Some results on the geometry of convex hulls in manifolds of pinched negative curvature

B. H. Bowditch
Faculty of Mathematical Studies, University of Southampton
Highfield, Southampton SO9 5NH, Great Britain.
Email: bhb@maths.soton.ac.uk

0. Introduction.

A "Hadamard manifold", X, is a complete simply-connected riemannian manifold of non-positive curvature. Such a manifold is diffeomorphic to $\mathbf{R}^\nu$, and can be naturally compactified to a closed ball $X_C = X \cup X_I$ on adjoining the "ideal sphere", X_I. We refer to [BaGS] for a general account of such manifolds.

In this paper we shall be assuming that X has pinched negative curvature, i.e. that all the sectional curvatures lie between two negative constants, which (on scaling the metric) we can take to be $-\kappa^2$ and -1, where $\kappa \geq 1$. In this case, X is a "visibility manifold", which means that any two points $x, y \in X_C$ are joined by a unique geodesic $[x, y]$, (where $[x, x] = \{x\}$). We say that a subset $A \subseteq X_C$ is *convex* if, for all $x, y \in A$, we have $[x, y] \subseteq A$. Given any closed subset $Q \subseteq X_C$, we define the (closed) convex hull, $\mathrm{hull}(Q)$, of Q to be the intersection of all the closed convex sets containing Q. Clearly, $\mathrm{hull}\{x, y\} = [x, y]$.

A major deficiency in the theory of Hadamard manifolds is the sparsity of good constructions of convex sets. In the general situation little seems to be known. The only obvious examples of convex sets are uniform neighbourhoods of points or of geodesic segments, and their intersections. We see, for example, that any three (non-ideal) points in a Hadamard manifold must lie in the boundary of their convex hull. Note that with variable curvature, one would expect generically for the convex hull of three points to have non-empty interior. It is by no means clear what the convex hull of three ideal points might look like, even when given an upper curvature bound away from 0.

In the special case of pinched curvature, there is a much more general construction due to Anderson [A]. Thus, for example, Anderson shows that if $Q \subseteq X_C$ is closed, then $X_I \cap \mathrm{hull}(Q) = X_I \cap Q$. In this paper, we aim to develop further the theory of convex sets in this context. Our paper splits into four sections.

The main result of Section 1 is that the map $[Q \mapsto \mathrm{hull}(Q)]$ which sends a closed set to its convex hull is continuous with respect to the Hausdorff topology (Theorem 1.1). The techniques employed in this section are rather different from the rest of the paper, although the results will be quoted later.

In later sections, we shall focus our attention mainly on convex hulls of finite sets of points. These play a central role in hyperbolic geometry as they are precisely the finite-sided finite-volume polyhedra. One would not expect such a nice picture in pinched variable curvature (for example a natural decomposition into faces), although many properties do generalise.

In Section 2, we describe how the convex hull of finite set $P \subseteq X_C$ is "tree-like", in that it approximates a certain spanning tree for P, in a manner that will be clarified later (Theorem 2.1). An analogous statement for hyperbolic polyhedra has been used [Be] to study the degeneration of discrete hyperbolic groups actions. The importance of generalising this fact is made apparent, for example, in [P].

In Section 3, we give generalisation of Anderson's construction. Specifically, we are aiming at Propositions 3.4 and 3.5.

In Section 4, we put together the ideas from the previous sections to give two new theorems. The first of these, Theorem 4.1, tells us that the volume of the convex hull of a set of n points of X_C is always finite, and in fact is bounded by some constant $C(\nu, \kappa, n)$, depending only on n, the dimension ν, and the pinching constant κ. It turns out that, for fixed ν and κ, $C(\nu, \kappa, n)$ is bounded by some polynomial in n. I suspect, in fact, that this could be improved to a linear function of n. In an appendix, I show that this is indeed the case in constant curvature. The second result of Section 4 (Theorem 4.2) tells us that the volume of the convex hull of a set of n points varies continuously in those points, provided that no two converge on the same ideal point.

The present paper combines two articles written at the University of Melbourne, under an Australian Research Council fellowship. I would like to thank Craig Hodgson for suggesting some of these questions to me. I am also endebted to the referee for many helpful comments.

1. Continuity of convex hulls.

In this section the main result will be Theorem 1.4. First we quote some basic results used through out this paper.

Notation

Recall, we are assuming that all the sectional curvatures of X lie in the interval $[-\kappa^2, -1]$. We write $T_x X$ for the tangent space of X at x. Given $\xi, \zeta \in T_x X$, we write $\langle \xi, \zeta \rangle$ and $|\xi| = \sqrt{\langle \xi, \xi \rangle}$, respectively, for the riemannian inner-product and norm on $T_x X$. Given $x \in X$ and $y \in X_C \setminus \{x\}$, write $\overrightarrow{xy} \in T_x X$ for the initial unit tangent vector of the geodesic from x to y, parameterised by arc length. If $z \in X_C \setminus \{y\}$, write $y\hat{x}z = \cos^{-1}\langle \overrightarrow{xy}, \overrightarrow{xz} \rangle \in [0, \pi]$ for the angle between $\overrightarrow{xy}$ and $\overrightarrow{xz}$. We write d for the induced path-metric on X_C. We shall sometimes refer to d as the "distance function" on X, to avoid any confusion with the riemannian inner-product.

Basic comparison theorems.

Given $\lambda \in (-\infty, 0)$, write $(\mathbf{H}^\nu(\lambda)$ for the ν-dimensional space of constant curvature $-\lambda^2$. We need the following variants of the Toponogov comparison theorems (see for example, [Sp] or [CE]). We write d_λ for the path-metric on $\mathbf{H}^\nu(\lambda)$.

Lemma 1.1 : *Suppose $x \in X$ and $y, z \in X \setminus \{x\}$. Choose points $x', y', z' \in \mathbf{H}^2(1)$ such that $d_1(x', y') = d(x, y)$, $d_1(x', z') = d(x, z)$ and $y'\hat{x}'z' = y\hat{x}z$. Then $d_1(y', z') \leq d(y, z)$.*

Lemma 1.2 : *Suppose $x \in X$ and $y, z \in X \setminus \{x\}$. Choose points $x', y', z' \in \mathbf{H}^2(\kappa)$ such that $d_\kappa(x', y') = d(x, y)$, $d_\kappa(x', z') = d(x, z)$ and $y'\hat{x}'z' = y\hat{x}z$. Then $d_\kappa(y', z') \geq d(y, z)$.*

Thus, the Rauch Comparison Theorem gives us the infinitesimal case with y close to z.

Another basic property of X is the convexity of the distance function, which is essentially Busemann's characterisation of non-positive curvature [Bu]:

Lemma 1.3 : *If $\alpha, \beta : [0,1] \longrightarrow X$ are geodesics parameterised proportionately to arc-length, then the map $[(t,u) \mapsto d(\alpha(t), \beta(u))]$ is convex on $[0,1]^2$.*

Discussion of the main result on continuity.

Let $\mathcal{C}(X_C)$ be the set of all closed subsets of X_C. Now, X_C is a topological ball, and hence metrisable. Choose a metric ρ on X_C. Given $P \in \mathcal{C}(X_C)$ and $r \geq 0$, we write $N(\rho)(P,r) = \{x \in X_C \mid \rho(x,P) \leq r\}$ for the uniform r-neighbourhood of P. Given $P, Q \in \mathcal{C}(X_C)$, write $\mathrm{hd}^+(\rho)(P,Q) \in [0,\infty)$ for the smallest $r \geq 0$ such that $P \subseteq N(\rho)(Q,r)$. Write

$$\mathrm{hd}(\rho)(P,Q) = \max(\mathrm{hd}^+(\rho)(P,Q), \mathrm{hd}^+(\rho)(Q,P)).$$

We call $\mathrm{hd}(\rho)(P,Q)$ the "Hausdorff distance" between P and Q, with respect to ρ. Thus, $\mathrm{hd}(\rho)$ is a metric on $\mathcal{C}(X_C)$. Since X_C is compact, it is easily verified that the induced topology on $\mathcal{C}(X_C)$ is independent of the choice of metric ρ. We refer to it as the *Hausdorff topology*. Thus $\mathcal{C}(X_C)$ is a compact hausdorff space in this topology.

We remark that a more natural approach would be note that since X_C is compact hausdorff, it admits a unique uniformity [K]. This naturally induces a uniformity, and hence a topology, on $\mathcal{C}(X_C)$.

Theorem 1.4 : *The map $[Q \mapsto \mathrm{hull}(Q)] : \mathcal{C}(X_C) \longrightarrow \mathcal{C}(X_C)$ is continuous, where $\mathcal{C}(X_C)$ is given the Hausdorff topology.*

In fact, we shall find a path-metric ρ on X_C such that $[Q \mapsto \mathrm{hull}(Q)]$ is distance non-increasing on $(\mathcal{C}(X_C), \mathrm{hd}(\rho))$.

Note that, clearly, the map $[(x,y) \mapsto \{x,y\}] : X_C \times X_C \longrightarrow \mathcal{C}(X_C)$ is continuous, and so as a special case we have that $[(x,y) \mapsto [x,y]] : X_C \times X_C \longrightarrow \mathcal{C}(X_C)$ is continuous. This is also a corollary of Proposition 1.5 below. However, this statement is easily verified directly, and we may leave it as an exercise. (Indeed, it is true without the lower curvature bound, $-\kappa^2$.)

Another consequence of the continuity of geodesics is that the convex hull map has to be "lower semicontinuous" in the following sense. Suppose ρ is a metric on X_C. Then, given $P \in \mathcal{C}(X_C)$ and $\epsilon > 0$, there is some $\delta > 0$ such that if $\mathrm{hd}^+(\rho)(P,Q) \leq \delta$, then $\mathrm{hd}^+(\rho)(\mathrm{hull}(P), \mathrm{hull}(Q)) \leq \epsilon$. (Note that $\mathrm{hd}^+(\rho)(P,Q) = \mathrm{hd}(\rho)(Q, P \cup Q)$, and so lower semicontinuity can be expressed in terms of the Hausdorff topology, and the partial order on $\mathcal{C}(X_C)$ by set inclusion.) To prove lower semicontinuity, suppose that P_n is any sequence with $\mathrm{hd}^+(\rho)(P,P_n) \to 0$. We claim $\mathrm{hd}^+(\rho)(\mathrm{hull}(P), \mathrm{hull}(P_n)) \to 0$. Let $H \in \mathcal{C}(X_C)$ be the set of all $y \in X_C$ such that $x_n \to y$ for some sequence (x_n) with $x_n \in \mathrm{hull}(P_n)$. From the continuity of geodesics, we see that H is convex. Clearly $P \subseteq H$, and $\mathrm{hull}(P) \subseteq H$. Now, since X_C is compact, we must have $\mathrm{hd}^+(\rho)(H, \mathrm{hull}(P_n)) \to 0$. Otherwise, we could find a sequence of points $y_n \in H$ with $\rho(y_n, \mathrm{hull}(P_n))$ bounded away from 0, and passing to a convergent subsequence would give a contradiction to the definition of H. It follows, then, that $\mathrm{hd}^+(\rho)(\mathrm{hull}(P), \mathrm{hull}(P_n)) \to 0$ as claimed.

