Boundaries of geometrically finite groups.

B. H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
bhb@maths.soton.ac.uk
[First draft: January 1998]

0. Introduction.

This paper is one of a series dealing with connectedness properties of boundaries of hyperbolic and relatively hyperbolic groups, and their connection with group splittings. In this paper, we prove two main results (Theorems 1.1 and 1.2) which, together with results appearing elsewhere, imply that the boundary of a relatively hyperbolic group is locally connected if it is connected, provided we place some mild constraints on the kinds of groups that can arise as peripheral (maximal parabolic) subgroups. Specifically, we show:

Theorem 0.1 : *Suppose that Γ is a relatively hyperbolic group, and that each peripheral subgroup is finitely presented, one-or-two-ended, and contains no infinite torsion subgroup. If the boundary of Γ is connected then it is also locally connected.*

It seems likely that at least the first two hypotheses on peripheral subgroups can be eliminated, or at least replaced by a hypothesis of finite generation. The proof of Theorem 0.1 will be described in Section 4.

The notion of a relatively hyperbolic group was outlined in Gromov's original paper on hyperbolic groups [Gr]. It has been elaborated on in [F] and [Sz]. Further discussion relevant to present paper is given in [Bo7] and [T].

A "relatively hyperbolic group", Γ, can be defined as one which admits a "geometrically finite" action on a proper hyperbolic space, X. We can assume that the action of Γ on ∂X is minimal, so that we can identify its boundary, $\partial\Gamma$, with ∂X. We refer to X as a "Cayley space" for Γ. (Thus, if Γ were hyperbolic, we could take the Cayley space to be a Cayley graph.) The results of [Bo7] allow us to make additional assumptions about our Cayley space, X, without loosing any generality. In particular, we can assume that X is "taut", and has "constant horospherical distorsion" about the parabolic points. This will be discussed in Section 1.

A familiar source of examples of relatively hyperbolic groups are geometrically finite kleinian groups in any dimension. In this case, the boundary can be identified with the limit set, and we can take the Cayley space to the convex hull of the limit set. Theorem 0.1 tells us that the such limit sets are connected if they are locally connected. This was already known for kleinian groups acting on 3-dimensional hyperbolic space (see [AnM]), but was open in higher dimensions. In this particular case of constant curvature, it seems that a number of technical complications can be avoided. An account of this case is given in [BoS]. Theorem 0.1 applies more generally to geometrically finite groups acting on pinched negatively curved spaces (Corollary 4.2). We also note that the case of hyperbolic group (with no parabolics) has already been dealt with from the results of [BeM,Bo3,Bo4,L,Swa,Bo5]. Some applications of local connectedness for hyperbolic groups are discussed in [Bo2]. Other implications and their generalisations to the relative case remain to be explored.

The two main results of the present paper both refer to the notion of a "separating horoball".

Definition : A horoball, B in a proper hyperbolic space, X, centred on a point, $x \in \partial X$ is "separating" if there exist two points $y, z \in \partial X \setminus \{x\}$ such that every path connecting y to z in X must meet B.

If ∂X is connected, then y and z lie in different quasicomponents of $\partial X \setminus \{x\}$. In particular, x is a global cut point of ∂X.

We shall speak (rather loosely) of a relatively hyperbolic group, Γ, having a separating horoball if there exists a (suitable) Cayley space for Γ with a separating horoball. Note that the centre of such a horoball is a well defined (global cut) point of $\partial\Gamma$. For all we know for the moment, the existence or otherwise of a separating horoball may depend on the Cayley space we choose. However, having established that the boundary is locally connected, we shall see, in retrospect, that the existence of a separating horoball in any Cayley space is equivalent to the existence of a global cut point in the boundary (Proposition 5.1).

For the first result we need to assume that our Cayley space is taut and has constant horospherical distorsion about each parabolic fixed point. In this case, if there are no separating horoballs (about any boundary point) then we deduce (Theorem 1.1) that $\partial\Gamma$ is locally connected. In particular, we see immediately:

Corollary 0.2 : *If the boundary of a relatively hyperbolic group is connected and has no global cut point, then it is locally connected.*

This generalises the result of Bestvina and Mess [BeM] in the case where Γ is hyperbolic. In fact, our proof of Theorem 1.1 can be viewed as a somewhat technical elaboration of the proof in [BeM].

The second result tells us what happens if Γ, in contrast, does have a separating horoball which is centred on a parabolic fixed point. In this case, we see (Theorem 1.2) that Γ admits a "proper peripheral splitting". This can be reinterpreted more intuitively by saying that Γ splits non-trivially over a parabolic subgroup relative to the set of all parabolic subgroups. For this result, we only need that our Cayley space is taut.

As discussed in Section 4, it is known from work elsewhere that, under the hypotheses of Theorem 0.1, every global cut point, and hence the centre of every separating horoball, must be a parabolic fixed point. Theorems 1.1 and 1.2 can therefore be viewed as telling us something about two opposing possibilities.

1. Review of relatively hyperbolic groups.

In this section, we give more precise statements of the main results of this paper. We begin by elaborating on the notion of a relatively hyperbolic group.

Let (X, ρ) be a proper (i.e. complete locally compact) hyperbolic space, with boundary ∂X. Thus, ∂X is compact metrisable.

As in [Bo7], we shall adopt certain conventions for dealing with (or rather, not bothering to deal with) additive constants. If $\xi, \zeta \in \mathbf{R}$, we write $\xi \preceq \zeta$ to mean that $\xi - \zeta$ is bounded above by some multiple of the hyperbolicity constant of X. The multiple involved might change during the course of an argument, but in principle is always recoverable, by tracing in detail all the steps of the argument back to the definitions. We similarly write $\xi \simeq \zeta$ to mean that $\xi \preceq \zeta$ and $\zeta \preceq \xi$. We write $\xi \ll \zeta$ to mean that $\zeta - \xi$ is greater than some suitably large multiple of the hyperbolicity constant. In this case, the multiple can, in principle, be determined by anticipating the remainder of an argument. If $x \in X$ and $Q \subseteq X$, we shall say that x "is close to" Q if $\rho(x, Q) \simeq 0$.

Given $x, y \in X \cup \partial X$, we write $[x, y]$ for some choice of geodesic from x to y. We shall say that X is *taut* if every point of X lies close to some biinfinite geodesic. It follows that if $p \in \partial X$ and $x \in X$, then there is some $q \in \partial X$ such that x lies close (though perhaps not quite as close) to $[p, q]$. If $x, y, z \in X \cup \partial X$, a *centre* of x, y, z is a point of X which is close to each of the geodesic $[x, y]$, $[y, z]$ and $[z, x]$. Provided that no two of x, y, z are both equal to same ideal point, then the set of centres is non-empty, compact, and of bounded diameter.

