alternative viewpoints. Since the first draft of this paper was written there have been numerous papers making use of relative hyperbolicity, and one could compile an extensive

Date: September 2011.

bibliography. Here, however, we will confine ourselves to discussing equivalent formulations of the notion. In that regard, we mention [Y, Da, Bu, DrS, O, GrovM, Gera]. We will say more about these later in the introduction. Of course, these papers include many other references to the subject.

We shall give two equivalent definitions of hyperbolicity for a group, Γ , relative to a set, \mathcal{G} , of infinite subgroups.

Definition 1. We say that Γ is *hyperbolic relative to \mathcal{G}* , if Γ admits a properly discontinuous isometric action on a path-metric space, X , with the following properties.

- (1) X is proper (i.e. complete and locally compact) and hyperbolic,
- (2) every point of the boundary of X is either a conical limit point or a bounded parabolic point,
- (3) the elements of \mathcal{G} are precisely the maximal parabolic subgroups of Γ , and
- (4) every element of \mathcal{G} is finitely generated.

Definition 2. We say that Γ is *hyperbolic relative to \mathcal{G}* , if Γ admits an action on a connected graph, K , with the following properties:

- (1) K is hyperbolic, and each edge of K is contained in only finitely many circuits of length n for any given integer, n ,
- (2) there are finitely many Γ -orbits of edges, and each edge stabiliser is finite,
- (3) the elements of \mathcal{G} are precisely the infinite vertex stabilisers of K , and
- (4) every element of \mathcal{G} is finitely generated.

The equivalence of these definitions will be proven in this paper (see Theorem 7.10). We refer to the elements of \mathcal{G} as “peripheral subgroups”. The insistence that they be finitely generated is rather artificial, though it seems to be needed for certain results.

The first definition could be viewed as a dynamical characterisation, except that we need to assume that we are dealing with the boundary of a proper hyperbolic space, and that the action arises from an isometric action on that space. In fact, it has been shown in [Y] that these last assumptions can be dispensed with. One therefore obtains a purely dynamical characterisation of a relatively hyperbolic group as a convergence group acting on a compact metrisable space, for which every point is either a conical limit point or a bounded parabolic point. This generalises the result of [Bo6], where it was assumed that there are no parabolics. (The argument of [Y] makes use of the fact, proven in [T3], that there are only finitely many conjugacy classes of maximal parabolic subgroups.)

The second definition (which we shall study first in this paper) will be phrased formally in terms of group actions on sets, which are “cofinite” (i.e. finitely many orbits). We keep much of the discussion fairly general. In principle, one could take a similar approach to studying other “geometric” properties of groups relative to a given set of subgroups. The standard (non-relative) case corresponds to those actions where all point stabilisers are finite. In this case, the interesting geometric properties seem to be quasiisometry invariant, though, in the more general setting, there are natural geometric properties which are not quasiisometry invariant, such as the property we call “finess” in Section 4. This approach is extended to group splittings in [Bo9].

Each of the various formulations of relative hyperbolicity have technical advantages in specific contexts, and each make apparent certain features not apparent from the others. We aim to explore such matters in this paper.

There is also a weaker notion, which we shall refer to here as “weak relative hyperbolicity”. This is a weakening of Definition 2 above. Let \mathcal{G} be a conjugacy invariant collection of infinite subgroups of a group Γ .

Definition. We say that Γ is *weakly hyperbolic relative to \mathcal{G}* if it admits an action on a connected graph, K , with the following properties:

- (1) K is hyperbolic,
- (2) there are finitely many Γ -orbits of edges, and
- (3) each element of \mathcal{G} fixes a vertex of K , and each vertex stabiliser of K contains an element of \mathcal{G} as a subgroup of finite index.

This gives rise to many more examples (see [Fa, Gers, Mas]). However, there seems to be a greater degree of arbitrariness about this definition. In this paper, we shall focus on the stronger definition. We briefly elaborate on these definitions, beginning with the dynamical formulation. Suppose that a group, Γ , acts by homeomorphism on a perfect metrisable compactum, M . We say that Γ is a “convergence group” (in the sense of [GehM]) if the induced action on the space of distinct triples is properly discontinuous. (For further discussion, see [GehM, T2, Fr1, Bo7].) We say that a convergence group is “(minimal) geometrically finite” if every point of M is either a conical limit point or a bounded parabolic point. (For definitions, see Section 6.) This is a natural generalisation of the Beardon-Maskit definition [BeM]. It has also been considered in this generality in [T3] and [Fr2]. It generalises the classical formulation of geometrical finiteness described by Ahlfors and Greenberg [A, Gre], see [Bo2, Bo4].

In this paper, we shall mostly confine ourselves to the case where M is the boundary, ∂X , of a complete, locally compact, Gromov hyperbolic space, X . If Γ acts properly discontinuously and isometrically on X , then the induced action on ∂X is a convergence action. We can thus define an action on X to be “geometrically finite” if the induced action on ∂X is geometrically finite in the dynamical sense. If $\Pi \subseteq \partial X$ is the set of parabolic points, then one can show that Π/Γ is finite. In fact, one can construct a strictly invariant system of horoballs, $(B(p))_{p \in \Pi}$, about Π , such that $(X \setminus \bigcup_{p \in \Pi} \text{int } B(p))/\Gamma$ is compact. In this way, we recover a generalisation of Marden’s definition of geometrical finiteness [Mar].

This gives us the first definition of relative hyperbolicity, namely that a group Γ is hyperbolic relative to a set of infinite finitely generated subgroups (the “peripheral subgroups”) if it admits a geometrically finite action on a proper hyperbolic space, where these subgroups are precisely the maximal parabolic subgroups. We note immediately that the intersection of any two peripheral subgroups is finite, and that there are finitely many conjugacy classes of peripheral subgroups. Moreover, each peripheral subgroup is equal to its normaliser.

Note that, in the case of a hyperbolic group (where $\mathcal{G} = \emptyset$) the quasiisometry class of the space X depends only on the group, Γ . This allows us to define the boundary, denoted $\partial\Gamma$, of Γ as ∂X . In the general situation, however, this is not the case. Indeed, there seems to be no particularly natural choice of quasiisometry class of space X . However, as we shall discuss in a moment, we can associate to a relatively hyperbolic group, (Γ, \mathcal{G}) , a space $X(\Gamma)$, together with a geometrically finite action of Γ , which is canonical up to Γ -equivariant quasiisometry, even though the construction is somewhat artificial. This allows us to give a formal definition of the boundary, $\partial(\Gamma, \mathcal{G})$ as $\partial X(\Gamma)$. One can go on to show that if X is any complete locally compact hyperbolic space admitting a geometrically finite action of Γ , then the limit set of ∂X is, in fact, Γ -equivariantly homeomorphic to $\partial(\Gamma, \mathcal{G})$. We thus see that $\partial(\Gamma, \mathcal{G})$ is a quite natural object to associate to (Γ, \mathcal{G}) , even if $X(\Gamma)$ is not. We frequently view \mathcal{G} as an intrinsic part of the structure of Γ , and abbreviate $\partial(\Gamma, \mathcal{G})$ to $\partial\Gamma$.

The proofs of the assertions of the last paragraph make use of the second formulation of relative hyperbolicity. We shall phrase this in terms of group actions on sets. To relate this to what we have said earlier, one can take this set to be the peripheral structure, with the group acting by conjugation.

Suppose we fix a group, Γ . A “ Γ -set” is simply a set, V , on which Γ acts. We refer to points of V as “vertices”. If $x \in V$, we write $\Gamma(x)$ for the vertex stabiliser. A “pair stabiliser” is a subgroup of the form $\Gamma(x) \cap \Gamma(y)$ for distinct $x, y \in V$. We shall usually assume that pair stabilisers are finite. A “ (Γ, V) -graph”, K , is a connected Γ -invariant graph with vertex set V , and with finitely many Γ -orbits of edges. Clearly any graph-theoretical property which is independent of the choice of K can be viewed as a property of a Γ -set, V . Since any two (Γ, V) -graphs are clearly quasiisometric, any quasiisometry invariant property, such as hyperbolicity, can be viewed as a property of a Γ -set. In view of this, we may define a Γ -set, V , to be “weakly hyperbolic” if some, and hence any, (Γ, V) -graph is hyperbolic.

This notion of weak hyperbolicity is equivalent to the formulation of relative hyperbolicity formulated in [Fa] (without the assumption of finite pair stabilisers) and used, for example, in [Mas] and [Gers]. However, it is weaker than that in which we are primarily interested here. For this we shall need an additional combinatorial hypothesis, which we shall call “finesness”. This is not a quasiisometry invariant. However, it turns out to be independent of the choice of (Γ, V) -graph.

In the same way that hyperbolicity can be viewed as a geometric weakening of certain properties of a simplicial tree (such as 0-thin triangles), the notion of finesness can be viewed as weakening of combinatorial properties. Thus, for example, a simplicial tree might be defined as a connected graph which contains no circuits. We say that a graph is “fine” if there are only finitely many circuits of a given length containing any given edge. This seems to be a fairly robust notion, and has several equivalent formulations, as we shall see in Section 2. We can thus define a Γ -set, V , to be “fine” if some (hence any) (Γ, V) -graph is fine.

We shall say that a Γ set, V , is “hyperbolic” if it has finite pair stabilisers, finitely generated point stabilisers, and is fine and weakly hyperbolic. This reduces to essentially two cases. If every vertex of V has finite stabiliser, then this is equivalent to saying that Γ is word hyperbolic in the sense of Gromov. On the other hand, if there exist vertices of infinite degree, then we can assume that every vertex has infinite degree. In the latter case, we get an equivalent definition of relatively hyperbolic group — if \mathcal{G} is a peripheral structure on Γ , then we can view \mathcal{G} as a Γ -set, where the action is by conjugation, and we can define (Γ, \mathcal{G}) to be hyperbolic if \mathcal{G} is hyperbolic as Γ -set. A more transparent reformulation of this definition is given by Proposition 4.9.

To relate this to the previous formulation of relative hyperbolicity, we use the complex given by Proposition 4.9. This is a Γ -invariant 2-dimensional simplicial complex, K , with 0-skeleton V , which has finite quotient under Γ , is locally finite away from V , and is “simplicially hyperbolic”. The last statement means that any cycle of length n in the 1-skeleton bounds a simplicial disc in K with the number of 2-simplices bounded above by a linear function of n . We may realise $K \setminus V$ as a complete locally compact path-metric space, $X(K)$, by giving each 2-simplex the structure of an ideal hyperbolic triangle, in such a way that the union of two adjacent triangles is isometric to an ideal hyperbolic square. It turns out that $X(K)$ is Gromov hyperbolic, and that the action of Γ on $X(K)$ is geometrically finite. This gives us our space, $X(\Gamma)$, referred to earlier.

Conversely, suppose we have a geometrically finite action of a group Γ on a complete locally compact hyperbolic space X . Let \mathcal{G} be the peripheral structure. We can reconstruct a fine hyperbolic graph, K , with vertex set \mathcal{G} , as a kind of “nerve” to an invariant set of horoballs. More precisely, we take any invariant set of horoballs for Γ , and we connect two peripheral subgroups by an edge if the distance between the corresponding horoballs is less than some sufficiently large constant. (Alternatively, if we choose the horoballs to overlap sufficiently, we can take it to be a nerve in the usual sense.) One verifies that K is indeed fine and hyperbolic, and that $\partial X(K)$ is Γ -equivariantly homeomorphic to ∂X .

There is a slight complicating factor, which we have alluded to. This concerns whether we should assume that the peripheral subgroups are all finitely generated. This is not needed in the discussion of geometrically finite groups in Section 6, but seems to be required, or at least desirable, for the constructions of Sections 3 and 4. We shall therefore take it as part of the hypotheses of a relatively hyperbolic group, but not of a geometrically finite group. This necessitates an extra clause in the statement of certain results (for example Proposition 7.9) which has no real mathematical content.

We should mention some other, and more recent, formulations of relative hyperbolicity. To set this in context we recall the case of a finite volume hyperbolic manifold (that is locally modelled on classical hyperbolic space). Such a manifold has finitely many (possibly zero) ends, each an exponentially tapering cusp. The manifold is therefore quasi-isometric to a finite wedge of rays (semi-infinite real intervals) connected at their respective endpoints. (One can also allow for convex boundaries — the geometrically finite case, or pinched negative

curvature.) This motivates Gromov’s original definition [Grom] — essentially a group acting properly on a proper hyperbolic space, whose quotient is quasi-isometric to a wedge of rays, each ray arising as the quotient of a horoball. Therefore the fundamental group of a finite volume (or geometrically finite) manifold is a basic example of a relatively hyperbolic group.

A different definition was given by Farb [Fa]. There a group is relatively hyperbolic if its “coned-off” Cayley graph is Gromov hyperbolic, and it satisfies a “bounded coset penetration” property (which plays the role of fineness in the present paper). This is also discussed in [Sz]. In this paper we give the two further definitions stated above. All the definitions mentioned are known to be equivalent, see [Da, Bu]. (In this discussion we are assuming peripheral subgroups to be finitely generated.)

Since the first draft of this paper was written, many other equivalent definitions have been formulated. We have already noted the dynamical description of Yaman [Y] as a geometrically finite convergence group acting on a compact metrisable space. We also note that Gerasimov [Gera] has shown that a convergence group is geometrically finite if and only if its induced action on the space of distinct pairs is cocompact. Given that a convergence action can be defined as one that is properly discontinuous on the space of distinct triples, this ties in with the characterisation of hyperbolic groups given in [Bo6].

In [DrS], Druţu, Sapir and Osin describe relatively hyperbolic groups in terms of their asymptotic cones — namely every asymptotic cone of the group should “tree graded” with respect to the peripheral subgroups. Osin [O] gives a more combinatorial description — a group is relatively hyperbolic if it satisfies a linear isoperimetric inequality “relative to” its peripheral subgroups. (His monograph includes a detailed discussion of relative hyperbolicity.) Groves and Manning [GroM] give another way coning off Cayley graphs, different to that of Farb, to give yet another characterisation.

The outline of this paper is as follows. In Section 2, we consider a few combinatorial properties of infinite graphs, in particular, introducing the notion of a “fine” graph. In Section 3, we construct a proper hyperbolic space starting with a fine hyperbolic graph, thereby giving us a notion of boundary for such a graph. In Section 4, we study actions of a group, Γ , on a set, allowing us to define the notion of a “hyperbolic Γ -set” and hence of a relatively hyperbolic group. In Section 5, we briefly review hyperbolic length spaces. In Section 6, we study geometrically finite actions on proper hyperbolic spaces, largely from a dynamical

point of view. In Section 7, we recover fine hyperbolic graphs from geometrically finite actions, thereby showing that a geometrically finite group with finitely generated peripheral subgroups is relatively hyperbolic by the combinatorial definition of Section 8. In Section 8, we shall explore properties of fine hyperbolic graphs. In Section 9, we show that the limit set of a geometrically finite group depends, up to equivariant homeomorphism, only on the group and the collection of peripheral subgroups. In Section 10, we describe some basic facts about splittings of relatively hyperbolic groups over finite subgroups, and how these are observed in the topology of the boundary.

I would like to thank Pekka Tukia and David Epstein for their comments on this paper. I am also indebted to Andrzej Szczepański for his ideas on relatively hyperbolic groups.

2. GRAPH THEORY

The main aim of this section will be to introduce the notion “finess” of a graph. We develop some basic properties of fine graphs, and describe some simple graph theoretical operations that preserve fineness. Additional properties of fine hyperbolic graphs will be explored in Section 8.

Let K be a graph with vertex set $V(K)$ and edge set $E(K)$. We can think of K as a 1-dimensional simplicial complex. (We are not allowing loops or multiple edges.) We write $V_0(K)$ and $V_\infty(K)$ respectively for the sets of vertices of finite and infinite degree. A *path* of length n connecting $x, y \in V$ is a sequence, $x_0x_1 \dots x_n$ of vertices, with $x_0 = x$ and $x_n = y$, and with each x_i equal to or adjacent to x_{i+1} . It is an *arc* if the x_i are all distinct. A *cycle* is a closed path ($x_0 = x_n$), and a *circuit* is a cycle with all vertices distinct. We regard two cycles as the same if their vertices are cyclically permuted (i.e. $x_0x_1 \dots x_{n-1}x_0 = x_kx_{k+1} \dots x_k$ for all k). We frequently regard arcs and circuits as subgraphs of K . Two arcs are *independent* if they meet only in their initial points. We speak of two paths or subgraphs as *distinct* simply to mean that they are not identical. The *length* of a subgraph is the cardinality of its edge set.

We put a *path metric*, d_K , on $V(K)$, where $d_K(x, y)$ is the length of the shortest path in K connecting x to y . (We set $d_K(x, y) = \infty$ if there is no such path.) When we speak of a subset of $V(K)$ as being *bounded* we mean with respect to this metric.

If $A \subseteq V(K)$, we write $K \setminus A$ for the full subgraph of K with vertex set $V(K) \setminus A$, i.e. the graph obtained by deleting the vertices A together with all incident edges. If $x \in K$, we abbreviate $K \setminus \{x\}$ to $K \setminus x$. We say that K is *n -vertex-connected* if $K \setminus A$ is connected for any subset,

A , of $V(K)$, of cardinality strictly less than n . (We are breaking with tradition slightly in that we are deeming the complete graph on n vertices to be n -vertex connected.)

Let K be a connected graph. A *block* of K is a maximal 2-vertex-connected subgraph. Two blocks meet, if at all, in a single vertex. The *block tree* of K is the bipartite graph whose vertex set is defined as the abstract union of the set, V , of vertices of K and the set, \mathcal{B} , of blocks of K , where a vertex, x , is adjacent to a block, B , in the block tree if $x \in B$ in K . It's not hard to see that the block tree is indeed a simplicial tree (i.e. a connected graph with no circuits).

So far, everything we have said is standard elementary graph theory. We now move on to some less standard definitions.

Definition. A collection, \mathcal{L} , of subgraphs of K is *edge-finite* if $\{L \in \mathcal{L} \mid e \in E(L)\}$ is finite for each edge $e \in E(K)$.

Definition. A subset $A \subseteq V(K)$ is *locally finite* in K if every bounded subset of A is finite.

If $x \in V(K)$, we write $V_K(x) \subseteq V(K)$ for the set of vertices adjacent to x .

We are now ready for the main result of this section:

Proposition 2.1. *Let K be a graph. The following are equivalent:*

(F1): *For each $n \in \mathbb{N}$, the set of circuits in K of length n is edge-finite.*

(F2): *For all $x, y \in V(K)$ and $n \in \mathbb{N}$, the set of arcs of length n connecting x to y is finite.*

(F3): *For any $x, y \in V(K)$ and $n \in \mathbb{N}$, there does not exist an infinite collection of pairwise independent arcs of length n connecting x to y .*

(F4): *Suppose $x, y \in V(K)$ are any pair of distinct vertices, and $n \in \mathbb{N}$. If \mathcal{L} is any edge-finite collection of connected subgraphs of K of length n , each containing both x and y , then \mathcal{L} is finite.*

(F5): *For each vertex $x \in K$, the set $V_K(x)$ is locally finite in $K \setminus x$.*

Proof. We prove (F2) \Rightarrow (F1) \Rightarrow (F3) \Rightarrow (F2), (F2) \Rightarrow (F4) \Rightarrow (F3) and (F2) \Rightarrow (F5) \Rightarrow (F3).

(F2) \Rightarrow (F1): Suppose $e \in E(K)$ has endpoints $x, y \in V(K)$. Each circuit, γ , of length n containing e gives us an arc, $\gamma \setminus e$, of length $n - 1$ connecting x and y . There can only be finitely many such arcs, and hence only finitely many such circuits.

(F1) \Rightarrow (F3): Suppose, for contradiction, that there is an infinite collection, $(\beta_i)_{i \in \mathbb{N}}$, of pairwise independent arcs of length n connecting

the same pair of distinct points. Let e be any edge of β_0 . Now, for each i , $\beta_0 \cup \beta_i$ is a circuit of length $2n$ containing e , contradicting (F1).

(F3) \Rightarrow (F2): Suppose, for contradiction, that (F2) fails. Let $n \in \mathbb{N}$ be minimal such that there exist distinct points, $x, y \in V(K)$ and an infinite collection, $(\beta_i)_{i \in \mathbb{N}}$, of distinct arcs of length at most n connecting x to y .

We claim that only finitely many of the β_i contain any given point, $z \in V(K) \setminus \{x, y\}$. To see this, note that z divides any such arc, β_i , into two subarcs, β_i^- and β_i^+ , each of length at most $n - 1$, and connecting x to z and z to y respectively. From the minimality of n , we see that there are only finitely many possibilities for β_i^- and β_i^+ , and hence for β_i , as claimed. It follows, more generally, that if $A \subseteq V(K) \setminus \{x, y\}$ is any finite set, then $\{i \in \mathbb{N} \mid \beta_i \cap A \neq \emptyset\}$ is finite.

We can thus, inductively, pass to a subsequence, $(\beta_{i_j})_j$, such that β_{i_j} meets no point of $\bigcup_{k < j} V(\beta_{i_k}) \setminus \{x, y\}$. In other words, $(\beta_{i_j})_j$ is an infinite collection of pairwise independent arcs of length at most n , connecting x to y . This contradicts (F3).

(F2) \Rightarrow (F4): Suppose, for contradiction, that $(L_i)_{i \in \mathbb{N}}$ is an edge-finite collection of connected subgraphs of k , each of length n , and each containing a pair of distinct points, $x, y \in V(K)$. Now, we can inductively pass to a subsequence, $(L_{i_j})_j$ of subgraphs which are pairwise edge-disjoint (since only finitely many graphs, L_i , contain any of the finite set, $\bigcup_{k < j} E(L_{i_k})$, of edges of K). Now, let β_j be any arc in L_{i_j} connecting x to y . These arcs are all edge-disjoint, and so, in particular, are distinct. Moreover, they all have length at most n , contradicting (F2).

(F4) \Rightarrow (F3): Clearly, any collection of independent arcs connecting a pair of points will be edge-disjoint, and hence finite by (F4).

(F2) \Rightarrow (F5): Suppose, for contradiction, that $x \in V$, and that $V_K(x)$ is not locally finite in $K \setminus x$. In other words, we can find some $y \in V_K(x)$, and some $n \in \mathbb{N}$, such that there is an infinite sequence, $(z_i)_{i \in \mathbb{N}}$, of distinct points in $V_K(x)$, each a distance at most n from y . Let α_i be a shortest path in $K \setminus x$ from y to z_i , and let β_i be α_i concatenated with the edge $z_i x$. We see easily that β_i is an arc of length at most $n + 1$ connecting y to x . Since $z_i \in V(\beta_i)$, these arcs are all distinct, contradicting (F2).

(F5) \Rightarrow (F3) Suppose, for contradiction, that there exist points $x, y \in V$, and an infinite sequence, $(\beta_i)_{i \in \mathbb{N}}$ of pairwise independent arcs of length n connecting x to y . Let z_i be the first vertex of the arc β_i after x , and let α_i be the subarc of β_i connecting z_i to y . Thus, $z_i \in V_K(x)$, and α_i has length $n - 1$ and lies in $K \setminus x$. We see that any pair of points, z_i and z_j are connected by a path of length $2(n - 1)$ in $K \setminus x$. Thus, $\{z_i \mid i \in \mathbb{N}\}$ is an infinite subset of $V_K(x)$, which is bounded in $K \setminus x$, contradicting (F5). \square

Definition. We say that a graph is *fine* if it satisfies one, hence all, of the properties (F1)–(F5) featuring in Proposition 2.1

(We remark that, for the applications we have in mind, one could equally well deal with “uniformly fine” graphs, where we assume a bound in terms of n on the number circuits in (F1). This seems to make little difference to the general theory, so we shall confine ourselves here to the weaker notion.)

We begin by making some trivial observations about fineness. Firstly:

Lemma 2.2. *Any subgraph of a fine graph is fine. Any locally finite graph is fine. A graph is fine if and only if each of its components is fine. A connected graph is fine if and only if each of its blocks is fine.*

Another point to note is that a 2-vertex-connected fine graph is countable. This follows, since in this case, every vertex has countable degree by property (F5).

We want to describe some slightly less trivial operations on graphs which preserve fineness. To this end, the following lemma will be useful.

Let K be a graph. Given an arc, α , in K , we write $e(\alpha)$ for the unordered pair of endpoints of α . If \mathcal{A} is a set of arcs, we write $K[\mathcal{A}]$ for the graph with vertex set $V(K[\mathcal{A}]) = V(K)$, and edge set, $E(K[\mathcal{A}]) = E(K) \cup \{e(\alpha) \mid \alpha \in \mathcal{A}\}$. (This union need not be disjoint — if α has length 1, then $e(\alpha)$ is already in $E(K)$.)

Lemma 2.3. *Suppose that K is a fine graph, and \mathcal{A} an edge-finite collection of arcs of bounded length in K . Then, $K[\mathcal{A}]$ is fine.*

Proof. Let k be the maximal length of any arc in \mathcal{A} . Suppose that L is a connected subgraph of $K[\mathcal{A}]$ of length n . We may associate to L , a connected subgraph, L' , of K , with $V(L) \subseteq V(L')$, and of length at most kn , as follows. For each edge, $e \in E(L)$, we choose some $\alpha \in \mathcal{A}$ with $e = e(\alpha)$, and let L' be the union of all such arcs. (That is the subgraph consisting of all vertices and all edges of such arcs. If $E(L)$ happens to be empty, then we set $L' = L$.)

We claim that if \mathcal{L} is an edge-finite collection of connected subgraphs of $K[\mathcal{A}]$, then $\mathcal{L}' = \{L' \mid L \in \mathcal{L}\}$ is edge-finite in K . To see this, fix

some edge $e \in K$. Now, $\mathcal{A}(e) = \{\alpha \in \mathcal{A} \mid e \in E(\alpha)\}$ is finite, and so $E(e) = \{e(\alpha) \in E(K[\mathcal{A}]) \mid \alpha \in \mathcal{A}(e)\}$ is finite. Now, if $e \in E(L)$ for some $L \in \mathcal{L}$, we see that $E(e) \cap E(L) \neq \emptyset$. Since $E(e)$ is finite, and \mathcal{L} is edge-finite, we see that the set of such $L \in \mathcal{L}$ is finite. This shows that \mathcal{L}' is edge-finite as claimed.