We thus see that the lower semicontinuity of convex hulls is fairly trivial. Achieving continuity in the pinched curvature case will involve us in a bit more work. The basic idea is as follows.

Given $Q \in \mathcal{C}(X_C)$, write

$$\mathrm{join}(Q) = \bigcup \{[x, y] \mid x, y \in Q\}.$$

Given the continuity of geodesics, we wee that $\mathrm{join}(Q)$ is closed in X_C. We define, inductively, $\mathrm{join}^{n+1}(Q) = \mathrm{join}(\mathrm{join}^n(Q))$ and $\mathrm{join}^\infty(Q) = \bigcup_{n=1}^\infty \mathrm{join}^n(Q)$. Clearly $\mathrm{join}^\infty(Q)$ is convex, and, again given the continuity of geodesics, we see that, if Q is closed, then $\mathrm{hull}(Q)$ is just the closure of $\mathrm{join}^\infty(Q)$. Our aim, then, will be to find a metric ρ on X_C such that the map $[Q \mapsto \mathrm{join}(Q)]$ is distance non-increasing on $(\mathcal{C}(X_C), \mathrm{hd}(\rho))$. It suffices therefore to show:

Proposition 1.5 : *There is some path-metric ρ on X_C such that if $x_0, y_0, x_1, y_1 \in X_C$, then*

$$\mathrm{hd}(\rho)([x_0, y_0], [x_1, y_1]) \leq \max(\rho(x_0, x_1), \rho(y_0, y_1)).$$

Some other observations about continuity.

Before we set about proving this, let us note another more trivial sense in which convex hulls vary continuously. We may define, in a similar fashion, a Hausdorff distance, $\mathrm{hd}(d)$, on the set $\mathcal{C}(X)$ of all closed subsets of X. In this case, the analogue of Proposition 1.5 follows directly from the convexity of the distance function (Lemma 1.3). We deduce:

Proposition 1.6 : *The map $[Q \mapsto \mathrm{hull}(Q)]$ is distance non-increasing on $(\mathcal{C}(X), \mathrm{hd}(d))$.*

On the subset of $\mathcal{C}(X_C)$ consisting of all compact subsets of X, the topologies given by $\mathrm{hd}(d)$ and $\mathrm{hd}(\rho)$ agree. However, in general, the topologies are quite different. For example, $(\mathcal{C}(X), \mathrm{hd}(d))$ has infinitely many components.

A related observation which will be used in Section 4 is:

Lemma 1.7 : *If $P, Q \subseteq X$ are convex, then $\mathrm{hd}(d)(\partial P, \partial Q) \leq \mathrm{hd}(d)(P, Q)$.*

Proof : Suppose for contradiction, that $\mathrm{hd}(d)(P, Q) = h$, and $\mathrm{hd}(\partial P, \partial Q) > h$. Without loss of generality, there is some $x \in \partial P$, with $d(x, \partial Q) = k > h$. Now $d(x, Q) \leq h$ so $N(d)(x, k) \subseteq Q$. Since P is convex, it's easy to see that there is some $y \in \partial N(d)(x, k)$ with $d(y, P) = k$, contradicting $\mathrm{hd}(d)(P, Q) < k$. ◊

Putting the last two results together, we see that the map $[Q \mapsto \partial \mathrm{hull}(Q)]$ is also distance nonincreasing on $(\mathcal{C}(X), \mathrm{hd}(d))$.

The metric ρ.

We next construct the metric ρ on X_C described by Proposition 1.5. In what follows, we shall write $|ds|$ for a riemannian norm defined pointwise on our space. This induces a path-metric, d, giving the distance between two points.

The metric ρ on X_C will arise from a construction of Floyd [F] (described originally in the context of discrete groups). We introduce this construction with reference to the Poincaré model for hyperbolic ν-space $\mathbf{H}^\nu = \mathbf{H}^\nu(1)$. Recall that $\mathbf{H}^\nu$ may be realised a conformal metric on the euclidean open unit ball, B, obtained by pointwise scaling the euclidean riemannian norm $|ds_{euc}|$. Thus, the hyperbolic norm, $|ds_{hyp}|$ is given at the point $x \in B$ by the formula $|ds_{hyp}| = \frac{2}{1-h^2}|ds_{euc}|$, where $h \in [0,1)$ is the euclidean distance $d_{euc}(o,x)$ from the origin $o \in B$. This induces the hyperbolic path-metric d_{hyp}. We may invert the process. To recover the euclidean ball, we fix a point $p \in \mathbf{H}^\nu$ and scale the riemannian norm at the point $x \in \mathbf{H}^\nu$ by a factor of $\frac{1}{2}\operatorname{sech}^2(r/2)$ where $r = d_{hyp}(x,p)$.

We can generalise this idea to our manifold X. Suppose that $f : [0,\infty) \longrightarrow (0,\infty)$ is a smooth function with $\int_0^\infty f(r)dr = R < \infty$. Fix any point $p \in X$, and set $\phi(x) = f(d(x,p))$ for $x \in X$. We now scale the riemannian norm $|ds|$ on X according to the function ϕ. Thus, the new norm, $|ds_f|$ is given at the point $x \in X$ by $|ds_f| = \phi(x)|ds|$. In this way, we get a riemannian metric (at least on $X \setminus \{p\}$), and we write d_f for the induced path-metric. In general, there may be a singularity at the point p. However, if f has the form $f(r) = f_0(r^2)$, where f_0 is smooth on a neighbourhood of 0, then the map $\phi : X \longrightarrow (0,\infty)$ will be smooth at p, and so we get a riemannian metric everywhere.

Now all d-geodesic rays emanating from p are also d_f-geodesic paths, each of which has d_f-length equal to R. (Note that if γ is a smooth curve joining p to some point q with $d(p,q) = k > 0$ and parameterised by arc length dt, then $\frac{d}{dt}d(p,\gamma(t)) \leq 1$, and so the d_f-length of γ is at least $\int_0^k f(r)dr$, with equality if and only if γ is a d-geodesic.) Also, if $s < R$, then $N(d_f,p,s) = N(d,p,r)$ where r is given by $\int_0^r f(t)dt = s$. In particular, each such ball is compact.

The idea, then, is to describe X_C as the metric completion of (X,d_f). However, we first need to ensure that f does not decay too fast. (For example, if we had $f(r) = O(e^{-\lambda r})$ with $\lambda > \kappa$, then we would just obtain the one-point compactification of X.) Suppose then that, for some $r_0 > 0$, we have $f(r) \geq \operatorname{cosech} r$ for all $r \geq r_0$. In this case, we have the following property. Suppose that β is a smooth path in $X \setminus N(d,p,r)$ joining points y and z. Then $y\hat{p}z$ is less than or equal to the d_f-length of β. This fact may be deduced from Lemma 1.1, or directly from its infinitesimal version (the Rauch Comparison Theorem). Now, we may use the d_f-exponential map based at p to identify X with a euclidean open metric ball B. It is easily checked that X_C is naturally identified with its closure N, so that the topologies agree. We thus need to verify that N is indeed the metric completion of $X \equiv B$ with respect to the metric d_f. To this end, we make the following simple observation:

Lemma 1.8 : *Suppose N is a compact, first countable topological space. Suppose $B \subseteq N$ is a dense subset which admits a metric ρ inducing the subspace topology on B. Then, N is (naturally homeomorphic to) the completion of B precisely if the following condition holds. Suppose (x_i) and (y_i) are sequences in B converging respectively to $x, y \in N$. Then*

$\rho(x_i, y_i) \to 0$ *if and only if* $x = y$. ♢

We apply this to our situation, with $\rho = d_f$. The "only if" part of the above criterion follows from the relation of d_f-length to visual distance at p already referred to. The "if" part is an exercise, on noting that euclidean distance along any ray emanating from p agrees with d_f-distance. (We remark that we do not need the lower curvature bound for this construction, unless we want explicit estimates for d_f.)

For definiteness, in the rest of this paper we shall set $f(r) = (\operatorname{sech} \kappa r)^\mu$ where $\mu > 0$ is sufficiently small. Specifically, we set $\mu = 1/4\kappa^2$. We choose this particular form for computational convenience. There is probably nothing very special about this formula, and I suspect that Proposition 1.5 is true much more generally.

We write $\rho = d_f$. Now, the completion of a path-metric space is a path-metric space, and so ρ is a path-metric on X_C. Suppose that $I \subseteq \mathrm{R}$ is some interval, and $\gamma : I \longrightarrow X$ is a smooth path. We may define the ρ-length of γ as

$$\mathrm{length}_\rho \gamma = \int_I \phi(\gamma(u)) \left| \frac{d\gamma}{du}(u) \right| du$$

where $d\gamma/du$ is shorthand for $\gamma_*(d/du)$. Clearly $\mathrm{length}_\rho \gamma$ agrees with the rectifiable length. Now, standard riemannian geometry allows us to approximate rectifiable paths by smooth paths of nowhere-vanishing derivative, and so:

Lemma 1.9 : *Suppose* $x, y \in X$ *and* $\epsilon > 0$. *Then there is a smooth path* $\gamma : [0,1] \longrightarrow X$ *such that* $\gamma(0) = x$, $\gamma(1) = y$ *and*

$$\phi(\gamma(u)) \left| \frac{d\gamma}{du}(u) \right| \leq \rho(x, y) + \epsilon$$

for all $u \in [0,1]$.

Proof of main theorem.

At last, we are ready to start on the proof of Proposition 1.5. To begin with, let us suppose that x_0, y_0, x_1, y_1 all lie in X. We shall describe later how to deal with ideal points. Set $l = \max(\rho(x_0, x_1), \rho(y_0, y_1))$. By Lemma 1.9, we can find paths $\gamma_i : [0,1] \longrightarrow X_C$ with $\gamma_i|(0,1)$ smooth, with $\gamma_0(0) = x_0$, $\gamma_0(1) = x_1$, $\gamma_1(0) = y_0$ and $\gamma_1(1) = y_1$, and so that

$$\phi(\gamma_i(u)) \left| \frac{d\gamma_i}{du}(u) \right| \leq l + \epsilon$$

for all $u \in [0,1]$ and $i = 0, 1$. It will be convenient to assume that $\gamma_0(u) \neq \gamma_1(u)$ for all $u \in (0,1)$. This can always be achieved by a small perturbation. (Alternatively, it will not be hard to see how to deal with a degenerate situation.)

Our first task is to span the rectangle $\gamma_0 \cup [x_0, y_0] \cup \gamma_1 \cup [x_1, y_1]$ by a ruled surface. More specifically, we are looking for a closed subset $S \subseteq \mathbf{R} \times [0,1]$ together with a smooth map $\beta : S \longrightarrow X$ with the following properties.