Suppose $p \in \partial X$. We say that $h : X \longrightarrow \mathbf{R}$ is a *horofunction* about p if, whenever $x, y, z \in X$ and z is a centre of (p, x, y), then $h(x) + \rho(x, z) \simeq h(y) + \rho(y, z)$. (We don't in general assume that h is continuous.) Note that a horofunction always exists, and any two horofunctions, h and h', differ by approximately an additive constant. More precisely, there is some $t \in \mathbf{R}$ such that $h'(x) \simeq h(x) + t$ for all $x \in X$. We say that a closed subset, $B \subseteq X$, is a *horoball about* p of there is a horofunction, h, about p such $h(x) \succeq 0$ for all $x \in B$, and $h(x) \preceq 0$ for all $x \in X \setminus B$. (In fact, we could choose h so that $B = h^{-1}[0, \infty)$.) We refer to p as the *centre* of B.

Suppose now that Γ acts isometrically and properly discontinuously on ∂X, and hence as a convergence group on ∂X. We say that the action of Γ is *geometrically finite* if every point of X is either a conical limit point or a bounded parabolic point and every peripheral (i.e. maximal parabolic) subgroup is finitely generated. (At this point, we are allowing the possibility of peripheral groups which are infinite torsion groups. For this reason, following [T], we speak of "parabolic points", rather than "parabolic fixed points" — a peripheral group need not contain a parabolic element. Note also that we don't demand that every parabolic group be finitely generated.) Let $\Pi \subseteq \partial X$ be the set of all parabolic points. It follows that Π/Γ is finite, and that each parabolic point is bounded. If $p \in \Pi$, then its stabiliser, $\Gamma(p)$ is peripheral. An *invariant system of horoballs* consists of a collection, $(B(p))_{p \in \Pi}$, of horoballs indexed by Π, such that $B(p)$ is a horoball about p for all $p \in \Pi$ and $B(\gamma p) = \gamma B(p)$ for all $\gamma \in \Gamma$. We say that $(B(p))_{p \in \Pi}$ is *r-separated* if $\rho(B(p), B(q)) \geq r$ whenever $p \neq q$. We say that $(B(p))_{p \in \Pi}$ is *strictly invariant* if it is r-separated for some $r \gg 0$. If Γ is geometrically finite, and $(B(p))_{p \in \Pi}$ is any strictly invariant collection of horoballs, then $(X \setminus \bigcup_{p \in \Pi} \operatorname{int} B(p))/\Gamma$ is compact.

It is not hard to see that we can always assume our space X to be taut. (This is necessarily the case if $\Pi = \emptyset$, provided we allow ourselves to increase the additive constant.)

Now an important feature of a hyperbolic space is the exponential distorsion of horospheres. To be more precise, suppose that h is a horofunction about a point $p \in \partial X$, and suppose $t \in \mathbf{R}$. We choose a constant, $k \simeq 0$, but sufficiently large, and define $S(t) = h^{-1}[t-k, t+k]$ to be the "horospherical shell" at level t. We write σ_t for the induced path metric on $S(t)$. Of course, σ_t might in general take infinite values, and may not behave particularly sensibly near the boundary of $S(t)$. However, if $x, y \in X$ with $h(x) \simeq h(y) \simeq t$, we have $\sigma_t(x,y) \succeq C_0 \omega^{-\rho(x,y)}$, where $\omega > 0$ and C_0 depend only on the hyperbolicity constant of (X, ρ). This is the "exponential distortion" referred to above. In general, there need be no exponential upper bound on σ_t.

We shall say that p has *constant horospherical distorsion* if there exist constants C_0, C_1 and $\omega = \omega(p)$ such that for all sufficiently large t, and all $x, y \in X$ with $h(x) \simeq h(y) \simeq t$, we have $C_0 \omega^{-\rho(x,y)} \preceq \sigma_t(x,y) \preceq C_1 \omega^{-\rho(x,y)}$. It's not hard to see that this is independent of the choice of horofunction about p. (Note that this is quite a strong assumption. It is true for constant negative curvature. For manifolds of pinched variable negative curvature, we get an upper exponential bound but with an exponent in general different from that of the lower bound. However this weaker condition does not seem to be adequate for the argument we present in Section 2.)

In [Bo7] it is shown that if a group Γ admits a geometrically finite action on a (taut) proper hyperbolic space, X, then one can always choose X so that all the parabolic points, Π, have constant horospherical distorsion. (In fact, the constant ω can be taken to be independent of $p \in \Pi$, though we don't need to know that here.)

Suppose $p \in \Pi$ is a parabolic point (which is necessarily bounded), and h is a horofunction about p. Let $\{\gamma_1, \ldots, \gamma_m\}$ be a finite set of generators for $\Gamma(p)$. It's easy to see that for all sufficiently large t, we have $\rho(x, \gamma_i x) \simeq 0$ for all $x \in X$ with $h(x) \simeq t$. ¿From this, it's not hard to see that the inclusion of the $\Gamma(p)$-orbit, $\Gamma(p)x$, of x in $S(t)$ is a quasiisometry from the word metric induced on $\Gamma(p)x$ to the metric σ_t on $S(t)$. Moreover, we can assume that every point of $S(t)$ is close to a point of $\Gamma(p)x$. We can thus choose a Γ-equivariant system of horofunctions, $(h_p)_{p \in \Pi}$, so that this holds whenever $h_p(x) \geq 0$. We can also assume that the lower bound on horospherical distorsion holds for all $t \geq 0$, and that $h_p(x) + h_q(x) \ll 0$ for all distinct $p, q \in \Pi$ and all $x \in X$.

Recall, from Section 0, the definition of a separating horoball. For future reference, we note that if B and B' are horoballs about the same point of ∂X with $B \subseteq B'$, then if B is separating, then so is B'. Also the centre of a separating horoball is a global cut point.

The first theorem says:

Theorem 1.1 : *Suppose that X is a taut proper hyperbolic space with no separating horoballs, and with ∂X connected. Suppose that Γ acts on X as a geometrically finite group in such a way that each parabolic point has constant horospherical distorsion. Then, ∂X is locally connected.*

Note that, in the hypotheses, we are disallowing any separating horoball, regardless of whether or not its centre lies in Π. However, in [Bo5] it is shown that, under mild additional hypotheses, any separating horoball must be centred on a parabolic point (see Section 4). The second theorem tells us what happens if such a separating horoball does indeed exist.

To make this more precise, we shall need the notion of a "peripheral splitting" as defined in [Bo8]. First, we recall the notion of a pretree as defined in [Bo3]. (See also the earlier work of [W] and [AdN].)

A *pretree* structure on a set, V, is a ternary relation, denoted xyz for $x, y, z \in V$, satisfying the following four axioms:

(T0) xyz implies $x \neq z$,

(T1) xyz holds if and only if zyx holds,

(T2) xyz and xzy cannot hold simultaneously, and

(T3) if xyz holds and $w \neq y$, then either xyw or wyz holds.

The intuitive interpretation of the statement xyz is that y lies "strictly between" x and z. The axioms express the "treelike" nature of this betweenness relation.

Given $x, y \in V$, we write $[x, y]$ for the *interval* $\{x, y\} \cup \{z \in V \mid xzy\}$. We say that V is *discrete* if every interval is finite.

Suppose now that V is a Γ-set, i.e. a set admitting an action of a group, Γ. Suppose also that V/Γ is finite. A *peripheral splitting* of V is (equivalent to) a Γ-invariant discrete pretree structure on V. We say that such a splitting is *proper* if there is no upper bound on the cardinality of intervals. (This is easily seen to be equivalent to the formulation given in [Bo8].) A proper peripheral splitting of V gives rise to a non-trivial action of Γ on a simplicial tree, and hence to a splitting of Γ.