Now, suppose that $x, y \in V(K[\mathcal{A}]) = V(K)$ are distinct, and that \mathcal{L} is an edge-finite collection of connected subgraphs of $K[\mathcal{A}]$ of bounded length containing x and y . We see that \mathcal{L}' is an edge-finite collection of subgraphs of K of bounded length containing x and y . By (F4), it follows that \mathcal{L}' is finite. But only finitely many graphs in \mathcal{L} can give rise to a given graph in \mathcal{L}' , so it follows that \mathcal{L} is finite. This verifies property (F4) for $K[\mathcal{A}]$. \square

More generally, suppose that \mathcal{L} is any edge-finite collection of subgraphs of bounded length. Let \mathcal{A} be the set of arcs which lie inside some graph in \mathcal{L} . Clearly this collection is also edge-finite. Let $K[\mathcal{L}] = K[\mathcal{A}]$. In other words, for every $L \in \mathcal{L}$, we span $V(L)$ by a complete graph. By Lemma 2.3, we get:

Lemma 2.4. *Suppose that K is fine, and that \mathcal{L} is an edge-finite collection of connected subgraphs of bounded length. Then, $K[\mathcal{L}]$ is fine.*

If $L \in \mathcal{L}$, we write $F(L)$ for the full subgraph of $K[\mathcal{L}]$ with vertex set $V(L)$. Thus, $F(L)$ is complete. We also see easily that the collection, $\{F(L) \mid L \in \mathcal{L}\}$ is edge finite (using property (F4)). Thus, for many purposes, we take an edge-finite collection of subgraphs of bounded length to consist entirely of complete subgraphs.

We shall use Lemma 2.4 in a number of constructions, three of which we describe below.

Let K be any graph, and $n \in \mathbb{N}$. We construct a graph K^n , with vertex set, $V(K^n) = V(K)$, by connecting distinct $x, y \in V(K^n)$ by an edge in K^n if and only if either $xy \in E(K)$ or x and y lie in some circuit in K of length at most n . Clearly, $K \subseteq K^n$.

If we write \mathcal{L} for the set of circuits of length at most n in K , we get that $K^n = K[\mathcal{L}]$. It thus follows from Lemma 2.4 that:

Lemma 2.5. *If K is fine, then K^n is fine for all $n \in \mathbb{N}$.*

The next operation we consider is that of “binary subdivision”. Suppose that K is a graph. If $e \in E(K)$, we can subdivide e into two edges by inserting an extra vertex at the “midpoint”, $m(e)$, of e . If we do this for each edge of K , we obtain a graph, L , whose realisation is identical to that of K . We now form a graph, K' , containing L , with vertex set $V(K') = V(L)$, by connecting two midpoints, $m(e_1)$ and

$m(e_2)$ by a new edge whenever there exists a third edge $e_3 \in E(K)$ such that $\{e_1, e_2, e_3\}$ forms a 3-circuit in K . We refer to K' as the *binary subdivision* of K .

Another way to express this construction is to define the “flag 2-complex”, $\Sigma(K)$, associated to a graph K . This is the 2-dimensional simplicial complex with 1-skeleton K , and with a set of three vertices of K spanning a simplex in $\Sigma(K)$ if and only if they form a 3-circuit in K . (Note that, if K is fine, then $\Sigma(K)$ is locally finite on the complement of the set of vertices.) Now, $\Sigma(K')$ is obtained from $\Sigma(K)$ by dividing each edge in two, and dividing each 2-simplex into four smaller simplices (that is, with one central simplex whose vertices are midpoints of the edges of the original, and the three others, each containing a vertex of the original). Note that the realisation of the flag 2-complex remains unchanged.

Lemma 2.6. *If K is fine, then so is the binary subdivision, K' .*

Proof. Let L be the graph obtained by dividing each edge of K in two. The circuits of L are of even length and correspond precisely to circuits of half their length in K . It follows that L is fine. Now, $K' \subseteq L^6$, and so, by Lemma 2.5, we see that K' is fine. \square

Suppose that $A \subseteq V(K)$ and $n \in \mathbb{N}$. We define a graph $K(A, n)$, with vertex set $V(K(A, n)) = A$, and with $x, y \in A$ joined by an edge in $K(A, n)$ if and only if there exists an arc, α , of length at most n in K , with $\alpha \cap A = \{x, y\}$. Recall that $V_\infty(K) \subseteq V(K)$ is the set of vertices of infinite degree in K .

Lemma 2.7. *Suppose that K is fine, $V_\infty(K) \subseteq A \subseteq V(K)$, and $n \in \mathbb{N}$. Then, $K(A, n)$ is fine.*

Proof. Let \mathcal{A} be the set of arcs of length at most n in K meeting A precisely in their endpoints. Since each vertex of $V(K) \setminus A$ has finite degree, we see easily that there are only finitely many possibilities for such arcs passing through any given edge. In other words, \mathcal{A} is edge-finite. Thus, by Lemma 2.4, $K([\mathcal{A}])$ is fine. Now, $K(A, n)$ is obtained from $K[\mathcal{A}]$ by deleting the isolated vertices $V(K) \setminus A$. It follows that $K(A, n)$ is fine. \square

In relation to the last construction, we should make a simple observation for future reference.

Definition. A subset, $A \subseteq V(K)$ is *r-dense* if every vertex of K is a distance at most r from some point of A .

The following is easily verified:

Lemma 2.8. *If K a connected graph, and $A \subseteq V(K)$ is r -dense, then $K(A, 2r + 1)$ is connected.*

We finish this section with another construction which preserves fineness. We first make a trivial observation:

Lemma 2.9. *Suppose that K is a graph and $A \subseteq V(K)$. Suppose that every point of A has degree 0 or 1 in K . If $K \setminus A$ is fine, then so is K .*

Proof. By property (F1), since every circuit of K lies in $K \setminus A$. \square

Suppose that $(L_i)_{i \in I}$ is a collection of subgraphs of a graph K , indexed by a set I . We say that $(L_i)_{i \in I}$ is “edge-finite” if the set $\{L_i \mid i \in I\}$ is edge-finite, and only finitely many L_i are equal to any given subgraph. In other words, $\{i \in I \mid e \in E(L_i)\}$ is finite for every edge $e \in E(K)$.

Lemma 2.10. *Suppose that K is a graph, $n \in \mathbb{N}$, and $A \subseteq V(K)$. Suppose that for each $x \in A$, there is a connected subgraph $L(x) \subseteq K \setminus A$, of length at most n , and containing every vertex adjacent to x in K . If $K \setminus A$ is fine, and the collection $(L(x))_{x \in A}$ is edge-finite, then K is fine.*

Proof. We may as well suppose that K is connected. Let M be a subgraph of K with $V(M) = V(K)$, which contains $K \setminus A$, and such that every element of A has degree 1 in M . Given $x \in A$, let $L'(x)$ be the subgraph of M consisting of $L(x)$ together with the unique edge of M incident on x . Thus, $L'(x)$ is connected, and has length at most $n + 1$, and vertex set $V(L(x)) \cup \{x\}$. Now, it’s clear that the collection $(L'(x))_{x \in A}$ is an edge-finite collection of subgraphs of M . Let $\mathcal{L} = \{L'(x) \mid x \in A\}$. By Lemma 2.4, $M[\mathcal{L}]$ is fine. But now, K is naturally embedded as a subgraph $M[\mathcal{L}]$, and so K is fine as claimed. \square

Fine graphs without finite order vertices have the property that geodesic arcs are extendible. More precisely:

Proposition 2.11. *Suppose that K is a fine graph with every vertex of infinite degree. Any finite geodesic arc in K lies inside an biinfinite geodesic arc.*

Proof. It’s enough to show that if $x, y \in V(K)$, then there is some $z \in V(K)$ adjacent to y such that $d_K(x, z) = d_K(x, y) + 1$. To see this, note that if z is adjacent to y , and $d_K(x, z) \leq d_K(x, y)$, then the edge xz together with any geodesic from x to z is an arc of length at most $d_K(x, y) + 1$ connecting x to y . There are only finitely many such arcs, and hence only finitely many such z . \square

As mentioned earlier, there is another way of formulating the notion of fineness as follows. Suppose K is a graph. Given $x \in V(K)$, write $E(x)$ for the set of edges incident on x . Suppose, to each $x \in V(K)$, we associate a locally finite function $L = L_x : E(x)^2 \rightarrow [0, \infty]$, which we can intuitively think of as measuring an “angle” between two adjacent edges. Here, “locally finite” means that for all $e \in E(x)$ and $r \geq 0$, the set $\{f \in E(x) \mid L(e, f) \leq r\}$ is finite. Given a path, $\beta = e_0, \dots, e_n$ in K , we define the “turning” of β as $L(\beta) = \sum_{i=1}^n L(e_{i-1}, e_i)$. Clearly only finitely many paths of a given turning can emanate from any given edge. We similarly define the turning of any cycle, where we sum over all the vertices of the cycle. We shall say that $L = [x \mapsto L_x]$ is an “angular structure” on K if there is a bound, in terms of n , on the turning of any circuit of length n . We see easily that K is fine if and only if it admits an angular structure. In fact, if $e, f \in E(x)$, we can define $L(e, f)$ to be the distance between the other endpoints of e and f in $K \setminus x$. In particular, we can always take L to be a metric on $E(x)$ for all x (allowing for infinite angles). Note that if K is 2-vertex connected, we can assume all angles to be finite.

Most of the arguments of later sections can be rephrased in terms of angles. For example, we note that in a fine hyperbolic graph, as discussed in Section 8, there is an upper bound on the angle between any two (quasi)geodesic with the same endpoints (one possibly ideal).

3. HYPERBOLIC GRAPHS AND COMPLEXES

In this section, we explore the notion of hyperbolicity of graphs and related spaces. One of the main objectives will be to associate to a graph, K , a space, $X(K)$, which will be hyperbolic under suitable hypotheses on K (see Theorem 3.8).

For the present purposes, we shall define hyperbolicity in terms of the linear isoperimetric inequality. Our discussion here is essentially combinatorial, though it will be convenient to phrase certain ideas in terms of CW-complexes. We shall look at other aspects of hyperbolicity in Section 5. Specific properties of fine hyperbolic graphs will be considered further in Section 8.

Suppose that K is a connected graph, and that $n \in \mathbb{N}$. We write $\Omega_n(K)$ for the CW-complex obtained by gluing a 2-cell to every circuit of length at most n in K . (Thus K is fine if and only if, for all $n \in \mathbb{N}$, $\Omega_n(K)$ is locally finite away from $V(K)$.)

Definition. We say that K is *n-simply connected* if $\Omega_n(K)$ is simply connected. We say that K is *coarsely simply connected* if it is *n-simply connected* for some n .

Clearly a graph is n -simply connected if and only if each of its blocks is.

We can define the notion of hyperbolicity in similar terms using “cellular discs”. By a *cellulation* of a space, we mean a representation of the space as a CW-complex. A *cellular map* between CW-complexes is one which sends cells into (possibly lower dimensional) cells. (This is more restrictive than the definition of “cellular map” commonly given in this context.) We can think of a cycle in K as consisting of a cellulation of the circle, S^1 , together with a cellular map of S^1 into K which sends each 1-cell homeomorphically onto an edge of K . We can define a *cellular disc*, (D, f) , (of *coarseness* n) as consisting of a cellulation of the disc, D , together with a map, f , of the 1-skeleton of D into K such that the boundary of each 2-cell in D has at most n 1-cells and gets mapped to a cycle (of length at most n) in K . We measure the *area* of D (or of (D, f)) as the number of 2-cells in D . We speak of (D, f) as a (*cellular*) *spanning disc* for the cycle $f|\partial D$, and of the cycle $f|\partial D$ as *bounding* (D, f) . Note that we can extend f to a cellular map $f : D \rightarrow \Omega_n(K)$, which gives us a more intuitive way of thinking about such an object. (There are a multitude of other ways of formulating the notion of “spanning disc” and “area”, and we shall see slight variations on these later. All we really require of such a notion is that it should satisfy a “rectangle inequality” similar to the coarea formula of Riemannian geometry — see, for example, [Bo1] for some discussion.)

There is another technical point we should make, namely that one can reformulate everything we say using circuits in place of cycles. This is based on the following simple observation. Suppose β is a cycle in K . Then, we can find a cellulation of the disc, D , with the number of 2-cells bounded by a fixed linear function of $\text{length}(\beta)$, and an extension of β to the 1-skeleton of D , such that the boundary of each 2-cell in D is either collapsed to an edge or point of K , or mapped onto a circuit. This can be achieved inductively. If two vertices of β get mapped to the same vertex of K , then we split β into two subcycles by connecting these two vertices by an edge which gets collapsed to a point in K . (Note that this process need not give us a spanning disc, since we have no control over the lengths of these circuits.)

We now equip our graph K with its path metric, d_K , obtained by assigning each edge a length 1. The following is one expression of the linear isoperimetric inequality:

Proposition 3.1. *A connected graph K is hyperbolic if and only if there is some $n \in \mathbb{N}$, and a fixed linear function, such that every cycle*

in K has a cellular spanning disc of coarseness n , whose area is bounded by the given linear function of the length of the cycle.

The number n together with the linear function will be referred to as “hyperbolicity parameters”.

From the earlier observation, we see that we could replace the word “cycle” by “circuit” in this proposition. An immediate consequence of this is that a connected graph is hyperbolic if and only its blocks are uniformly hyperbolic (i.e. with fixed hyperbolicity parameters). We also note that hyperbolic graph is coarsely simply connected.

For some purposes, it is better to define hyperbolicity in terms of simplicial complexes, rather than cell complexes.

By a *flag 2-complex*, we shall mean a 2-dimensional simplicial complex with the property that every 3-circuit in the 1-skeleton bounds a 2-simplex. Prior to the statement of Lemma 2.6, we defined the flag complex, $\Sigma(K)$, associated to a graph, K , which we may identify with $\Omega_3(K)$. Thus, $\Sigma(K)$ is simply connected if and only if K is 3-simply connected. We note:

Lemma 3.2. *If K is n -simply connected, then K^n is 3-simply connected.*

(Here K^n is the graph defined in Section 2.)

We also get a combinatorial notion of hyperbolicity.

Definition. A *simplicial disc*, (D, f) , in K , consists of a triangulation of the disc D , together with a simplicial map, f , of D into $\Sigma(K)$. We speak of (D, f) as a (*simplicial*) *spanning disc* bounding the cycle $f|\partial D$ in K . We measure the *area* of (D, f) as the number of 2-simplices in the triangulation of D .

Definition. We say that a connected graph, K , is *simplicially hyperbolic* if there is a fixed linear function such that every cycle in K bounds a simplicial disc, whose area is bounded by this linear function of the length of the cycle.

As before, we can replace the word “cycle” by “circuit”, without changing the definition. A graph is simplicially hyperbolic if and only if its blocks are uniformly simplicially hyperbolic. A simplicially hyperbolic graph is obviously 3-simply connected.

Note that if β is a circuit, then we can assume that a simplicial spanning disc, (D, f) , maps ∂D homeomorphically to β .

Clearly this notion has combinatorial as well as geometric content — it is certainly not a quasiisometry invariant. However, we note:

Lemma 3.3. *If K is hyperbolic, then K^n is simplicially hyperbolic for some n .*

The following discussion of links in K is only indirectly relevant to the main objective of this section (the construction of the space $X(K)$). However it will help to clarify certain points arising.

Let K be any graph, and $x \in V(K)$. Recall that $V_K(x) \subseteq V(K)$ is the set of vertices adjacent to x . Let $L_K(x)$ be the full subgraph of K with vertex set $V_K(x)$. Thus $L_K(x)$ is precisely the link of x in $\Sigma(K)$. The following is a simple exercise:

Lemma 3.4. *If K is 2-vertex-connected and 3-simply connected, then $L_K(x)$ is connected for each $x \in V(K)$.*

(It follows that the space $\Sigma(K) \setminus V(K)$ is connected.) The proof of Lemma 3.4 is essentially the same as that of the following result.

If L is a connected subgraph of K , we say that L is *linearly distorted* if given any $x, y \in V(L)$, $d_L(x, y)$ is bounded above by some fixed linear function (the *distortion bound*) of $d_K(x, y)$.

Lemma 3.5. *If K is 2-vertex-connected and simplicially hyperbolic, then $L_K(x)$ is linearly distorted in $K \setminus x$ for all $x \in V(K)$. Moreover, the distortion bound is independent of $x \in V(K)$.*

Proof. Suppose $y, z \in L = L_K(x)$ are distinct. Let β be an arc of minimal length, say n , connecting y to z . We see that $\beta \cup yxz$ is a circuit of length $n + 2$ in K . It thus bounds a simplicial disc, $f : D \rightarrow \Sigma(K)$, whose area is linearly bounded in terms of n . We identify ∂D with its image in K . Let S be the star of x . Consider the preimage of the component of $f^{-1}S$ which contains x . Its boundary lies in $f^{-1}L$, is connected, and contains $\partial D \cap f^{-1}y$ and $\partial D \cap f^{-1}z$. It follows that there is an arc in $f^{-1}L$ connecting $\partial D \cap f^{-1}y$ to $\partial D \cap f^{-1}z$. The length of this arc is linearly bounded in terms of n , and its image in K connects y to z in L . This shows that L is linearly distorted in K (and indeed that L is connected, as claimed by Lemma 3.4. Moreover, the distortion bound is uniform. \square

Note that if $\Sigma(K)$ is locally finite away from $V(K)$, then $L_K(x)$ is locally finite for all $x \in V(K)$. It is now a simple consequence of Lemma 3.5 that:

Lemma 3.6. *If K is simplicially hyperbolic and $\Sigma(K)$ is locally finite away from $V(K)$, then K is fine.*

Proof. Note that each block of K satisfies the same hypotheses, and so, using Lemma 3.5 satisfies condition (F5) of fineness. By Lemma 2.2 we see that K is fine. \square

We remark that the proofs of Lemmas 3.5 and 3.6 have not really used the fact that $\Sigma(K)$ is a flag 2-complex. Both lemmas hold if we take Σ to be any simplicially hyperbolic 2-complex with 1-skeleton K .

(An indirect proof of Lemma 3.6 can also be given from the construction of the space, $X(K)$, together with the results of Section 7.)

We now assume that K is 2-vertex-connected, and set about the construction of the space $X(K)$. We can view this as a geometric realisation of the space, $\Sigma(K) \setminus V(K)$, which, as we have already observed, is connected. To do this, we take, for each 2-simplex in $\Sigma(K)$, an ideal hyperbolic triangle. We glue these triangles together by isometry along their edges, with the same combinatorial pattern as $\Sigma(K)$. There is a canonical way of performing these identifications, which may be characterised by saying that the union of two adjacent triangles is isometric to an ideal hyperbolic square. The following is a simple observation:

Lemma 3.7. *If K is fine, then $X(K)$ is complete and locally compact.*

This will turn out to be the only case of interest to us, though we have no need to impose the assumption of fineness in what follows.

Our main result will be:

Theorem 3.8. *If K is 2-vertex-connected and simplicially hyperbolic, then $X(K)$ is hyperbolic.*

Here, of course, we use “hyperbolic” in the sense of Gromov [Grom], which we shall elaborate on shortly. (See also Section 5.)

In fact we will draw further consequences from our construction, notably that there is a canonical embedding of $V(K)$ in the ideal boundary, $\partial X(K)$, of $X(K)$. Moreover, a point of $V(K)$ is isolated in $\partial X(K)$ if and only if it has finite degree in K .

Let’s focus for the moment on showing that $X(K)$ is hyperbolic. We shall need another notion of spanning disc, appropriate to the present context. The idea is to take a riemannian metric on the disc D and a lipschitz map from D into our space X , where of course, we need to bound the lipschitz constant. However, we need to make allowance for the fact that X is not simply connected. We thus make do with a riemannian metric on a planar surface $H \subseteq D$, bounded by ∂D and a finite number of circles S_1, \dots, S_k in D . We insist that the lengths of each of these circles is bounded above by some fixed constant. We measure the area of D as the riemannian area plus the sum of the lengths of the curves S_i . A “perforated” spanning disc thus consists of a lipschitz map, f , of H into X . We speak of (D, f) as “spanning” the loop $f|\partial D$. It is easily checked that this notion satisfies the rectangle inequality as described in [Bo1]. It follows that X is hyperbolic if every

loop, β , in X bounds a spanning disc whose area is linearly bounded by $\text{length}(\beta)$.

The following construction is a special case of a more general procedure of taking “cusps” over metric spaces, described in [Grom].

By a *spike* in the hyperbolic plane, \mathbb{H}^2 , we mean a closed region bounded by two asymptotic geodesic rays and a horocyclic arc of length 1. More precisely, we can define a spike, Y , as the region $[0, 1] \times [1, \infty)$ in the upper half-space model. We write $Y_t \subseteq Y$ for the region $[0, 1] \times [e^t, \infty)$ (so that $Y = Y_0$). Note that Y_t is the intersection of Y with a horodisc of hyperbolic height t above the horocyclic boundary of Y .

Suppose now that L is a connected graph. We construct a space by taking a spike for every edge of L , and gluing them together, by isometry along the bounding rays, in the pattern prescribed by L . This gives us a hyperbolic 2-complex denoted $\text{cusp}(L)$, with a copy of L embedded in the 1-skeleton as the union of all the horocyclic edges. (Here we use “hyperbolic” in the sense of hyperbolic geometry.) The remainder of the 1-skeleton consists of a set of geodesic rays, one for each vertex of L . These rays are asymptotic, in the sense that the distance between any two of them is bounded. We can thus embed $\text{cusp}(L)$ as a Hausdorff topological space, $\text{cusp}(L) \cup \{p\}$, by adjoining an “ideal point”, p , which compactifies each of these rays. A neighbourhood base of p is given as the complements of bounded subsets of $\text{cusp}(L)$. (In the case of particular interest, where L is locally finite, $\text{cusp}(L)$ will be complete and locally compact, and we obtain the one-point compactification of $\text{cusp}(L)$.)

We observe that $\text{cusp}(L)$ is hyperbolic (in the sense of Gromov). This is readily seen by spanning a loop, β , in $\text{cusp}(L)$ simply by coning over the ideal point p . The area of the disc thus obtained is at most $\text{length}(\beta)$. We have the technical detail that our disc is not entirely contained in our space. However we can easily put this right by making a small hole in the disc at p , and pushing it into $\text{cusp}(L)$. Thus gives us a perforated spanning disc, in the sense described earlier. This shows that $\text{cusp}(L)$ is hyperbolic.

It is now a simple matter to show that the ideal boundary of $\text{cusp}(L)$ consists of a single point which we may identify with our point p . Moreover, L is a horocycle about p . It is also worth observing that the closed subset, $B_t \subseteq \text{cusp}(L)$ obtained as the union of the sets Y_t in the construction of $\text{cusp}(L)$, is a horoball about p .

Suppose we subdivide an ideal triangle into four pieces using three horocyclic arcs which touch pairwise at points on the edges of the triangle. (This is combinatorially equivalent to the subdivision of a simplex into four smaller simplices, as described in Section 2 in getting from the

complex $\Sigma(K)$ to $\Sigma(K')$.) We have the happy coincidence that each horocyclic arc has length 1, so that each of the three unbounded pieces is isometric to a spike, as described earlier. (This saves us the bother of rescaling.) Now, subdividing each triangle of $X(K)$ in this way, we obtain a realisation of $\Sigma(K') \setminus V(K)$ as a complex with 2-cells locally modelled on the hyperbolic plane. Here K' is the binary subdivision of K . Now, if $x \in V(K)$, we write $S(x)$ for the star of x in $\Sigma(K')$. Thus, $L(x) = L_{K'}(x)$ is the boundary of $S(x)$ in $\Sigma(K')$. Clearly $S(x) \setminus \{x\}$ is isometric to $\text{cusp}(L(x))$. Note that the interiors of these stars are all disjoint.

Given $A \subseteq V(K)$, recall that $K' \setminus A$ is defined as the full subgraph of K with vertex set $V(K') \setminus A$. We see that $\Sigma(K' \setminus A) = \Sigma(K') \setminus \bigcup_{x \in A} \text{int } S(x)$. Let $P = \Sigma(K' \setminus V(K))$. Thus, $V(P) = V(K') \setminus V(K)$.

We can now give the proof of Theorem 3.8.

Proof. Let β be a loop (i.e. a closed rectifiable path) in $X(K)$. We suppose that β consists of an alternating sequence of paths, $\alpha_1, \gamma_1, \alpha_2, \gamma_2, \dots, \alpha_k, \gamma_k$, where each α_i consists of a sequence of edges in the 1-skeleton of P , and each γ_i is a path of length at least 2 in lying in $S(x_i)$ for some $x_i \in V(K)$. (More precisely, we can homotop β to such a path, only increasing its length by an amount which depends linearly on the original length, and such that the area of the homotopy is similarly bounded.) Let $y_i, z_i \in V(P)$ be the endpoints of the path γ_i . Let γ'_i be the path $y_i x_i z_i$ in K' . Let β' be the path in K' obtained from β by replacing each path γ_i by γ'_i . Thus $\text{length}(\beta') \leq \text{length}(\beta)$. Now, by hypothesis, K and hence K' is simplicially hyperbolic, and hence bounds a simplicial disc, D_0 , whose area (as measured by the number of 2-cells) is linearly bounded in terms of $\text{length}(\beta')$. Note that the riemannian area of its image in $X(K)$ (punctured at the vertices $V(K)$) is certainly less than π times this combinatorial area.

Now, consider the closed path $\gamma_i \cup \gamma'_i$ in $S(x_i) = \text{cusp}(L(x_i)) \cup \{x_i\}$. As in our earlier discussion, we see that this spans a disc, D_i , of area at most $\text{length}(\gamma_i)$, obtained by coning over the point x_i . We perform this construction for each $i = 1, \dots, k$. In this way, we obtain a spanning disc, D , for β , by joining each D_i to D_0 along the arc γ'_i . Note that $\sum_{i=1}^k \text{area}(D_i) \leq \sum_{i=1}^k \text{length}(\gamma_i) \leq \text{length}(\beta)$. We conclude that the total hyperbolic area of D is linearly bounded in terms of $\text{length}(\beta)$.

As before, we are left with the technical point that our disc may pass (finitely often) through points of $V(K)$. However, we can simply puncture the disc at these points and push it slightly into $X(K)$. Since everything is locally just a cone, we can control the length of the new boundary curves arising. In this way we get a perforated spanning

disc in $X(K)$, of the type described earlier. It follows that $X(K)$ is hyperbolic as claimed. \square

In fact, we see that the hyperbolicity parameters of $X(K)$, depend only on those of K .

Let $\partial X(K)$ be the ideal boundary of $X(K)$. This is a metrisable topological space. In the case of real interest to us, where $X(K)$ is locally compact, $\partial X(K)$ will be compact, though we have no reason to assume this at present.