(1) There are smooth functions $q_0, q_1 : [0,1] \longrightarrow \mathbf{R}$ such that $q_0(u) < q_1(u)$ for all $u \in (0,1)$, and so that $S = \{(t,u) \in \mathbf{R} \times [0,1] \mid q_0(u) < t < q_1(u)\}$ (Figure 1a).

(2) $\gamma_i = \beta \circ \sigma_i$, where $\sigma_i : [0,1] \longrightarrow \mathbf{R} \times [0,1]$ is given by $\sigma_i(u) = (q_i(u), u)$ for $i = 0, 1$.

(3) The map $\alpha_u = [t \mapsto \beta(t,u)] : [q_0(u), q_1(u)] \longrightarrow X$ is a d-geodesic parameterised with respect to arc length, for all $u \in [0,1]$.

(4) $\left\langle \frac{\partial \beta}{\partial t}(t,u), \frac{\partial \beta}{\partial u}(t,u) \right\rangle = 0$ for all $(t,u) \in S$.

Note, in property (4), that the vectors $\frac{\partial \beta}{\partial t} = \beta_*(\partial/\partial t)$ and $\frac{\partial \beta}{\partial u} = \beta_*(\partial/\partial u)$ are well-defined over the whole of S.

Now,

$$\frac{d\sigma_i}{du} = \frac{dq_i}{du}\frac{\partial}{\partial t} + \frac{\partial}{\partial u}$$

and so

$$\frac{d\gamma_i}{du}(u) = \frac{dq_i}{du}(u)\frac{\partial \beta}{\partial t}(\sigma_i(u)) + \frac{\partial \beta}{\partial u}(\sigma_i(u)).$$

Thus,

$$\frac{dq_i}{du}(u) = \left\langle \frac{d\gamma_i}{du}(u), \frac{\partial \beta}{\partial t}(\sigma_i(u)) \right\rangle.$$

Note that $\xi_i(u) = \frac{\partial \beta}{\partial t}(\sigma_i(u))$ is determined by the points $x = \sigma_0(u)$ and $y = \sigma_1(u)$. Thus $\xi_0(u) = \overrightarrow{xy}$ and $\xi_1(u) = -\overrightarrow{yx}$.

Suppose, then, that we have γ_0 and γ_1, and want to construct β. We can obtain the functions q_i, up to an additive constant, by integrating the quantity $\left\langle \frac{d\gamma_i}{du}(u), \xi_i(u) \right\rangle$. We see easily that $\frac{d}{du}(q_1(u) - q_0(u)) = \frac{d}{du}(d(\gamma_0(u), \gamma_1(u)))$, and so we can arrange that $q_1(u) - q_0(u) = d(\gamma_0(u), \gamma_1(u))$ for all $u \in [0,1]$. Now, let $\alpha_u : [q_0(u), q_1(u)] \longrightarrow X$ be the geodesic joining $\gamma_0(u)$ to $\gamma_1(u)$, parameterised with respect to arc-length. Define $\beta : S \longrightarrow X$ by $\beta(t,u) = \alpha_u(t)$. Thus, $\beta \circ \sigma_i = \gamma_i$ and $\frac{\partial \beta}{\partial t}(\sigma_i(u)) = \xi_i(u)$ for $u \in (0,1)$. Now (from the Implicit Function Theorem), we know that $\xi_0(u)$ varies smoothly in u. It follows that β is smooth. We need finally to verify property (4). From the formula for $\frac{dq_0}{du}$, we find that $\left\langle \frac{\partial \beta}{\partial t}(\sigma_0(u)), \frac{\partial \beta}{\partial u}(\sigma_0(u)) \right\rangle = 0$ for all $u \in (0,1)$. Now the vector field $\left[t \mapsto \frac{\partial \beta}{\partial u}(t,u)\right]$ along α_u is the first variation of a geodesic, and so its component parallel to α_u is constant, and thus equal to 0, i.e. $\left\langle \frac{\partial \beta}{\partial t}(t,u), \frac{\partial \beta}{\partial u}(t,u) \right\rangle = 0$ for all $(t,u) \in S \cap (\mathbf{R} \times (0,1))$ and so, by continuity, for all $(t,u) \in S$. We have thus constructed β.

We now claim:

Lemma 1.10 : *For all $(t,u) \in S$, we have*

$$\phi(\beta(t,u)) \left| \frac{\partial \beta}{\partial u}(t,u) \right| \leq l + \epsilon.$$

Given this lemma, we may complete the proof of Proposition 1.5 as follows:

Suppose $x_0, y_0, x_1, y_1 \in X$, and S, β are as above. Given $t \in [q_0(0), q_0(1)]$, let $\tau : [0,1] \longrightarrow S$ be the path defined as follows. If $q_0(u) < t < q_1(u)$ for all $u \in (0,1)$, we set $\tau = [u \mapsto (t,u)]$. Otherwise, we let τ begin as the path $[u \mapsto (t,u)]$ and continue until it runs into either σ_0 or σ_1. We then continue along either σ_0 or σ_1 until we arrive at $\sigma_0(1)$ or $\sigma_1(1)$.

Now, let $\delta = \beta \circ \tau : [0,1] \longrightarrow X$. Thus δ is a path joining $\alpha_0(t) \in [x_0, y_0]$ to $\alpha_1(\delta(1)) \in [x_1, y_1]$. Moreover, $\frac{d\delta}{du}(u)$ is either $\frac{\partial \beta}{\partial u}(t,u)$ or $\frac{d\gamma_i}{du}(u)$. In any case, we have $\phi(\delta(u)) \left| \frac{d\delta}{du}(u) \right| \leq l + \epsilon$, and so $\text{length}_\rho \delta \leq l + \epsilon$. Thus $\rho(\alpha_0(t), [x_1, y_1]) \leq l + \epsilon$. But t and $\epsilon > 0$ were arbitrary, and we may also invert the roles of $[x_0, y_0]$ and $[x_1, y_1]$, and conclude that $\text{hd}(\rho)([x_0, y_0], [x_1, y_1]) \leq l$.

Now suppose that $x_0, y_0, x_1, y_1 \in X_C$ are arbitrary. Choose $\epsilon > 0$. If $x_0 \neq y_0$, then we can find $x_0', y_0' \in [x_0, y_0] \cap X$ so that $[x_0, x_0'] \subseteq N(\rho)(x_0, \epsilon)$ and $[y_0, y_0'] \subseteq N(\rho)(y_0, \epsilon)$. (This is trivial given that ρ induces the usual topology on X_C.) If $x_0 = y_0$, we find $x_0' = y_0' \in X$ so that $d(x_0, x_0') \leq \epsilon$. In either case, we have $\text{hd}(\rho)([x_0, y_0], [x_0', y_0']) \leq \epsilon$. We can similarly find $x_1', y_1' \in X$ with $\text{hd}(\rho)([x_1, y_1], [x_1', y_1']) \leq \epsilon$. The general case of Proposition 1.5 now follows by applying the first part, and letting ϵ tend to 0.

Proof of Lemma 1.10 : Fix $u \in (0,1)$ and write $q_0 = q_0(u)$ and $q_1 = q_1(u)$. For $t \in [q_0, q_1]$, set

$$g(t) = \phi(\beta(t,u)),$$

$$j(t) = \left| \frac{\partial \beta}{\partial u}(t,u) \right|$$

and

$$G(t) = g(t) j(t).$$

We want to show that $G(t) \leq l + \epsilon$.

Now $j(q_i) = \left| \frac{\partial \beta}{\partial u}(\sigma_i(u)) \right| \leq \left| \frac{d\gamma_i}{du}(u) \right|$, and so $G(q_i) \leq \phi(\gamma_i(u)) \left| \frac{d\gamma_i}{du}(u) \right| \leq l + \epsilon$. It thus suffices to see that G cannot attain a maximum in the open interval (q_0, q_1).

We shall use primes and double primes, G', G'' etc., to denote the first and second derivatives with respect to t.

Write $\alpha = \alpha_u$ for the geodesic $[t \mapsto \beta(t,u)]$. Now, $\left[t \mapsto \frac{\partial \beta}{\partial u}(t,u) \right]$ is a Jacobi field along α. Thus, except where it vanishes, j is smooth in t. Moreover, from the Jacobi equation and the upper curvature bound (see for example [CE]), we have that $j''(t) \geq j(t)$.

We shall want to bound the first and second derivatives of g. Now,

$$g(t) = \phi(\alpha(t)) = f(h(t)) = (\text{sech}\, \kappa h(t))^\mu$$

where $h(t) = d(p, \alpha(t))$, and $\mu = 1/4\kappa^2$. Thus, $g(t) = (H(t))^{-\mu}$ where $H(t) = \cosh \kappa h(t)$. We claim that $|H'(t)| \leq \kappa H(t)$ and $|H''(t)| \leq \kappa^2 H(t)$. Note that H is smooth, even in the case where $\alpha(t) = p$. In this special case, the inequalities are easily verified, so we shall assume that $\alpha(t) \neq p$. Let $r(x) = d(x,p)$, so $h(t) = r(\alpha(t))$.

Now, $H'(t) = \kappa dr(\alpha'(t)) \sinh \kappa h(t)$. Since $|dr| \leq 1$, the first inequality follows.

For the second inequality, write $D^2 r$ for the second derivative of r at the point $x = \alpha(t)$. Thus $D^2 r$ restricted to $\ker dr$ is the second fundamental form of the sphere of radius

$r(x) = h(t)$ at x. From the lower curvature bound, the principal curvatures of such a sphere are at most $\kappa \coth \kappa(h(t))$ (i.e. that of a sphere of radius $h(t)$ in $\mathbf{H}^\nu(\kappa)$). We see that $|D^2 r(\alpha'(t), \alpha'(t))| \leq \kappa \coth \kappa h(t)(1 - dr(\alpha'(t))^2)$. Now,

$$H''(t) = \kappa \sinh(\kappa h(t)) D^2 r(\alpha'(t), \alpha'(t)) + \kappa^2 \cosh(\kappa h(t))(dr(\alpha'(t)))^2,$$

from which we deduce that $0 \leq H''(t) \leq \kappa^2 H(t)$, as required. This proves the claim.

Now, recall that $g(t) = (H(t))^{-\mu}$. Thus $g'(t) = -\mu H'(t)(H(t))^{-1-\mu}$ and $g''(t) = -\mu H''(t)(H(t))^{-1-\mu} + \mu(1+\mu)(H'(t))^2(H(t))^{-2-\mu}$. We see that

$$|g'(t)| \leq \kappa\mu g(t)$$

and

$$|g''(t)| \leq \kappa^2 \mu(2+\mu) g(t).$$

Now, finally, suppose for contradiction, that $G(t) = g(t)j(t)$ attains a maximum a some point $t \in (q_0, q_1)$. Thus $G'(t) = g'(t)j(t) + g(t)j'(t) = 0$ and so

$$\begin{aligned} \frac{G''(t)}{G(t)} &= \frac{j''(t)}{j(t)} - 2\left(\frac{g'(t)}{g(t)}\right)^2 + \frac{g''(t)}{g(t)} \\ &\geq 1 - 2(\kappa\mu)^2 - \kappa^2\mu(2+\mu) \\ &= 1 - \kappa^2\mu(2+3\mu) \\ &\geq 1 - 3\kappa^2\mu \geq 1/4. \end{aligned}$$

Thus $G''(t) > 0$ contradicting the existence of such a t.