Now, suppose that Γ is a relatively hyperbolic group, and that X is a Cayley space for Γ. Let $\Pi \subseteq \partial\Gamma \equiv \partial X$ be the set of parabolic points. Thus Π is a Γ-set with Π/Γ finite. It turns out that a peripheral splitting of Π is essentially the same as a splitting of Γ over a parabolic subgroup relative to the set of all parabolic subgroups. We shall only concern ourselves here with the formulation in terms of peripheral splittings. We shall show:

Theorem 1.2 : *Suppose that Γ acts as a geometrically finite group on a taut proper hyperbolic space X, with ∂X connected. Let $\Pi \subseteq \partial X$ be the set of parabolic points. Suppose that X has a separating horoball centred on a point of Π. Then, Π admits a proper peripheral splitting.*

We note that Π also admits a Γ-invariant pretree structure arising directly from the separation properties of ∂X as a continuum. This pretree structure is a refinement of the peripheral splitting we obtain in Theorem 1.2.

2. Proof of Theorem 1.1.

We begin the proof of Theorem 1.1 by giving a criterion for a metric space to be locally connected.

We should first recall the usual definitions (see for example [K,HY]). A hausdorff topological space is said to be *locally connected* at a point $x \in M$ if x admits a base of open connected neighbourhoods. (In general, it is not sufficient for x to have a base of connected neighbourhoods.) We say that M is *locally connected* if it is locally connected at every point. It's not hard to see that M is locally connected if and only if every point has a base of connected neighbourhoods. (If the latter condition holds, one verifies that each connected component of any open set is open, and this implies local connectedness everywhere.) It is the latter statement that we shall verify here.

The criterion we shall use is:

Lemma 2.1 : *Suppose that (M, d) is a complete metric space. Suppose that there are constants $\alpha, \epsilon, K, \mu > 0$ and with $\mu < 1$, such that for all $x, y \in M$ with $d(x, y) < \epsilon$, there is a finite sequence $x = x_0, x_1, \ldots, x_n = y$ of points of M with $d(x_i, x_{i+1}) \leq \mu d(x, y)$ for all $i \in \{0, \ldots, n-1\}$ and $d(x_i, x_j) \leq K d(x, y)^\alpha$ for all $i, j \in \{0, \ldots, n\}$. Then, M is locally connected.*

Proof : Given $x, y \in M$ with $d(x, y) \leq \epsilon$, we shall construct a path from x to y in M by iterated interpolation, using the completeness of M.

Given $n \in \mathbf{N}$, we write $I_n = \{p/n \mid p \in \{0, 1, \ldots, n\}\} \subseteq [0, 1]$. Thus, if $m|n$, then $I_m \subseteq I_n$. If $n(1), n(2), n(3), \ldots$ is a sequence of natural numbers with $n(i)|n(j)$ whenever $i < j$, then we write $I_\infty = \bigcup_{i=1}^\infty I_{n(i)}$. If $n(i) \to \infty$, then I_∞ is dense in $[0, 1]$.

Now, suppose $x, y \in M$. We iterate the interpolation given by the hypotheses of the lemma. At the ith stage, we may as well assume that the number of new points introduced between any consecutive pair at the $(i-1)$th stage is constant, and equal to $m(i) - 1$, say. Let $n(i) = \prod_{j=1}^i m(j)$. We thus get a sequence of maps from $I_{n(i)}$ to M each extending the previous one, and hence a map $\beta : I_\infty \longrightarrow M$, with $\beta(0) = x$ and $\beta(1) = y$.

Now, for each $i \in \mathbf{N}$ and $p, q \in \{0, \ldots, n(i)\}$, we have $d(\beta(p/n(i)), \beta((p+1)/n(i))) \leq \mu^i d(x, y)$ and $d(\beta(p/n(i)), \beta(q/n(i))) \leq K d(x, y)^\alpha (1 + \mu^\alpha + \mu^{2\alpha} + \cdots + \mu^{(i-1)\alpha}) \leq A d(x, y)^\alpha$, where $A = K/(1 - \mu^\alpha)$. In particular, the map β is continuous, and so, by completeness, extends to a continuous map $\beta : [0, 1] \longrightarrow M$, with $\mathrm{diam}(\beta([0, 1])) \leq A d(x, y)^\alpha$. In summary we have shown that any pair of points x, y with $d(x, y) \leq \epsilon$ can be connected by a path of diameter at most $A d(x, y)^\alpha$.

Now, suppose that $x \in M$ and that $0 < \eta < \epsilon$. Let H be the connected component of the closed ball $N(x, \eta)$ with $x \in H$. Let $\delta = (\eta/A)^{1/\alpha}$. Suppose that $y \in N(x, \delta)$. Then y is connected to x by a path of diameter at most $A\delta^\alpha = \eta$, and hence lying in $N(x, \eta)$. It follows that $y \in H$. We have shown that $N(x, \delta) \subseteq H$, and so H is a connected neighbourhood of x contained in $N(x, \eta)$.

Now, as observed above, any space with a base of connected neighbourhoods for any point is locally connected. Thus, M is locally connected. ◊

Now let's return to the boundary, ∂X, of our hyperbolic space X. We shall effectively show that ∂X is locally connected by showing that $\partial X \setminus \{p\}$ is locally connected for each $p \in \partial X$. For convenience (in the case where $\Pi \neq \emptyset$), we shall take $p \in \Pi$. There is a fairly standard way of metrising $\partial X \setminus \{p\}$, which we now go on to describe. We shall need the notion of a "quasiultrametric", as discussed in [GhH].

Suppose $\lambda \geq 1$. A λ-*quasiultrametric*, σ, on a set, M, is a symmetric binary map $\sigma : M \times M \longrightarrow [0, \infty)$ such that for all $x, y, z \in M$, $\sigma(x,y) \leq \lambda \max\{\sigma(x,z), \sigma(y,z)\}$. Clearly, if σ is a λ-quasiultrametric, and $\alpha > 0$, then σ^α is λ^α-quasiultrametric. Thus, by taking a suitable power, α, we can convert any quasiultrametric to a λ-quasiultrametric with λ arbitrarily close to 1, in particular, $\lambda < \sqrt{2}$. (Note that a quasiultrametric need not be continuous.)

Now suppose that σ is a λ-quasiultrametric with $\lambda < \sqrt{2}$. Given $x, y \in M$, define $d(x,y)$ to be the infimum of the sums $\sum_{i=0}^{n-1} d(x_i, x_{i+1})$ as $x_0, \ldots, x_n$ ranges over all finite sequences with $x = x_0$ and $y = x_n$. It is shown in [GhH] that d is a metric and that $(3 - 2\lambda)\sigma \leq d \leq \sigma$.

Note that it follows that, in the hypotheses of Lemma 2.1, we can work with a λ-quasiultrametric in place of a metric, provided we make sure that μ is sufficiently small (depending only on λ). The constants α, ϵ, K do not require any dependence on λ.