Suppose $x \in V(K)$. Write $S'(x) = S(x) \setminus \{x\}$, so that $S'(x)$ is isometric to $\text{cusp}(L(x))$. Suppose $y \in V_{K'}(x)$. The edge yx (minus the point x) gives us a ray, β , in $S'(x) \subseteq X(K)$. We parameterise β by arc-length such that $\beta(0) = y$. From the intrinsic geometry of $S'(x) = \text{cusp}(L(x))$, we see easily that β is a geodesic in $\text{cusp}(L(x))$, and that the distance of $\beta(t)$ from $L(x)$ is equal to t . Now, $L(x)$ is the boundary of $S'(x)$ in $X(K)$, and so it follows easily that β is, in fact, a geodesic ray in $X(K)$. Now, any two such rays in $S'(x)$ are asymptotic and so define an ideal point, $p(x) \in \partial X(K)$. Moreover, we see that $S'(x)$ is a horoball about $p(x)$ in $X(K)$. In particular it follows (see Section 5) that $S'(x)$ is quasiconvex in $X(K)$. Moreover the constant of quasiconvexity is independent of x . (Another proof of this can be given via Lemma 3.5) Note that if $x \neq y$, then it's easily seen that $p(x) \neq p(y)$, and so p gives us a canonical embedding of $V(K)$ in $\partial X(K)$.

Let $B(x, t) \subseteq S(x)$ be the subset corresponding to B_t in $\text{cusp}(L(x))$. We see that $B(x, t)$ is also a horoball about x . Moreover, from the quasiconvexity of $S(x)$, it's not hard to see that $B(x, t)$ is actually convex for all sufficiently large t (independently of x).

For future reference, it's worth noting that, given the space $X(K)$ and the collection of horoballs, $\mathcal{B} = \{B(x, t) \mid x \in V(K)\}$ for a fixed t , we can recover the graph K geometrically as a graph with vertex set \mathcal{B} , and with two vertices connected by an edge if the distance between them in $X(K)$ is at most (or in this case equal to) $2t$. This construction makes sense for any hyperbolic space with a collection of horoballs — indeed in any metric space with a collection of preferred subsets. We shall make use of this construction in Section 7.

We finish this section with another result which may be proven along similar lines to Theorem 3.8

Lemma 3.9. *Suppose K is simplicially hyperbolic and $A \subseteq V(K)$. Suppose that the links $L_K(x)$ are uniformly hyperbolic for $x \in A$. Then, $K' \setminus A$ is hyperbolic.*

Proof. Suppose β is a cycle in $K' \setminus A$. Now β bounds a simplicial disc, $f : D \rightarrow \Sigma(K)'$, whose area is linearly bounded in terms of $\text{length}(\beta)$. Suppose that U is a component of $f^{-1}(\text{int } S(x))$ for some $x \in A$. The closure of U in D is a subcomplex which is a planar surface. The outer boundary of U is a closed curve whose length is at most the area of U . Now f maps this boundary component to a cycle in $L(x)$. Since $L(x)$ is hyperbolic, this cycle bounds a cellular disc in $L(x)$, whose area is, in turn, linearly bounded by area of U . Replacing $f|U$ by this cellular disc, and performing this construction for each such U , we get a cellular disc in $K' \setminus A$ spanning β . We have increased the area by an amount linearly bounded in terms of the area of D . Thus its total area is still linearly bounded in terms of $\text{length}(\beta)$. We have thus verified a form of the linear isoperimetric inequality for $K' \setminus A$. \square

4. GROUPS ACTING ON SETS

In this section, we extend the notions of fineness and hyperbolicity to group actions on sets. This will lead to the first definition of relative hyperbolicity, given at the end of this section.

As mentioned in the introduction, the approach of this section could be used to study other “relative” geometric properties of groups. It is based on the construction of graphs which play the role of Cayley graphs in the non-relative case (where all point stabilisers are finite).

Certain kinds of relative splittings of groups can be phrased naturally in these terms. This gives a convenient formal setting in which to explore splittings of relatively hyperbolic groups. For some applications of this, see [Bo9].

Suppose Γ is a group. By a Γ -set, we mean a set, V , on which Γ acts. We shall refer to the points of V as *vertices*. Thus, if $x \in V$, we refer to the group, $\Gamma(x) = \{g \in \Gamma \mid gx = x\}$ as a *vertex stabiliser*. A *pair stabiliser* is a subgroup of the form $\Gamma(x) \cap \Gamma(y)$, where $x, y \in V$ are distinct. We write $V = V_0 \sqcup V_\infty$, where V_0 and V_∞ are, respectively, the sets of vertices with finite and infinite stabilisers. Clearly these are Γ -invariant. We say that V is *cofinite* if V/Γ is finite.

Definition. Given a Γ -set, V , a (Γ, V) -graph is a connected Γ -invariant graph with vertex set V and with finite quotient under Γ .

By having “finite quotient” we mean that there are finitely many Γ -orbits of vertices and edges. (We do not take this to imply that there are no edge-inversions — so there might be no well-defined quotient graph.)

We shall say that V is *connected* if it admits a (Γ, V) -graph. Clearly, a connected Γ -set is cofinite.

(Note that Γ is itself a Γ -set under left multiplication. In this case, a (Γ, Γ) -graph is precisely a Cayley graph. Clearly, Γ is connected as a Γ -set if and only if it is finitely generated as a group.)

Lemma 4.1. *Suppose V is a Γ -set and $W \subseteq V$ is Γ -invariant. If V is cofinite, and W is connected and non-empty, then V is connected. Conversely, if V is connected and $V_\infty \subseteq W$, then W is connected.*

Proof. For the first statement, let K be a (Γ, W) -graph, and let $V' \subseteq V \setminus W$ be a (finite) Γ -orbit transversal of $V \setminus W$. For each $x \in V'$, connect x to any point of W by an edge, $e(x)$. Let L be the graph with vertex set V and edge set $E(K) \cup \bigcup_{x \in V'} \Gamma e(x)$, where $\Gamma e(x)$ is the Γ -orbit of $e(x)$. We see easily that L is a (Γ, V) -graph.

The second statement follows by a similar argument to Proposition 4.10. \square

Definition. We shall say that a Γ -set is of *finite type* if it is connected and has all pair stabilisers finite, and all vertex stabilisers finitely generated.

Although we are ultimately only interested in Γ -sets of finite type, we shall only introduce these hypotheses as we need them.

We continue our study of Γ -sets with a simple observation.

Lemma 4.2. *Suppose V is a connected Γ -set. Any two (Γ, V) -graphs are quasiisometric via a Γ -equivariant quasiisometry which restricts to the identity on V .*

Thus any quasiisometry invariant property of graphs gives rise to a property of connected Γ -sets. An obvious example of this is hyperbolicity, and we shall return to this later. We shall also see that certain combinatorial properties, notably fineness, are also independent of the choice of graph, despite not being quasiisometry invariant. We begin with some general observations.

Lemma 4.3. *Suppose that V is a connected Γ -set, and K is a fine (Γ, V) -graph. Then all edge stabilisers of K are finite if and only if all pair stabilisers of V are finite.*

Proof. The “only if” bit is trivial. Conversely, suppose that each edge stabiliser of K is finite. Clearly the stabiliser of any non-trivial arc is finite. Now, suppose $x, y \in V$ are distinct. We connect x and y by an arc, α , in K . By property (F2) of fineness, the set of $(\Gamma(x) \cap \Gamma(y))$ -images of α is finite. It follows that $\Gamma(x) \cap \Gamma(y)$ is finite, as required. \square

Lemma 4.4. *Suppose that V is a connected Γ -set with finite pair stabilisers, and K is a (Γ, V) -graph. If \mathcal{L} is a Γ -invariant collection of finite subgraphs of K with \mathcal{L}/Γ finite, then \mathcal{L} is edge-finite.*

Proof. If $L \in \mathcal{L}$ and $e \in E(K)$, then $\{g \in \Gamma \mid ge = e\}$ is finite. Thus, only finitely many Γ -images of L contain the edge e . The result follows since \mathcal{L}/Γ is finite. \square

Lemma 4.5. *Suppose that V is a connected Γ -set with finite pair stabilisers, and K and L are (Γ, V) -graphs. If K is fine, then so is L .*

Proof. Suppose $e \in E(L)$. Since K is connected, we can find an arc, α , in K whose endpoints coincide with those of e , i.e. $e = e(\alpha)$ in the notation of Lemma 2.3. We can perform this construction Γ -equivariantly, giving us a Γ -invariant set, \mathcal{A} , of arcs in K with $L \subseteq K[\mathcal{A}]$. Since $E(L)/\Gamma$ is finite, we see that \mathcal{A}/Γ is finite, so \mathcal{A} is edge-finite by Lemma 4.4. By Lemma 2.3, we see that $K[\mathcal{A}]$ and hence also L is fine. \square

Definition. We say that a Γ -set is *fine* if it is connected, has finite pair stabilisers, and if some (and hence every) (Γ, V) -graph is fine.

An alternative way of formulating this is given by the following lemma:

Lemma 4.6. *A connected Γ -set, V is fine if and only if a (Γ, V) -graph has finite edge stabilisers and finitely many Γ -orbits of n -circuits for any n .*

(Here, “a” could be interpreted either as “some” or “every”.)

Proof. Let K be a (Γ, V) -graph. The “if” statement follows from Lemmas 4.3 and 4.4. Conversely, suppose K is fine. Let E_0 be a finite set of edges containing an edge from each Γ -orbit. Now, any n -circuit must have some Γ -image meeting E_0 . By fineness, there are only finitely many possibilities for such images. \square

For many purposes it is convenient (if not essential) to restrict to (Γ, V) -graphs which are 2-vertex-connected. There are several ways, according to context, one can reduce to this case. In general, note that if K is a (Γ, V) -graph, then Γ acts on the block tree, T , of K . We thus get a splitting of Γ whose vertex groups are either vertex stabilisers of K , or setwise block stabilisers. We also note that if B is a block of K , and $\Gamma(B)$ is the setwise stabiliser of B , then $V(B)$ is a connected $\Gamma(B)$ -set, and B is a $(\Gamma(B), V(B))$ -graph. Thus, for many purposes, we don’t lose much by restricting to a 2-vertex-connected graph.

Another observation which helps to justify this liberty is the following:

Lemma 4.7. *Suppose that V is a Γ -set of finite type. Then V admits a 2-vertex-connected (Γ, V) -graph.*

Proof. Let K be any (Γ, V) -graph. Suppose $x \in V$. Recall that $V_K(x)$ is the set of adjacent vertices. Now $\Gamma(x)$ acts on $V_K(x)$ with finite vertex stabilisers and finite quotient. Since $\Gamma(x)$ is finitely generated, we can find a connected $\Gamma(x)$ -invariant graph, $H(x)$, with vertex set $V_K(x)$, and with $E(H(x))/\Gamma(x)$ finite. We can perform this construction Γ -equivariantly for each $x \in V$. Let L be the graph with vertex set V , and with edge set $E(L) = E(K) \cup \bigcup_{x \in V} E(H(x))$. Thus, $E(L)/\Gamma$ is finite. We see that L is a 2-vertex-connected (Γ, V) -graph as required. \square

We also have a converse to Lemma 4.7, though it requires additional hypotheses. Note that since coarse simple connectedness is a quasiisometry invariant, we can speak of a connected Γ -set as being “coarsely simply connected”.

Lemma 4.8. *Suppose a Γ -set, V , is fine and coarsely simply connected. If V admits a 2-vertex-connected (Γ, V) -graph, then it is of finite type.*

Proof. In other words, we want to show that all vertex stabilisers are finitely generated.

Let K be a 2-vertex-connected (Γ, V) -graph. Since K is coarsely simply connected, there is some n such that $\Sigma(K^n)$ is simply connected (Lemma 3.2). Suppose $x \in V$. By Lemma 3.4, the link, $L_{K^n}(x)$, of x in $\Sigma(K^n)$ is connected. Since K^n is fine (Lemma 2.5), $L_{K^n}(x)$ is locally finite. Now, $\Gamma(x)$ acts on $L_{K^n}(x)$ with finite edge stabilisers and finite quotient. It follows that $\Gamma(x)$ is finitely generated. \square

We are now ready for the main definition of this section.

Definition. We say that a connected Γ -set, V , is *weakly hyperbolic* if some (hence every) (Γ, V) -graph is hyperbolic (in the sense of Gromov, as described in the Section 3).

Definition. A Γ -set is *hyperbolic* if it has finite type, and is fine and weakly hyperbolic.

In other words, it admits some fine weakly hyperbolic (Γ, V) -graph with all edge stabilisers finite and all vertex stabilisers finitely generated.

Since this last definition is a bit cryptic, we give a reformulation in the form of the following proposition:

Proposition 4.9. *A Γ -set, V , is hyperbolic if and only if we can represent V as the vertex set, $V = V(\Sigma)$, of a Γ -invariant simplicial complex, Σ , such that Γ acts on Σ with finite edge stabilisers and finite quotient, and such that Σ is simplicially hyperbolic and has no cut vertex.*

From the fact that Σ has finite quotient and finite edge stabilisers, we see immediately that it is locally finite away from $V(\Sigma)$. In the definition, we can replace “simplicially hyperbolic” by “hyperbolic” (in the usual geometric sense of Gromov), together with the additional assumption that there are finitely many Γ -orbits of n -circuits in the 1-skeleton of Σ for any $n \in \mathbb{N}$. (In fact, for all n sufficiently large in relation to the hyperbolicity parameters will do.) This latter assumption is, in turn, equivalent to saying that, for any n , we can place some bound on the area of a simplicial disc spanning any circuit (or cycle) of length n in the 1-skeleton of Σ . It turns out that there is no loss in assuming that Σ is a flag 2-complex, i.e. every 3-circuit in the 1-skeleton of Σ bounds a 2-simplex. The assumption that Σ has no cut vertex corresponds to assuming that all vertex stabilisers are finitely generated. It is questionable how natural this assumption is, though without it, we would be lead into a number of complications. This point was also discussed in the introduction.

We can now prove Proposition 4.9:

Proof. Suppose V is hyperbolic. By Lemma 3.6, V admits a 2-vertex-connected (Γ, V) -graph, K . Now, K is hyperbolic, so by Lemma 3.3, K^n is simplicially hyperbolic for some n . Let $\Sigma = \Sigma(K^n)$. Since K^n is fine and has finite quotient, we see that that Σ has finite quotient. Since, V had finite pair stabilisers, Σ has finite edge stabilisers.

Conversely, suppose V is the vertex set of a 2-complex, Σ , with the properties stated. Let K be the 1-skeleton of Σ . We have already observed that Σ is locally finite away from V . Thus, by Lemma 3.6, and the subsequent remark, we see that K is fine. By Lemma 4.6, V has finite pair stabilisers. By Lemma 4.8, V has finitely generated vertex stabilisers. We see that V is of finite type, fine and weakly hyperbolic, as required. \square

We remark that, from the construction, it is easily seen that the complex, K , given by Proposition 4.9 can be assumed to contain any given Γ -invariant 2-complex with finite quotient.

We refer to a vertex $x \in V$ as having “finite” or “infinite degree” according to whether $\Gamma(x)$ is finite or infinite. Recall that V_0 and V_∞ are, respectively, the sets of vertices of finite and infinite degree.

Proposition 4.10. *Suppose that V is a hyperbolic Γ -set, and that $W \subseteq V$ is a Γ -invariant subset with $V_\infty \subseteq W$. Then, W is a hyperbolic Γ -set.*

Proof. Let K be a fine hyperbolic (Γ, V) -graph. We can assume that $W \neq \emptyset$, so that W is r -dense in K . Let $L = K(W, 2r + 1)$. By Lemma 2.8, L is connected, and by Lemma 2.7, it is fine. Also, from

its description as $K[\mathcal{A}]$ for an edge-finite arc system, \mathcal{A} , (after deleting the vertices of $V \setminus W$), we see that it has finite quotient. Thus, L is a (Γ, W) -graph. We see that W is fine. Since K and L are quasiisometric, we see that L is hyperbolic. It follows that W is hyperbolic. \square

Note that the argument shows that Proposition 4.10 in fact holds for the properties of fineness and weak hyperbolicity individually. We saw in Lemma 4.1 that it also holds for the property of connectedness.

We can go on to give a stronger result:

Proposition 4.11. *Suppose that V is a hyperbolic Γ -set, and $W \subseteq V$ is Γ -invariant, and that $\Gamma(x)$ is hyperbolic for each $x \in V \setminus W$. Then W is a hyperbolic Γ -set.*

Proof. Let K be a 2-vertex-connected (Γ, V) -graph, and K' be its binary subdivision. Let $A = V \setminus W$. Now, for each $x \in A$, $\Gamma(x)$ acts with finite quotient on the link $L_K(x)$. Thus $L_K(x)$ is hyperbolic. Since V/Γ is finite, the hyperbolicity parameters are uniform over A . By Lemma 2.6, K' and hence $K' \setminus A$ is fine. By Lemma 3.9 $K' \setminus A$ is hyperbolic. Now $K' \setminus A$ is a $(\Gamma, V(K') \setminus A)$ -graph for the Γ -set $V(K') \setminus A$. Thus $V(K') \setminus A$ is hyperbolic. Now, $W \subseteq V(K') \setminus A$ and every point of $(V(K') \setminus A) \setminus W = V(K') \setminus V$ has finite degree. Thus, by Proposition 4.10, W is hyperbolic. \square

We also have the following converse of Proposition 4.10

Lemma 4.12. *Suppose that V is a cofinite Γ -set and that $W \subseteq V$ is a non-empty Γ -invariant subset with $V_\infty \subseteq W$. If W is hyperbolic, then V is hyperbolic.*

Proof. Clearly V is of finite type. Let K be a (Γ, V) -graph such that $K \setminus W$ is connected. Thus $K \setminus W$ is a (Γ, W) graph, and hence fine and hyperbolic. We can clearly construct K so that no two vertices of W are adjacent. Given $x \in V \setminus W$, since $\Gamma(x)$ is finite and $K \setminus W$ is connected, we can find a finite connected $\Gamma(x)$ -invariant subgraph, $L(x)$ of $K \setminus W$. We perform this construction Γ -equivariantly, as x ranges over $V \setminus W$. By Lemma 4.4, the set $\{L(x) \mid x \in V \setminus W\}$ is edge-finite. In fact, the argument of Lemma 4.4, tells us $(L(x))_{x \in V \setminus W}$ is edge-finite as an indexed collection. By Lemma 2.10, K is fine. Now, K and $K \setminus V$ are quasiisometric, and so K is also hyperbolic. It follows that V is hyperbolic. \square

Suppose, now, that V is a hyperbolic Γ -set. Let $\Sigma = \Sigma(K)$ be a flag 2-complex of the type described by Proposition 4.9. Let $X(\Sigma)$ be the space obtained by realising each 2-simplex of $\Sigma \setminus V$ as an ideal

hyperbolic triangle, as described in Section 3. Thus, Γ acts by isometry on Σ . In this case, $X(\Sigma)$ is complete and locally compact. By Theorem 3.8, it is (Gromov) hyperbolic. Its boundary, $\partial X(\Sigma)$, is a compact metrisable topological space. We can construct a Γ -invariant system of horoballs about points of V , such that the closure of the complement is compact. In the terminology of Section 6, we shall say that the action of Γ on $X(\Sigma)$ is “geometrically finite”.

Clearly, this construction depends on the choice of flag 2-complex, Σ . It turns out, however, that its quasiisometry class does not.

To see this, consider two 2-vertex-connected 3-simply connected (Γ, V) -graphs K_1 and K_2 , so that the 2-complexes $\Sigma(K_1)$ and $\Sigma(K_2)$ are simplicially hyperbolic. Let $K_1 \cup K_2$ be the graph with vertex set V , and edge set $E(K_1 \cup K_2) = E(K_1) \cup E(K_2)$. Clearly $K_1 \cup K_2$ is also a (Γ, V) -graph. Since $K_1 \cup K_2$ is hyperbolic, $(K_1 \cup K_2)^n$ is 3-vertex-connected for some n . Let $L = (K_1 \cup K_2)^n$, so that L is (Γ, V) -graph. Thus $\Sigma(L)$ has the property described by Proposition 4.9. Now, $\Sigma(K_1)$ and $\Sigma(K_2)$ are both subcomplexes of $\Sigma(L)$, and so $X(K_1)$ and $X(K_2)$ are both subspaces of $X(L)$. It's easily seen that each $X(K_i)$ is quasidense in $X(L)$, and that the inclusion is a quasiisometry. Moreover, it clearly respects geodesic rays tending to any point of V . We thus get a quasiisometry of $X(K_1)$ to $X(K_2)$. This extends to a homeomorphism of the boundaries which respects the canonical inclusion of V . Clearly everything can be taken to be Γ -equivariant.

In summary, we have shown:

Proposition 4.13. *Suppose V is a hyperbolic Γ -set, and K_1 and K_2 are 2-vertex-connected 3-simply connected (Γ, V) -sets. Then, there is a Γ -equivariant quasiisometry from $X(K_1)$ to $X(K_2)$. Moreover, the induced Γ -equivariant homeomorphism from $\partial X(K_1)$ to $\partial X(K_2)$ respects the natural inclusion of V into these spaces.*

We may write $X(V)$ and $\partial V = \partial X(V)$ for the spaces constructed in this manner. They are well defined up to Γ -equivariant quasiisometry, and Γ -equivariant homeomorphism, respectively. We may identify V as a subset of ∂V . A point of V is isolated in ∂V if and only if it has finite degree (unless V consists of a single point). In fact it's not hard to see that every isolated point of ∂V lies in V .

It is natural to define the *restricted boundary*, $\partial_0 V$, of V to be equal to $\partial V \setminus V_0$. (Recall that V_0 and V_∞ are, respectively, the subsets of V with finite and infinite stabilisers.) Thus, $\partial_0 V$ is a closed Γ -invariant subset of ∂V . Intrinsically it is a metrisable continuum, and is either perfect, or else consists of at most two points. This is probably best seen by splitting into different cases.

Firstly, if $V_\infty = \emptyset$, then Γ is word hyperbolic, and $\partial_0 V$ is naturally homeomorphic to $\partial\Gamma$. In this case, $\partial_0 V$ is the set of accumulation point of ∂V . If V_∞ is finite, then it must consist of a single point, and $\partial_0 V$ consist precisely of this point. The interesting case for us is when V_∞ is infinite. In this case, V_∞ is itself hyperbolic (Proposition 4.10), and we may identify $\partial_0 V$ with $\partial_0 V_\infty = \partial V_\infty$. Again, $\partial_0 V$ is the set of accumulation points of ∂V .

It finally remains to define a “relatively hyperbolic group” in these terms. Briefly, if Γ admits a hyperbolic Γ -set V , we say that Γ is “hyperbolic relative to” the collection of infinite vertex stabilisers, $\{\Gamma(x) \mid x \in V_\infty\}$. If $V_\infty = \emptyset$, we clearly get the standard notion of a (Gromov) hyperbolic group. (In this case the hypothesis of fineness is superfluous.) We can thus assume that $V_\infty \neq \emptyset$, and rephrase the definition in terms of a “peripheral structure”.

Definition. A *peripheral structure* on a group Γ , consists of a set, \mathcal{G} , of infinite subgroups of Γ , that each $G \in \mathcal{G}$ is equal to its normaliser in Γ , and each Γ -conjugate of G lies in \mathcal{G} .

We refer to an element of \mathcal{G} as a *peripheral subgroup*.

We can think of \mathcal{G} as a Γ -set, where Γ acts by conjugation. Note that the stabiliser of a peripheral subgroup under this action is precisely the peripheral subgroup itself. We can formally rephrase Definition 2 of the introduction in these terms.

Definition. We say that Γ is *hyperbolic relative to \mathcal{G}* if either $\mathcal{G} = \emptyset$ and Γ is hyperbolic, or if $\mathcal{G} \neq \emptyset$ and \mathcal{G} is a hyperbolic Γ -set.

We speak of (Γ, \mathcal{G}) as a *relatively hyperbolic group*.

We note that in such a case, each peripheral subgroup is finitely generated, and any two intersect in a finite group.

The results of this section translate into results about relatively hyperbolic groups. For example, we can associate to such a group, (Γ, \mathcal{G}) , a canonical boundary, $\partial(\Gamma, \mathcal{G})$, which is a compact metrisable space on which Γ acts by homeomorphism. From the description of the canonical inclusion of the Γ -set in the boundary, we see that each peripheral subgroup is the stabiliser of a point of $\partial(\Gamma, \mathcal{G})$. We shall see that the action is, in fact, a “geometrically finite” convergence action, and the peripheral groups are precisely the maximal parabolic subgroups. We shall elaborate on this in Section 6.

We remark that we could take a more unified approach to the above definitions, for example, by taking V to be the Γ -set $\Gamma \sqcup \mathcal{G}$, where Γ acts by left multiplication on Γ and by conjugation on \mathcal{G} . We could then

define (Γ, \mathcal{G}) to be hyperbolic if V is hyperbolic, and define $\partial(\Gamma, \mathcal{G}) = \partial_0 V$.

5. HYPERBOLIC PATH-METRIC SPACES

So far, our treatment of (Gromov) hyperbolic spaces has been primarily in terms of isoperimetric inequalities. In this section, we summarise other basic geometric facts which will be needed in the remainder of this paper. For simplicity, we shall deal mainly with length spaces, though the discussion will be seen to generalise easily to path-metric spaces, with minor modifications. It also extends to path-pseudometric spaces. We shall finish with a brief discussion of the special case of proper (i.e. complete locally compact) path-metric spaces. Such spaces are necessarily length spaces.

The notion of a hyperbolic space was defined by Gromov in [Grom]. Further expositions can be found in [Bo1, CDP, GhH, Sh].

Let (X, ρ) be a metric space. Given $x \in X$ and $r \geq 0$, we write $N(x, r) = \{y \in X \mid \rho(x, y) \leq r\}$. Given $Q \subseteq X$, write $N(Q, r) = \bigcup_{x \in Q} N(x, r)$. Suppose that β is a path connecting x to y in X . We say that β is a *geodesic* if $\text{length } \beta = \rho(x, y)$. (We won't worry about parameterisations in this paper.) We say that X is a *length space* if every pair of points are connected by a geodesic. A *triangle*, (α, β, γ) , consists of three geodesic segments cyclically connecting three points of X . We say that a point, $x \in X$ is a *k-centre* of (α, β, γ) if $\max\{\rho(x, \alpha), \rho(x, \beta), \rho(x, \gamma)\} \leq k$. We say that (X, ρ) is *hyperbolic* if there is some k such that every triangle has a k -centre. We refer to k as the *hyperbolicity constant*.