In summary, there is no maximum of G on (q_0, q_1) and so $G(t) \leq \max(G(q_0), G(q_1)) \leq l + \epsilon$ as required. ◇

2. Spanning trees.

In this section, we describe the treelike nature of convex hulls. First, we introduce some terminology and notation.

Notation.

From now on, we deal with only one metric on X, namely d, the path-metric induced from the riemannian metric on X. If $Q \subseteq X_C$ is closed, we write $N(Q, r) = Q \cup \{x \in X \mid d(x, Q \cap X) \leq r\}$ for the uniform r-neighbourhood of Q. Thus, $N(Q, r)$ is closed in X_C.

By a *(combinatorial) tree*, T, we mean a simply-connected finite 1-complex, with vertex set $V(T)$, and edge set $E(T)$. We write $V_0(T) \subseteq V(T)$ for the set of *extreme points* of T, i.e. the those vertices which have degree 1. We demand that each vertex of $V_1(T) = V(T) \setminus V_0(T)$ should have degree at least 3. It follows that $|V(T)| \leq 2|V_0(T)| - 2$, and so there are only a finite number of combinatorial types of trees with a given number of extreme points. Given $s, t \in T$, we shall write $\alpha(s, t)$ for the arc in T joining s to t.

Suppose that $P \subseteq X_C$ is finite. By a *(geodesic) spanning tree*, (T, f) for P, we mean a tree T together with a map $f : T \longrightarrow X_C$, such that:

(a) $f|V_0(T)$ is a bijection from $V_0(T)$ to P,

(b) $f(V_1(T)) \subseteq X$, and

(c) if $e \in E(T)$, then $f(e) = [f(v), f(w)]$ where $v, w \in V(T)$ are the endpoints of e.

It will be convenient to allow for the possibility that $v = w$ so that $f(e)$ is a single point. Otherwise, we shall assume that $f|e$ is injective. Note that, up to isotopy along the edges, f is determined by its restriction to $V(T)$. Note also that $f(T) \cap X_I = P \cap X_I$.

The main theorem on spanning trees.

Theorem 2.1 : *(Figure 2a.) Suppose that $P \subseteq X_C$ is a set of n points. Then, there is a spanning tree (T, f) for P such that*

(1) hull(P) *contains $f(T)$ and lies inside an r_1-neighbourhood of $f(T)$, and*

(2) Suppose $s, t \in T$ and u lies in the arc $\alpha(s, t) \subseteq T$ joining s to t. If β is any path from $f(s)$ to $f(t)$ lying in hull(P)*, then $f(u)$ lies a distance at most r_2 from β.*

Here $r_i = r_i(\kappa, n)$ are functions only of n and κ, which have the form $r_i(\kappa, n) = \lambda(\kappa) + \mu_i(n)$. Moreover, we can arrange that $\mu_1(n) = O(\log \log n)$ and $\mu_2(n) = O(\log n)$.

In most (if not all) cases, one can take f to be injective, so that we get an embedded tree. If the dimension ν is at least 3, this can always be acheived by a small perturbation. It seems more natural, however, to speak in terms of immersed trees.

Note that property (1), alone, is not sufficient to capture the treelike nature of hull(P). Without property (2), we could form a spanning tree simply by choosing any point $a \in$ hull(P), and joining it to each point of P by a geodesic path. In this way, r_1 would be independent of n.

Even if property (2) is added, I suspect that r_1 can be made independent of n, i.e. that we should be able to get rid of the term $\mu_1(n)$. However, $\mu_1(n) = O(\log \log n)$ is the best I can do. On the other hand, $\mu_2(n) = O(\log n)$ is the best possible, as can be seen by considering a set of n points evenly spaced about a circle of radius r in the hyperbolic plane. In this case, the convex hull is a regular polygon with n vertices. It is not hard to see that the best spanning tree (in the sense of minimising $\mu_2(n)$) is obtained by joining each vertex to the centre by a geodesic segment of length r. Now, as r tends to infinity, $2r$ minus the length of a side of the polygon tends to $-\log \sin(\pi/n) = O(\log n)$. I make no attempt here to find the best multiplicative constant.

There are several ways one might attempt to refine this result. One of these will be relevant to the proof of Theorem 4.1 in Section 4. Note that the term $\mu_2 = O(\log n)$ only really enters when we have a cluster of $O(n)$ vertices of $f(V_1(T))$ in a small region of X. Thus, if we have a long edge $f(e)$ in our spanning tree, we would expect that hull(P) should have small cross-section along most of $f(e)$. In other words, hull(P) separates into two pieces joined by a long thin tube, which we can imagine as a tubular neighbourhood of $f(e)$. Such tubes have bounded volume, as will be explained in Sections 3 and 4.

It is by no means clear that the lower curvature bound $-\kappa^2$ is necessary. Perhaps the term $\lambda(\kappa)$ can be removed. However, Anderson's construction gives $\lambda(\kappa) \to \infty$ as $\kappa \to \infty$. For this reason, we do not bother to estimate $\lambda(\kappa)$ here. The reader can obtain such an estimate by referring to [A] and [Bo2]. We note however that λ can be assumed continuous in κ.

A basic geometric lemma.

To study the geometry of spanning trees, we shall need a simple result (Lemma 2.3) related to well-known facts about the approximation of quasigeodesics by geodesics in hyperbolic space. The argument we apply is a standard one. First, we note the following simple consequence of Toponogov's comparison theorem (Lemma 1.1), and some hyperbolic trigonometry:

Lemma 2.2 : *Suppose that $a, b \in X$, and $p \in [a, b]$ is the midpoint of $[a, b]$. Set $r = d(a, p) = \frac{1}{2} d(a, b)$. Suppose that β is a path from a to b with $d(p, \beta) \geq r$. Then* $\operatorname{length} \beta \geq \pi \sinh r$. ♢

Lemma 2.3 : *Suppose the points $x, y \in X$ are joined by a path β of length at most $d(x, y) + h$, where $h \geq 0$. Then, β lies inside a $\phi(h)$-neighbourhood of the geodesic $[x, y]$. Conversely, $[x, y]$ lies inside a $\theta(h)$-neighbourhood of β. Here $\theta(h) = O(\log h)$ and $\phi(h) = O(h)$ are universal functions of h.*

Proof : Choose $p \in [x, y]$ so as to maximise $d(p, \beta)$. Let $r = d(p, \beta)$. Let $a \in [x, p]$ and $b \in [y, p]$ satisfy $d(a, p) = d(b, p) = r$. If $d(x, p) \geq 2r$, let $a' \in [x, p]$ be the point with $d(a', p) = 2r$, and choose $z \in \beta$ with $d(z, a') \leq r$. If $d(a', p) < 2r$, set $a' = z = x$. Note that $d(p, [a', z]) \geq r$. Similarly choose a point $b' \in [y, p]$ and $w \in \beta$ with $d(w, b') \leq r$ and $d(p, [b', w]) \geq r$ (Figure 2b). It will not matter to us in what order the points z and w occur along β. Let γ be the segment of β lying between z and w. Then, by Lemma 2.2, we have

$$d(a, a') + d(a', z) + \operatorname{length} \gamma + d(w, b') + d(b', b) \geq \pi \sinh r,$$

and so

$$\operatorname{length} \gamma \geq \pi \sinh r - 4r.$$

Let β' be the path obtained from β by replacing γ with the path $[z, a'] \cup [a', b'] \cup [b', w]$. We have

$$\begin{aligned} h = \operatorname{length} \beta - d(x, y) &\geq \operatorname{length} \beta - \operatorname{length} \beta' \\ &\geq \operatorname{length} \gamma - 6r \geq \pi \sinh r - 10r. \end{aligned}$$

Thus $r \leq \theta(h)$ where $\theta(h) = O(\log h)$, and so

$$[x, y] \subseteq N(\beta, \theta(h)).$$

Now suppose that $q \in \beta$. The point q divides β into two subpaths β_1 and β_2. By continuity, we can find some $s \in [x, y]$ with $d(s, \beta_1) \leq \theta(h)$ and $d(s, \beta_2) \leq \theta(h)$. Since $\operatorname{length} \beta \leq d(x, y) + h$, it follows easily that $d(p, s) \leq \phi(h) = 2\theta(h) + h/2 = O(h)$. Thus $\beta \subseteq N([x, y], \phi(h))$. ♢

Spanning trees.

Next, we describe the spanning tree construction. Given the upper curvature bound, we see that X is k-hyperbolic in the sense of Gromov [Gr], for some fixed parameter, k. (Here k depends only on the precise formulation of hyperbolicity we choose to use.) In [Gr, Section 3.2], Gromov outlines a method of constructing spanning trees of finite sets in such spaces. We quote the following refinement of this result [Bo2, Theorem 7.6.1]:

Lemma 2.4 : *Suppose $P \subseteq X$ is a set of n points. Then, there is a geodesic spanning tree, (T, f) for B, with the property that if $v, w \in V_0(T)$, then*

$$\operatorname{length} f(\alpha(v,w)) \leq d(f(v), f(w)) + h(n)$$

where $h(n) = O(\log n)$. ◊

Here, $\operatorname{length} f(\alpha(v,w))$ is equal to $\sum_{i=1}^{p} d(f(v_i), f(v_{i-1}))$ where $v_0 = v$, $v_p = w$ and $v_1, v_2, \ldots, v_{p-1}$ are the successive points of $V(T)$ along the arc $\alpha(v,w)$. Note that it follows that for arbitrary $s, t \in T$, then $\operatorname{length} f(\alpha(s,t)) \leq d(f(s), f(t)) + h(n)$. Inspection of the construction of [Bo1, Chapter 7], shows that $f(T) \subseteq \operatorname{hull}(P)$.

Most of the work in proving this lemma is involved in obtaining the logarithmic bound on $h(n)$ (which gives us the polynomial bound on volume in Section 3). If one is unconcerned about this, it is possible to give an elementary argument as follows. We choose an arbitrary order on the set of n points, and construct an embedded spanning tree $f(T)$ inductively by joining the $(i+1)$th point by a geodesic arc to the nearest point on the spanning tree of the first i points (see [Bo1, Lemma 3.3.1]). We easily see the existence of some bound $h(n)$. With some work, it turns out to be linear in n. (Unfortunately, the argument of [Bo1, Chapter 7] is not guaranteed to give us an embedded tree in the case where X has dimension 2, though I suspect this ought to be possible.)

We want a version of Lemma 2.4 which allows for the possibility of P containing some ideal points:

Lemma 2.5 : *Suppose $P \subseteq X_C$ is a set of n points. Then, there is a spanning tree (T, f) for P such that if $s, t \in T$ and $\alpha(s,t)$ is the arc joining them, then*

$$\operatorname{length} f(\alpha(s,t)) \leq d(f(s), f(t)) + h(n),$$

where $h(n) = O(\log n)$ is the same constant as in Lemma 2.4.