Now, let X be any proper hyperbolic space. Suppose that $p \in \partial X$, and that h is a horofunction about p. Write $M = \partial X \setminus \{p\}$. Given $x, y \in X \cup M$, not both equal to the same ideal point, we write $\langle x, y \rangle = \langle x, y \rangle_h \simeq -h(z)$, where z is a centre of (p, x, y). (For definiteness, we can take it to the infimum of possible values of $-h(z)$ as z varies over all possible centres, relative to some fixed hyperbolicity constant.)

Now fix some $C > 1$ and let $\sigma(x,y) = C^{-\langle x,y \rangle}$. Thus, σ is a λ-quasiultrametric, where λ depends only on the hyperbolicity constant of X and on C. By taking C to be sufficiently close to 1, we can arrange that $\lambda < \sqrt{2}$, and so σ is bilipschitz related to a metric, d, as described above.

Suppose, now, that Γ acts on X, satisfying the hypotheses of Theorem 1.1. Suppose that $p \in \Pi$ (or any point of ∂X if $\Pi = \emptyset$). Let h be a horofunction about p. We aim to show that the metric, d, as defined in the previous paragraph, satisfies the hypotheses of Lemma 2.1. As observed earlier, we can work instead with the quasiultrametric σ, provided we take μ to be sufficiently small. Reinterpreting this in terms of the bracket $\langle ., . \rangle$, we shall show that given any fixed $r > 0$, if $x, y \in M$ with $\langle x, y \rangle$ sufficiently large, we can find a finite sequence $x = x_0, \ldots, x_n = y$, such that for all i, j, $\langle x_i, x_{i+1} \rangle \geq \langle x, y \rangle + r$ and $\langle x_i, x_j \rangle \geq \nu \langle x, y \rangle - s$, where $\nu > 0$ and s are some constants independent of x and y (in fact, depending only on r and the hyperbolicity constants). We shall denote the condition of the previous sentence by $(*)$. We shall split the proof of $(*)$ into two cases, depending on whether or not some centre of (p, x, y) lies inside one of a suitably chosen invariant system of horoballs.

To deal with these cases, we need some more definitions. Let $(h_p)_{p \in \Pi}$ be a Γ-equivariant set of horofunctions, with the properties described in Section 1 in relation to constant horospherical distorsion, and with $h_p(x) + h_q(x) \ll 0$ for all $x \in X$ and distinct $p, q \in \Pi$. Let $(B_0(p))_{p \in \Pi}$ be a strictly invariant collection of standard horoballs for X, so that $h_p(x) \preceq 0$ for all $x \in X \setminus B_0(p)$ and $h_p(x) \succeq 0$ for all $x \in B_0(p)$. Similarly, we let $(B(p))_{p \in \Pi}$ be another invariant system of horoballs for X, so that $h_p(x) \preceq t$ for all $x \in X \setminus B(p)$ and $h_p(x) \succeq t$ for all $x \in B(p)$, where $t \gg r$ is a constant chosen in a manner to be disussed later. Note that the system $(B(p))_{p \in \Pi}$ is $(2t)$-separated.

Let $\Theta = \Theta(\partial X)$ be the space of distinct triples of ∂X. Then, ∂X is locally compact, and Γ acts properly discontinuously on Θ. Let $\Theta_0 \subseteq \Theta$ be the subset of triples which have a centre in $\partial X \setminus \bigcup_{p\in\Pi} \operatorname{int} B(p)$. Thus, Θ_0 is closed and Γ-invariant. Since the action of Γ is geometrically finite, we see easily that Θ_0/Γ is compact (cf. [T,Bo7]). We shall split the proof of $(*)$ into two parts, according to whether $(p,x,y)\in\Theta_0$ or $(p,x,y)\notin\Theta_0$.

The first case will follow along the lines of [BeM]. (As in that paper, this is sufficient to deal with the case where $\Pi=\emptyset$ so that Γ is a hyperbolic group.) It is here that we use the assumption that there is no separating horoball. (For the proof of the first case, the choice of the strictly invariant system of horoballs is irrelevant. All we really care about is the fact that Θ_0/Γ is compact.) In the second case, we see that each centre of (p,x,q) lies in $B(q)$ for some $q\in\Pi$. The argument relies on a slightly delicate study of the geometry of $B(q)$. This is where we use the assumption of constant horospherical distorsion.

Let's begin with the case where $(p,x,y)\in\Theta_0$. In fact, we shall prove a stronger result than $(*)$. We first note that if $a,b,c,d\in\partial X\setminus\{p\}$, then the quantity $\langle a,b\rangle_h-\langle c,d\rangle_h$ does not depend on the choice of horofunction about p, at least up to a small additive constant. It can thus be viewed as a function of a, b, c, d and p alone. By slight abuse of notation, we shall denote it by $\langle a,b\rangle_p-\langle c,d\rangle_p$. (In fact it can be defined in terms of the "crossratios" involving the five points a,b,c,d,p (cf. [Bo6]).) Note that we can find a neighbourhood, U, of p in ∂X, and neighbourhoods of a, b, c and d in $\partial X\setminus U$, such that for all $q\in U$ and all a', b', c' and d' in the respective neighbourhoods, we have $\langle a',b'\rangle_q-\langle c',d'\rangle_q\simeq\langle a,b\rangle_p-\langle c,d\rangle_p$.

Using the compactness of Θ_0/Γ, we shall prove a strong uniform version of the statement $(*)$, where we take $\nu=1$:

Lemma 2.2 : *Suppose that X, Γ satisfy the hypotheses of Theorem 1.1 (except that we don't need constant horospherical distorsion at this point). Let Θ_0 be as defined above. Suppose $t\in\mathbf{R}$. Then, there is some $l\in\mathbf{R}$ with the following property. Suppose $p,x,y\in\partial X$ are distinct, and $(p,x,y)\in\Theta_0$. Then, there is a finite sequence, $x=x_0,x_1,\ldots,x_n=y$ with $\langle x_i,x_{i+1}\rangle_p-\langle x,y\rangle_p\geq t$ for all $i\in\{0,\ldots,n-1\}$ and $\langle x_i,x_j\rangle_p-\langle x,y\rangle_p\geq -l$ for all $i,j\in\{0,\ldots,n\}$.*

Proof : Suppose $\theta=(p,x,y)\in\Theta_0$. For notational convenience, we choose a horofunction, h, about p so that $\langle x,y\rangle_h=0$. Choose $u\gg t$, and let B be a horoball about p such that $h(x)\leq -u$ for all $x\in X\setminus B$. Now, since no horoball separates, we can connect x to y by a path in $X\setminus B$. By choosing enough points on this path, we can find a finite sequence of points, $a_1,a_2,\ldots,a_{n-1}$, in X with $h(a_i)\leq -u$ for all $i\in\{1,\ldots,n-1\}$, and close enough together so that $\langle a_i,a_{i+1}\rangle_h\succeq u$ for all $i\in\{0,\ldots,n-1\}$. Since X is taut, we can find, for each $i\in\{1,\ldots,n-1\}$, a point $x_i\in\partial X\setminus\{p\}$ such that a_i lies close to a geodesic connecting p to x_i. Let $x_0=x$ and $x_n=y$. We thus get that $\langle x_i,x_{i+1}\rangle_h\succeq u$ for all i. Now, let $l(\theta)\gg-\min\{\langle x_i,x_j\rangle_h\mid i\neq j\}$.