As in [Bo1], we adopt the following convention regarding ‘‘approximate’’ inequalities. Suppose $\zeta, \xi \in \mathbb{R}$. We write $\zeta \preceq \xi + c$, where c is some multiple of the hyperbolicity constant. We behave as though the relation \preceq were transitive, so the constant, c , may change during the course of an argument. However, at any given stage, it is completely determined. Its value can, in principle, be ascertained by retracing the steps of a given proof back to the definitions, though we won't bother to keep track of this. We write $\zeta \succeq \xi$ to mean $\xi \preceq \zeta$, and $\xi \simeq \zeta$ to mean $\xi \preceq \zeta$ and $\zeta \preceq \xi$. We write $\zeta \ll \xi$ to mean that $\zeta + C \leq \xi$, where C is some sufficiently large multiple of the hyperbolicity constant, (which can be ascertained, generally by anticipating the remaining steps of an argument). We view the relations $\alpha \ll \xi$ and $\xi \preceq \zeta$ as mutually exclusive. If $x \in X$ and $Q \subseteq X$, we say that x is *close to* Q if $\rho(x, Q) \simeq 0$.

Given $x, y \in X$, we write $[x, y]$ for a geodesic from x to y . Any two geodesics are close (in the sense that every point of one is close to the

other), so this choice will not matter to us (except where we need to ensure that a construction is canonical). We can also define a *centre* of a triple (x, y, z) as a point close to each of $[x, y]$, $[y, z]$ and $[z, x]$. Such centres always exist, and any two centres are close to each other.

We say that a subset, Q , is λ -*quasiconvex* if $[x, y] \subseteq N(Q, \lambda)$ for all $x, y \in Q$. We shall here always take λ to be some fixed multiple of the hyperbolicity constant, and refer to Q as simply being *quasiconvex*. Note that in such a case, $N(Q, \lambda)$ is intrinsically a path-metric space, and is intrinsically hyperbolic in the induced path-metric. Note also that every uniform neighbourhood of a quasiconvex set is quasiconvex.

Suppose that $Q \subseteq X$ is quasiconvex. Given $x \in X$, we may associate a subset $P_Q(x) \subseteq Q$ by $P_Q(x) = \{y \in Q \mid \rho(x, y) \simeq \rho(x, Q)\}$. Thus, $P_Q(x)$ is always nonempty and $\text{diam } P_Q(x) \simeq 0$. It will be notationally convenient to single out an arbitrary point $\text{proj}_Q(x) \in P_Q(x)$. This gives us a “projection map” $\text{proj}_Q : X \rightarrow Q$, which we may as well take to be the identity on Q . Other choices will not matter to us. Now, proj_Q need not be continuous. However, it is approximately distance decreasing:

Lemma 5.1. *If $Q \subseteq X$ is quasiconvex, and $x, y \in X$, then $\rho(\text{proj}_Q x, \text{proj}_Q y) \preceq \rho(x, y)$.*

Another point to note is:

Lemma 5.2. *Suppose $Q, Q' \subseteq X$ are both quasiconvex, and $\rho(Q, Q') \gg 0$. Then, $\text{diam}(\text{proj}_Q Q') \simeq 0$.*

The proofs of these results are elementary (cf. [Bo1]).

Suppose $Q, Q' \subseteq X$ are quasiconvex and $\rho(Q, Q') \gg 0$. Let β be a path in X connecting Q to Q' with length $\beta \simeq \rho(Q, Q')$. Any two such paths will remain close to β . We shall write $[Q, Q']$ for some choice of such path. Similarly, given $x \in X$ and $Q \subseteq X$ quasiconvex, we shall write $[x, Q]$ for some choice of path from x to Q with length $[x, Q] \simeq \rho(x, Q)$.

As usual, we define the ideal boundary, ∂X , of X as the set of parallel classes of geodesic rays in X . (Two rays are “parallel” if they remain a bounded distance apart.) Given $x \in X \cup \partial X$ and $y \in \partial X \setminus \{x\}$, we can connect x to y by a geodesic, and we denote by $[x, y]$ one choice of such. We refer to it as a *geodesic ray* or *biinfinite geodesic* depending on whether $x \in X$ or $x \in \partial X$. We define a topology on ∂X as follows. Given $p \in \partial X$, $x \in X$ and $r \geq 0$, we let $D(p, x, r) = \{y \in X \mid \rho(x, [p, y]) \geq r\}$. We fix some $x \in X$ and let $\{D(p, x, r) \mid r \geq 0\}$ be a neighbourhood base of p . Changing x gives an equivalent

neighbourhood base. We deem the subspace topology on X to be the metric topology. It turns out that $X \cup \partial X$ is metrisable.

We may extend the notion of a centre to define the centre of a triple, x, y, z , of points of $X \cup \partial X$, provided no two of x, y, z are equal to the same ideal point.

Given a subset $A \subseteq \partial X$, we may define *join* A to be the set of points $x \in X$ such that $\rho(x, \beta) \simeq 0$ for some biinfinite geodesic, β , with both endpoints in A . For any A , *join* A is quasiconvex.

Definition. We say that X is *taut* if $X = \text{join } \partial X$.

(Of course, there is an additive constant $\simeq 0$ implicit in this definition.)

Lemma 5.3. *Suppose X is taut, and $p \in \partial X$. Given any $x \in X$, there is some $q \in \partial X$ such that $\rho(x, [p, q]) \simeq 0$.*

(Note that the implied additive constant be greater than that featuring in the definition of tautness.) The proof is elementary.

We next want to define horofunctions and horoballs.

Definition. Suppose $p \in \partial X$. A *horofunction* about p is a function $h : X \rightarrow \mathbb{R}$ such that if $x, a \in X$ and $\rho(a, [x, p]) \simeq 0$, then $h(a) \simeq h(x) + \rho(x, a)$.

(Note that we not assuming horofunctions to be continuous.) It's easily seen that if $x, y \in X$ and a is a centre of x, y, p , then $h(x) + \rho(x, a) \simeq h(y) + \rho(y, a)$. (This might serve as an alternative equivalent definition of a horofunction.) By standard arguments, we can always construct a horofunction about any ideal point. Moreover, any two horofunctions about the same point differ approximately by an additive constant. More precisely, if h and h' are horofunctions about the same point, then there is some $t \in \mathbb{R}$, such that $h'(x) \simeq h(x) + t$ for all $x \in X$.

Lemma 5.4. *Suppose $p, q \in \partial X$ with $p \neq q$. Suppose h_p and h_q are horofunctions about p and q respectively. Then, there is some $t \in \mathbb{R}$ such that $h_p(x) + h_q(x) + 2\rho(x, [p, q]) \simeq t$ for all $x \in X$.*

Proof. Suppose $x, y \in [p, q]$. Without loss of generality, we have $\rho(y, [x, p]) \simeq 0$ and $\rho(x, [y, q]) \simeq 0$. Thus, $h_p(y) \simeq h_p(x) + \rho(x, y)$. Similarly, $h_q(x) \simeq h_q(y) + \rho(x, y)$. Thus, $h_p(x) + h_q(x) \simeq h_p(y) + h_q(y)$. In other words, $h_p + h_q$ is approximately constant, equal to t , say, along $[p, q]$.

Now, suppose $x \in X$. Let $z \in [p, q]$ be a centre of (p, q, x) . Then $\rho(z, [x, p]) \simeq 0$, so $h_p(z) \simeq h_p(x) + \rho(x, z)$. Similarly, $h_q(z) \simeq h_q(x) +$

$\rho(x, z)$. Now $\rho(x, z) \simeq \rho(x, [p, q])$. Thus, we get that $h_p(x) + h_q(x) + 2\rho(x, [p, q]) \simeq h_p(z) + h_q(z) \simeq t$ as required. \square

Definition. A *horoball* about p is a closed subset $B \subseteq X$ such that there is a horofunction about p such that $h(x) \succeq 0$ for all $x \in B$, and $h(x) \preceq 0$ for all $x \in X \setminus B$.

(In fact, we could choose the horofunction, h , so that $B = h^{-1}[0, \infty)$.)

Note that p is uniquely determined by B , and we refer to it as the *centre* of the horoball. Note also that a horoball is quasiconvex. Moreover, if B is a horoball about p , then so is $N(B, r)$ for all $r \geq 0$.

We next want to consider quasigeodesics. Suppose that $f : [0, \infty) \rightarrow [0, \infty)$ is a non-decreasing function. Suppose β is a path in X . (As with geodesics, we are unconcerned about parameterisations.) We say that β is an $O(f)$ -path if, given any subpath, α , of β , we have $\text{length } \alpha \preceq f(\rho(x, y))$ where x and y are the endpoints of α . If f is linear, we refer to β as a *quasigeodesic*. In this case, we refer to f as the *linear bound* on β . The following is well-known:

Lemma 5.5. *Suppose β is a quasigeodesic connecting points $x, y \in X \cup \partial X$. Then there is some $r \geq 0$ such that $\beta \subseteq N([x, y], r)$ and $[x, y] \subseteq N(\beta, r)$. Moreover, r depends only on the hyperbolicity constant and the linear bound of β .*

We shall more or less prove this in the course of proving the next lemma, which appears to be folklore, but less standard:

Lemma 5.6. *Suppose f is a subexponential function. Then any $O(f)$ -path is quasigeodesic. Moreover the linear bound depends only on f and the hyperbolicity constant.*

In fact, it's enough to assume that $f(t) = o(\lambda^t)$, where λ is small enough in relation to the hyperbolicity constant. Since I don't have an explicit reference for Lemma 5.6, we outline a proof below. This essentially just amounts to one of the standard arguments for proving Lemma 5.5 (cf. [GhH]).

We can now prove Lemma 5.6 :

Proof. Let β be an $O(f)$ -path connecting x_1 to x_2 in X . Choose $a \in [x_1, x_2]$ so as to maximise $\rho(a, \beta)$. Let $t = \rho(a, \beta)$. Thus $\rho(a, x_i) \geq t$ for $i = 1, 2$. Let $y_i, z_i \in [x_i, a]$ be points with $\rho(a, y_i) = t$ and $\rho(a, z_i) = 2t$. If $\rho(a, x_i) < 2t$, we take $z_i = x_i$. Let w_i be a point of β with $\rho(z_i, w_i) \leq t$. If $z = x_i$, we take $w_i = x_i$. Let α be the subpath of β connecting w_1 to w_2 . Note that $\rho(w_1, w_2) \leq 6t$, so $\text{length } \alpha \simeq f(6t)$. Let γ be the path $[y_1, z_1] \cup [z_1, w_1] \cup \alpha \cup [w_2, z_2] \cup [z_2, y_2]$. Thus, $\text{length } \gamma \preceq 4t + f(6t)$. Also,

$\rho(a, \gamma) \geq t$, and $\rho(y_1, y_2) = 2t$. In other words, γ connects two points (namely y_1 and y_2) in the boundary of $N(x, t)$, while not entering the interior of $N(x, t)$. Now, a standard result about hyperbolic spaces tells us that the length of such a path is bounded below (up to an additive constant) by an exponential function, say $[u \mapsto \omega^u]$, of the distance between its endpoints. Here, ω , depends only on the hyperbolicity constant. Thus, $\omega^t \leq 4t + f(6t)$. Since $f(6t) = o(\omega^{6t})$, we get an upper bound, say s , on possible values of t . (Here s depends only on f and the hyperbolicity constant.) We have thus shown that $[x, y] \subseteq N(\beta, s)$.

We now apply the standard argument (subdividing x, y into $O(\rho(x, y))$ subpaths of bounded length) to show that β lies inside a uniform neighbourhood of $[x, y]$. With just a little more work, we conclude that length β is bounded by a linear function of $\rho(x, y)$.

Since we can apply exactly the same argument to any subpath of β , we conclude that β is quasigeodesic. \square

A result of a similar but more combinatorial flavour was proven in [Bo1]. This concerns “broken geodesics” i.e. piecewise geodesic paths (where the “break points” are regarded as part of the structure). This time we put upper bounds on the number of geodesic segments and on the pairwise “backtracking” between distinct segments. First we give some definitions.

We shall consider “directed” geodesic segments, i.e. with an ordering on their endpoints (though a weaker statement for non-directed segments would suffice for this paper). We say that two directed geodesics, $[x, y]$ and $[z, w]$, *backtrack* a distance $l \gg 0$ if $\rho(x, y) + \rho(w, z) \simeq \rho(x, w) + \rho(y, z) + 2l$. Less formally, this means that they remain close but directed in opposite directions for a distance approximately l . We say that a broken geodesic backtracks at most l if any pair of distinct directed geodesic segments thereof backtrack at most l .

The following is given as Proposition 7.3.4 of [Bo1]. (It is really a statement about a finite sequence of points, viewed as the break points of a broken geodesic.)

Proposition 5.7. *Let β be a broken geodesic path consisting with n geodesic segments with backtracking at most $l \geq 0$. Let x, y be the endpoints of β . Then, $\text{length } \beta \leq \rho(x, y) + h(n)$, where h is a function depending on l and on the hyperbolicity constant, k , and which is linear in n .*

(In fact, h has the form $h(n) = (f(k) + g(l))j(n)$ where f, g and j are fixed linear functions.)

Note, in particular, that a closed broken geodesic with n -segments (i.e. an “ n -gon”) which has bounded backtracking has total length bounded by linear function of n .

One can give a fairly elementary argument to obtain a quadratic bound in n (along the lines of Corollary 7.2 of this paper). The linear bound appears to be more subtle. All we require in this paper is a subexponential bound.

Proposition 5.7 is used in the proof of Theorem 7.11, but is not required elsewhere in this paper.

This more or less concludes our general discussion of hyperbolic length spaces. We note that our arguments through with only slight modification to path-metric spaces. In this case, we should redefine a geodesic connecting x to y to be a path, β , with length $\beta \simeq \rho(x, y)$. In fact, we can generalise still further to path-pseudometric spaces. A path-pseudometric space is hyperbolic if and only if its hausdorffification (as a path-metric space) is hyperbolic. Also, the fact that hyperbolicity is a quasiisometry invariant remains true for path-pseudometrics.

In the other direction, one might want to specialise to proper (i.e. complete locally compact) hyperbolic path-metric spaces, X . Such spaces are necessarily length spaces, and all closed bounded subsets are compact. Moreover, the ideal boundary, ∂X , is compact. If $A \subseteq \partial X$ is compact, then $Q = \text{join}(A)$ is closed, quasiconvex and intrinsically hyperbolic in the induced path-metric. We may identify ∂Q with A . Moreover $Q \cup A$ is precisely the closure of Q in $X \cup \partial X$.

6. GEOMETRICALLY FINITE GROUPS

In this section, we discuss the notion of “geometrical finiteness” for groups acting on proper (i.e. complete locally compact) Gromov hyperbolic spaces. This notion has its roots in work of Ahlfors, Greenberg and others in the context of 3-dimensional kleinian groups [A, Gre]. Since then, other formulations have been given in increasingly general contexts. The cases of constant and variable negative curvature are discussed respectively in [Bo2] and [Bo4]. A purely dynamical study of this notion has recently been made by Tukia [T3] (see also [Fr2, Y]). Of course, at each stage of the generalisation process, certain features are lost. However, even in the purely dynamical set-up there remain many interesting results and questions, and it seems to be the most natural context in which to express many aspects of the subject. Our account in this section, can be thought of falling somewhere between that in [Bo4] and that in [T3] in generality.

For a proper Gromov hyperbolic space, X , we shall see essentially two ways of formulating geometrical finiteness, based on the those of Beardon and Maskit [BeM] and Marden [Mar] respectively (“GF2” and “GF1” in [Bo2]). The Beardon-Marden definition is intrinsic to ∂X , and is one of the definitions considered in [T3]. Here, we shall interpret the Marden definition in terms of systems of horoballs, which is appropriate to the discussion of Section 7. The second dynamical definition given by [T3] can also be viewed as a further generalisation of the Marden definition to case of a convergence action on a perfect compact metrisable space M . Tukia [T3] shows the equivalence of these definitions in the dynamical setting. In the special case where M is the boundary of a Gromov hyperbolic space, it can probably also be deduced from the results of this section, though we shall not pursue this matter here. In fact, the dynamical description of geometrical finiteness gives an alternative formulation of a relatively hyperbolic group [Y]. This generalises the result of [Bo6] in the case when the peripheral structure is empty.

We shall begin by recalling briefly some basic facts about convergence groups. These were introduced by Gehring and Martin [GehM]. For further discussion, see for example, [Fr1, T2, Bo7].

Suppose that M is a compact metrisable topological space. Suppose that a group, Γ , acts by homeomorphism on M . We say that Γ is a *convergence group* if the induced action on the space of distinct triples of M is properly discontinuous. In such a case, we write $\Lambda\Gamma \subseteq M$ for the *limit set* of Γ , which (unless Γ is virtually cyclic) is the unique minimal closed non-empty Γ -invariant subset. We say that the action is *minimal* if $\Lambda\Gamma = M$. We say that a subgroup $G \leq \Gamma$ is *parabolic* if it is infinite, fixes some point of M , and contains no loxodromics. (Here we shall allow for the possibility of an infinite torsion parabolic group). In this case, the fixed point of G is unique. We refer to it as a *parabolic point*. (Tukia has suggested that one should reserve the term “parabolic fixed point” for one which is fixed by a parabolic element, which need not be the case here. In [T3], Tukia shows that any point fixed by infinitely many non-loxodromic elements is a parabolic point.) The stabiliser, $\text{stab}_\Gamma(p)$ of a parabolic point is necessarily a parabolic group. There is thus a natural bijective correspondence between parabolic points in M and maximal parabolic subgroups of Γ . We say that a parabolic group, G , with fixed point p , is *bounded* if the quotient $(M \setminus \{p\})/G$ is compact. (It is necessarily hausdorff.) We say that p is a *bounded parabolic point* if $\text{stab}_\Gamma(p)$ is bounded. A *conical limit point* is a point $y \in M$ such that there exists a sequence $(\gamma_i)_{i \in \mathbb{N}}$ in Γ , and distinct points $a, b \in M$, such that $\gamma_i y \rightarrow a$ and $\gamma_i x \rightarrow b$ for all $x \in M \setminus \{y\}$.

There are a number of alternative formulations of conical limit point, see [Bo7] and [T3]. It is shown in [T3] that:

Proposition 6.1. *A conical limit point cannot be a parabolic point.*

The notion of a conical limit point has quite a long history in the theory of fuchsian and kleinian groups (see [He]). They have also been called “radial limit points” or “points of approximation”.

If Γ is a non-elementary group (i.e. not finite, virtually cyclic or parabolic), then $\Lambda\Gamma$ is perfect (i.e. has no isolated points). Moreover, $y \in \Lambda\Gamma$ is a (bounded) parabolic point with respect to the action of Γ on $\Lambda\Gamma$ if and only if it is a (bounded) parabolic point with respect to the action on M . The same goes for conical limit points. This is the reason, it is convenient to restrict to minimal actions of Γ (so that M is perfect). In the context of (boundaries of) Gromov hyperbolic spaces, there is no essential loss of generality in doing this. (Since we can replace X by $\text{join } \Lambda\Gamma$, which is intrinsically taut and hyperbolic with ideal boundary $\Lambda\Gamma$ — see Section 5.) However, as observed by Tukia [T3], in the (possibly) more general setting of convergence actions, the issue appears to be more subtle. We shall not pursue this general question here.

We are now ready for the definition of geometrical finiteness based on that of Beardon and Maskit [BeM]:

Definition. Let M be a compact metrisable space, and suppose that Γ is a convergence group acting on M . We say that (the action of) Γ is *geometrically finite* if every point of M is either a conical limit point or a bounded parabolic point.

It follows that M is perfect, and that the action of Γ on M is minimal.

Let \mathcal{G} be the set of maximal parabolic subgroups of Γ . It’s easy to see that \mathcal{G} is a peripheral structure on \mathcal{G} as defined at the end of Section 4. Moreover, the intersection of any two peripheral subgroups is finite. Tukia [T3] shows that there are finitely many conjugacy classes of peripheral subgroups. (We give a proof below, in the case of Gromov hyperbolic spaces.)

In general a maximal parabolic subgroup of a convergence group need not be finitely generated. It’s unclear what additional hypotheses, geometric or dynamical, are needed to force maximal parabolic subgroups to be finitely generated. (For example, it is true in a riemannian manifold of pinched negative curvature [Bo3].) We shall have no reason to assume this in this section.

We now explore what this means in the case where M is the boundary of a proper hyperbolic length space, X . Most of the arguments are

variations of those that can be found in [Bo2] and [Bo4]. We shall need to find substitutes for certain results and constructions. For example, in the context of pinched Hadamard manifolds, the Margulis Lemma ensures that horoballs can be chosen to be disjoint. Here, this follows from Lemma 6.8, which rests ultimately on the dynamical result of Proposition 6.1.

Suppose then that X is a proper hyperbolic length space with boundary, ∂X . Thus, ∂X is a compact metrisable topological space. It will be convenient to assume that X is taut, i.e. every point of X is a bounded distance from a biinfinite geodesic, though for most purposes, this will not be necessary.

If Γ acts properly discontinuously and isometrically on X , then we get an induced action of Γ on ∂X , which is a convergence action (see, for example, [Fr1, T2, Bo7]). In this context, a point $y \in \partial X$ is a conical limit point if and only if for some (hence every) point $x \in X$, there is a sequence of Γ -images of x which converges to y while remaining within bounded distance of some geodesic ray with endpoint at y .

Suppose that G is a discrete parabolic group acting on X with fixed point $p \in \partial X$. If $h : X \rightarrow \mathbb{R}$ is a horofunction about p , then so is $[x \mapsto \inf\{h(gx) \mid g \in G\}]$. (Recall that we are not assuming horofunctions to be continuous.) We see that:

Lemma 6.2. *If G is a parabolic group with fixed point p , then there is a G -invariant horofunction about p .*

Similarly, if $B \subseteq X$ is a horoball about p , then so is the closure, B' , of $\bigcup_{g \in G} gB$. Note that $B \subseteq B'$, and that B' lies in a bounded neighbourhood of B . We see that there is no loss in assuming horofunctions and horoballs are G -invariant if necessary.

We need to consider the geometric meaning of a ‘‘bounded parabolic point’’. Given a horoball B , we write $\text{fr } B$ for the frontier of B in X .

Lemma 6.3. *Suppose that G is a parabolic group with fixed point p . Suppose that B is a G -invariant horoball about p . Then G is a bounded parabolic group if and only if $(\text{fr } B)/G$ is compact.*

Proof. Let F be the subset of $(\text{fr } B) \times (\partial X \setminus \{p\})$ consisting of pairs (x, q) such that x lies within a some fixed distance, $r \gg 0$, of some biinfinite geodesic connecting p to q . Thus, F is closed and G -invariant, and the projection to $\partial X \setminus \{p\}$ is surjective and proper. The same is true of the projection to $\text{fr } B$, provided we choose r large enough so that every point of X is within a distance r of some biinfinite geodesic with one endpoint at p . (This is possible by Lemma 5.3, since we are assuming

that X is taut.) It follows that G acts cocompactly on $\partial X \setminus \{p\}$ if and only if it does so on F and hence, in turn, on $\text{fr } B$. \square

Lemma 6.4. *Suppose that G is a bounded parabolic group with fixed point p . If B, B' are G -invariant horoballs about p , then $(B' \setminus \text{int } B)/G$ is compact.*

Proof. This is equivalent to showing that if h is a horofunction about p , and $a, b \in \mathbb{R}$, then the closure of $h^{-1}[a, b]$ in X has compact quotient under G . The argument is the same as that of Lemma 6.3, replacing $\text{fr } B$ by the closure of $h^{-1}[a, b]$. \square

(In this case, we don't know that the projection of F to the first coordinate is necessarily surjective, unless $b - a \gg 0$, in which case, we can deduce a converse to Lemma 6.4.)

Although we shall not formally be using the fact, it is worth noting that:

Proposition 6.5. *If G is a bounded parabolic group with fixed point p , then the space $((X \cup \partial X) \setminus \{p\})/G$ has precisely one topological end.*

In fact, a neighbourhood base of the end is given by the collection of sets of the form B/G as B ranges over the set of G -invariant horoballs about p . The proof is along similar lines to the previous two lemmas.

Definition. If G is a bounded parabolic group, then a *cuspidal region* for G is a set of the form B/G where B is a G -invariant horoball.

We note that every point of B is a bounded distance from some G -image of a fixed geodesic ray tending to p . This bound can be assumed to be uniform (i.e. depending only on the hyperbolicity parameters) provided we take B sufficiently small. We also note that B/G is quasi-isometric to a ray.

Now let us return to our group Γ acting properly discontinuously isometrically on X . We suppose that the action is minimal and that X is taut.

Lemma 6.6. *Suppose that p is a bounded parabolic fixed point, and that $h : X \rightarrow \mathbb{R}$ is a $\text{stab}_\Gamma(p)$ -invariant horofunction about p . If $x \in X$, then $\{h(\gamma x) \mid \gamma \in \Gamma\}$ is bounded away from ∞ .*

Proof. Suppose, to the contrary, that we can find a sequence γ_i in Γ , with $h(\gamma_i x) \rightarrow \infty$. Since p is bounded, we can suppose that each point $\gamma_i x$ is a bounded distance from a geodesic ray tending to p . We see that p is a conical limit point, contradicting Proposition 6.1. \square

We can thus find a G -invariant horoball about p , which misses the orbit Γx .

Now, suppose that $\Pi \subseteq \partial X$ is a Γ -invariant set of parabolic points.

Definition. An *invariant system of horofunctions* for Π consists of a collection, $(h_p)_{p \in \Pi}$, of horofunctions indexed by p , such that h_p is a horofunction about p , and such that $h_{\gamma p}(x) = h_p(\gamma^{-1}x)$ for all $x \in X$ and $p \in \Pi$.

Lemma 6.7. *Suppose that $\Pi \subseteq \partial X$ is a Γ -invariant set of bounded parabolic points with Π/Γ finite. Suppose $x \in X$. Then, $h_p(x) \rightarrow -\infty$ as p varies over Π (i.e. $\{p \in \Pi \mid h_p(x) \geq t\}$ is finite for all $t \in \mathbb{R}$).*

Proof. Since Π/Γ is finite, it's enough to prove the result in the case where Π is the orbit of a single point. By Lemma 6.6, there is some $a \geq 0$ such that $h_p(x) \leq a$ for all $p \in \Pi$. Also, by Lemma 6.4, and the fact that Γx is discrete, we see that $\{p \in \Pi \mid h_p(x) \in [b, c]\}$ is finite for all $b, c \in \mathbb{R}$. \square

Given $p, q \in \Pi$, with $p \neq q$, and $x \in X$, we write $H_{pq}(x) = h_p(x) + h_q(x)$. Also we write $[p, q]$ for some choice of geodesic connecting p and q . Thus, Lemma 5.4 tells us that $H_{pq}(x) + 2\rho(x, [p, q])$ is bounded above and below (by constants whose difference depends only on the hyperbolicity parameters). Under the same hypotheses as Lemma 6.7, we see that these functions are uniformly bounded away from ∞ .