Proof : As remarked after Lemma 2.4, the case where $P \subseteq X$ is already dealt with.

For a general $P \subseteq X_C$, we choose a sequence (P_i) of subsets of X, each with n points, and with P_i tending to P. From the first part, we obtain a spanning tree (T_i, f_i) for each P_i. We can imagine $V_0 = V_0(T)$ as a fixed set, with $f_i(v)$ tending to a certain element $f(v) \in P$, for all $v \in V_0$. Thus, $f : V_0 \longrightarrow P$ is a bijection. Now there are only finitely many possibilities for combinatorial trees with extreme points V_0. Thus, passing to a subsequence, we can take $T_i = T$ to be a fixed tree. It now suffices to define $f(u)$ for all $u \in V_1(T) = V(T) \setminus V_0$. We shall take $f(u)$ to be a limit point of the sequence $f_i(u)$. However we do not want $f(u)$ to be an ideal point, so we have to rule out this possibility.

Suppose then, that $u \in V_1(T)$. By definition, u has degree at least 3. Choose $v_1, v_2, v_3 \in V_0(T)$ so that no two lie in the same component of $T \setminus \{u\}$. In other words, $u \in \alpha_1 \cap \alpha_2 \cap \alpha_3$ where $\alpha_j = \alpha(v_j, v_{j+1})$ and $3 + 1 = 1$. From the construction, and applying Lemma 1.3, we have $f_i(\alpha_j) \subseteq N([f_i(v_j), f_i(v_{j+1})], \rho)$ for all $i \in \mathbf{N}$ and $j \in \{1, 2, 3\}$, where $\rho = \phi(h(n))$. In particular, $f_i(u) \in \bigcap_{j=1}^{3} N([f_i(v_j), f_i(v_{j+1})], \rho)$. Now, as $i \to \infty$, we have $f_i(v_j) \to f(v_j)$ and so the geodesic $[f_i(v_j), f_i(v_{j+1})]$ converges to $[f(v_j), f(v_{j+1})]$. In particular, given any $\epsilon > 0$, then for all sufficiently large i, we have $f_i(u) \in N = \bigcap_{j=1}^{3} N([f(v_j), f(v_{j+1})], \rho + \epsilon)$. Now, this intersection, N, is a compact subset of X (see the discussion of "centres" in [Bo1, Chapter 3].) Thus, passing to a subsequence, we have that $f_i(u)$ converges to a point $f(u) \in X$.

We have thus defined $f : V(T) \longrightarrow X_C$. We may extend f over T by sending each edge $e \in E(T)$ to the geodesic segment $[f(t), f(u)]$ where $t, u \in V(T)$ are the endpoints of e. Note that $f_i(e)$ converges to $f(e)$, so the conclusion of the lemma may be verified. $\diamondsuit$

Note that, in the above proof, we have $f_i(u) \in \mathrm{hull}(P_i)$ for all i, and for all $u \in V_1(T)$. It follows, by Theorem 1.5, that $f(u) \in \mathrm{hull}(P)$. Thus, $f(T) \subseteq \mathrm{hull}(P)$.

Proof of the main theorem.

From now on, we assume that X as curvature pinched between $-\kappa^2$ and -1. The proof of Theorem 2.1 will combine the results of the last section with the convex hull construction of Anderson [A]. The ideas behind this construction will be described in Section 3. For the present section we just need to quote one direct consequence, which is described in [Bo2].

We say that a closed set $Q \subseteq X_C$ is K-quasiconvex if a geodesic joining any two points x, y of Q remains within a distance K of Q, i.e. $[x, y] \subseteq N(Q, K)$. In [Bo2] it was shown that, in such a case, $\mathrm{hull}(Q)$ lies in a uniform R-neighbourhood of Q, where R depends only on K and κ. The idea is that if we are sufficiently far away from a quasiconvex set, it will appear "small" as measured by the maximal angle subtended by two points in the set. Now Anderson's construction may be used to find a convex surface separating us from the set.

Now suppose $Q \subseteq X_C$ is an arbitrary closed set. Recall the definition, $\mathrm{join}(Q) = \bigcup\{[x, y] \subseteq X_C \mid x, y \in Q\}$, thought of as a first approximation to the convex hull. Now any two points of $\mathrm{join}(Q)$ can be joined by a piecewise geodesic path in $\mathrm{join}(Q)$ with at most 3 geodesic segments. It follows that $\mathrm{join}(Q)$ is $(2 \cosh^{-1} \sqrt{2})$-quasiconvex. We arrive at the following (described in [Bo2]):

Lemma 2.6 : *If $Q \subseteq X_C$ is closed, then* $\mathrm{hull}(Q)$ *lies inside a σ-neighbourhood of* $\mathrm{join}(Q)$ *where $\sigma = \sigma(\kappa)$ is some fixed function of the pinching constant κ.*

(Although it is not explicitly stated in [Bo2], it is apparent from the construction that σ is independent of the dimension ν.)

Proof of Theorem 2.1 : Suppose that (T, f) is the spanning tree given by Lemma 2.5. Thus $f(T) \subseteq \mathrm{hull}(P)$. Applying Lemma 2.3, we see that if $s, t \in T$, we have $[f(s), f(t)] \subseteq N(f(\alpha(s, t)), \mu_1(n))$ where $\mu_1(n) = \theta(h(n)) = O(\log \log n)$. In particular, we have $\mathrm{join}(P) \subseteq N(f(T), \mu_1(n))$. By Lemma 2.6, it follows that $\mathrm{hull}(P) \subseteq$

$N(\mathrm{join}(P), \sigma(\kappa)) \subseteq N(f(T), r_1)$, where $r_1 = \sigma(\kappa) + \mu_1(n)$. This proves property (1).

To see Property (2), suppose $s, t \in T$ and $u \in \alpha(s,t)$. We can suppose that $u \notin V_0(T)$, and so $T \setminus \{u\}$ is disconnected. Thus we can write $T = T_1 \cup T_2$ with $s \in T_1$ and $t \in T_2$ and such that $u \in \alpha(x,y)$ whenever $x \in T_1$, and $y \in T_2$. (Thus $T_1 \cap T_2 = \{u\}$.) Now, let β be any path joining $f(s)$ to $f(t)$ in $\mathrm{hull}(P)$. By continuity and using Property (1), we can find some $b \in \beta$ with $d(b, f(T_1)) \leq r_1$ and $d(b, f(T_2)) \leq r_1$. Thus, we can find $x \in T_1$ and $y \in T_2$ with $d(b, f(x)) \leq r_1$ and $d(b, f(y)) \leq r_1$ (Figure 2c). By the construction of (T, f), we have that

$$\mathrm{length}\, f(\alpha(x,y)) \leq d(f(x), f(y)) + h(n) \leq 2r_1 + h(n).$$

Since $u \in \alpha(x,y)$, we have, without loss of generality, that $d(f(u), f(x)) \leq \frac{1}{2}(2r_1 + h(n))$. It follows that $d(f(u), \beta) \leq r_1 + \frac{1}{2}(2r_1 + h(n)) = 2r_1 + \frac{1}{2}h(n) = \lambda(\kappa) + \mu_2(n)$, where $\lambda(\kappa) = 2\sigma(\kappa)$ and $\mu_2(n) = 2\mu_1(n) + \frac{1}{2}h(n) = O(\log n)$. ◊

3. Tubular neighbourhoods of geodesics.

In this section, we describe a variation of Anderson's construction of convex sets. Specifically we are aiming at Propositions 3.4 and 3.5. These will be used in the proof of Theorems 4.1 and 4.2.

As remarked in the introduction, a uniform neighbourhood of a geodesic segment is always convex (by the convexity of the distance function, Lemma 1.3). The problem for us is that, given a fixed radius, there is no upper bound on the volumes of such neighbourhoods. Indeed the volume will be infinite if one of the endpoints is ideal. To deal with this problem we will need to vary the radius along the tube in such a way that convexity is preserved. Our basic building blocks will be called "joints". They are convex pieces used to connect together pieces of tube of different radii. By choosing these radii appropriately we arrange that total volumes remain bounded.

Basic observations.

Recall that X has dimension ν and curvature pinched between $-\kappa^2$ and -1. Given a closed convex set $Q \subseteq X_C$, we shall write $\pi = \pi_Q : X_C \longrightarrow Q$ for the nearest point retraction. This map is continuous (see for example [Bo2]). We shall write vol_ν for the ν-dimensional volume. For $m \geq 0$, we write $\Delta(m)$ for the m-volume of the unit sphere in euclidean $(m+1)$-space (so that $\Delta(0) = 2$).

Let us begin by recalling some basic facts about hyperbolic ν-space, $\mathbf{H}^\nu$. The volume of a uniform r-ball in $\mathbf{H}^\nu$ equals $\Delta(\nu-1)\int_0^r \sinh^{\nu-1} x dx \leq \frac{\Delta(\nu-1)}{\nu-1} e^{(\nu-1)r}$. The boundary of the r-ball is a totally umbilic surface with principal curvatures equal to $\coth r$. Suppose x, y are distinct points of $\mathbf{H}^\nu_I$. Let π be the nearest point retraction of $\mathbf{H}^\nu_C$ to $[x,y]$. Suppose $a, b \in [x,y] \cap X$, and let $l = d(a,b)$. Then, for all $r > 0$,

$$\begin{aligned}\mathrm{vol}_\nu(N([x,y], r) \cap \pi^{-1}[a,b]) &= l\Delta(\nu-2)\int_0^r \sinh^{\nu-2} x \cosh x dx \\ &= \frac{l\Delta(\nu-2)}{\nu-1}\sinh^{\nu-1} r.\end{aligned}$$

The boundary $\partial N([x,y],r)$ has one (longitudinal) principal curvature equal to $\tanh r$, and all the remaining principal curvatures (in the radial directions) equal to $\coth r$.

From these observations, we obtain bounds on the corresponding quantities in X. These may be proven by standard arguments, using Jacobi fields and the Rauch Comparison theorem (see [CE]). Thus, the volume of a uniform r-ball in X is at most $\frac{\Delta(\nu-1)}{\kappa^\nu(\nu-1)}e^{\kappa(\nu-1)r}$. Also the principal curvatures of a sphere if radius r lie between $\coth r$ and $\kappa\coth\kappa r$. Suppose that $x,y\in X_I$, and $\pi : X_C \longrightarrow [x,y]$ is the nearest point retraction. Suppose that $a,b\in[x,y]\cap X$ and $l=d(a,b)$ and $r>0$. Then,

$$\mathrm{vol}_\nu(N([x,y],r)\cap\pi^{-1}[a,b]) \leq \frac{l\Delta(\nu-2)}{\kappa^{\nu-1}(\nu-1)}\sinh^{\nu-1}\kappa r.$$

Also, the principal curvatures of $\partial N([x,y],r)$ all lie between $\tanh r$ and $\kappa\coth\kappa r$. Note that for any $a\in[x,y]\cap X$, the preimage $\pi^{-1}(a)$ is a properly embedded codimension-1 submanifold—the image of a subspace under the exponential map based at a.