Now choose small enough open sets, $O(\theta),U(\theta),V(\theta)$ containing p,x,y respectively such that if $q\in O(\theta)$, $x'\in U(\theta)$ and $y'\in V(\theta)$, then $\langle x_i',x_j'\rangle_q-\langle x',y'\rangle_q\simeq\langle x_i,x_j\rangle_p-\langle x,y\rangle_p=\langle x_i,x_j\rangle_h$ for all $i,j\in\{0,\ldots,n\}$, where $x_0'=x$, $x_n'=y$ and $x_i'=x_i$ for all $i\in\{1,\ldots,n-1\}$. In particular, we see (since $t\ll u$) that $\langle x_i',x_{i+1}'\rangle_q-\langle x',y'\rangle_q\geq t$ and $\langle x_i',x_j'\rangle_q-\langle x',y'\rangle_q\geq -l(\theta)$.

Since Θ_0/Γ is compact, we can find a finite set, $\theta_1,\theta_2,\ldots,\theta_m\in\Theta_0$ such that $\bigcup_{j=1}^m(O(\theta_j)\times U(\theta_j)\times V(\theta_j))/\Gamma=\Theta_0/\Gamma$. The result now follows, with $l=\max\{l(\theta_j)\mid 1\leq j\leq m\}$. ◊

To deal with the second case, namely when $(p, x, y) \notin \Theta_0$, we need a general lemma about Cayley graphs. Suppose G is a group. A *Cayley graph* of G is a (locally finite) connected graph, K, on which G acts freely and cocompactly without edge inversions and with precisely one orbit of vertices (so that the quotient, K/G, is a finite wedge of circles). Such a graph exists if and only if G is finitely generated. We give K the combinatorial path-metric, σ, so that the distance between adjacent vertices equals 1.

Lemma 2.3 : *Let G be a finitely generated infinite group, and let (K, σ) be a Cayley graph for G. Suppose x and y are vertices of K. Then, there is a ray, β, from y to infinity in K with $\sigma(x, \beta) \geq r/2$.*

By a "ray" we mean a (properly embedded) semi-infinite arc. In fact, the ray we construct will be geodesic in K.

Proof : Since K is locally finite and unbounded, and admits a cocompact action, we can certainly find a biinfinite geodesic in K. Since the G-orbit of y is the whole vertex set, we can find rays, β_0 and β_1 based at y, such that $\beta_0 \cup \beta_1$ is a biinfinite geodesic. Suppose, for contradiction, that we can find points $z_i \in \beta_i$ with $\sigma(x, z_i) < r/2$. We see that $\sigma(z_0, z_1) < r$. But now, $\sigma(y, z_i) \geq r - \sigma(x, z_i) > r/2$. Since y lies on a geodesic segment from z_0 to z_1, we arrive at the contradiction $\sigma(z_0, z_1) > r$. This shows that either $\sigma(x, \beta_0) \geq r/2$ or $\sigma(x, \beta_1) \geq r/2$. $\diamondsuit$

Question : Can one improve on $r/2$? If not in general, under what circumstances? For example, does there exist a constant r_0 such that $r/2$ can be replaced by $r - r_0$?

We gave the above result for Cayley graphs since it has a clean statement in this form. The argument clearly works for any path-metric space, which admits a cocompact (not necessarily free) action of G. In this case we should replace $r/2$ by $r/2 - r_0$ for some constant r_0. In fact, all we really need to know here is that there is a lower bound which is linear in r.

We now return to the objective of proving $(*)$ in the case where $p \in \Pi$, and $x, y \in \partial X \setminus \{p\}$ with $(p, x, y) \notin \Theta_0$ and $\langle x, y \rangle_{h_p}$ sufficiently large. In fact, from the choice we have made of h_p, it will be enough to assume that $\langle x, y \rangle_{h_p} \geq 0$. This ensures that no centre of (p, x, y) can lie in $B(p)$, and so it follows that each centre of (p, x, y) must lie in $B(q)$ for some $q \in \Pi$.

For the rest of the argument, the points p, q, x and y will be held fixed. We shall abbreviate $\langle ., . \rangle_{h_p}$ to $\langle ., . \rangle_p$ and $\langle ., . \rangle_{h_q}$ to $\langle ., . \rangle_q$. Suppose $z, w \in \partial X \setminus \{p\}$ and $\rho(B_0(p), [z, w]) \gg 0$. ¿From the construction of $B_0(p)$, we have the approximate inequalities $\langle z, w \rangle_p \simeq \rho(B_0(p), [z, w]) \simeq \rho(B_0(p), b)$, where b is any centre of (p, z, w).

Let A be the $\Gamma(q)$-orbit of some point in X with $h_q(z) \simeq 0$ for all $z \in A$. Given any $a \in \partial X \setminus \{p\}$, we write a' for some point of A close to a geodesic $[a, q]$. (Such a point exists from the choice of the horoball $B_0(q)$. Any two such points will be close, so the choice involved doesn't really matter to us.) We write $s = h_p(p')$. ¿From the construction of the horoballs $B_0(p)$ and $B_0(q)$, we have $s \simeq \rho(B_0(p), p') \simeq \rho(B_0(p), B_0(q))$.

Thus, in particular, $s \succeq 0$. (This is where we use the assumption that $p \in \Pi$.) Given any $a, b \in X \cup \partial X \setminus \{q\}$, we write $H(a,b) = -\langle a,b\rangle_q$. Thus, if $a, b \in \partial X \setminus \{q\}$, we get that $H(a,b) \preceq H(a',b')$. If $H(a,b) \gg 0$, then $H(a,b) \simeq H(a',b')$ and $\rho(a',b') \simeq 2H(a,b)$. (In other words, we think of $H(a,b)$ as approximately half the length of the segment of $[a,b]$ that lies inside $B_0(q)$.) Figure 1 is meant to suggest the upper half-space model of hyperbolic space, with the ideal point, q, at infinity. Points and distances illustrated are meant to be accurate up to an additive constant.

Let c be a centre of (p,x,y). Now $c \in B(q)$, so $h_q(c) \succeq t$, and so $H(x,y) \succeq t$. Since c lies close to the geodesic $[p,x]$, we see that $H(p,x) \succeq t$. Let $u = H(p,x) - t \succeq 0$. Since $H(p,x) \simeq t + u \gg 0$, we have $\rho(p',x') \simeq 2(t+u)$. Now, $c \in B(q)$. Also $[p,x]$ intersects (or lies close to) $B(q)$ in a segment of length approximately $2u$, starting at a distance approximately t from p'. Since c lies close to this geodesic, we have $\rho(p',c) \preceq t+2u$. Thus, $\langle x,y\rangle_p \simeq \rho(B_0(p),c) \simeq \rho(B_0(p),p') + \rho(p',c) \preceq s+t+2u$.