Lemma 6.8. *If Π is a Γ -invariant set of bounded parabolic points with Π/Γ finite, then there is some $t \in \mathbb{R}$ such that $h_p(x) + h_q(x) \leq t$ for all distinct $p, q \in \Pi$ and for all $x \in X$.*

Proof. Suppose not. Then we can find sequences, $(x_i)_i$ in X and sequences $(p_i)_i$ and $(q_i)_i$ in Π such that $H_{p_i q_i}(x) \rightarrow \infty$. Since Π/Γ is finite, we can suppose that $p_i = p$ is constant. Lemma 5.4 tells us that, up to an additive constant, the maximum values of H_{pq_i} are attained along the geodesic $[p, q_i]$. Thus, we can choose the points x_i so that $h_p(x_i)$ is bounded above and below. Now, by Lemma 6.4, and the $\text{stab}_\Gamma(p)$ -invariance of h_p , we can assume that the x_i all lie in compact set. In fact, we can assume that they converge on some $x \in X$. Now, for all $q \in \Pi$, $|h_q(x_i) - h_q(x)| \leq \rho(x, x_i)$ is bounded, and so $H_{pq_i}(x_i) \simeq H_{pq_i}(x)$. We see that $H_{pq_i}(x) \rightarrow \infty$ and so $h_{q_i}(x) \rightarrow \infty$ contradicting Lemma 6.7. \square

We can assume that, in Lemma 6.8, $t \geq 0$. Thus, replacing each horofunction h_p by the horofunction $h_p - t$, we have shown that, under the hypotheses of Lemma 6.8, we can construct an invariant system,

$(h_p)_{p \in \Pi}$, of horofunctions, such that $h_p(x) + h_q(x) \leq 0$ for all $x \in X$ and distinct $p, q \in \Pi$. In fact, we can eliminate the hypothesis that Π/Γ is finite from this statement. To do this, we first note:

Lemma 6.9. *The set of all bounded parabolic points is countable.*

Proof. This is shown in [T3] for convergence groups. It is based on the observation that if p is a bounded parabolic fixed point, then there is a finite subset, $S(p) \subseteq \text{stab}_\Gamma(p)$, such that p is the unique common fixed point of the elements of $S(p)$. Since Γ is countable, there are countably many such subsets. \square

Lemma 6.10. *Suppose $\Pi \subseteq \partial X$ is a Γ -invariant set of bounded parabolic fixed points. Then we can find an invariant system of horofunctions, $(h_p)_{p \in \Pi}$, such that $h_p(x) + h_q(x) \leq 0$ for all $x \in X$ and distinct $p, q \in \Pi$.*

Proof. By Lemma 6.9, we can write Π as an increasing union, $\Pi = \bigcup_n \Pi_n$, where $\Pi_0 = \emptyset$, $\Pi_n \subseteq \Pi_{n+1}$ and $\Pi_{n+1} \setminus \Pi_n$ is a Γ -orbit for each n . (Thus, Π_n/Γ has cardinality n .) Suppose, inductively that we have constructed h_p for each $p \in \Pi_n$. Let $(h'_p)_{p \in \Pi_{n+1}}$ be an invariant system of horofunctions for Π_{n+1} such that $h'_p = h_p$ for all $p \in \Pi_n$. By Lemma 6.8, there is some $t \geq 0$ such that $h'_p + h'_q \leq t$ for all distinct $p, q \in \Pi_{n+1}$. Given $p \in \Pi_{n+1}$, let h_p be the horofunction $h'_p - t$. This extends our system of horofunctions (h_p) to Π_{n+1} in such a way that $h_p + h_q \leq 0$ for all distinct p, q . We now continue inductively. \square

In particular, the sets $h_p^{-1}(0, \infty)$ are pairwise disjoint. We can rephrase this in terms of horoballs.

Let Π be a Γ -invariant set of parabolic points.

Definition. An *invariant system of horoballs* is a collection, $(B(p))_{p \in \Pi}$, indexed by Π , such that $B(p)$ is a horoball about p and such that $B(\gamma p) = \gamma B(p)$ for all $\gamma \in \Gamma$ and all $p \in \Pi$.

Such a system is *r-separated* if $\rho(B(p), B(q)) \geq r$ for all distinct $p, q \in \Pi$.

It is *strictly invariant* if it is *r-separated* for some $r > 0$.

Proposition 6.11. *If Π is a Γ -invariant set of bounded parabolic points and $r \geq 0$, then there exists an *r-separated invariant system of horoballs* for Π .*

Proof. Let $(h_p)_{p \in \Pi}$ be the system of horofunctions given by Lemma 6.10. Let $s \gg r$. Choose horoballs $(B(p))_{p \in \Pi}$ equivariantly so that $B(p) \subseteq h_p^{-1}[s, \infty)$ for all $p \in \Pi$. \square

We now set about describing a geometric formulation of geometrical finiteness, based on that of Marden [Mar].

Given an invariant system, $\mathcal{B} = (B(p))_{p \in \Pi}$ of horoballs, we write $Y(\mathcal{B}) = X \setminus \bigcup_{p \in \Pi} \text{int } B(p)$. Thus, $Y(\mathcal{B})$ is a closed Γ -invariant subset of X .

Proposition 6.12. *Suppose that $\Pi \subseteq \partial X$ is a Γ -invariant set of parabolic points. Suppose we can find a strictly invariant system, \mathcal{B} , of horoballs for Π such that $Y(\mathcal{B})/\Gamma$ is compact. Then Π/Γ is finite, each point of Π is a bounded parabolic point, and each point of $\partial X \setminus \Pi$ is a conical limit point.*

Proof. We first note that if $p \in \Pi$, then $(\text{fr } B(p))/\text{stab}_\Gamma(p)$ can be identified as a closed subset, $F(p)$, of $Y(\mathcal{B})/\Gamma$. Since $Y(\mathcal{B})/\Gamma$ is compact, it follows that $F(p)$ is compact, and so p is bounded by Lemma 6.3. From the strict invariance of the horoball system, we see that the sets $F(p)$ cannot accumulate in $Y(\mathcal{B})/\Gamma$, as p ranges over an orbit transversal of Π . It follows that Π/Γ is finite.

Now choose any $x \in Y(\mathcal{B})$. There is some $r > 0$ such that $Y(\mathcal{B}) \subseteq N(\Gamma x, r)$. Suppose $y \in \partial X \setminus \Pi$. Let α be a geodesic ray tending to y . Now α cannot eventually remain within any one horoball, and thus meets $Y(\mathcal{B})$ in an unbounded set. It follows that we can find an infinite sequence $\gamma_i \in \Gamma$ with $\gamma_i x \rightarrow y$ and with $\rho(\gamma_i x, \alpha) \leq r$. Thus, y is a conical limit point. \square

By Lemma 6.1, it follows that Π is, in fact, the set of all parabolic points.

Proposition 6.13. *Let Π be the set of all bounded parabolic points. Suppose that every point of $\partial X \setminus \Pi$ is a conical limit point. Then Π/Γ is finite. Moreover, for any $r \geq 0$, we can find an r -separated collection, \mathcal{B} , of horoballs such that $Y(\mathcal{B})/\Gamma$ is compact.*

Proof. By Proposition 6.11, we can find an r -separated collection $\mathcal{B} = (B(p))_{p \in \Pi}$ of horoballs for Π .

We claim $Y(\mathcal{B})/\Gamma$ is compact. Since $Y(\mathcal{B})$ is closed, it is enough to show that $Y(\mathcal{B}) \subseteq N(\Gamma x, r)$ for some $x \in Y(\mathcal{B})$ and $r \geq 0$. Suppose not. Fix $x \in Y(\mathcal{B})$, and let $(x_i)_{i \in \mathbb{N}}$ be a sequence in $Y(\mathcal{B})$ with $\rho(x_i, \Gamma x) \rightarrow \infty$. After translation by a suitable element of Γ we can suppose that $\rho(x_i, x) = \rho(x_i, \Gamma x)$ for all i . After passing to a subsequence, we can suppose that x_i tends to some $y \in \partial X$.

If $y \notin \Pi$, then y is a conical limit point. We can thus find a sequence, $(\gamma_j)_{j \in \mathbb{N}}$, in Γ with $\gamma_j x \rightarrow y$ and $\gamma_j x$ remaining within a bounded distance of a geodesic ray. Now, a simple geometric argument shows that we can find some $j \in \mathbb{N}$ such that for all $z \in X$,

sufficiently close to y , we have $\rho(z, \gamma_j x) < \rho(z, x)$. In particular, $\rho(x_i, \Gamma x) \leq \rho(x_i, \gamma_j x) < \rho(x_i, x) = \rho(x_i, \Gamma x)$, thereby contradicting the choice of x_i .

We can thus suppose that $y \in \Pi$. Now, $(\text{fr } B)/\text{stab}_\Gamma(p)$ is compact, so there exists $t \geq 0$ such that $\text{fr } B$ lies inside the t -neighbourhood of the $\text{stab}_\Gamma(p)$ -orbit of x . In particular, $\text{fr } B \subseteq N(\Gamma x, t)$. Now, since $x_i \rightarrow y$, and $x_i \notin \text{int } B(y)$, a simple geometric argument shows that we can find $z_i \in [x, y_i] \cap \text{fr } B$ with $z_i \rightarrow y$, and $\rho(x, z_i) \rightarrow \infty$. Now, we can find $\gamma_i \in \Gamma$, with $\rho(z_i, \gamma_i x) \leq t$. Since $z_i \in [x, y_i]$, we see that $\rho(x_i, x) - \rho(x_i, \gamma_i x) \geq \rho(z_i, x) - t \rightarrow \infty$. In particular, $\rho(x_i, \gamma_i x) < \rho(x_i, x)$ for all sufficiently large i , contradicting the choice of x_i as before.

We have thus proven the claim that $Y(\mathcal{B})/\Gamma$ is compact. Choosing any $r > 0$, we see finally, by Proposition 6.12, that Π/Γ is finite. \square

Proposition 6.14. *Suppose $\Pi \subseteq \partial X$ is a Γ -invariant set of bounded parabolic points with Π/Γ finite and \mathcal{B} an invariant system of horoballs for Π with $Y(\mathcal{B})/\Gamma$ compact. Then, every point of $\partial X \setminus \Pi$ is a conical limit point.*

Proof. Let $\mathcal{B} = (B(p))_{p \in \Pi}$. Given $r > 0$, we can find an r -separated collection, $\mathcal{B}' = (B'(p))_{p \in \Pi}$ of horoballs for Π . Without loss of generality, we can suppose that $B'(p) \subseteq B(p)$. By Lemma 6.4, $(B(p) \setminus \text{int } B'(p))/\text{stab}_\Gamma(p)$ is compact for each $p \in \Pi$. Now $Y(\mathcal{B}')/\Gamma$ is a union of $Y(\mathcal{B})/\Gamma$ and the projection to X/Γ of a finite number of sets of this form. It follows that $Y(\mathcal{B}')/\Gamma$ is compact. The result follows by Proposition 6.12. \square

Propositions 6.12, 6.13 and 6.14, together give us a number of equivalent ways of describing geometrical finiteness.

To relate this to Marden's definition [Mar], recall that a cusp region is a set of the form $B(p)/\text{stab}_\Gamma(p)$, where p is a bounded parabolic point in ∂X . We thus see that Γ is geometrically finite if and only if X/Γ can be written a union of a compact set together with a finite number of pairwise disjoint cusp regions. In particular, X/Γ has finitely many ends, and is quasiisometric to a wedge of rays.

Note also, that we have shown:

Proposition 6.15. *If Γ is geometrically finite, then there are finitely many orbits of parabolic points, and each parabolic point is bounded.*

This is also shown in [T3] for general geometrically finite convergence groups.

Suppose that Γ is geometrically finite. Let \mathcal{G} be the set of maximal parabolic subgroups. We have already observed that \mathcal{G} is peripheral

structure. We shall see in Section 9 that (Γ, \mathcal{G}) is a relatively hyperbolic group, and that ∂X is Γ -equivariantly homeomorphic to $\partial(\Gamma, \mathcal{G})$, as defined in Section 4.

7. SYSTEMS OF QUASICONVEX SETS

In this section, we describe a construction which, in some sense, is an inverse to that described in Section 4. There, we constructed a hyperbolic space with a preferred set of horoballs, starting with a fine hyperbolic graph. In this section, we start with a hyperbolic space, with a preferred system of quasiconvex sets (for example horoballs) and construct a graph as a kind of “nerve” of this system. More precisely, we construct a graph whose vertex set is the set of quasiconvex sets, and regard two such sets as adjacent in the graph if they are a bounded distance apart in X . Under appropriate hypotheses (that the system is “sufficiently separated”, “quasidense” and “locally finite”) we verify that this graph is fine and hyperbolic (Proposition 7.8). These conditions are automatic for a standard system of horoballs for a minimal geometrically finite action of a group on a proper hyperbolic space. As a corollary, we deduce that such a group is hyperbolic relative to the set of maximal parabolic subgroups (by the definitions of Section 4). In Section 9, we shall carry this analysis further to the study of the boundaries of such spaces.

In fact, with a more careful analysis, we can drop the assumption of sufficient separability (Proposition 7.12). This has an application to proving that a hyperbolic group is hyperbolic relative to quasiconvex subgroups (Theorem 7.11), though it will not be needed elsewhere.

As an intermediate step towards proving Proposition 7.8, we shall see that if we take any sufficiently separated collection of quasiconvex sets in a hyperbolic length space, and collapse each quasiconvex set to a point, then the resulting space is also hyperbolic. This is the geometric interpretation of Proposition 7.4.

We begin by giving the main definitions of this section. Let (X, ρ) be a hyperbolic length space. Let $(Q(p))_{p \in \Pi}$ be a collection of subsets of X , indexed by a set Π .

Definition. We say that $(Q(p))_{p \in \Pi}$ is *r-separated* if $\rho(Q(p), Q(q)) \geq r$ for all distinct $p, q \in \Pi$.

We say that $(Q(p))_{p \in \Pi}$ is *t-dense* if $X = N(\bigcup_{p \in \Pi} Q(p), t)$.

We say that $(Q(p))_{p \in \Pi}$ is *locally finite* if $\{p \in \Pi \mid \rho(x, Q(p)) \leq u\}$ is finite for all $x \in X$ and $u \geq 0$.

We say that $(Q(p))_{p \in \Pi}$ is *sufficiently separated* if it is *r-separated* for some $r \gg 0$ (sufficiently large in relation to the hyperbolicity constant).

We say that $(Q(p))_{p \in \Pi}$ is *quasidense* if it is t -dense for some $t \geq 0$.

For the moment, we shall just assume that $(Q(p))_{p \in \Pi}$ is a sufficiently separated system of closed (uniformly) quasiconvex sets. In particular, we require that $\text{diam}(\text{proj}_{Q(p)} Q(q)) \simeq 0$ whenever $p \neq q$ (see Lemma 5.2). Given $p \neq q$, we shall write $[Q(p), Q(q)]$ for some choice of path connecting $Q(p)$ to $Q(q)$ with $\text{length}[Q(p), Q(q)] \simeq \rho(Q(p), Q(q))$. As discussed in Section 5, any two such paths are close, so this choice will not matter to us.

We may define a path-pseudometric, σ , on X , by modifying ρ so that it is identically zero on each of the sets $Q(p)$. The hausdorffification of (X, σ) can be thought of as the path-metric space obtained by collapsing each of the sets $Q(p)$ to a point. Our first objective will be to show that (X, σ) is hyperbolic. We begin by making some more precise definitions.

In what follows, we shall abuse notation slightly by regarding a path in X as a subset of X , and denoting concatenation of paths as unions etc. Given a path, β , we shall write $\partial\beta \subseteq X$ for the unordered pair of endpoints of β .

By a *relative path*, (γ, α) , we mean a path γ , with a subset $\alpha \subseteq \gamma$ consisting of a disjoint union of a finite number of subpaths, $\alpha_1, \alpha_2, \dots, \alpha_n$, of γ , occurring in this order along γ , such that for each $i \in \{1, \dots, n\}$, there is some $p_i \in \Pi$ such that $\partial\alpha_i \subseteq Q(p_i)$. Moreover, we assume that $p_{i+1} \neq p_i$ for each i . (Note that $\partial\gamma$ may or may not meet α .) We write $n(\gamma, \alpha) = n$ and $L(\gamma, \alpha) = \text{length}(\gamma \setminus \alpha)$. We can write $\gamma = \beta_0 \cup \alpha_1 \cup \beta_1 \cup \dots \cup \alpha_n \cup \beta_n$, where β_0, \dots, β_n are the complementary paths of α in γ . Note that one or both of β_0 and β_n may be empty. For each $i \in \{1, \dots, n-1\}$, the path β_i connects $Q(p_i)$ to $Q(p_{i+1})$. In particular, we have $\text{length}\beta_i \geq \rho(Q(p_i), Q(p_{i+1})) \geq r$, where $r \gg 0$ is the constant featuring in the hypothesis that $(Q(p))_{p \in \Pi}$ is sufficiently separated. We thus have $L(\gamma, \alpha) = \sum_{i=0}^n \text{length}\beta_i \geq (n-1)r$. In particular, we see that $L(\gamma, \alpha)$ is bounded below by a fixed linear function of $n(\gamma, \alpha)$.

We say that a relative path (γ, α) is *efficient* if $p_i \neq p_j$ whenever $i \neq j$. We say that (γ, α) is *semipolygonal* if each arc α_i is a geodesic segment. We say that it is *polygonal* if it is semipolygonal and $\beta_i = [Q(p_i), Q(p_{i+1})]$ for all $i \in \{1, \dots, n-1\}$ and $\beta_0 = [x, Q(p_1)]$ and $\beta_n = [y, Q(p_n)]$ where x, y are the endpoints of γ , if either β_0 or β_n is non-empty.

We can also speak about a closed relative path, (γ, α) , where γ is a closed path in X , and α is finite disjoint union of subpaths. Thus the

components, $\alpha_1, \dots, \alpha_n$, are cyclically ordered and $\gamma = \alpha_1 \cup \beta_1 \cup \alpha_2 \cup \dots \cup \alpha_n \cup \beta_n$. We similarly define $n(\gamma, \alpha)$, $L(\gamma, \alpha)$, efficient closed paths, etc.

Given $x, y \in X$, we let $\sigma(x, y) = \inf\{L(\gamma, \alpha)\}$, where (γ, α) ranges over all relative paths from x to y . Of course, we could restrict to efficient polygonal paths. It's easily seen that σ is a path-pseudometric. (A σ -continuous path consists essentially of a sequence of ρ -continuous paths connecting different sets $Q(p)$. We can fill in the gaps with ρ -geodesics to give us a relative path, say (γ, α) . The σ -length of the original path equals $L(\gamma, \alpha)$.)

We aim to prove that (X, σ) is hyperbolic. To this end, we need some lemmas about relative paths.

Lemma 7.1. *There is a fixed linear function, H , such that if (γ, α) is an efficient closed semipolygonal path, then $\text{length } \alpha_i \leq H(L(\gamma, \alpha))$ for each component, α_i , of α .*

Proof. Let $\gamma = \alpha_1 \cup \beta_1 \cup \alpha_2 \cup \dots \cup \alpha_n \cup \beta_n$, so that $\partial\alpha_i \subseteq Q(p_i)$. To simplify notation, we shall verify the statement for α_1 .

Given $i \in \{1, \dots, n\}$, let x_i and y_i , be the endpoints of β_i . Thus, $x_i \in Q(p_i)$ and $y_i \in Q(p_{i+1})$. Let $x'_i = \text{proj}_{Q(p_1)} x_i$ and $y'_i = \text{proj}_{Q(p_1)} y_i$. Thus, $x'_1 = x_1 \in Q(p_1)$ and $y'_n = y_n \in Q(p_1)$. By Lemma 5.1, we have $\rho(x'_i, y'_i) \leq \rho(x_i, y_i) \leq \text{length } \beta_i$. Let $\beta'_i = [x'_i, y'_i]$. Then $\text{length } \beta'_i \leq \text{length } \beta_i$. Given each $i \in \{2, \dots, n\}$ we similarly define α'_i to be $[y_{i-1}, x_i]$. Since $y_{i-1}, x_i \in Q(p_i)$ and $p_i \neq p_1$ (since (γ, α) is efficient), we have $\text{length } \alpha'_i \simeq 0$, by Lemma 5.2. Now let $\delta = \beta'_1 \cup \alpha'_2 \cup \beta'_2 \cup \dots \cup \alpha'_n \cup \beta'_n$. Thus $\text{length } \delta - L(\gamma, \alpha) = \sum_{i=2}^n \text{length } \alpha'_i + \sum_{i=1}^n (\text{length } \beta'_i - \text{length } \beta_i)$ is bounded above by a linear function of $n(\gamma, \alpha)$ and hence by a linear function of $L(\gamma, \alpha)$. We deduce that $\text{length } \delta \leq H(L(\gamma, \alpha))$ for some fixed linear function, H . Now, α_1 and δ both connect x_0 to y_n and α_1 is geodesic. Thus, $\text{length } \alpha_1 \leq \text{length } \delta \leq H(L(\gamma, \alpha))$ as required. \square

In particular, it follows that $\text{length } \gamma \leq L(\gamma, \alpha) + n(\gamma, \alpha)H(L(\gamma, \alpha))$. Now, $n(\gamma, \alpha)$ is bounded above by a linear function of $L(\gamma, \alpha)$. This gives us a quadratic bound on $\text{length } \gamma$. In other words, we have shown:

Corollary 7.2. *There is a fixed quadratic function, S , such that if (γ, α) is an efficient closed semipolygonal path, then $\text{length } \gamma \leq S(L(\gamma, \alpha))$.*

(In fact, one could obtain a linear bound using Proposition 5.7. All we shall need here is a subexponential bound.)

Note that, if (γ, α) is relative path, and $\delta \subseteq \gamma$, then $(\delta, \delta \cap \alpha)$ is also a relative path. If (γ, α) is efficient, semipolygonal or polygonal, then so is $(\delta, \delta \cap \alpha)$. The appropriate way of defining quasigeodesics with respect

to the path-pseudometric, σ , is as follows. We say that a relative path, (γ, α) , is a σ -quasigeodesic if there is a linear function, say J , such that if $\delta \subseteq \gamma$ is a subpath, then $L(\delta, \delta \cap \alpha) \preceq J(\sigma(x, y))$, where x, y are the endpoints of δ . Note that $\sigma(x, y) \leq \rho(x, y)$. For our purposes, it will, in fact, be enough to assume that $L(\delta, \delta \cap \alpha) \preceq J(\rho(x, y))$.

Lemma 7.3. *Suppose that (γ, α) is an efficient semipolygonal path in X . If (γ, α) is a σ -quasigeodesic, then γ is a quasigeodesic in X with respect to the path-metric ρ . Moreover, the linear bound of the latter depends only on the hyperbolicity constant and the linear bound of the former.*

Proof. It's enough to verify that length γ is bounded by a linear function of $\rho(x, y)$, where x, y are the endpoints of γ (since we can apply the same argument to any subpath of γ). By Lemma 5.6, it's enough, in fact, to find a subexponential bound. Here, we shall give a quadratic bound.

To this end, we connect x to y by a geodesic $[x, y]$, and let δ be the closed path $\gamma \cup [x, y]$. Now, (δ, α) is a closed efficient semipolygonal path. Also, $L(\delta, \alpha) = L(\gamma, \alpha) + \rho(x, y)$, which is bounded by a linear function of $\rho(x, y)$ (since (γ, α) is a σ -quasigeodesic). Now, length $\gamma \leq \text{length } \delta \leq S(L(\delta, \alpha))$, where S is the quadratic function given by Corollary ???. Thus, length γ is quadratically bounded in terms of $L(\gamma, \alpha)$ and hence in terms of $\rho(x, y)$ as claimed. \square

We can use this to verify that (X, σ) is hyperbolic. We should first convince ourselves that every geodesic in (X, σ) can indeed be viewed as a σ -quasigeodesic, as we have defined it. Since we don't know that (X, σ) is a length space, we should view a geodesic in (X, σ) as a σ -continuous path whose σ -length is equal to the σ -distance between its endpoints up to a small additive constant. As discussed earlier, we can, in turn, view this as a sequence of paths, β_i , with each β_i connecting distinct sets $Q(p_i)$ and $Q(p_{i+1})$. The fact that is almost distance minimising means that it is necessarily the case that the p_i are all distinct (since we could simply omit any subpath with returns to the same set $Q(p_i)$, significantly reducing the σ -length). Now, filling in the gaps by ρ -geodesics, α_i , we obtain a relative path, (γ, α) , with the same endpoints, and with $L(\gamma, \alpha)$ equal to the σ -length of the original σ -geodesic. (In fact, if an endpoint of the path lies in a set $Q(p)$, we are free to attach a ρ -geodesic arc to any other point of $Q(p)$, so we can take any point $Q(p)$ as an endpoint of the path γ .) We see that (γ, α) is an efficient semipolygonal σ -quasigeodesic, by the earlier definitions. By Lemma 7.3, we see that γ is a ρ -quasigeodesic. We can now prove:

Proposition 7.4. *Let (X, ρ) be hyperbolic, and let $(Q(p))_{p \in \Pi}$ be a sufficiently separated system of uniformly quasiconvex sets, and let σ be the path pseudometric on X obtained by modifying the metric on each set $Q(p)$ to be zero. Then, the path-pseudometric space, (X, σ) , is hyperbolic.*

Proof. Suppose that $(\gamma_1, \gamma_2, \gamma_3)$ is a σ -geodesic triangle. From the preceding discussion, we can assume that each γ_i is a ρ -quasigeodesic, and so, by Lemma 5.5 remains within a ρ -bounded distance of a ρ -geodesic, δ_i , with the same endpoints (with the extra freedom to determine an endpoint lying in one of the sets $Q(p)$ as noted above). In this way, we obtain ρ -geodesic triangle, $(\delta_1, \delta_2, \delta_3)$. Now, any ρ -centre for $(\delta_1, \delta_2, \delta_3)$ will be a σ -centre for $(\gamma_1, \gamma_2, \gamma_3)$. This proves that (X, σ) is hyperbolic. \square

Suppose now that $(Q(p))_{p \in \Pi}$ is quasidense, i.e. it is t -dense for some $t > 0$. It's easily seen that $t \geq r/2$, where $r \gg 0$ is the constant featuring in the hypothesis that $(Q(p))_{p \in \Pi}$ is sufficiently separated. Let u be a constant greater than $2t$. We can construct a graph, K , with vertex set Π , by deeming two points $p, q \in \Pi$ to be adjacent if $\rho(Q(p), Q(q)) \leq u$.