The following may also be proven by comparison with hyperbolic space.

Lemma 3.1 : *Given $K>0$, there is some $l=l(K)>0$ so that the following holds. If $x,y\in X_I$ are distinct, and $\pi : X_C \longrightarrow [x,y]$ is the nearest point retraction, then for all $p,q\in X$ with $d(p,q)\leq l$, we have that $d(\pi(p),\pi(q))\leq Ke^{-r}$, where $r=\min(d(p,[x,y]),d(q,[x,y]))$.* $\diamondsuit$

A variation on Anderson's construction.

We now describe the idea behind Anderson's construction. Given $x\in X$, we write T_xX for the tangent space to X at x. We write $T^1_xX\subseteq T_xX$ for the unit tangent space at x. Given $\xi,\zeta\in T_xX$, we write $\langle\xi,\zeta\rangle$ and $|\xi|$ for the riemannian inner-product and norm respectively.

Given a smooth function $\phi : X\longrightarrow \mathbf{R}$, we write $\mathrm{grad}\,\phi$ for the gradient vector field, and write $D^2\phi$ for the second derivative of ϕ. Thus, if $\xi,\zeta\in T_xX$, we have $D^2\phi(\xi,\zeta) = D^2\phi(\zeta,\xi) = \langle\nabla_\xi\mathrm{grad}\,\phi,\zeta\rangle$. We write

$$|D^2\phi(x)| = \max\{|D^2\phi(x)(\xi,\xi)| \mid \xi\in T^1_xX\}.$$

Suppose that $Q\subseteq X$ is closed and convex. Define $\rho=\rho_Q : X\longrightarrow[0,\infty)$ by $\rho(x) = d(x,Q)$. Thus ρ is C^1, and $|\mathrm{grad}\,\rho|=1$, on $X\setminus Q$ (see [BaGS]). Let us assume that ρ is smooth on $X\setminus Q$. (This is always true in the cases that interest us, for example if Q is a single point or a bi-infinite geodesic. In fact it is enough to assume that ρ is C^2.) The boundaries of uniform neighbourhoods of Q are level sets of ρ. We aim to join together pieces of such level sets by convex surfaces, obtained from perturbations of ρ. Our goal, in this regard, is Lemma 3.3.

Now, $D^2\rho(x)(\xi,\mathrm{grad}\,\rho)=0$ for all $\xi\in T_xX$, and $D^2\rho(x)$ restricted to the subspace $(\mathrm{grad}\,\rho(x))^\perp=\ker d\rho(x)$ gives us the second fundamental form of the surface $\partial N(Q,\rho(x))$ at x. Since $\partial N(Q,\rho(x))$ is strictly convex, the second fundamental form is positive definite. It follows that if $\xi\in T^1_xX$, then $D^2\rho(x)(\xi,\xi)\geq(1-\langle\xi,\mathrm{grad}\,\rho\rangle^2)m(x)$, where $m(x)$ is the minimal principal curvature of $\partial N(Q,\rho(x))$ at x. In fact, using the Jacobi field equation, we find that always $m(x) \geq \tanh \rho(x)$. We shall only need this result here in the case where Q is a bi-infinite geodesic, which we described above.

Now, suppose that we have a map $\psi : X \longrightarrow \mathbf{R}$ which is continuous on X, and smooth on $X \setminus Q$. Suppose that $\psi(x) \leq 0$ for all $x \in Q$, and that $\langle \mathrm{grad}\, \psi, \mathrm{grad}\, \rho \rangle > 0$ everywhere on $X \setminus Q$. Given $r > 0$, let $M(r) = \psi^{-1}(-\infty, r]$. Then $M(r)$ is a connected submanifold of X with smooth boundary $\partial M(r) = \psi^{-1}(r)$, and containing Q in its interior. We may compute the second fundamental form of $\partial M(r)$ at $x \in \partial M(r)$ as $\frac{1}{|\mathrm{grad}\, \psi(x)|} D^2\psi(x)$ restricted to $\ker d\psi(x)$. Thus, $M(r)$ will be convex if $D^2\psi(x)$ is positive definite on $\ker d\psi(x)$.

We shall take ψ to be a perturbation of ρ. Thus $\psi = \rho - \epsilon\phi$ where $\epsilon \geq 0$, and $\phi : X \longrightarrow [0,1]$ is smooth, and satisfies $|\mathrm{grad}\, \phi| \leq c_1$ and $|D^2\phi| \leq c_2$ where c_1 and c_2 are constants. If $\epsilon < 1/c_1$, then $\langle \mathrm{grad}\, \psi, \mathrm{grad}\, \rho \rangle \geq 1 - c_1\epsilon > 0$ on $X \setminus Q$. Suppose $r > 0$, and $x \in \partial M(r)$. Then $\rho(x) \geq \psi(x) = r$. If $\xi \in \ker d\psi(x) \cap T^1_x X$, then $|\langle \xi, \mathrm{grad}\, \rho \rangle| \leq c_1\epsilon < 1$, and so $D^2\rho(x)(\xi,\xi) \geq (1 - (c_1\epsilon)^2)m(x)$. Thus $D^2\psi(x)(\xi,\xi) \geq (1-(c_1\epsilon)^2)m(x) - c_2\epsilon$. Therefore, given that $m(x) \geq \tanh \rho(x) \geq \tanh r$, the manifold $M(r)$ will be convex provided that $c_2\epsilon \leq (1 - (c_1\epsilon)^2) \tanh r$. Note that

$$N(Q,r) \subseteq M(r) \subseteq N(Q, r+\epsilon),$$

and that

$$\partial M(r) \cap \phi^{-1}(0) = \partial N(Q,r) \cap \phi^{-1}(0)$$
$$\partial M(r) \cap \phi^{-1}(1) = \partial N(Q,r+\epsilon) \cap \phi^{-1}(1).$$

The following lemma gives us a suitable perturbation, ϕ.

Lemma 3.2 : *Given any $l > 0$, there exist constants $c_1, c_2, \eta > 0$, depending on l and κ, such that for all $p \in X$, there is a smooth map $\phi = \phi_p : X \longrightarrow [0,1]$ such that $|\mathrm{grad}\, \phi| \leq c_1$ and $|D^2\phi| \leq c_2$ everywhere, and such that $\phi(x) = 0$ if $d(x,p) \leq \eta$ and $\phi(x) = 1$ if $d(x,p) \geq l - \eta$.*

Proof : Let ρ_p be defined by $\rho_p(x) = d(p,x)$. Thus ρ_p is smooth on $X \setminus \{p\}$, and we know from the above discussion that $|D^2\rho_p(x)| \leq \kappa \coth \kappa\rho_p(x)$. Choose any $\eta < l/2$, and some smooth function $g : [0,\infty) \longrightarrow [0,1]$ such that $g|[0,\eta] \equiv 0$, $g|[l-\eta,\infty) \equiv 1$ and such that, for all $r \geq 0$, $|g'(r)| \leq c_1 \tanh \kappa r$ and $|g''(r)| \leq c_3$, where c_1 and c_3 depend only on κ and l. Now let $\phi = \phi_p = g \circ \rho_p$. Then $|\mathrm{grad}\, \phi(x)| \leq |g'(\rho_p(x))| \leq c_1$ and $|D^2\phi(x)| \leq |g''(\rho_p(x))| + |g'(\rho_p(x))||D^2\rho_p(x)| \leq c_3 + \kappa c_1 = c_2$. ♢

Let's return to our discussion with $Q \subseteq X$ closed and convex, and with $\rho(x) = \rho(Q,x)$ smooth on $X \setminus Q$. Given $r > 0$, we choose $\epsilon \geq 0$ so that $c_2 \leq (1 - (c_1\epsilon)^2) \tanh r$. Given $p \in X$, write $\psi_p = \rho - \epsilon\phi_p$. We see that $M_p(r,\epsilon) = \psi_p^{-1}(-\infty, r]$ is convex. Suppose we have $A \subseteq \partial N(Q, r+\epsilon)$ and $B \subseteq \partial N(Q,r)$ both closed, and such that $d(A,B) \geq l$. Set $M_B(r,\epsilon) = \bigcap_{p \in B} M_p(r,\epsilon)$. Then $M_B(r,\epsilon)$ convex, and

$$N(Q,r) \subseteq M_B(\rho,\epsilon) \subseteq N(Q, r+\epsilon),$$

and

$$\partial M_B(r,\epsilon) \cap N(A,\eta) = \partial N(Q,r+\epsilon) \cap N(A,\eta)$$
$$\partial M_B(r,\epsilon) \cap N(B,\eta) = \partial N(Q,r) \cap N(B,\eta).$$

The construction given in [A] takes Q to be a single point. Here, we take Q to be a bi-infinite geodesic.

Lemma 3.3 : *For all $k > 0$ there is some $\delta = \delta(\kappa, k)$ such that the following holds (Figure 3a).*

Suppose $x, y \in X_I$ are distinct. Let $\pi : X_C \longrightarrow [x, y]$ be the nearest point retraction. Suppose $r > 0$, and that $a, b \in [x, y] \cap X$ with $b \in [a, y]$ and $d(a, b) \geq ke^{-r}$. Suppose that $r \leq R \leq r + \delta \tanh r$. Then there is a convex set $M \subseteq X$ such that

$$N([x, y], r) \subseteq M \subseteq N([x, y], R),$$

and

$$\partial M \cap U = \partial N([x, y], R) \cap U$$
$$\partial M \cap V = \partial N([x, y], r) \cap V,$$

where U, V are, respectively, neighbourhoods in X of the sets $\partial N([x, y], R) \cap \pi^{-1}[x, a]$ and $\partial N([x, y], r) \cap \pi^{-1}[y, b]$.

Proof : Given $k > 0$, let $l = l(k)$ be the constant of Lemma 3.1. Given this, let c_1 and c_2 be the constants of Lemma 3.2. Choose $\delta > 0$ so that $c_2\delta \leq 1 - (c_1\delta)^2$. Thus, δ depends only on k and κ. Now suppose that x, y, a, b, r, R are as in the hypotheses. Let $\epsilon = R - r \leq \delta \tanh r \leq \delta$. Thus $c_2\epsilon \leq (1 - (c_1\epsilon)^2) \tanh r$. Set $Q = [x, y]$ and let $A = \partial N(Q, R) \cap \pi^{-1}[x, a]$ and $B = \partial N(Q, r) \cap \pi^{-1}[y, b]$. Thus, by Lemma 3.1, we have $d(A, B) \geq l$. Set $M = M_B(r, \epsilon)$, $U = N(A, \eta)$ and $V = N(B, \eta)$, where η comes from Lemma 3.2. The result follows from the above discussion. $\diamondsuit$

Given x, y, a, b, r, R as in the hypotheses of Lemma 3.3, we shall write $J(a, b, R, r) = M \cap \pi^{-1}[a, b]$, where M is the convex set thus constructed. We may think of $J(a, b, R, r)$ as a "joint" used to connect two tubes of unequal radii. Write $\partial_0 J(a, b, R, r) = \partial M \cap \pi^{-1}[a, b]$. Since $J(a, b, R, r) \subseteq N([x, y], R) \cap \pi^{-1}[a, b]$, we have

$$\mathrm{vol}_\nu J(a, b, R, r) \leq d(a, b) \frac{\Delta(\nu - 2)}{\kappa^{\nu-1}(\nu - 1)} \sinh^{\nu-1} \kappa R.$$

(Recall that $\Delta(\nu - 2)$ is the volume of $(\nu - 2)$-sphere.)