Now, we can $\Gamma(q)$-equivariantly identify A with the vertex set of a Cayley graph of $\Gamma(q)$. By Lemma 2.3, we can find an infinite path, $z'_0, z'_1, z'_2, \ldots$ in A, with $z'_0 = x'$ and with z'_i and z'_{i+1} adjacent in the Cayley graph, and which satisfies the conclusion of the lemma in that it does not approach too closely the point p'. Thus, $\rho(z'_i, z'_{i+1}) \simeq \sigma_0(z'_i, z'_{i+1}) \simeq 0$ for all i, where σ_0 is the path-metric in the horospherical shell at level 0 with respect to h_q. Now, the word metric on A agrees with the metric σ_0 to within linear bounds. In particular (from the construction of the path $(z'_i)_i$) there is some constant, $\lambda > 0$, such that $\sigma_0(p',z'_i) \succeq \lambda\sigma_0(p',x')$ for all i. Now, since q has constant horospherical distortion, we have $\sigma_0(p',z'_i) \preceq C_1\omega^{\rho(p',z'_i)}$ and $\sigma_0(p',x') \succeq C_0\omega^{\rho(p',x')}$, where $\omega = \omega(q)$ is constant. Thus, $C_1\omega^{\rho(p',z'_i)} \succeq \lambda C_0\omega^{\rho(p',x')}$, so $\rho(p',x') \preceq \rho(p',z'_i) + C_2$, where $C_2 = \log(C_1/\lambda C_0)/\log\omega$. By incorporating the constant C_2 into the notation $\preceq$, we may as well write $\rho(p',x') \preceq \rho(p',z'_i)$. (Now, the additive constant may depend on the geometry of horoballs as well as the hyperbolicity constant, but this will not matter to us.) We thus have $\rho(p',z'_i) \succeq 2(t+u) \gg 0$.

Let $z_i \in \partial X$ be such that z'_i is close to the geodesic $[q,z_i]$. Thus $z_i \to q$ in ∂X. We see that $H(p,z_i) \simeq H(p',z'_i) \simeq \frac{1}{2}\rho(p',z'_i) \succeq t+u$. Let b_i be a centre of (p,q,z_i). Then $\rho(p',b_i) \simeq h_q(b_i) \simeq H(p,z_i) \simeq t+u$. ¿From this, we obtain $\langle q,z_i\rangle_p \simeq \rho(B_0(p),p') + \rho(p',b_i) \succeq s + (t+u) \geq \frac{1}{2}(s+t+2u) \succeq \frac{1}{2}\langle x,y\rangle_p$.

Let c_i be a centre of (q,z_i,z_{i+1}). Since $\rho(z_i,z_{i+1}) \simeq 0$, we have $H(z_i,z_{i+1}) \simeq 0$, so that c_i lies close to the geodesics $[z_i,z'_i]$ and $[z_{i+1},z'_{i+1}]$. Since $H(p,z_i) \gg 0$, the z'_i lies close to $[p,z_i]$, and so c_i lies close to $[p,z_i]$. Similarly c_i lies close to $[p,z_{i+1}]$. In other words, c_i is also a centre of (p,z_i,z_{i+1}). Moreover, we see that $\rho(p',c_i) \succeq \rho(p',z_i) \simeq 2(t+u)$, so $\langle z_i,z_{i+1}\rangle_p \simeq \rho(B_0(p),[z_i,z_{i+1}]) \simeq \rho(B_0(p),p') + \rho(p',c_i) \succeq s+2(t+u)$. Thus, $\langle z_i,z_{i+1}\rangle_p - \langle x,y\rangle_p \succeq (s+2t+2u) - (s+t+2u) \succeq t$, so $\langle z_i,z_{i+1}\rangle_p - \langle x,y\rangle_p \geq t/2$, provided we choose t large enough (in relation to each of the constants $\lambda, C_0, C_1, \omega$ as q ranges over a Γ-transversal of Π, as well as the hyperbolicity constant of X).

We now perform exactly the same construction starting with the point y in place of x. This gives us a sequence $w_i \in \partial X \setminus \{p\}$ with $w_i \to q$, and with $w_0 = y$, $\langle w_i,w_{i+1}\rangle_p - \langle x,y\rangle_p \geq t/2$, and $\langle w_i,q\rangle_p \succeq \frac{1}{2}\langle x,y\rangle_p$ for all i.

Now, since $z_i \to q$ and $w_i \to q$, we get that $\langle z_i,q\rangle_p \to \infty$ and $\langle w_i,q\rangle_p \to \infty$, and so $\langle z_i,w_i\rangle_p \succeq \min\{\langle z_i,q\rangle_p, \langle w_i,q\rangle_p\} \to \infty$. Thus, we can find some m with $\langle z_m,w_m\rangle_p \geq$

$\langle x,y\rangle_p + (t/2)$. Now let $n = 2m+2$ and let $x_0, \ldots, x_n$ be the finite sequence $z_0, z_1, \ldots, z_{m-1}, z_m, w_m, w_{m-1}, \ldots, w_1, w_0$. Then, $x = x_0$ and $y = x_n$. Moreover, $\langle x_i, x_{i+1}\rangle_p - \langle x,y\rangle_p \geq t/2$ for all i. Also, if $i, j \in \{0, \ldots, n\}$, we get $\langle x_i, x_j\rangle_p \succeq \min\{\langle x_i, q\rangle_p, \langle x_j, q\rangle_p\} \succeq \frac{1}{2}\langle x,y\rangle_p$.

This proves condition $(*)$ for the case where $p \in \Pi$ and $(p,x,y) \notin \Theta_0$. (Note that this case is vacuous if $\Pi = \emptyset$.)

Thus, translating this back in terms of quasiultrametrics, and thence to metrics, we have verified the hypotheses of Lemma 2.1, with $M = \partial X \setminus \{p\}$, for all $p \in \Pi$ (or for all $p \in \partial X$ if $\Pi = \emptyset$). If Π is nonempty then it has at least two elements (in fact, it is dense in ∂X). This shows that ∂X is locally connected. We have thus proven Theorem 1.1, from which Corollary 0.2 immediately follows.

3. Proof of Theorem 1.2.

We begin in a very general context. Let X be a proper hyperbolic space, and let $P \subseteq \partial X$ be any subset of the boundary. Suppose that $(B(p))_{p\in P}$ is a set of horoballs in X so that $B(p)$ is centred on p for all $p \in P$.

Fix for the moment some $p \in P$. Now, $B(p) \cup \{p\}$ is closed in $X \cup \partial X$. For convenience, we can assume that each connected component of $X \setminus B(p)$ is path-connected (for example, by taking X to be some kind of simplicial complex, and B to be a subcomplex). Now, it's easy to see that the closure of each component of $X \setminus B(p)$ in $(X \cup \partial X) \setminus (B(p) \cup \{p\})$ is also path-connected. Moreover, we can assume that these closures are pairwise disjoint. Let $\mathcal{C}(p)$ be the set of sets of the form $C \cap \partial X$ as C runs over all such closures where $C \cap \partial X$ is non-empty. It's easily seen that $\mathcal{C}(p)$ partitions $\partial X \setminus \{p\}$ into a disjoint union of non-empty clopen subsets. (For the moment, it is possible that $\mathcal{C}(p)$ consists of a single element.) Note that two points of $\partial X \setminus \{p\}$ lie in the same element of $\mathcal{C}(p)$ if and only if they are connected by a path in $(X \cup \partial X) \setminus (B(p) \cup \{p\})$. (We also see easily that $\mathcal{C}(p)/\Gamma(p)$ is finite, though we won't explicitly use this fact.)