Definition. We refer to K as the u -nerve of $(Q(p))_{p \in \Pi}$.

In fact, for the purposes of Section 9, we shall need a slightly more general notion. For this, we allow ourselves two constants, $u_2 \geq u_1 > 2t$, and say that K is a *nerve* for $(Q(p))_{p \in \Pi}$ if it contains the u_1 -nerve and is contained in the u_2 -nerve. When we speak about nerves, it will always be understood that the constant u_1 is sufficiently large (i.e. greater than $2t \gg 0$).

Now, it's easily seen that a nerve, K , is connected, and so has a combinatorial path metric, d_K . Moreover, by standard elementary arguments, we see that (K, d_K) is quasiisometric to (X, σ) . We observed in Section 5 that hyperbolicity is a quasiisometry invariant for path-pseudometric spaces. We can thus conclude immediately that:

Lemma 7.5. *Suppose $(Q(p))_{p \in \Pi}$ is a quasidense sufficiently separated system of quasiconvex sets. Then, any nerve for $(Q(p))_{p \in \Pi}$ is hyperbolic.*

(One could give a direct proof of this, without explicitly introducing the pseudometric, σ , though the argument becomes more complicated. A more general argument is given for Proposition 7.12.)

Suppose that π is a path in a nerve, K , of $(Q(p))_{p \in \Pi}$. For the moment, let's assume that $\pi = p_0 p_1 \dots p_n$ is finite, where $p_i \in \Pi$. For

each $i \in \{1, \dots, n-1\}$, we let β_i be the path $[Q(p_i), Q(p_{i+1})]$ from $Q(p_i)$ to $Q(p_{i+1})$. Thus, $\text{length } \beta_i \simeq \rho(Q(p_i), Q(p_{i+1})) \leq u_2$. As usual, we connect the paths, β_i , together by geodesic segments to give us a path $\gamma = \gamma(\pi) = \beta_0 \cup \alpha_1 \cup \beta_1 \cup \dots \cup \alpha_{n-1} \cup \beta_{n-1}$. Let $\alpha \subseteq \gamma$ be the union of the α_i . Thus, (γ, α) is a polygonal path. Note that $L(\gamma, \alpha) = \sum_{i=0}^{n-1} \text{length } \beta_i$ is bounded above by a linear function of n . If π is an arc, the (γ, α) is efficient. If π is also quasigeodesic in K , then (γ, α) is a σ -quasigeodesic. It follows that γ is a ρ -quasigeodesic. Note that we are free to attach geodesic arcs α_0 and α_n to either end of γ , with $\partial\alpha_0 \subseteq Q(p_0)$ and $\partial\alpha_n \subseteq Q(p_n)$, and the same argument shows that $\alpha_0 \cup \gamma \cup \alpha_n$ is a ρ -quasigeodesic, joining an arbitrary point of $Q(p_0)$ to an arbitrary point of $Q(p_n)$. As a result, we may conclude (via Lemma 5.5) that γ is a bounded distance from the path $[Q(p_0), Q(p_n)]$.

We can apply these constructions to rays and biinfinite paths in K as well as to cycles and circuits.

In summary, to each path π in K , we have associated a polygonal path $\gamma(\pi)$, or more properly $(\gamma(\pi), \alpha(\pi))$, in X with the following properties:

Lemma 7.6. *If π' is a subpath of π then $\gamma(\pi')$ is a subpath of $\gamma(\pi)$. If π is an arc, then $\gamma(\pi)$ is efficient. If π is a quasigeodesic arc, then $\gamma(\pi)$ is a quasigeodesic in (X, ρ) . If π is a finite quasigeodesic arc connecting p to q in Π , the $\gamma(\pi)$ is a bounded distance from $[Q(p), Q(q)]$. (Here the bounds given by the conclusions, depend only on the bounds implicit in the hypotheses.)*

We now finally introduce the assumption of local finiteness into the proceedings:

Lemma 7.7. *Suppose that $(Q(p))_{p \in \Pi}$ is a locally finite sufficiently separated system of quasiconvex sets. Then, any nerve of $(Q(p))_{p \in \Pi}$ is fine.*

(We have only formally defined “nerve” under the assumption that $(Q(p))_{p \in \Pi}$ is quasidense, though here this hypothesis is irrelevant.)

Proof. Suppose π is a circuit in K . Let $(\gamma(\pi), \alpha(\pi))$ be the efficient closed polygonal path given by Lemma 7.6. Now, there is a bound (linear in n) on the quantity $L(\gamma(\pi), \alpha(\pi))$, and hence, by Corollary 7.2, a bound (quadratic in n) on $\text{length } \gamma(\pi)$. This bound, say $C(n)$, depends only on n and the hyperbolicity constant.

Suppose $p_0, p_1 \in \Pi = V(K)$ are adjacent in K . Let $\delta = [Q(p_0), Q(p_1)]$. If π is a circuit of length n containing the edge p_0p_1 , then $\delta \subseteq \gamma(\pi)$.

Thus, $\gamma \subseteq N(\delta, C(n))$. If $q \in \pi \cap \Pi$, then $\rho(\delta, Q(q)) \leq C(n)$. By local finiteness of $(Q(p))_{p \in \Pi}$, there are only finitely many possibilities for q , and hence for the circuit π . In other words, there are only finitely many circuits of a given length in K containing any given edge. This shows that K is fine. \square

Combining Lemmas 7.5 and 7.7 we get:

Proposition 7.8. *Suppose that $(Q(p))_{p \in \Pi}$ is a locally finite quasidense sufficiently separated system of quasiconvex subsets of X . Then any nerve of $(Q(p))_{p \in \Pi}$ is fine and hyperbolic.*

Note that if X is proper (i.e. complete and locally compact), then the hypotheses that $(Q(p))_{p \in \Pi}$ is locally finite is redundant.

We now want to apply this to geometrically finite actions on a proper hyperbolic space.

Suppose then that (X, ρ) is a taut proper hyperbolic length space, and that Γ acts as a minimal geometrically finite group on X . Let $\Pi \subseteq \partial X$ be the set of (bounded) parabolic points. We can view Π as a Γ -set, which we have already seen is cofinite (Proposition 6.15). The vertex stabilisers are precisely the maximal parabolic subgroups.

By Proposition 6.13, we can find a sufficiently separated system, $(B(p))_{p \in \Pi}$, of horoballs for Γ . Assuming that $\Pi \neq \emptyset$, this system is necessarily quasidense. Let K be a Γ -invariant nerve for the system $(B(p))_{p \in \Pi}$. (For example, take K to be the u -nerve for some sufficiently large u .) Since $(X \setminus \bigcup_{p \in \Pi} \text{int } B(p))$ is compact, it is easily seen that K/Γ is finite. Thus, K is a (Γ, Π) -graph. By Proposition 7.8, K is fine and hyperbolic. If, in addition, we assume that all peripheral subgroups are finitely generated, then we conclude that Π is a hyperbolic Γ -set. In summary, we have shown:

Proposition 7.9. *Suppose that Γ acts as a (minimal) geometrically finite group on a (taut) proper hyperbolic length space, X . Suppose that all maximal parabolic subgroups are finitely generated. Let $\Pi \subseteq \partial X$ be the set of (bounded) parabolic points. Then, Π is a hyperbolic Γ -set.*

(From the discussion of Section 6, we see that the hypotheses of minimality and tautness are superfluous here.)

In other words, we have shown that a geometrically finite group is hyperbolic relative to its maximal parabolic subgroups.

Putting these various results together, we can now conclude:

Theorem 7.10. *Definition 1 and Definition 2 of the Introduction are equivalent.*

Proof. To deduce Definition 1 from Definition 2, put together Proposition 4.9, Lemma 3.7 and Theorem 3.8, as discussed at the end of Section 4. For the converse, use Proposition 7.9 \square

It is possible for the same group, Γ , to act as a minimal geometrically finite group, on two taut proper hyperbolic length spaces, say X and X' , with same set, \mathcal{G} , of peripheral (i.e. maximal parabolic) subgroups. The spaces, X and X' , need not be quasiisometric. However, we shall see in Section 9, that the boundaries ∂X and $\partial X'$ are Γ -equivariantly homeomorphic.

Note that one of these spaces, X' , could be taken to be the space arising from the construction of Section 4 (for example, starting with the nerve of X). Now, $\partial X'$ is, by definition, the boundary, $\partial(\Gamma, \mathcal{G})$, of the relatively hyperbolic group (Γ, \mathcal{G}) . We shall therefore conclude that ∂X is Γ -equivariantly homeomorphic to $\partial(\Gamma, \mathcal{G})$.

We shall finish this section with a proof of the following result which is described in [Gers] (although the argument presented there contains a gap, of no consequence to the remainder of that paper).

Theorem 7.11. *Suppose that Γ is a hyperbolic group, and that \mathcal{G} is a collection of infinite quasiconvex subgroups. Suppose that \mathcal{G} is invariant under conjugacy, and that there are finitely many conjugacy classes. Then, Γ is weakly hyperbolic relative to \mathcal{G} . If, in addition, each element of \mathcal{G} is equal to its normaliser, and the intersection of any distinct pair of elements of \mathcal{G} is finite, then Γ is hyperbolic relative to \mathcal{G} .*

The proof amounts to a variation on Proposition 7.8 where we drop the assumption of sufficient separation. Of course, this gives an alternative route to Proposition 7.8 from that already described. However, since the additional techniques are not required elsewhere, we prefer to view it as a supplement to the main theme of this paper.

Let us return, for the moment, to the general setting of a hyperbolic space (not necessarily proper). We shall show:

Proposition 7.12. *Let (X, ρ) be a hyperbolic space, and let $(Q(p))_{p \in \Pi}$ be a quasidense collection of uniformly quasiconvex subsets of X . Then, any nerve of $(Q(p))_{p \in \Pi}$ is hyperbolic.*

Recall what this means. We are assuming that $(Q(p))_{p \in \Pi}$ is t -dense for some $t \geq 0$. By definition, a nerve, K , is a graph with vertex set Π , and with $p, q \in \Pi$ adjacent if $\rho(Q(p), Q(q)) \leq u_1$ and not adjacent if $\rho(Q(p), Q(q)) > u_2$ where $u_2 \geq u_1 > 2t$ are fixed. For Proposition 7.12, we shall need to insist that $u_1 \gg 0$. This condition was automatically satisfied in the case where the quasiconvex sets were sufficiently separated.

The argument is similar to that of the previous results, except that we need to bring Proposition 5.7 into play. As usual, we shall verify the thin triangle property.

Suppose that $\pi = p_0 p_1 \cdots p_n$ is a geodesic arc of length n in K . Thus, $p_0, p_1, \dots, p_n \in \Pi$ are distinct. Let $\gamma = \beta_0 \cup \alpha_1 \cup \beta_1 \cup \cdots \cup \alpha_{n-1} \cup \beta_{n-1}$ be a broken geodesic, where β_i is a geodesic of length $\leq u_2$ connecting $Q(p_i)$ to $Q(p_{i+1})$, and α_i is a geodesic with both endpoints in $Q(p_i)$. Among all such broken geodesics, we choose γ to have length close to minimal.

We claim that γ has bounded backtracking. In other words, there is some $l \geq 0$, depending only on the hyperbolicity constant (and quasiconvexity constant) and the quasidensity constant, t , such that the backtracking of γ , in the sense defined in Section 5, is at most l . Since the geodesics β_i all have bounded length, we just want to check if $i < j$ then α_i and α_j don't backtrack a distance more than l .

Suppose they did. Choose points $z \in \alpha_i$ and $w \in \alpha_j$ with $\rho(z, w) \simeq 0$, and such that z is closest to the initial endpoint of α_i with the property that such a w exists. By quasiconvexity, we can find $z' \in Q(p_i)$ and $w' \in Q(p_j)$ close to z and w respectively. Since $\rho(z', w') \simeq 0 \ll u_1$, we have $\rho(z', w') \leq u_1$, so that p_i and p_j are adjacent in K . Since π is geodesic, it follows that $j = i + 1$. We now replace β_i with $[z', w']$ and adjust α_i and α_{i+1} accordingly. By choosing l sufficiently large (which we can do a-priori), we see that this significantly reduces the length of γ , contradicting approximate minimality. This proves the claim that γ has bounded backtracking.

Suppose that γ connects $x \in Q(p_0)$ to $y \in Q(p_n)$. Since π is geodesic, and $(Q(p))_{p \in \Pi}$ is quasidense, we see that n is bounded above by a linear function of $\rho(x, y)$. Also, by Proposition 5.7, it follows that $\text{length } \gamma - \rho(x, y)$ is bounded by a linear function of n . Thus, $\text{length } \gamma$ is bounded by a linear function of $\rho(x, y)$.

Now, this argument applies equally well to any subpath of γ , and so we see that γ is, by definition, a quasigeodesic. (In fact, by Lemma 5.6, it we see that any subexponential bound in Proposition 5.7 would suffice.)

From here, the argument proceeds as before. Given a geodesic triangle, (π_1, π_2, π_3) in K , we get a quasigeodesic triangle, $(\gamma_1, \gamma_2, \gamma_3)$ in X . By Lemma 5.5, this has a centre in X . By quasidensity of $(Q(p))_{p \in \Pi}$, this gives rise to a centre of (π_1, π_2, π_3) in K as required.

We finally note that since $u_1 > 2t$, by quasidensity, K , is connected. This proves Proposition 7.12.

In general, of course, the nerve need not be fine. For this, we need an additional hypothesis on our collection of quasiconvex sets.

We shall say that $(Q(p))_{p \in \Pi}$ has *bounded penetration* (cf. [Fa]) if, given any $r \geq 0$, there is some $D \geq 0$ such that for all distinct $p, q \in \Pi$, the set $N(Q(p), r) \cap N(Q(q), r)$ has diameter at most D . In fact, it is enough to assume this for some fixed $r \gg 0$ — the statement for all r then follows. Another way to view this condition is that the “projection” of any set $Q(p)$ onto any other set $Q(q)$ has bounded diameter. Note that bounded penetration is clearly implied by sufficient separation.

Under this additional hypothesis, it is clear that if π is any arc in K , and γ is an associated broken geodesic in X as constructed above, then γ has bounded backtracking. By applying the same argument to a circuit π , we see that there is a (linear) bound on the length of γ in terms of the length of π .

The following follows exactly as in Lemma 7.7:

Lemma 7.13. *Suppose that $(Q(p))_{p \in \Pi}$ is a locally finite collection of uniformly quasiconvex sets with bounded penetration. Then any nerve of K is fine.*

We can now prove Theorem 7.11:

Proof. Suppose that \mathcal{G} is a disjoint union of conjugates H_1, \dots, H_k . Any quasiconvex subgroup of a hyperbolic group is of finite index in its commensurator (in this case, the unique maximal commensurable subgroup, and the setwise stabiliser of its limit set). This commensurator is itself quasiconvex, and equal to its normaliser. Thus, without loss of generality, we can assume that each element of \mathcal{G} is equal to its normaliser. Let \mathcal{G} be a disjoint union of conjugates of $H_1, \dots, H_k \in \mathcal{G}$.

Let (X, ρ) be a graph on which Γ acts freely with finite quotient, X/Γ . We can assume that the graph X/Γ has at least k vertices. For each $i \in \{1, \dots, k\}$, choose a vertex, v_i , of X , so that v_1, \dots, v_k project to distinct vertices of X/Γ .

Now, if $G \in \mathcal{G}$, then $G = gH_i g^{-1}$ for some $i \in \{1, \dots, k\}$ and $g \in \Gamma$. Let $Q(G) = gH_i v_i$. Since H_i is equal to its normaliser, this is independent of the choice of g . Moreover this construction is Γ -equivariant, i.e. $Q(gGg^{-1}) = gQ(G)$ for all $G \in \mathcal{G}$ and $g \in \Gamma$. Also, if $gQ(G) \cap Q(G) \neq \emptyset$, then $g \in G$.

The sets $Q(G)$ are uniformly quasiconvex. Moreover, they are disjoint, and so form a locally finite collection. Let K be a nerve of $(Q(G))_{G \in \mathcal{G}}$. Thus, \mathcal{G} is precisely the set of vertex stabilisers of K . By Proposition 7.12, K is hyperbolic. By local finiteness of $(Q(G))_{G \in \mathcal{G}}$, there are only finitely many Γ -orbits of edges. Thus, by definition, Γ is weakly hyperbolic relative to \mathcal{G} .

For the second statement, we want to verify that $(Q(G))_{G \in \mathcal{G}}$ has bounded penetration. Suppose, to the contrary, that there are sequences of vertices, $(x_1)_i, (y_1)_i, (z_1)_i$ and $(w_1)_i$, and groups $G_i, G'_i \in \mathcal{G}$, with $G_i \neq G'_i$, $x_i, y_i \in Q(G_i)$, $x'_i, y'_i \in Q(G'_i)$, with $\rho(x_i, x'_i)$ and $\rho(y_i, y'_i)$ bounded, and with $\rho(x_i, y_i) \rightarrow \infty$. By local finiteness of X , we can assume that y_i and y'_i are constant, hence $G_i = G$ and $G'_i = G'$ are constant. In summary, we have sequences $x_i \in Q(G)$ and $x'_i \in Q(G')$, with $\rho(x_i, x'_i)$ bounded, but with $\rho(x_i, x_j)$ unbounded. Since X is locally finite, it follows that on passing to a subsequence, we can find $g_i \in \Gamma$ such that $x_i = g_i x_0$ and $x'_i = g_i x'_0$. But now it follows that $g_i \in G \cap G'$. By hypothesis there are only finitely many such g_i giving us a contradiction.

Finally, by Lemma 7.13, it follows that K is fine. Every element of \mathcal{G} , being quasiconvex, is finitely generated. Thus, Γ is hyperbolic relative to \mathcal{G} . \square

8. FINE HYPERBOLIC GRAPHS.

In this section, we shall explore further the properties of fine hyperbolic graphs. This seems to be a fairly robust notion, and such graphs share many of the properties of locally finite hyperbolic graphs. The discussion will set the scene for the description of boundaries of relatively hyperbolic groups in Section 9.

In the case of 2-vertex connected graphs, we can construct the space $X(K^n)$ where n is any sufficiently large number so that K^n is simplicially hyperbolic, as discussed in Section 3. Indeed this space is potentially a useful tool in studying such graphs, as we shall see in Section 9. In particular, the boundary we shall define for K can be canonically identified with the usual Gromov boundary of $X(K^n)$.

Recall that a graph is fine if every edge is contained in finitely many circuits of length at most n . We do not place any bound on this number, though it is clear that one can easily reinterpret many of the results here for “uniformly fine” hyperbolic graphs, where this number is controlled as a function of n . Thus, various sets which are proven to be finite here will be seen to have bounded cardinality in that case. All the fine graphs of interest in studying relatively hyperbolic groups are uniformly fine in this sense.

Another point worth noting is that for hyperbolic graphs (i.e. graphs which are Gromov hyperbolic as path metric spaces), it is sufficient to check fineness for circuits of a certain bounded length:

Proposition 8.1. *Given $k \in \mathbb{N}$ there is a number $n(k)$ such that if K is a hyperbolic graph of hyperbolicity constant k with the property that*

any set of circuits of length at most $n(k)$ containing any given edge is finite, then K is fine. (In fact, $n(k)$ can be taken to be a linear function of k .)

Proof. The argument is a variation one of the standard arguments for deriving a linear isoperimetric inequality from the “thin triangles” property (cf. [Bo1]). In general, the idea is to inductively cut any cycle into (a linearly bounded number of) cycles of bounded length. Here, we perform essentially the same construction, but taking care to ensure that if we start with a circuit, then we end up only with circuits. Only the inductive step is relevant here. We prove, by induction on n starting at $n = n(k)$, that there are only finitely many circuits of length n containing any given edge.

Suppose then, that γ is a circuit in K of length greater than $n(k)$. By a standard argument (see for example, Lemma 6.1.4 of [Bo1]) we can cut γ into two subarcs, γ_1 and γ_2 , and find a geodesic arc, α , in K , such that γ_1 , γ_2 and α all have the same pair of endpoints, and such that $\text{length } \alpha < \min\{\text{length } \gamma_1, \text{length } \gamma_2\}$. (It is from this result that the constant $n(k)$ arises.) Moreover, we can suppose that α meets γ only in these endpoints. For, in the general case, let $\alpha_1, \dots, \alpha_m$ be the closures of the components of $\alpha \setminus \gamma$. By considering the successive distances between the endpoints of these arcs in the induced path-metric on γ , we see easily that there is some i such that the endpoints of α_i cut γ into two subarcs each of length greater than $\text{length } \alpha_i$. We can then replace α by α_i .

Now, let $\beta_i = \alpha \cup \gamma_i$. Thus, β_1 and β_2 are circuits with $\gamma \subseteq \beta_1 \cup \beta_2$, and $\text{length } \beta_i < \text{length } \gamma$. Suppose we fix an edge $e \in E(\gamma)$. Without loss of generality, $e \in E(\beta_1)$. By the induction hypothesis, there are only finitely many possibilities for β_1 . Since $\beta_1 \cap \beta_2$ contains at least one edge, there are thus only finitely many possibilities for β_2 . Since $\gamma \subseteq \beta_1 \cup \beta_2$, there are finitely many possibilities for γ , as required. \square

From now on, we shall assume that K is a fine hyperbolic graph with combinatorial metric $d = d_K$. Let ∂K be the usual Gromov boundary of K (as discussed in Section 5). This carries a natural topology, though for the moment, we just regard it as a set.

Suppose $f, g : \mathbb{N} \rightarrow \mathbb{N}$ are functions. We write $f \leq g$ to mean that $f(n) \leq g(n)$ for all $n \in \mathbb{N}$. We write $1_{\mathbb{N}}$ for the identity function. We are interested here only in linear functions, or function bounded above by linear functions. If $f \geq 1_{\mathbb{N}}$ is such a function, an “ f -quasigeodesic” arc is an arc β such that $\text{length } \alpha \leq f(d(x, y))$ for any subarc, α of β , where x, y are the endpoints of α . Thus, a $1_{\mathbb{N}}$ -quasigeodesic arc is a geodesic.

We stress that throughout this section we are only interested in quasi-geodesic arcs, i.e. paths which never return to the same vertex. This is essential to the application of the fineness assumption.

Given $a, b \in V(K) \cup \partial K$, we write $\mathcal{F}_f(a, b)$ for the set of f -quasigeodesic arcs from a to b , and write $F_f(a, b) = \bigcup \mathcal{F}_f(a, b)$. Thus, $F_f(a, b)$ is a connected subgraph of K . We shall abbreviate $\mathcal{F}_{1_{\mathbb{N}}}(a, b)$ and $F_{1_{\mathbb{N}}}(a, b)$ to $\mathcal{F}(a, b)$ and $F(a, b)$ respectively. Clearly, if $f \leq g$, then $F_f(a, b) \subseteq F_g(a, b)$. Also, if $c \in F(a, b)$, then $F(a, c) \subseteq F(a, b)$. For any $a \in \partial K$, we set $F_f(a, a) = \mathcal{F}_f(a, a) = \emptyset$.

If $\alpha, \beta \in \mathcal{F}_f(a, b)$, then there is a constant $r \geq 0$, depending only on the hyperbolicity constant and f , such that $\beta \subseteq N(\alpha, r)$.

Note that if $a, b \in V(K)$, then $F(a, b)$ is finite, by fineness. More generally, we have:

Lemma 8.2. *For any $a, b \in V(K) \cup \partial K$ and any linear function, f , the graph $F_f(a, b)$ is locally finite.*

Proof. Let α be any geodesic from a to b . Since $F_f(a, b) \subseteq N(\alpha, r)$, it's enough to show that $F_f(a, b) \cap N(x, r)$ is finite for all $x \in V(\alpha)$. For simplicity, we shall assume that $a, b \in \partial K$, or at least, that $d(x, a) \geq 4r$ and $d(x, b) \geq 4r$. The remaining cases are dealt with similarly.

Let $y_1, y_2 \in V(\alpha)$ be vertices on either side of α such that $d(x, y_1) = d(x, y_2) = 4r$. Suppose that $z \in V(F_f(a, b)) \cap N(x, r)$. Then, $z \in V(\beta)$ for some $\beta \in \mathcal{F}_f(a, b)$. The point z divides β into two subarcs β_1 and β_2 . Since $z \in N(x, r)$ and $\beta \subseteq N(\alpha, r)$, it is easily seen that $d(y_1, \beta_1) \leq 2r$ and $d(y_2, \beta_2) \leq 2r$. Let δ_i be a geodesic connecting y_i to the nearest point, z_i , in β_i . Let ϵ be subarc of β connecting w_1 to w_2 , and let $\gamma = \delta_1 \cup \epsilon \cup \delta_2$. Thus, γ is an arc of bounded length connecting y_1 and y_2 and containing z . By fineness, there are only finitely many possibilities for γ and hence for z . \square

Using Lemma 8.2, it's not hard to see that if $a, b, c \in V(K) \cup \partial K$, then the set $V(F_f(a, b)) \cap V(F_f(b, c)) \cap V(F_f(c, a))$ is finite. Also, for suitable f , depending only on the hyperbolicity constant, this set is non-empty, provided that no two of a, b, c are both equal to the same point of ∂K . This leads naturally to a strong notion of a centre of a triple of three points. We shall not pursue this notion further here, so some of the subsequent constructions could be rephrased in these terms.

Lemma 8.3. *Suppose $f \geq 1_{\mathbb{N}}$ is a linear function, and $e \in E(K)$. Then there is a finite set $E_f(e) \subseteq E(K)$ with the property that if α and β are f -quasigeodesic arcs connecting the same pair of points of $V(K) \cup \partial K$ and with $e \in E(\alpha)$, then $E_f(e) \cap E(\beta) \neq \emptyset$.*

Proof. Suppose $a_1, a_2 \in V(K) \cup \partial K$ and $\alpha, \beta \in \mathcal{F}_f(a_1, a_2)$, with $e \in E(\alpha)$. We can suppose that $e \notin E(\beta)$. We aim to construct a circuit, γ , of bounded length, with $e \in E(\gamma)$ and $E(\beta) \cap E(\gamma) \neq \emptyset$. By fineness, there are only finitely many possibilities for γ . We can thus set $E_f(e)$ to be the union of all sets $E(\gamma)$ as γ ranges over all such possibilities.

Let α_1, α_2 be the components of $\alpha \setminus e$ with endpoints a_1, a_2 respectively. Thus, $\alpha = \alpha_1 \cup e \cup \alpha_2$. Let ϵ be any geodesic connecting a_1 and a_2 . Thus, $\alpha, \beta \in N(\epsilon, r)$ for some fixed $r \geq 0$.