Application to the construction of long thin tubes.

Suppose that $x, y \in X_I$, and $p \in [x, y] \cap X$. Let $H = \pi^{-1}[x, p]$. (Thus H is the image of a half-space under the exponential map based at p.) We construct a convex set containing $H \cup [x, y]$ by stringing together a bi-infinite sequence of joints as follows.

For convenience, set $k = 1$, as let $\delta = \delta(\kappa, 1)$ be the constant given by Lemma 3.3. Let $c = \tanh 1$, and $\eta = c\delta$. Let $L = 1/(1 - e^{-\eta})$. Note that $\tanh r \geq cr$ for $r \in [0, 1]$ and $\tanh r \geq c$ for $r \in [1, \infty)$.

We form a bi-infinite sequence $(a_i)_{i=-\infty}^{\infty}$ of points of $[p, y]$, with $a_{i+1} \in [a_i, y]$ for all i, as follows. We set $a_0 \in [p, y]$ to be the point such that $d(a_0, p) = L$, and demand, for all $i \geq 0$, that $d(a_i, a_{i+1}) = 1$ and $d(a_{-(i+1)}, a_{-i}) = e^{-\eta i}$. Note that as $i \to \infty$, we have $a_i \to y$, and, since $L = \sum_{i=0}^{\infty} d(a_{-i}, a_{-(i+1)})$, we have that $a_{-i} \to p$.

For $i \geq 0$, set $r_i = (1+\eta)^{-i}$ and $r_{-i} = 1+\eta i$. Thus, $r_{i+1} < r_i$ for all i. If $i \geq 0$, then $r_i - r_{i+1} = \eta(1+\eta)^{-(i+1)} = \eta r_{i+1} = \delta(c r_{i+1}) \leq \delta \tanh r_{i+1}$ and $1 = d(a_{i+1}, a_i) \geq e^{-r_{i+1}}$. Thus, by Lemma 3.3, we can construct the joint $J_i = J(a_i, a_{i+1}, r_i, r_{i+1})$. We also have that $r_{-(i+1)} - r_{-i} = \eta = \delta c \leq \delta \tanh r_{-i}$ and $d(a_{-(i+1)}, a_{-i}) = e^{-\eta i} \geq e^{-(1+\eta i)} = e^{-r_{-i}}$. Again, by Lemma 3.3, we construct $J_{-(i+1)} = J(a_{-(i+1)}, a_{-i}, r_{-(i+1)}, r_{-i})$.

Let $J = H \cup \bigcup_{i=-\infty}^{\infty} J_i$. (Figure 3b.) Thus, J is connected, with boundary $\partial J = \bigcup_{i=-\infty}^{\infty} \partial_0 J_i$. We see that, for all i, the boundary ∂J agrees with $\partial N([x,y], r_i)$ on some neighbourhood, U, of $\partial N([x,y], r_i) \cap \pi^{-1}(a_i)$, i.e. $\partial J \cap U = (\partial_0 J_i \cup \partial_0 J_{i-1}) \cap U = \partial N([x,y], r_i) \cap U$. Since convexity for a connected set is a local property, we see that J is convex. Clearly $H \cup [x,y] \subseteq J$.

For $i \geq 0$, we have

$$\mathrm{vol}_\nu J_i \leq \frac{\Delta(\nu-2)}{\kappa^{\nu-1}(\nu-1)} \sinh^{\nu-1} \kappa r_i.$$

Now, $r_i = (1+\eta)^{-i} \leq 1$, and so $\sinh^{\nu-1} \kappa r_i \leq r_i^{\nu-1} \sinh^{\nu-1} \kappa = (1+\eta)^{-i(\nu-1)} \sinh^{\nu-1} \kappa$. Thus

$$\mathrm{vol}_\nu J_i \leq \frac{\Delta(\nu-2)}{\nu-1} \left(\frac{\sinh \kappa}{\kappa}\right)^{\nu-1} (1+\eta)^{-(\nu-1)i},$$

and so

$$\begin{aligned} \mathrm{vol}_\nu(J \cap \pi^{-1}[a_0, y]) &= \sum_{i=0} \mathrm{vol}_\nu J_i \\ &\leq \frac{\Delta(\nu-2)}{\nu-1} \left(\frac{\sinh \kappa}{\kappa}\right)^{\nu-1} \left(\frac{1}{1-(1+\eta)^{-(\nu-1)}}\right) \end{aligned}$$

which is finite, and a function only of ν and κ.

Similarly, for any fixed $i_0 \geq 0$, we have that $\mathrm{vol}_\nu(J \cap \pi^{-1}[a_{-i_0}, a_0]) = \sum_{i=1}^{i_0} \mathrm{vol}_\nu J_{-i}$, which is bounded by some function of ν, κ and i_0. Note that given any $q \in [p,y] \setminus \{p\}$, we can find some i_0, such that $q \in [p, a_{i_0}]$. This i_0 depends only on κ and $d(p,q)$. We conclude:

Proposition 3.4 : *(Figure 3c.) Given any $\zeta > 0$, there is some constant $K(\nu, \kappa, \zeta)$ such that the following holds. Suppose that $x, y \in X_I$ are distinct points, and that $p, q \in [x,y] \cap X$ with $q \in [p,y]$ and $d(p,q) \geq \zeta$. Let $\pi : X_C \longrightarrow [x,y]$ be the nearest point retraction, and let $H = \pi^{-1}[x,p]$ and $H_0 = \pi^{-1}[q,y]$. Then,*

$$\mathrm{vol}_\nu(H_0 \cap \mathrm{hull}(H \cup \{y\})) \leq K(\nu, \kappa, \zeta).$$

$\diamondsuit$

In fact, we see that $K(\nu, \kappa, \zeta) \to 0$ as $\zeta \to \infty$.

By a similar argument, we arrive also at the following:

Proposition 3.5 : *(Figure 3d.) Given any $\zeta > 0$, there is some constant $K' = K'(\nu, \kappa, \zeta)$ such that the following holds. Suppose that $x, x' \in X_I$ and $p, q, p', q' \in [x, x'] \cap X$ are points occurring in the order $xpqq'p'y$ along $[x, x']$, so that $d(p, q) \geq \zeta$ and $d(p', q') \geq \zeta$. Let $H = \pi^{-1}[x, p]$, $H' = \pi^{-1}[x', p']$ and $H_0 = \pi^{-1}[q, q']$. Then,*

$$\mathrm{vol}_\nu(H_0 \cap \mathrm{hull}(H \cup H')) \leq K'(\nu, \kappa, \zeta).$$

$\diamondsuit$

For notational convenience, we set $K'(\nu, \kappa, \zeta) = K(\nu, \kappa, \zeta)$. (Thus, Proposition 3.4 may be regarded as a corollary of Proposition 3.5.)

4. Boundedness and continuity of volume.

The first result of this section is the fact that convex hull of of finite sets have finite, indeed bounded volume:

Theorem 4.1 : *Given $n \in \mathbf{N}$, there is some constant $C(\nu, \kappa, n)$ such that if $P \subseteq X_C$ is a set of n points, then $\mathrm{vol}_\nu \mathrm{hull}(P) \leq C(\nu, \kappa, n)$. Moreover, for fixed ν and κ, $C(\nu, \kappa, n)$ is bounded by some polynomial in n.*

We also note that, for fixed ν and n, $C(\nu, \kappa, n)$ can be assumed continuous in κ. As far as I know, it may be possible to remove dependence on κ altogether, though I suspect not.

The second result of this section shows how these volumes vary continuously. Let $P = \{p_1, \ldots, p_n\}$. Thus P, and hence $\mathrm{hull}(P)$ vary continuously in $(p_1, \ldots, p_n) \in X_C^n$. In proving Theorem 4.1, we will effectively show that most of the volume of $\mathrm{hull}(P)$ lies inside a certain compact convex set. Usually this set can be chosen to be locally constant. The only problem arises if two vertices p_i and p_j converge on the same ideal point. Let Λ be the set of $(p_1, \ldots, p_n) \in X_C^n$ such that for two distinct $i, j \in \{1, \ldots, n\}$, we have $p_i = p_j \in X_I$.

Theorem 4.2 : *The map from X_C^n to $[0, \infty)$ which sends $(p_1, \ldots, p_n)$ to $\mathrm{vol}_\nu \mathrm{hull}\{p_1, \ldots, p_n\}$ is continuous on $X_C^n \setminus \Lambda$.*

Proof of boundedness of volume.

The ingredients we use for Theorem 4.1 are the existence of a spanning tree (T, f) with the property that $\mathrm{length}\, f(\alpha(s, t)) \leq d(f(s), f(t)) + h(n)$ for all $s, t \in T$ (Lemma 2.5), together with the fact that for such a tree we have $\mathrm{hull}(P) \subseteq N(f(T), r_1(\kappa, n))$ (Theorem 2.1). If we want the polynomial bound, we need that $h(n) = O(\log n)$ and that $r_1(\kappa, n) = \lambda(\kappa) + \mu_1(n)$ where $\mu_1(n) = O(\log n)$. (We know that $\mu_1(n) = O(\log \log n)$.) I suspect that, in fact, $C(\nu, \kappa, n)$ is always bounded by a linear function of n.

Given such a spanning tree, (T, f), we write $V(T) = V_0(T) \sqcup V_1(T)$, where $V_0(T)$ is the set of extremal vertices, and $V_1(T)$ is the set of internal vertices. Thus, $f(V_0(T)) = P$. We write $E_0(T)$ for the set of *extremal* edges, i.e. those incident on some vertex of $V_0(T)$. We write $E_1(T) = E(T) \setminus E_0(T)$ for the set of *internal* edges. We have $|V_0(T)| = |E_0(T)| = n$ and $|V_1(T)| \leq n-2$ and $E_1(T) \leq n-3$.

The proof of Theorem 4.1 is based on the observation (Lemma 4.5) that hull(P) lies inside a certain neighbourhood of $f(T)$ which consists of uniform balls about each internal vertex, together with tubes along each of the edges. The volumes of these tubes are bounded by the results of Section 3. The balls about the vertices can be taken to have radii $O(\log n)$ which gives us our polynomial bound on $C(\nu, \kappa, n)$.

We assume that $n \geq 3$. Suppose that $e \in E_1(T)$ with endpoints $v_0, v_1 \in V_1(T)$. Any point in the interior of e divides T into two components, T_1 and T_2, with $v_i \in T_i$. Let $W_i = T_i \cap V_0(T)$ and $P_i = f(W_i)$. Thus $P = P_0 \sqcup P_1$. Let $\pi : X_C \longrightarrow f(e)$ be the nearest point retraction to $f(e) = [f(v_0), f(v_1)]$.