Suppose now that ∂X is connected. We define a ternary relation on ∂X as follows. Given $x, y, p \in \partial X$, we shall write xpy if and only if $p \in P$ and x and y lie in different elements of $\mathcal{C}(p)$. This relation is clearly coarser that induced from the standard pretree structure on a continuum. (In other words, if xpy holds, then x any y lie in different quasicomponents of $\partial X \setminus \{p\}$ — see [Bo3].) We note:

Lemma 3.1 : *The ternary relation thus defined on ∂X satisfies the pretree axioms.*

Proof : Axioms (T0), (T1) and (T3) are immediate. Axiom (T2) follows from the observation made above that this relation is coarser than one we already know to be a pretree structure. ◊

Suppose now that $\Pi \subseteq \partial X$ is any subset with $P \subseteq \Pi$. Suppose that $(B_0(p))_{p\in\Pi}$ is a system of pairwise disjoint horoballs about points of Π, with $B_0(p) \subseteq B(p)$ for all $p \in P$. Suppose, moreover, that if $p \in P$ and $q \in \Pi \setminus P$, then $B(p) \cap B_0(q) = \emptyset$. Finally we suppose that the collection $(B(p))_{p\in P}$ is locally finite (i.e. only finitely many horoballs $B(p)$ meet any given compact subset of X). We restrict the pretree structure defined above to the set Π. In this case, we have:

Lemma 3.2 : *The pretree structure on* Π *is discrete.*

Proof : Suppose that $p, p' \in \Pi$. Connect $B_0(p)$ to $B_0(p')$ by a compact path α in X. We extend α to a path, β, connecting p to p' in X with $\beta \setminus \alpha \subseteq B_0(p) \cup B_0(p')$. Thus, $\beta \cap B(q) \subseteq \alpha$ for all $q \in P \setminus \{p, p'\}$. We see that if pqp', then $B(q) \cap \alpha \neq \emptyset$. By local finiteness, there are only finitely many such q. This shows that the interval $[p, p']$ is finite. $\diamondsuit$

Suppose now that Γ acts a geometrically finite group on a proper taut hyperbolic space, X. Let $\Pi \subseteq \partial\Gamma \equiv \partial X$ be the set of parabolic points. Thus, Π/Γ us finite. As discussed in [Bo7], 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 $p, q \in \Pi$. Moreover, we have $h_p(x) \to -\infty$ for any fixed x, as p varies over Π.

Suppose, now, that X a has separating horoball, B_0, centred on a point $p_0 \in \Pi$. Let $B(p_0)$ be the closure of the union of all images of B_0 under the stabiliser of p_0. Thus $B(p_0)$ is also a horoball about p_0. Since it contains B_0, it is also separating. Taking the set of images of $B(p_0)$ under Γ, we obtain a Γ-invariant system of separating horoballs, $(B(p))_{p\in P}$, where P is the Γ-orbit of p_0. We can suppose that $h_p(x) \geq t$ for all $p \in P$ and $x \in B(p)$, for some fixed $t \in \mathbf{R}$. We can assume that $t \ll 0$. Now, choose a Γ-invariant system of horoballs, $(B_0(p))_{p\in\Pi}$, such that $h_p(x) \geq -t$ for all $p \in \Pi$ and $x \in B_0(p)$. We therefore see that the systems $(B(p))_{p\in P}$ and $(B_0(p))_{p\in\Pi}$ satisfy the hypotheses of Lemma 3.2. We thus obtain a discrete Γ-invariant pretree structure on Π. In other words, we have a peripheral splitting of Π thought of as a Γ-set.

We finally need to verify that this is a proper splitting — in other words that there are intervals of arbitrarily large finite cardinality. To this end, we may as well restrict our pretree structure to P, giving us a peripheral splitting of P.

Now, since P is dense as a subset of $\partial\Gamma$ (in the usual topological sense), we see easily that no point of P in terminal in the pretree structure we have defined. (In other words, for all $p \in P$, there exist $q, r \in P$ such that qpr holds.) But now suppose we have $p, p' \in P$ with $[p, p']$ of maximal cardinality. There exist $q, r \in P$ with qpr, so without loss of generality (by axiom (T3)) we have qpp'. But now we get the contradiction that $[p, p']$ is a proper subset of the interval $[q, p']$. (Indeed, continuing in this way, we obtain an infinite arc in P, i.e. an infinite linearly ordered subset.)

We have shown that Π admits a proper peripheral splitting. This concludes the proof of Theorem 1.2.

4. Proof of Theorem 0.1.

In this section, we explain how Theorems 1.1 and 1.2 together with some results from elsewhere imply Theorem 0.1. For further discussion of this, see [Bo8].

The following result was proven in [Bo5] using results and ideas from [Bo3,L,Bo4,Swa].

Theorem 4.1 : *Suppose that* Γ *is a relatively hyperbolic group, and that each peripheral subgroup is finitely presented, one-or-two ended, and contains no infinite torsion subgroup. If* $\partial\Gamma$ *is connected, then every global cut point of* $\partial\Gamma$ *is a parabolic point.* $\diamondsuit$

This represents the main additional ingredient of Theorem 0.1, and is where the hypotheses on peripheral subgroups arise.

Suppose then, that Γ satisfies these hypotheses. Let $\Pi \subseteq \partial\Gamma$ be the set of parabolic points. In [Bo8], it was shown that Π admits a maximal peripheral splitting. Thus, Γ can be split into "components" each of which is itself relatively hyperbolic, and whose boundary can be identified as a closed subset of $\partial\Gamma$. These subsets are glued together in a treelike fashion along parabolic points. Each of these component boundaries is connected. Moreover, if they are all locally connected, we can deduce that $\partial\Gamma$ is locally connected.

Suppose G is a component, so that $\partial G \subseteq \partial\Gamma$. We note that each point of $\Pi \cap \partial G$ is a parabolic point for the action of G on ∂G. (For if the G-stabiliser of p were finite, the construction of [Bo8] would give a Cayley space for Γ with disconnected boundary. In that paper, this is expressed more formally in term of actions on sets. It is implicit in the assertion that the component boundaries are connected, and so, in particular have no isolated points.)

Suppose, for contradiction, that $\partial\Gamma$ is not locally connected. Then, the boundary, ∂G, of some component, $G \leq \Gamma$, is not locally connected. Using the constructions of [Bo7] (as discussed in Section 1), we can find a Cayley space for G which is taut, and which has constant horospherical distorsion about each parabolic point. By Theorem 1.1, this Cayley space admits a separating horoball. The centre, $p \in \partial G$, of this horoball is a global cut point of ∂G, and thus, by an observation in [Bo8], is also a global cut point of $\partial\Gamma$. By Theorem 4.1, p is a parabolic point for the action of Γ on $\partial\Gamma$. In other words, $p \in \Pi$. But now it follows from the observation made earlier, that p is also a parabolic for the action of G on ∂G. In other words, the Cayley space for G admits a separating horoball centred on a parabolic point, and so satisfies the hypotheses of Theorem 1.2 (with G in place of Γ). We deduce that the G-set $\Pi \cap \partial G$ admits a proper peripheral splitting. This splitting gives us a refined peripheral splitting for the Γ-set Π. This contradicts maximality of that splitting, and therefore proves Theorem 0.1. (For more details of the statements concerning peripheral splittings, see [Bo8].)