Suppose first that α and β meet only in their endpoints (i.e. $\alpha_i \cap \beta = \{a_i\}$ if $a_i \in V(K)$ and $\alpha_i \cap \beta = \emptyset$ if $a_i \in \partial K$). We construct arcs δ_i connecting α_i to β as follows.

Let $x \in V(\epsilon)$ be a point with $d(x, e) \leq r$. If $d(x, a_i) \leq 4r$, set $\delta_i = \{a_i\}$. If $d(x, a_i) > 4r$, let $x_i \in V(\epsilon)$ be the vertex of ϵ a distance $4r$ from x in the direction of a_i . Since $\alpha_i \subseteq N(\epsilon, r)$ and $d(x_i, \alpha_i) \leq r$, we see easily that $d(x_i, \alpha_i) \leq 2r$. Let $y_i \in V(\alpha_i)$ be a vertex with $d(x_i, y_i) \leq 2r$. We also choose $z_i \in V(\beta)$ with $d(x_i, z_i) \leq 2r$. Let γ_i be any geodesic from y_i to z_i . We can find a subarc, δ_i , of γ_i which meets α_i in a single point, y'_i , and meets β in a single point, z'_i . By construction, we see that δ_1, δ_2 and e are all disjoint. Let α' be the subarc of α between y'_1 and y'_2 , and let β' be the subarc of β between z'_1 and z'_2 . Note that $e \in E(\alpha')$ and $E(\beta') \neq \emptyset$. Let γ be the circuit $\alpha' \cup \delta_1 \cup \beta' \cup \delta_2$. We see that γ has bounded length, $e \in E(\gamma)$ and $E(\gamma) \cap E(\beta) \neq \emptyset$ as required.

To deal with the general case, note that some subarc, β_0 , of β must connect a vertex of α_1 to a vertex of α_2 (where we allow for the possibility that $\beta_0 \cap \alpha_i = \emptyset$ and β_0 has endpoint $a_i \in \partial K$). By choosing a minimal such subarc, we can suppose that β_0 meets α only in these points. Let α_0 be the subarc of α connecting these points. Thus, $e \in \alpha_0$. By replacing α by α_0 and β by β_0 , we are reduced to the situation already dealt with. \square

The next step will be to define a boundary of K . The construction is analogous to that described in [Bo5]. This boundary will be a compact hausdorff space, incorporating $V(K)$ as a dense subset. Vertices of finite degree in K will be isolated. In the case where K is 2-vertex connected, this boundary can be identified with the usual Gromov boundary of the complex $X(K^n)$ for sufficiently large n , as constructed in Section 4. Details of this will be discussed in Section 9.

Let $\Delta K = V(K) \cup \partial K$. We put a topology on ΔK as follows. Given $a \in \Delta K$ and $A \subseteq V(K)$, we write

$$M(a, A) = \{b \in \Delta K \mid A \cap V(F(a, b)) \subseteq \{a\}\}.$$

If $A, B \subseteq M$, then $M(a, A) \cap M(a, B) = M(a, A \cup B)$. In particular, if $A \subseteq B$, then $M(a, B) \subseteq M(a, A)$. Note also that $a \in M(a, A)$.

We shall take the collection $\{M(a, A)\}_A$ as A ranges over all finite subsets of $V(K)$, as a basis for the topology about a . In other words, we define a set $U \subseteq \Delta K$ to be open if for all $a \in U$ there is a finite subset $A \subseteq V(K)$ such that $M(a, A) \subseteq U$. The above observations show that this is indeed a topology. We shall eventually see that $M(a, A)$ is a neighbourhood of a in this topology, so the collection $\{M(a, A)\}_A$ is a neighbourhood base for a . First, however, we describe a number of equivalent bases for a .

Note that $M(a, A)$ is the set of points $b \in \Delta K$ such that each geodesic from a to b meets A , if at all, in a . We could similarly use quasigeodesics. Thus, given a function, $f : \mathbb{N} \rightarrow \mathbb{N}$, with $f \geq 1_{\mathbb{N}}$ and f bounded by a linear function, we define

$$M_f(a, A) = \{b \in \Delta K \mid A \cap V(F_f(a, b)) \subseteq \{a\}\}.$$

Thus, if $f = 1_{\mathbb{N}}$, we recover $M(a, A)$. We similarly define $M'_f(a, A)$ to be the set of points $b \in \Delta K$ such that there exists an f -quasigeodesic arc, β , from a to b with $V(\beta) \cap A \subseteq \{a\}$.

We can do a similar thing with edges. Given a subset $I \subseteq E(K)$, we write

$$P_f(a, I) = \{b \in \Delta K \mid I \cap E(F_f(a, b)) = \emptyset\}.$$

We set $P'_f(a, I)$ to be the set of $b \in \Delta K$ such that there exists an f -quasigeodesic arc from b to a with $E(\beta) \cap I = \emptyset$. In each of the above definitions we shall drop the subscript f if $f = 1_{\mathbb{N}}$, i.e. if we are dealing with geodesics.

We claim that we can obtain equivalent bases for the point $a \in \Delta K$ by taking any of the collections $\{M_f(a, A)\}_A$, $\{M'_f(a, A)\}_A$, $\{P_f(a, I)\}_I$ or $\{P'_f(a, I)\}_I$, where $f \geq 1_{\mathbb{N}}$ is any fixed function bounded above by a linear function, and where A or I ranges over all finite subsets of $V(K)$ or $E(K)$ respectively. We shall also see that if $a \in V(K)$, then it is sufficient to allow I to range only over finite subsets of edges incident on a .

To prove these facts, we first make some trivial observations. Firstly, note that $M_f(a, A) \subseteq M'_f(a, A)$ and $P_f(a, I) \subseteq P'_f(a, I)$. Also, if $f \leq g$, then

$$M_f(a, A) \supseteq M_g(a, A)$$

$$M'_f(a, A) \subseteq M'_g(a, A)$$

$$P_f(a, I) \supseteq P_g(a, I)$$

$$P'_f(a, I) \subseteq P'_g(a, I).$$

Now, suppose $I \subseteq E(K)$ is finite and $f \geq 1_{\mathbb{N}}$. Let $J = \bigcup_{e \in I} E_f(e)$, where $E_f(e)$ is the finite set of edges associated to e as given by Lemma 8.3. We see that $P'_f(a, J) \subseteq P_f(a, I)$. Since $P_f(a, I) \subseteq P'_f(a, I)$, we deduce the equivalence of $\{P_f(a, I)\}_I$ and $\{P'_f(a, I)\}_I$ as I ranges over all finite subsets of $E(K)$.

Given $f, g \geq 1_{\mathbb{N}}$, let $h = \max\{f, g\}$. Putting the above facts together, we derive the equivalence of $\{P_f(a, I)\}_I$ with $\{P_h(a, I)\}_I$ and hence with $\{P_g(a, I)\}_I$. Thus, up to equivalence, these bases are independent of f .

Now, given any finite $I \subseteq E(K)$, let $A(I) \subseteq V(K)$ be the set of vertices incident on some element of I . Thus, $M_f(a, A(I)) \subseteq P_f(a, I)$, and $M'_f(a, A(I)) \subseteq P'_f(a, I)$.

Conversely, given any finite set $A \subseteq V(K)$, let $I(A)$ be the set of edges of $\bigcup_{b \in A} E(F_f(a, b))$ which are incident on some element of A . By Lemma 8.2, $I(A)$ is finite. We now have $P_f(a, I(A)) \subseteq M_f(a, A)$ and $P'_f(a, I(A)) \subseteq M'_f(a, A)$.

We have thus derived the equivalence with the bases $\{M_f(a, A)\}_A$ and $\{M'_f(a, A)\}_A$ as A ranges over all finite subsets of $V(K)$.

Finally, note that if $a \in V(K)$, then in the last but one paragraph, we could instead have taken $I(A)$ to be the set of edges of $\bigcup_{b \in A} E(F_f(a, b))$ which are incident to a . We again get $P_f(a, I(A)) \subseteq M_f(a, A)$ and $P'_f(a, I(A)) \subseteq M'_f(a, A)$. We thus derive the equivalence with the bases $\{P_f(a, I)\}_I$ and $\{P'_f(a, I)\}_I$ as I ranges over finite subsets of edges incident on a .

This proves all the equivalences we previously asserted.

In applications, a certain class of quasigeodesic arc frequently arises. Suppose α and β are geodesics meeting only at a common endpoint, $a \in V(K)$. Suppose, moreover, that a is the nearest point of α to each point of β . Then $\alpha \cup \beta$ is an f_0 -quasigeodesic, where we can take f_0 to be a fixed linear function (with multiplicative constant 1). We shall abbreviate $M_{f_0}(a, A)$ to $M_0(a, A)$ etc. Note that if γ is any geodesic, and $x \in \Delta K$, we can “drop a perpendicular”, β , from x to γ at some point $y \in V(\gamma)$, which we call “foot” of the perpendicular. Now, y separates γ into two (possibly trivial) subarcs, α_1 and α_2 , so that $\beta \cup \alpha_1$ and $\beta \cup \alpha_2$ satisfy the above conditions, and so are f_0 -quasigeodesic arcs. In the case where $x \in \partial K$, the existence of a perpendicular calls for a diagonal sequence argument, noting that there are only finitely many possibilities for y , and hence for each of the successive vertices of β , using Lemma 8.2. The same argument allows us to drop a perpendicular from any element of ΔK to any finite subset

of $V(K)$. This last observation will be used in the proof of Proposition 8.6.

We are finally ready to verify that $\{M(a, A)\}_A$ is indeed a neighbourhood base for any $a \in \Delta K$. In other words, given any finite $A \subseteq V(K)$, we can find a finite $B \subseteq V(K)$ such that if $b \in M(a, B)$ then there exists a finite $C \subseteq V(K)$ such that $M(b, C) \subseteq M(a, A)$. In fact, from the equivalence of the bases we have described, it is enough to verify the following statement. Given any finite $A \subseteq V(K)$ and $b \in M_0(a, A)$ we have $M_0(b, A) \subseteq M(a, A)$.

We can suppose that $b \neq a$. Given any $c \in M_0(b, A)$, let α be any geodesic from a to c . We drop a perpendicular, β from b to α , whose foot divides α into two subarcs α_1 and α_2 so that $\gamma_1 = \beta \cup \alpha_1$ and $\gamma_2 = \beta \cup \alpha_2$ are both f_0 -quasigeodesic arcs. Since $b \in M_0(a, A)$, we have $\gamma_1 \cap A \subseteq \{a\}$. Since $c \in M_0(b, A)$, we have $\gamma_2 \cap A \subseteq \{b\}$. Since $\alpha \subseteq \gamma_1 \cup \gamma_2$, we conclude that $\alpha \cap A \subseteq \{a\}$. Since α was arbitrary, it follows that $c \in M(a, A)$. Thus $M_0(b, A) \subseteq M(a, A)$ as claimed.

We remark that one can go on to show fairly easily that the collection $\{M(a, A)\}$ as a ranges over V and A ranges over all finite subsets of V , forms a base for the topology on ΔK . It thus follows that ΔK is second countable.

We now move on to study properties of the boundary, ΔK . Firstly, we note:

Lemma 8.4. *ΔK is hausdorff.*

Proof. Suppose $a, b \in \Delta K$ are distinct. Choose an edge $e \in E(F(a, b))$. Thus e lies on a geodesic α from a to b . By dropping a perpendicular to α as in the previous argument, we see that $P_0(a, \{e\}) \cap P_0(b, \{e\}) = \emptyset$. \square

Proposition 8.5. *The subspace topology on ∂K induced from ΔK agrees with the usual boundary topology on ∂K .*

Proof. In the boundary topology, a neighbourhood base for a point $a \in \partial K$ is given by the collection $\{M(a, A)\}_A$ as A ranges over all bounded subsets of $V(K)$. We therefore need to replace bounded subsets by finite subsets. In other words, given a bounded subset $A \subseteq V(K)$ we want to find a finite subset $B \subseteq V(K)$ such that $M(a, B) \subseteq M(a, A)$.

Fix any vertex $c \in V(K)$. By Lemma 8.2, the set $V(F_0(a, c))$ is locally finite. Thus, for any $r \in \mathbb{N}$, the set $B(r) = N(c, r) \cap V(F_0(a, c))$ is finite. We claim that $M(a, B(r)) = M(a, N(c, r))$. To see this, suppose that $b \notin M(a, N(c, r))$. Then there is some geodesic α from a to b with $d(c, \alpha) \leq r$. Let β be a geodesic from c to a nearest

point, $x \in V(\alpha)$. Concatenating β with a subarc of α , we get an f_0 -quasigeodesic arc connecting c to a and containing the point x . Thus, $x \in B(r)$, and so $\alpha \cap B(r) \neq \emptyset$. It follows that $b \notin M(a, B(r))$. Since $M(a, N(c, r)) \subseteq M(a, B(r))$, the claim follows.

Now, any bounded set, $A \subseteq V(K)$, lies inside $N(c, r)$ for some $r \geq 0$, and so $M(a, B(r)) \subseteq M(a, A)$ as required. \square

We note that $V(K)$ is dense in ΔK . This is easily seen directly from the definition of the topology. If x is any point of ∂K , choose any geodesic ray in K converging on x . The sequence of vertices of this geodesic tends to x in the topology of ΔK .

We also remark that if $x \in V(K)$, and $n \in \mathbb{N}$, then the n -ball, $N(x, r) \subseteq V(K)$ is closed in ΔK . In fact, if $V_0(K)$ is infinite, then x is the unique accumulation point of depth n in $N(x, r)$.

For the next proof, we shall need the following ‘‘shortcutting’’ procedure.

Suppose that $a, b \in \Delta K$ are not both equal to the same point of ∂K , and that $c \in V(K)$. Suppose that α and β are geodesics connecting a to c and c to b respectively. We want to shortcut c by jumping from α to β near a centre of the the points a, b, c . (Here we use ‘‘centre’’ in the sense of a Gromov hyperbolic space, as discussed in Section 5. It is well defined up to a bounded set.) If α and β meet at some point near the centre, then we choose the nearest intersection point on α to a . We then jump directly from α to β at this point. If there is no such intersection point, choose any points $x \in V(\alpha)$ and $y \in V(\beta)$ near the centre, and let δ be a geodesic from x to y . By replacing δ by a suitable subarc, if necessary, we can suppose that $\alpha \cap \delta = \{x\}$ and $\beta \cap \delta = \{y\}$. Let α' be the subarc of α between a and x , and let β' be the subarc of β between b and y . Let $\gamma = \alpha' \cup \delta \cup \beta'$. Thus, γ is an f_1 -quasigeodesic arc from a to b , where f_1 depends only on the hyperbolicity constant. We shall abbreviate $M_{f_1}(a, A)$ to $M_1(a, A)$ etc. Note that since $\text{length } \delta$ is bounded, this shortcutting only changes things in a bounded neighbourhood of α , where the bound depends only on the hyperbolicity constant. Note also that if $a \in V(K)$ and $c \neq a$, we can easily arrange that the edge of α incident to a also lies in γ .

We are now ready for:

Proposition 8.6. *ΔK is compact.*

Proof. We begin with the observation that if $A \subseteq V(K)$ is finite, then $\Delta K = \bigcup_{a \in A} M'(a, A)$. This can be seen by dropping a perpendicular from any point of ΔK to A .

Now, let $(A_n)_{n \in \mathbb{N}}$ be an increasing sequence of finite subsets of $V(K)$ with $V(K) = \bigcup_{n \in \mathbb{N}} A_n$. Suppose that \mathcal{U} is an open cover of ΔK . We claim that there is some n such that for all $a \in A_n$, there is some $U \in \mathcal{U}$ such that $M'(a, A_n) \subseteq U$. This gives rise to a finite subcover of ΔK by the observation of the first paragraph.

Suppose to the contrary, there for each n there is some point $a_n \in A_n$ such that $M'(a_n, A_n)$ is not contained in any element of \mathcal{U} .

Suppose first that $a_n = a$ is constant. Now, $a \in U$ for some $U \in \mathcal{U}$. Thus, there is some finite set $B \subseteq V(K)$ such that $M'(a, B) \subseteq U$. Now $B \subseteq A_n$ for all sufficiently large n , and so we get $M'(a, A_n) \subseteq M'(a, B)$ contrary to our assumption.

We can therefore assume that the a_n are all distinct. Fix any $x_0 \in V(K)$. For each n , we choose some geodesic, α_n , connecting x_0 to a_n .

Now $x_0 \in U$ for some $U \in \mathcal{U}$, and so we can find a finite set, I , of edges incident on x_0 such that $P'_1(x_0, I) \subseteq U$. Moreover, we can suppose that each endpoint of each edge of I lies in A_n for all n .

Suppose that $I \cap E(\alpha_n) = \emptyset$. In this case, we claim that $M'(a_n, A_n) \subseteq P'_1(x_0, I)$. To see this, suppose $b \in M'(a_n, A_n)$. We connect b to a_n by a geodesic, β , with $V(\beta) \cap A_n = \{a_n\}$. In particular, we note that $x_0 \notin V(\beta)$. By the shortcutting procedure described earlier, we can find an f_1 -quasigeodesic arc, γ , connecting b to x_0 , with $e \in E(\gamma)$. Thus, $I \cap E(\gamma) = \emptyset$. This shows that $b \in P'_1(x_0, I)$, as claimed. We therefore see that $M'(a_n, A_n) \subseteq U$, again contrary to our assumption. We deduce that $I \cap E(\alpha_n) \neq \emptyset$.

Since I is finite, after passing to a subsequence, we can assume that the first edge of the geodesics α_n is fixed. Let $x_1 \in V(K)$ be the other endpoint of this edge. Now, by applying exactly the same argument, with x_1 replacing x_0 (and ignoring the first edge of the arcs α_n), we can assume that second edge of each of the geodesics α_n is also fixed. Continuing in this way, and using a diagonal sequence argument, we can find a fixed geodesic ray, ϵ , emanating from x_0 , such that for each n , the geodesic α_n agrees with ϵ for at least the first n edges.

Let $y \in \partial K$ be the endpoint of ϵ . There is some $U \in \mathcal{U}$ such that $y \in U$, and hence some finite set $B \subseteq V(K)$ such that $M'_1(y, B) \subseteq U$. We can suppose that $B \subseteq A_n$ for all n . Now, $B \subseteq N(x_0, r)$ for some $r \geq 0$. Let n be some number much bigger than r . We claim that $M'(a_n, A_n) \subseteq M'_1(y, B)$.

To see this, connect y to a_n by any geodesic α . Since ϵ and α_n agree on an initial segment of length much bigger than r , we see (by the ‘‘thin triangle’’ property) that $d(x_0, \alpha)$ is also much bigger than r . Suppose that $z \in M'(a_n, A_n)$. Then there is a geodesic, β , connecting z to a_n meeting A_n only in the point a_n . In particular, $V(\beta) \cap B = \emptyset$, since

$B \subseteq A_n$. We now carry out the shortcutting procedure on $\alpha \cup \beta$, to obtain an f_1 -quasigeodesic, γ , from z to y . Since the shortcutting only changes things on a bounded neighbourhood of α , and since $d(x_0, \alpha)$ is much bigger than r , we get that $V(\gamma) \cap B = \emptyset$. Thus, $z \in M'_1(y, B)$ as claimed.

This shows that $M'(a_n, A_n) \subseteq U$ contrary to our assumption. We have thus derived a contradiction, proving the result. \square

We remark that since ΔK is hausdorff and second countable, it is also metrisable. In the case where K is 2-vertex-connected, this fact also follows from the results of Section 9 (since ΔK can be identified with $\partial X(K)$).

We finally note that the boundary remains unchanged by some of the operations described in Section 2.

Suppose that K and L are fine hyperbolic graphs with the same vertex set, $V = V(K) = V(L)$, and with $E(K) \subseteq E(L)$. Suppose moreover that if $x, y \in V$ are adjacent in L then they are connected by an arc of length at most p in L . In other words, the metrics d_K and d_L are related by $d_L \leq d_K \leq pd_L$. Then, the embedding of K in L is a quasiisometry, so we get a natural identification of the sets ∂K and ∂L , and hence of ΔK and ΔL .

Now, any geodesic in K is an f -quasigeodesic arc in L , where f is the map $[n \mapsto pn]$. Thus, if $a \in \Delta K \equiv \Delta L$, and $A \subseteq V$, we see that the set $M'(a, A) \subseteq \Delta K$ as defined in K is a subset of the set $M'_f(a, A) \subseteq \Delta L$ as defined in L . This shows that the identification map from ΔK to ΔL is continuous. Since ΔK is compact and ΔL is hausdorff, it must be a homeomorphism. We have shown:

Proposition 8.7. *If $K \subseteq L$ is a quasiisometric embedding of fine hyperbolic graphs with the same vertex set V , then there is a natural homeomorphism from ΔK to ΔL which is the identity on V .*

In particular, in the notation of Section 2, we see that ΔK can be naturally identified with ΔK^n for any $n \in \mathbb{N}$, and indeed with ΔL for any graph L with $K \subseteq L \subseteq K^n$.

If K is the graph featuring in the second definition of a relatively hyperbolic group, (Γ, \mathcal{G}) , in the introduction, then ΔK can be taken to be the boundary of (Γ, \mathcal{G}) . This will be elaborated on in the next section.

We remark that one can give a unified approach to boundaries of hyperbolic and relatively hyperbolic groups. If K is a fine hyperbolic graph, then it is easily verified that the isolated points of ΔK are precisely the finite degree vertices of K . Now if K is any (Γ, V) -graph,

we can take the boundary of (Γ, \mathcal{G}) to be ΔK minus the set of isolated points.

9. BOUNDARIES

This section can be viewed a continuation of Section 7. The main goal will be to show that the boundary of a relatively hyperbolic group is well defined. We already seen a number of ways of associating a boundary to such a group, either via geometrically finite groups, as suggested in Section 6, or via fine hyperbolic graphs, as suggested in Section 8. We therefore need to relate these constructions. Most of this can be done without introducing group actions.

We discussed fine hyperbolic graphs in some detail in Section 8. Most of this is not logically required for the results of this section, though it helps to set them in context. We shall just need to recall the definition of the boundary.

Let K be a fine hyperbolic graph, with vertex set $V(K)$ and Gromov boundary ∂K . Let $\Delta K = V(K) \cup \partial K$. Given $a \in \Delta K$ and $A \subseteq V(K)$, let $M(a, A)$ be the set of points $b \in \Delta K$ such that any geodesic from b to a meets A , if at all, only in the point a . We define a set $U \subseteq \Delta K$ to be open if for all $a \in U$, there is a finite set $A \subseteq V(K)$ such that $M(a, A) \subseteq U$.

Now, suppose that (X, ρ) is a proper hyperbolic length space. The (Gromov) boundary, ∂X , is compact metrisable. Suppose Π is a subset of ∂X , and that $(B(p))_{p \in \Pi}$ is a sufficiently separated quasidense system of horoballs. It follows easily that Π is a countable dense subset of ∂X . Let K be a nerve of $(B(p))_{p \in \Pi}$. By definition, the vertex set, $V(K)$, is precisely the indexing set Π . We can thus identify Π both as a subset of ∂X and as a subset of ΔK . We aim to show:

Proposition 9.1. *There is a homeomorphism from ΔK to ∂X which fixes Π .*

Since Π is dense in ∂X , it follows that this homeomorphism must be unique.

Before we set about proving this, let's recall more precisely the relevant definitions. We are assuming that there are constants, r_1, r_2, r_3, r_4 , such that $(B(p))_{p \in \Pi}$ is r_1 -separated and r_2 -dense, and such that K contains the r_3 -nerve of $(B(p))_{p \in \Pi}$ and is contained in the r_4 -nerve. We necessarily have that $r_1 \leq r_2$ and that $r_3 \leq r_4$. We need in addition that $2r_2 < r_3$ and that r_1 is sufficiently large in relation to the hyperbolicity constant. In other words, $r_1 \leq 2r_2 < r_3 \leq r_4$, and r_1 is large.

For the purposes of proving Proposition 9.1, it will be convenient to alter the meaning of the notations \simeq , \simeq , \ll , etc. and the terminology “close”, to take account of these new constants. In other words, we write $\zeta \preceq \xi$ to mean that $\zeta \leq \xi + c$, where c depends on r_1, r_2, r_3, r_4 as well as the hyperbolicity constant. As usual, the value of c at any given point in an argument is completely determined as a function of these basic constants, though we won't bother to explicitly keep track of it. With this new convention, we can write, for example $\rho(B(p), B(q)) \simeq 0$ whenever $p, q \in \Pi$ are adjacent in K .

Suppose π is a path in K . As described by Lemma 7.6 we may construct a polygonal path, $\gamma(\pi)$, in X . If $p \in \pi \cap \Pi$, then $\gamma(\pi) \cap B(p) \neq \emptyset$. Moreover, any point of $\gamma(\pi)$ lies close to $B(p)$ for some $p \in \pi \cap \Pi$. If π is a geodesic, then $\gamma(\pi)$ is a quasigeodesic. If $\pi' \subseteq \pi$ connects $p, q \in \pi \cap \Pi$, then $\gamma(\pi') \subseteq \gamma(\pi)$ and $\gamma(\pi')$ lies a bounded distance from $[B(p), B(q)]$. (In this section, we may as well take $[B(p), B(q)]$ to be a shortest geodesic connecting $B(p)$ to $B(q)$.) Note that $[B(p), B(q)]$ lies close to $[p, q] \setminus (B(p) \cup B(q))$.

We want to associate a geodesic, $\delta(\pi)$, in X to a geodesic path, π , in K . We consider three possibilities. Firstly, if π connects $p, q \in \Pi$, we set $\delta(\pi) = [p, q]$. Secondly, suppose π is a geodesic ray based at $p \in \Pi$. Now, $\gamma(\pi)$ is a quasigeodesic ray in X , and so converges to some $y \in \partial X$. Since $\gamma(\pi)$ meets infinitely many horoballs in X , we see easily that $y \notin \Pi$. We set $\delta(\pi) = [p, y]$. Finally, if π is a biinfinite geodesic, then $\gamma(\pi)$ is a biinfinite quasigeodesic with endpoints $x, y \in \partial X \setminus \Pi$. We set $\delta(\pi) = [x, y]$.

Note that, in the first case, $\gamma(\pi)$ lies close to $\delta(\pi) \setminus (B(p) \cup B(q))$. In the second case, it lies close to $\delta(\pi) \setminus B(p)$. In the final case, it lies close to $\delta(\pi)$. Here “close to” means that each lies in a bounded neighbourhood of the other.

To construct our homeomorphism, we shall also need to associate geodesics in K with geodesics in X .

Lemma 9.2. *Suppose α is a geodesic ray in X , tending to a point, y , of $\partial X \setminus \Pi$. Then there is a geodesic ray, π , in K such that $\gamma(\pi)$ lies a bounded distance from α .*

Proof. Let $(x_n)_{n=1}^\infty$ be a sequence of points of α with x_0 an endpoint of α , and with $\rho(x_0, x_n)$ tending monotonically to ∞ (so that $x_n \rightarrow y$). For each n , we choose $q_n \in \Pi$ so that $\rho(x_n, B(q_n)) \simeq 0$. Passing to a subsequence if necessary, we can suppose that the q_n are all distinct, and that $\rho(x_m, B(q_n)) \gg 0$ whenever $m \neq n$. (Here we use the fact that $y \notin \Pi$.) It follows easily that if $0 < m < n$, then $\rho(x_m, [B(p_0), B(p_n)]) \simeq 0$.

Now, for each n , we choose a geodesic, π_n , in K connecting q_0 to q_n . Since $\gamma(\pi_n)$ runs close to $[B(q_0), B(q_n)]$, we see that if $0 < m < n$, then x_m lies close to $B(p_{m,n})$ for some $p_{m,n} \in \pi_n \cap \Pi$. Since $(B(p))_{p \in \Pi}$ is locally finite, we see that for a given m , there are only finitely many possibilities for $p_{m,n}$. By a diagonal sequence argument, we can therefore suppose that $p_{m,n} = p_m$ is independent of $n > m$. In other words, $p_m \in \pi_n$ for all $n > m$. We set $p_0 = q_0$. It's now easy to see that the points p_m all lie on a geodesic, π , in K emanating from p_0 . Since $\rho(x_m, B(p_m)) \simeq 0$ and $x_m \rightarrow y$, we see that π must be infinite, i.e. a geodesic ray. Moreover, $\gamma(\pi)$ lies a bounded distance from α as required. \square

By a slight variation on this argument, we can also deduce:

Lemma 9.3. *Suppose that α is a biinfinite geodesic in X with both endpoints in $\partial X \setminus \Pi$. Then, there is a biinfinite geodesic, π , in K such that $\gamma(\pi)$ lies a bounded distance from α .*

We are now in a position to define the map $f : \Delta K \rightarrow \partial X$, which we shall eventually prove to be a homeomorphism. We set $f|_{\Pi}$ to be the inclusion of Π in ∂X . Suppose $a \in \partial K$. Let π be a geodesic in K tending to a . We define $f(a)$ to be the endpoint of the quasigeodesic $\gamma(\pi)$ in X . We have already observed that $f(a) \notin \Pi$. By Lemma 9.2 f is surjective. To see that f is injective, it's enough to show that if $a, b \in \partial K$ are distinct, then $f(a) \neq f(b)$. But we can find rays π and π' tending to a and b respectively such that $\pi \cup \pi'$ is a biinfinite geodesic in K . Now $\gamma(\pi \cup \pi')$ is a biinfinite geodesic in X , and its endpoints, by definition, are $f(a)$ and $f(b)$. Therefore $f(a) \neq f(b)$ as claimed.

Of course, we still have to verify that f is well defined, i.e. does not depend on the choice of ray π . Suppose that π and π' are both rays in K tending to $a \in \partial K$. Let $y, y' \in \partial X \setminus \Pi$ be the endpoints of $\gamma(\pi), \gamma(\pi')$ respectively. Suppose, for contradiction, that $y \neq y'$. Using Lemma 9.3, we can find rays τ, τ' , in K such that $\tau \cup \tau'$ is a biinfinite geodesic in K , and such that $\gamma(\tau)$ and $\gamma(\tau')$ tend to y and y' respectively. Let b, b' be the endpoints of τ, τ' in ∂K . Thus, $b \neq b'$. By the injectivity part of the argument, we see that $b = a = b'$. This contradiction shows that $y = y'$ as claimed.

So far, we have constructed a bijective map $f : \Delta K \rightarrow \partial X$. To show that it is a homeomorphism, we need some preliminary observations.

Fix a point $w \in X$. Given $r \geq 0$, let $\Pi(r) = \{p \in \Pi \mid \rho(w, B(p)) \leq r\}$. Thus, $\Pi(r)$ is finite.

Suppose that π is a geodesic in K with $\pi \cap \Pi(r) = \emptyset$. Since $\delta(\pi)$ is a bounded distance from $\bigcup_{p \in \pi \cap \Pi(r)} B(p)$, we see that $\rho(w, \delta(\pi)) \succeq r$.

Now suppose π is a geodesic ray or segment in K with an endpoint at $p \in \Pi$. Suppose that $\pi \cap \Pi(r) \subseteq \{p\}$. In this case, we see that $\rho(w, \delta(\pi) \setminus B(p)) \succeq r$. Now a simple geometric argument shows that if δ is any biinfinite geodesic with endpoint at p , then (since $w \notin B(p)$) we have $\rho(w, \delta \setminus B(p)) \preceq 2\rho(w, \delta)$. In particular, putting $\delta = \delta(\pi)$, we see that $\rho(w, \delta(\pi)) \succeq r/2$.

Conversely, suppose that $A \subseteq \Pi$ is finite, and that π is a geodesic in K satisfying $\rho(\delta(\pi), \bigcup_{p \in A} B(p)) \gg 0$. Now, $\gamma(\pi)$ is close to δ , and meets $B(p)$ for each $p \in \pi \cap \Pi$. We thus conclude that $\pi \cap A = \emptyset$.

We are now ready for the proof of Proposition 9.1:

Proof. We have constructed a bijective map, $f : \Delta K \longrightarrow \partial X$. We need to show that it is a homeomorphism. We first show that it is continuous. We fix a basepoint, $w \in X$.

Suppose $a \in \partial K$ and $r \geq 0$. Suppose $b \in M(a, \Pi(r))$. Let π be a geodesic connecting a to b in K . Thus $\delta(\pi)$ connects $f(a)$ to $f(b)$ in X . Now, by the definition of $M(a, \Pi(r))$, we have that $\pi \cap \Pi(r) = \emptyset$. From the previous discussion, we have that $\rho(w, \delta(\pi)) \succeq r$. In other words, we can ensure that $\rho(w, [f(a), f(b)])$ is arbitrarily large, by choosing r large enough. This shows that f is continuous at a .

Suppose $a \in V(K)$ and $r \geq 0$. Suppose $b \in M(a, \Pi(r))$. This time, we have $\pi \cap \Pi(r) \subseteq \{a\}$ for any geodesic, π , from a to b in K . From the earlier discussion, we deduce that $\rho(a, \delta(\pi)) \succeq r/2$, so we can again ensure that $\rho(w, [f(a), f(b)])$ is arbitrarily large. This shows that f is continuous at a .

We have shown that f is continuous everywhere. It remains to show that f^{-1} is continuous.

Suppose, then, that $x \in \partial X \setminus \Pi$, and that $A \subseteq V(K)$ is finite. Now, $A \cup \bigcup_{p \in A} N(B(p), C) = \bigcup_{p \in A} (\{p\} \cup N(B(p), C))$ is closed in $X \cup \partial X$ for all $C \geq 0$. Taking $C \gg 0$, we see that for all $y \in \partial X$ sufficiently close to x , we get that $\rho(\delta, \bigcup_{p \in A} B(p)) \gg 0$ for all geodesics δ connecting x to y . Now, suppose that π is a geodesic in K connecting $f^{-1}(x)$ to $f^{-1}(y)$. We have $\rho(\delta(\pi), \bigcup_{p \in A} B(p)) \gg 0$, so $\pi \cap A = \emptyset$. Thus, $f^{-1}(y) \in M(f^{-1}(x), A)$. Since A is arbitrary, this shows that f^{-1} is continuous at x .

Finally, suppose $x \in \Pi$, so that $f^{-1}(x) = x$. Suppose $A \subseteq V(K) \setminus \{x\}$ is finite. The same argument shows that $f^{-1}(y) \in M(x, A)$ for all $y \in \partial X$ sufficiently close to x . This shows that f^{-1} is continuous at x .

We have shown that f^{-1} is continuous everywhere, thus concluding the proof of Proposition 9.1. \square

We remark that we could have bypassed some of the above proof, using the fact that a continuous bijection from a compact space to a hausdorff space is continuous (using Proposition 8.4). The above argument is perhaps more generally applicable in that we have made no use of the local compactness of X , other than the fact that the collection $(B(p))_{p \in \Pi}$ is locally finite.

Note that, if K is 2-vertex connected, we could apply Proposition 9.1. to the proper hyperbolic space $X(K)$ constructed in Section 3. We conclude that, in this case, the space, ΔK , we have defined is really the same as the boundary $\partial X(K)$. This gives another proof of the fact that ΔK is compact and metrisable in this case.

We now apply this result to geometrically finite groups. Suppose that Γ acts a geometrically finite group on a taut proper hyperbolic space, X . Let $\Pi \subseteq \partial X$ be the set of (bounded) parabolic points, which we assume to be non-empty. As discussed in Section 7, we can construct a sufficiently separated Γ -invariant system of horoballs, $(B(p))_{p \in \Pi}$, for Γ . Let $r \gg 0$, and let K be the r -nerve of $(B(p))_{p \in \Pi}$. Thus, K is a (Γ, Π) -graph. By Proposition 7.8, K is fine and hyperbolic. From the naturality of the construction, Γ acts by homeomorphism on ΔK . By Proposition 9.1, there is a homeomorphism from ΔK to ∂X fixing Π . This homeomorphism is unique, and hence Γ -equivariant.

We can now prove:

Theorem 9.4. *Suppose that X and X' are taut proper hyperbolic length spaces, and that Γ acts as a minimal geometrically finite group on both X and X' . Let $\Pi \subseteq \partial X$ and $\Pi' \subseteq \partial X'$ be the respective sets of (bounded) parabolic points. Suppose that there is a Γ -equivariant bijection, $f : \Pi \longrightarrow \Pi'$ such that $\Gamma(p) = \Gamma(f(p))$ for all $p \in \Pi$. Then we can extend f to a Γ -equivariant homeomorphism $f : \partial X \longrightarrow \partial X'$.*

Proof. We identify Π and Π' via f . Thus, Π is a Γ -set. Let L and L' be Γ -equivariant nerves for the systems of horoballs for the actions of Γ on X and X' respectively. Let K be a (Γ, Π) -graph containing both L and L' . Thus, K is a nerve for both systems of horoballs. By Proposition 9.1 and the above discussion, we get Γ -equivariant homeomorphisms, $g : \Delta K \longrightarrow \partial X$ and $g' : \Delta K \longrightarrow \partial X'$, both fixing Π . Now let $f = g' \circ g^{-1}$. This gives the desired homeomorphism. \square

We have been considering only minimal geometrically finite actions and taut hyperbolic spaces. However, as observed in Section 6, there is no essential loss of generality in doing this. If Γ acts as a geometrically finite group on a proper hyperbolic space, X , then it acts as a minimal

geometrically finite group on the taut proper hyperbolic space, join $\Lambda\Gamma$, where $\Lambda\Gamma \subseteq \partial X$ is the limit set of Γ .

Translating Theorem 9.4 into these terms, we see that if we have geometrically finite action of a group Γ on proper hyperbolic spaces, X and X' , with $f : \Pi \rightarrow \Pi'$ as in the hypotheses, then we can extend f to a Γ -equivariant homeomorphism between the corresponding limit sets. For geometrically finite kleinian groups (constant negative curvature of any dimension) this is proven in [Tu].

Theorem 9.4 enables us to make simplifying assumptions about a geometrically finite action when we are interested in relating properties of the boundary to properties of the group. Thus, for example, we can assume that X has the form of a complex built out of ideal hyperbolic triangles in the manner described in Section 3 (at least, if all peripheral subgroups are finitely generated). In particular, this puts a strong control on the geometry of horoballs. This is needed, for example, in [Bo9].

10. CONNECTEDNESS OF BOUNDARIES

In this section, we make a few observations about splittings of relatively hyperbolic groups over finite subgroups. We express such splittings in terms of actions of the group on simplicial trees with finite edge stabilisers and without edge inversions. A splitting is said to be *relative* to a certain collection of subgroups if every subgroup in this collection fixes a vertex of the tree. It is *non-trivial* if no vertex of the tree is fixed by Γ . The quotient of the tree by the group gives a representation of the group as the fundamental group of a graph of groups.

The first result tells us when a boundary is connected:

Proposition 10.1. *The boundary, $\partial\Gamma$, of a relatively hyperbolic group, Γ , is connected if and only if Γ does not split non-trivially over any finite subgroup relative to peripheral subgroups.*

There is also an accessibility result:

Proposition 10.2. *Any relatively hyperbolic group can be expressed as the fundamental group of a finite graph of groups with finite edge groups and with every peripheral subgroup conjugate into a vertex group, and with the property that no vertex group splits non-trivially over any finite subgroup relative to the peripheral subgroups.*

We also note:

Proposition 10.3. *Suppose a relatively hyperbolic group splits as a graph of groups with finite edge groups and relative to the peripheral*

subgroups. Then each vertex group is hyperbolic relative to the peripheral subgroups that it contains.

In the last result, we need to admit the possibility that a vertex group might be finite, or might consist of just a single peripheral subgroup. In all cases, the boundary of a vertex group is naturally identified as a closed subset of the boundary, $\partial\Gamma$, of the whole group, Γ .

If the splitting is maximal the sense of Proposition 10.2, then the boundary of each vertex group is connected. In fact, each component of $\partial\Gamma$ is either a single point or the boundary of an infinite non-peripheral vertex group. A singleton component is either fixed point of a vertex group that is also a peripheral group, or else corresponds to one of the ideal points of the tree corresponding to the splitting.

In the context of hyperbolic groups, Propositions 10.2 and 10.2 are simple consequences respectively of Stallings's theorem on ends of finitely generated groups and Dunwoody's accessibility theorem for finitely presented groups, and Proposition 10.3 is a simple exercise. One can easily adapt these arguments to deal with the relative case, though since there seems to be no explicit reference, it is appropriate to give some commentary.

In view of the fact that the hyperbolic case is well-known, it is convenient to restrict to the non-hyperbolic case, i.e. where the set, \mathcal{G} of peripheral subgroups of Γ is non-empty. Thus, Γ acts on a Γ -set V , such that Γ is precisely the set of point stabilisers of V , so that we can identify $\partial\Gamma$ with ∂V . Now, Γ acts on a simplicially hyperbolic 2-complex, Σ , in the manner described by Proposition 4.9. Let K be the 1-skeleton of Σ , and let $X = X(K)$ be the hyperbolic complex constructed in Section 3. Thus we can identify $\partial\Gamma \equiv \partial V$ with $\Delta K \equiv \partial X$.

We shall say that K is *multiended* if and only if we can partition $V(K)$ into two disjoint infinite subsets $A \sqcup B$ such that only finitely many edges of K connect A to B . Note that if K is fine and every vertex has infinite degree (as in the case above) then the same holds for the full subgraphs, K_A and K_B , on the sets A and B . Thus, by Lemma 2.11 it's enough to suppose that A and B are non-empty. Indeed, this implies that K_A and K_B must be intrinsically unbounded, from which it follows easily that A and B must be unbounded as subsets of K .

Returning the case where K is the 1-skeleton of our complex, Σ , we note:

Lemma 10.4. *The space ΔK is disconnected if and only if K is multiended.*

Proof. This is probably most easily seen using the complex X .

Suppose that $\Delta K \equiv \partial X$ is disconnected. Write $\partial X = P \sqcup Q$, where P and Q are disjoint closed and non-empty. Let $A = P \cap V$ and $B = Q \cap V$. Since V is dense in ∂X , A and B are non-empty. Now, $X \cup \partial X$ is compact hausdorff, so we can find disjoint open subsets, $O \supseteq P$ and $U \supseteq Q$, of $X \cup \partial X$. Now $(X \cup \partial X) \setminus (O \cup U)$ is a compact subset of X which meets every edge of K connecting A to B . Thus, the set of such edges is finite.

Conversely, if only finitely many edges of K connect A to B , then there is a compact subset of $\Sigma \setminus V \equiv X$ which separates A from B . It is now easily seen that the closures of A and B in ∂X are disjoint. Since V is dense in ∂X , these closures partition ∂X . \square

Propositions 10.1 and 10.2 now follow by combinatorial arguments. We only require that Σ is simply connected and locally finite away from V . Following Dunwoody [Du], we define a *track* on Σ to be a connected subset $\Sigma \setminus V$ which meets each edge of Σ in a finite set, and each 2-simplex of Σ in a finite set of disjoint arcs, each arc connecting two different faces of the simplex. If $A \subseteq B$ are connected by a finite set of edges, then we can find a track, t , which meets an edge of Σ (in one point) if and only if this edge connects A and B . Thus, t separates A from B . In other words, K is multiended if and only if Σ admits a finite separating track.

The proof of Proposition 10.1 is now identical to that given in [Du]. If Σ admits a finite separating track, then one can find such a track, t , such that any Γ -image of t is either equal to t or disjoint from it. This gives rise to an action of Γ on a simplicial tree, T , whose edges are in bijective correspondence to the set, Γt , of Γ -images of t , and whose vertices are in bijective correspondence with the connected components of $\Sigma \setminus \bigcup \Gamma t$. In particular, each vertex stabiliser of Σ fixes a vertex of T . In other words, we get a splitting relative to \mathcal{G} . (In [Du] it is assumed that the complex, Σ , is locally finite, though only the fact that it is locally finite away from V is used.)

The converse of this result is easier. If Γ acts on a simplicial tree, T , with each element of \mathcal{G} fixing a vertex, then there is an equivariant map of V to the vertex set, $V(T)$, of T . We can assume that the action on T is non-trivial. Let K be a any (Γ, V) -graph. If e is any edge of T , then it is easily seen that only finitely many edges of K connect vertices of K which map to vertices of T on different sides of e . This finite set of edges separates K into (at least) two unbounded components. Thus, K is multiended.

Putting the results of the previous two paragraphs together with Lemma 10.4, we deduce Proposition 10.1.

The accessibility result of Proposition 10.2 also follows immediately from [Du]. There is a bound on the number of non-parallel tracks one embed in the quotient complex, Σ/Γ . (We may to subdivide Σ in order to take the quotient.) Thus there is a bound on the complexity of any representation of Γ as a reduced graph of groups. The result follows by taking such a splitting of maximal complexity.

To prove Proposition 10.3, suppose that Γ acts on a simplicial tree, T , such that each element of \mathcal{G} fixes a (unique) element vertex of T . We can assume that this action is minimal. By collapsing down the irrelevant bits of the graph of groups, we can suppose that the quotient is a finite graph. Moreover, after subdividing the edges of T , we can assume that the vertex set, $V(T)$, can be partitioned as $V(T) = V_0(T) \sqcup V_1(T)$, such that each element of $V_0(T)$ has finite degree and finite stabiliser, and such that each edge of T has one endpoint in each of these sets. We obtain an equivariant map, $\phi : V \rightarrow V_1(T)$. We can extend this to an equivariant map of a (Γ, V) -graph, L , into T . After subdividing the edges of L , we obtain a graph, M , whose vertex set can be written as a disjoint union, $V \sqcup W$, where W is a cofinite Γ -set with finite point stabilisers, such that each edge of M is either mapped to a vertex or to an edge of T . Now the preimage in W of each element of $V_0(T)$ is finite. By connecting every pair of points in such a preimage by an edge, we obtain a new $(\Gamma, V \cup W)$ -graph, $K \supseteq M$, and an extension $\phi : K \rightarrow T$, with the property that $\phi^{-1}x$ is a finite connected subgraph of K for each $x \in V_0(T)$. Note that K is fine and hyperbolic.

Now suppose that $x \in V_1(T)$. Let $T(x)$ be the star of x in T , i.e. the union of all incident edges. Let $K(x) = \phi^{-1}(T(x)) \subseteq K$. Thus, $K(x)$ is a subgraph of K . Moreover, since K is connected, and the preimage of each vertex of T adjacent to x is connected, we see easily that $K(x)$ is connected. Now, the vertex stabiliser, $\Gamma(x)$, acts on $K(x)$ with finite quotient. Thus, $K(x)$ is a $(\Gamma(x), V(K(x)))$ -graph. Since it is a subgraph of K , it must be fine. The infinite degree vertices of $K(x)$ are precisely the vertices of V which correspond to those elements of \mathcal{G} which lie inside $\Gamma(x)$. Thus, to prove Proposition 10.3, it remains only to check that the graph, $K(x)$, is hyperbolic.

There are several ways to see this. Probably the simplest is to note that the diameters of the subgraphs $\phi^{-1}y$ where y is adjacent to x , are bounded. Thus if α is any geodesic connecting two points of $K(x)$, there is a bound on the length of the components of $\alpha \setminus K(x)$. It follows that $K(x)$ is quasiconvex in K , and hence intrinsically hyperbolic.

This proves proposition 10.3.

The results of this section represent only the first step towards studying topological properties of boundaries. The next natural things to consider are splittings over parabolic subgroups. Splitting of this type are studied in [Bo8]. Topologically they can be recognised from global cut points in the boundary. There is an accessibility result for such splittings, as well as an analogue of Proposition 10.3.

These results are used in [Bo9] to show that the boundary of a relatively hyperbolic group is locally connected if it is connected — at least provided one assumes that each peripheral subgroup is finitely presented, one-or-two ended and contains no infinite torsion subgroup.

REFERENCES

- [A] L.V.Ahlfors, *Fundamental polyhedrons and limit sets for Kleinian groups* : Proc. Nat. Acad. Sci. U.S.A. **55** (1966) 251–254.
- [BeM] A.F.Beardon, A.Maskit, *Limit points of kleinian groups and finite sided fundamental polyhedra* : Acta Math. **132** (1974) 1–12.
- [Bo1] B.H.Bowditch, *Notes on Gromov’s hyperbolicity criterion for path-metric spaces* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 64–167.
- [Bo2] B.H.Bowditch, *Geometrical finiteness for hyperbolic groups* : J. Functional Anal. **113** (1993) 245–317.
- [Bo3] B.H.Bowditch, *Discrete parabolic groups* : J. Differential Geom. **38** (1993) 559–583.
- [Bo4] B.H.Bowditch, *Geometrical finiteness with variable negative curvature* : Duke Math. J. **77** (1995) 229–274.
- [Bo5] B.H.Bowditch, *A class of incomplete non-positively curved manifolds* : Pacific J. Math. **172** (1996) 1–39.
- [Bo6] B.H.Bowditch, *A topological characterisation of hyperbolic groups* : J. Amer. Math. Soc. **11** (1998) 643–667.
- [Bo7] B.H.Bowditch, *Convergence groups and configuration spaces* : in “Group theory down under” (ed. J.Cossey, C.F.Miller, W.D.Neumann, M.Shapiro), de Gruyter (1999), 23–54.
- [Bo8] B.H.Bowditch, *Boundaries of relatively hyperbolic groups* : Math. Z. **230** (1999) 509–527.
- [Bo9] B.H.Bowditch, *Peripheral splittings of groups* : Trans. Amer. Math. Soc. **353** (2001) 4057–4082.
- [Bu] I.Bumagin, *On definitions of relatively hyperbolic groups* : in “Geometric methods in group theory”, Contemporary Mathematics Vol. 372, Amer. Math. Soc. (2005) 189–196.
- [CDP] M.Coornaert, T.Delzant, A.Papadopoulos, *Les groupes hyperboliques de Gromov* : Lecture Notes in Mathematics No. 1441, Springer Verlag (1990).
- [Da] F.Dahmani, *Les groupes relativement hyperboliques et leurs bords* : Thesis, Université Louis Pasteur (2003).
- [DrS] C.Druţu, M.Sapir, *Tree-graded spaces and asymptotic cones of groups* : Topology **44** (2005) 959–1058.

- [Du] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449–457.
- [Fa] B.Farb, *Relatively hyperbolic groups* : Geom. Funct. Anal. **8** (1998) 810–840.
- [Fr1] E.M.Freden, *Negatively curved groups have the convergence property* : Ann. Acad. Sci. Fenn. Ser. A Math. **20** (1995) 333–348.
- [Fr2] E.M.Freden, *Properties of convergence groups and spaces* : Conform. Geom. Dynam. **1** (1997) 13–23.
- [GehM] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I* : Proc. London Math. Soc. **55** (1987) 331–358.
- [Gera] V.Gerasimov, *Expansive convergence groups are relatively hyperbolic* : Geom. Funct. Anal. **19** (2009) 137–169.
- [Gers] S.M.Gersten, *Subgroups of word hyperbolic groups in dimension 2* : J. London Math. Soc. **54** (1996) 261–283.
- [GhH] E.Ghys, P.de la Harpe (eds.), *Sur les groupes hyperboliques d’après Mikhael Gromov* : Progress in Maths. 83, Birkhäuser (1990).
- [Gre] L.Greenberg, *Fundamental polyhedra for Kleinian groups* : Ann. Math. **84** (1966) 433–441.
- [Grom] M.Gromov, *Hyperbolic groups* : in “Essays in Group Theory” (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263.
- [GroVM] D.Groves, J.F.Manning, *Dehn filling in relatively hyperbolic groups* : Israel. J. Math. **168** (2008) 317–429.
- [He] G.A.Heglund, *Fuchsian groups and transitive horocycles* : Duke Math. J. **2** (1936) 530–542.
- [Mar] A.Marden, *The geometry of finitely generated Kleinian groups* : Ann. Math. **99** (1974) 383–462.
- [Mas] H.A.Masur, Y.M.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103–149.
- [O] D.V.Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems* : Mem. Amer. Math. Soc. **179** (2006).
- [Sh] H.Short et al. *Notes on word hyperbolic groups* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 3–63.
- [Sz] A.Szczepański, *Relatively hyperbolic groups* : Michigan Math. J. **45** (1998) 611–618.
- [Tu] P.Tukia, *On isomorphisms of geometrically finite Möbius groups* : Publ. Math. Inst. Hautes Étud. Sci. **61** (1985) 171–214.
- [T2] P.Tukia, *Convergence groups and Gromov’s metric hyperbolic spaces* : New Zealand J. Math. **23** (1994) 157–187.
- [T3] P.Tukia, *Conical limit points and uniform convergence groups* : J. Reine. Angew. Math. **501** (1998) 71–98.
- [Y] A.Yaman, *A topological characterisation of relatively hyperbolic groups* : preprint, Southampton 2001.

MATHEMATICS INSTITUTE, UNIVERSITY OF WARWICK, COVENTRY CV4 7AL,
GREAT BRITAIN

E-mail address: B.H.Bowditch@warwick.ac.uk