It seems likely that the hypotheses of Theorem 4.1, and thus of Theorem 0.1, can be weakened. The result probably remains valid if we just assume that the peripheral subgroups are finitely generated, and contain no infinite torsion subgroup. Some variations on the approach to Theorem 4.1 have been suggested by Swenson [Swe], but more work needs to be done to sort out the details.

We note that Theorem 0.1 applies to the case of a geometrically finite group acting on a pinched hadamard manifold (see [Bo1]). In such a case, the peripheral subgroups are necessarily finitely generated nilpotent, and hence satisfy the hypotheses. We conclude:

Corollary 4.2 : *Suppose* Γ *is a geometrically finite group acting on a pinched hadamard manifold. If its limit set is connected, then it is locally connected.* $\diamondsuit$

(Note that the closed convex hull of the limit set is a Cayley space for X. However it will not usually satisfy the constant horospherical distorsion criterion. We therefore need to go through the rigmarole described in [Bo7] to find one that does.)

Note that Corollary 4.2 deals with geometrically finite kleinian groups in any dimension. In this case, some of the technical complications of our proof, such as that mentioned in the last paragraph, can be avoided. See [BoS] for discussion of this case.

5. Cut points in locally connected boundaries.

In this section, we consider one corollary of local connectedness, namely the equivalence of global cut points and separating horoballs (Proposition 5.1). Although this is of no consequence to the general logic of the argument, it rounds off the main discussion of this paper. We shall finish with some brief speculation on more significant potential applications of local connectedness.

Suppose that Γ is a relatively hyperbolic group and that $\partial\Gamma$ is a Peano continuum, i.e. connected and locally connected. Let $\Pi \subseteq \partial\Gamma$ be the set of parabolic points. Suppose that $p \in \Pi$, and let $G \leq \Gamma$ be the stabiliser of p. Now, each parabolic point is "bounded", so $(\partial\Gamma \setminus \{p\})/G$ is compact hausdorff. Note also that by local connectedness, each component of $\partial\Gamma \setminus \{p\}$ is open in $\partial\Gamma$. In particular, components and quasicomponents of $\partial\Gamma \setminus \{p\}$ coincide.

Proposition 5.1 : *Suppose that Γ acts as a geometrically finite group on a taut proper hyperbolic space, X (so that $\partial\Gamma \equiv \partial X$). Suppose ∂X is a Peano continuum. If $p \in \partial X$ is a parabolic point and a global cut point of ∂X, then X has a separating horoball about p.*

Proof : Let C be a component of $\partial X \setminus \{p\}$. We claim that there is a horoball, B, centred on p such that any path in X connecting C to $X \setminus (C \cup \{p\})$ must meet B.

Suppose not. Let $G \leq \Gamma$ be the stabiliser of p. Let $(B_i)_{i\in\mathbf{N}}$ be a sequence of G-invariant horoballs about p with $\bigcup_{i=0}^{\infty} B_i = X$. Now a simple geometric exercise shows that for each i, we can find a geodesic, α_i, connecting points $x_i \in C$ and $y_i \in X \setminus (C \cup \{p\})$ and with $\alpha_i \cap B_i = \emptyset$. Since $(\partial X \setminus \{p\})/G$ is compact, we can suppose, after passing to a subsequence, that x_i converges to some point, $z \in \partial X$. Since $\alpha_i \cap B_i = \emptyset$, we see that y_i also converges on z. Let U be a connected neighbourhood of z in $\partial X \setminus \{p\}$. For all sufficiently large i, we have $x_i, y_i \in U$, contradicting the assumption that they lie in different components. $\diamondsuit$

In retrospect, it is probably possible to greatly simplify the proof of Theorem 4.1 under the assumption of local connectedness. In the case of a hyperbolic groups, an elementary proof of the non-existence of global cut points was given in [Bo2], again under the assumption that the boundary is locally connected.

It is interesting to speculate on what new results about relatively hyperbolic groups could be obtained knowing that the boundary is a locally connected. For example, is it true that every local cut point is either a parabolic point or else plays a role in the JSJ splitting, similar to that described in [Bo2]?

References.

[AdN] S.A.Adeleke, P.M.Neumann, *Relations related to betweenness: their structure and automorphisms* : to appear in Memoirs Amer. Math. Soc.

[AnM] J.W.Anderson, B.Maskit, *On the local connectivity of limit sets of kleinian groups* : Complex Variables **31** (1996) 177–183.

[BeM] M.Bestvina, G.Mess, *The boundary of negatively curved groups* : J. Amer. Math. Soc. **4** (1991) 469–481.

[Bo1] B.H.Bowditch, *Geometrical finiteness with variable negative curvature* : Duke Math. J. **77** (1995) 229–274.

[Bo2] B.H.Bowditch, *Cut points and canonical splittings of hyperbolic groups* : to appear in Acta. Math.

[Bo3] B.H.Bowditch, *Treelike structures arising from continua and convergence groups* : to appear in Memoirs Amer. Math. Soc.

[Bo4] B.H.Bowditch, *Group actions on trees and dendrons* : to appear in Topology.

[Bo5] B.H.Bowditch, *Connectedness properties of limit sets* : to appear in Trans. Amer. Math. Soc.

[Bo6] B.H.Bowditch, *A topological characterisation of hyperbolic groups* : to appear in J. Amer. Math. Soc.

[Bo7] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton (1997).

[Bo8] B.H.Bowditch, *Peripheral splittings of groups* : preprint, Southampton (1997).

[BoS] B.H.Bowditch, G.A.Swarup, *Cut points in the boundaries of hyperbolic groups* : in preparation.

[F] B.Farb, *Relatively hyperbolic groups and automatic groups with applications to negatively curved manifolds* : thesis, Princeton (1994).

[GhH] E.Ghys, P.de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Mathematics, Vol. 83, Birkhäuser (1990).

[Gr] 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.

[HY] J.G.Hocking, G.S.Young, *Topology* : Addison-Wesley (1961).

[K] J.L.Kelley, *General topology* : Graduate Texts in Mathematics, 21, Springer-Verlag (reprint of Van Nostrand edition 1955).

[L] G.Levitt, *Non-nesting actions on actions on real trees* : Bull. London Math. Soc. **30** (1998) 46–54.

[Swa] G.A.Swarup, *On the cut point conjecture* : Electron. Res. Announc. Amer. Math. Soc. **2** (1996) 98–100, (Electronic).

[Swe] E.L.Swenson, *A cut point tree for a continuum* : preprint (1997).

[Sz] A.Szczepański, *Relatively hyperbolic groups* : preprint, Gdansk (1995).

[T] P.Tukia, *Conical limit points and uniform convergence groups* : preprint, Helsinki (1996).

[W] L.E.Ward, *Axioms for cutpoints* : in "General topology and modern analysis", Proceedings, University of California, Riverside (ed. L.F.McAuley, M.M.Rao) Academic Press (1980) 327–336.