Lemma 4.3 : *If $e \in E_1(T)$, and v_0, v_1, P_0, P_1, π are as above, then $d(f(v_i), \pi(p)) \leq h(n)$ for all $p \in P_i$, and $i = 0, 1$.*

Proof : Let $p = f(w)$ where $w \in W_i$. Let $\pi(p) = f(u)$ where $u \in e$. Suppose first, that $p \in X$. By the definition of π, we have $d(f(w), f(u)) = d(p, \pi(p)) \leq d(p, f(v_i)) \leq$ length $f(\alpha(w, v_i))$, By the construction of (T, f) (Lemma 1.5), we have length $f(\alpha(w, u)) \leq d(f(w), f(u)) + h(n)$. Thus $d(f(v_i), \pi(p)) = d(f(v_i), f(u)) =$ length $f(\alpha(w, u)) -$ length $f(\alpha(w, v_i))$ $\leq h(n)$. The case where $p = f(w) \in X_I$ can be dealt with by taking a sequence of points $w_j \in T \setminus \{w\}$ tending to w. ♢

By a similar argument, we have:

Lemma 4.4 : *Suppose $e \in E_0(T)$ is incident on $v \in V_1(T)$ and $w \in V_0(T)$. Then $d(f(v), \pi(p)) \leq h(n)$ for all $p \in f(V_0(T) \setminus \{w\})$.* ♢

Now suppose $e \in E(T)$. For any $\zeta \geq 0$, we define $S(e, \zeta)$ to be a (possibly empty) closed segment of $f(e)$ as follows. If $e \in E_1(T)$, incident on $v, w \in V_1(T)$, let $S(e, \zeta) = \{x \in f(e) \mid d(x, \{f(v), f(w)\}) \geq h(n) + \zeta\}$. Thus, by Lemma 4.3, $d(S(e, \zeta), \pi(p)) \geq \zeta$ for all $p \in P$. If $e \in E_0(T)$, incident on $v \in V_1(T)$ and $w \in V_0(T)$, let $S(e, \zeta) = \{x \in f(e) \mid d(x, f(v)) \geq h(n) + \zeta\}$. Thus, by Lemma 4.4, $d(S(e, \zeta), \pi(p)) \geq \zeta$ for all $p \in P \setminus \{f(w)\}$. In either case, set $G(e, \zeta) = \text{hull}(P) \cap \pi^{-1}(S(e, \zeta))$. Applying Propositions 3.4 and 3.5, we find that

$$\text{vol}_\nu G(e, \zeta) \leq K(\nu, \kappa, \zeta)$$

for all $e \in E(T)$.

For the proof of Theorem 4.2, we will need to note that, given any $c > 0$, we can assume that $G(e, \zeta)$ lies inside a c-neighbourhood of $f(e)$, provided ζ is sufficiently large depending on c and κ.

We now come to the result that confines the convex hull to a union of balls and thin tubes. Let $R = R(\kappa, n) = r_1(\kappa, n) + h(n) = \lambda(\kappa) + \mu_1(n) + h(n) = \lambda(\kappa) + O(\log n)$.

Lemma 4.5 : *For any $\zeta > 0$,*

$$\text{hull}(P) \subseteq \bigcup_{v \in V_1(T)} N(f(v), R + \zeta) \cup \bigcup_{e \in E(T)} G(e, \zeta).$$

Proof : Suppose $x \in \mathrm{hull}(P)$. Let $y \in f(T)$ be a nearest point in $f(T)$ to x. We thus have $d(x, y) \leq r_1$. Now, $y \in e$ for some $e \in E(T)$. If $y \in S(e, \zeta)$, then $x \in G(e, \zeta)$. If $y \in f(e) \setminus S(e, \zeta)$, then, by definition of $S(e, \zeta)$, there is some $v \in V_1(T)$, incident on e, so that $d(f(v), y) \leq h(n) + \zeta$. Thus $x \in N(f(v), R + \zeta)$. ♢

Proof of Theorem 4.1 : For the proof, we take $\zeta = 1$.

In Section 3, we gave an upper bound for the volume of a uniform ball. Thus,

$$\mathrm{vol}_\nu N(f(v), R + 1) \leq \frac{\Delta(\nu - 1)}{\kappa^\nu(\nu - 1)} e^{\kappa(\nu-1)(R(\kappa,n)+1)} = B(\nu, \kappa, n).$$

From the form of $R(\kappa, n)$, we see that, for fixed κ and ν, $B(\nu, \kappa, n)$ is bounded by some polynomial in n. By Lemma 4.5, we have that

$$\begin{aligned} \mathrm{vol}_\nu \, \mathrm{hull}(P) &\leq |V_1(T)| B(\nu, \kappa, n) + |E(T)| K(\nu, \kappa, 1) \\ &\leq (n - 2) B(\nu, \kappa, n) + (2n - 3) K(\nu, \kappa, 1) \\ &= C(\nu, \kappa, n). \end{aligned}$$

For fixed ν, κ, we see that $C(\nu, \kappa, n)$ is bounded by a polynomial in n. This concludes the proof of Theorem 4.1. ♢

Proof of continuity of volume.

To prove Theorem 4.2, we need to observe that the boundary, ∂Q of a convex subset $Q \subseteq X$ has zero Lebesgue measure. (Note, for example, that the Lebesgue density of Q at any point of ∂Q is at most $\frac{1}{2}$.) Thus, if Q is compact, we can choose $\eta > 0$, to make $\mathrm{vol}_\nu N(\partial Q, \eta)$ arbitrarily small.

We shall also need the following lemma, which will confine most of the volume of a convex hull to a certain bounded set.

Lemma 4.6 : *Suppose $A_1, \ldots, A_n$ are closed subsets of X_C satisfying $X_I \cap A_i \cap A_j = \emptyset$ if $i \neq j$. Then there is a compact convex set $M \subseteq X$ with the following property. Suppose $P = \{p_1, \ldots, p_n\}$, with $p_i \in A_i$ for all i, and suppose (T, f) is a spanning tree for P satisfying the same criterion as that of Lemma 2.5, (namely* $\mathrm{length}\, f(\alpha(s, t)) \leq d(f(s), f(t)) + h(n)$ *for all $s, t \in T$). Then, $f(v) \in M$ for each internal vertex $v \in V_1(T)$.*

Proof : Suppose v separates the three extremal vertices $v_i, v_j, v_k \in V_0(T)$, so that $p_\alpha = f(v_\alpha) \in A_\alpha$ for $\alpha \in \{i, j, k\}$. As in the proof of Lemma 1.5, we see that $f(v) \in N([p_i, p_j], \rho) \cap N([p_j, p_k], \rho) \cap N([p_k, p_i], \rho)$ for some fixed $\rho > 0$. Now this intersection is bounded. Moreover, as p_i, p_j, and p_k vary in A_i, A_j and A_k respectively, these intersection are all contained in some bounded subset, $D(i, j, k)$ of X. (This is an elementary consequence of Gromov hyperbolicity of X—see [Gr] or [Bo1].) Now choose some compact ball M, which contains the sets $D(i, j, k)$ for all distinct $i, j, k \in \{1, \ldots, n\}$. ♢

Proof of Theorem 4.2 : Suppose $(p_1, \dots, p_n) \in X_C^n \setminus \Lambda$, and $\epsilon > 0$. Choose $\zeta > 0$ so that $K(\nu, \kappa, \zeta) < \epsilon/4n$, where $K(\nu, \kappa, \zeta)$ is the constant in Proposition 3.4. Choose neighbourhoods A_i of p_i so that $X_I \cap A_i \cap A_j = \emptyset$ if $i \neq j$. Let $M \subseteq X$ be the compact convex set given by Lemma 4.6, and let $M' = N(M, R + \zeta)$ where $R = R(\kappa, n)$ is the constant of Lemma 4.5. Choose $\eta > 0$ so that $\mathrm{vol}_\nu N(\partial(M' \cap \mathrm{hull}(P)), \eta) < \epsilon/2$. By continuity in the Hausdorff topology (Theorem 1.5), we can assume (shrinking the A_i if necessary) that if $q_i \in A_i$ for $i = 1, \dots, n$, then $\mathrm{hd}(d)(M' \cap \mathrm{hull}(P), M' \cap \mathrm{hull}(Q)) \leq \eta$, where $Q = \{q_i, \dots, q_n\}$. (Note that Theorem 1.5, refers to a different metric on X, so we need to observe that any two metrics induce the same uniformity on the compact set M'.) So, by Lemma 1.7, $\mathrm{hd}(d)(\partial(M' \cap \mathrm{hull}(P)), \partial(M' \cap \mathrm{hull}(Q))) \leq \eta$, and so $|\mathrm{vol}_\nu(M' \cap \mathrm{hull}(P)) - \mathrm{vol}_\nu(M' \cap \mathrm{hull}(Q))| \leq \epsilon/2$.

Now, let (T, f) be a spanning tree for P. By Lemma 4.5, we have $\mathrm{hull}(P) \subseteq \bigcup_{v \in V_1(T)} N(f(v), R + \zeta) \cup \bigcup_{e \in E(T)} G(e, \zeta)$. By Lemma 4.6, if $v \in V_1(T)$, then $f(v) \in M$, so $N(f(v), R + \zeta) \subseteq M'$. If e is an internal edge of T, then it follows that $f(e) \subseteq M$, so, from the remarks following Lemma 4.4, we can assume that $G(e, \zeta) \subseteq M'$. Thus, $\mathrm{hull}(P) \setminus M' \subseteq \bigcup_{e \in E_0(T)} G(e, \zeta)$, where $E_0(T)$ is the set of extremal edges of T. So, $\mathrm{vol}_\nu(\mathrm{hull}(P) \setminus M') \leq nK(\nu, \kappa, \zeta) \leq n(\epsilon/4n) = \epsilon/4$.

Now exactly the same argument shows that $\mathrm{vol}_\nu(\mathrm{hull}(Q) \setminus M') \leq \epsilon/4$. Putting these facts together, we see that $|\mathrm{vol}_\nu \mathrm{hull}(P) - \mathrm{vol}_\nu \mathrm{hull}(Q)| \leq \epsilon$. ♢

5. Appendix.

In this appendix, we give a brief discussion of the case of constant negative curvature. In this case, we can use a different technique to obtain a linear upper bound on volumes.

Let $\mathbf{H}^\nu$ be ν-dimensional hyperbolic space (of constant curvature -1). We can define a (closed, convex, finite volume) polytope in $\mathbf{H}^\nu_C$ as the convex hull of a finite set of points. Given such a polytope, Π, there is a unique minimal such finite set, which we refer to as the set of *vertices*, $\mathrm{vert}(\Pi)$, of Π. Thus $\mathrm{vert}(\Pi)$ is the union of $\Pi \cap \mathbf{H}^\nu_I$ and the set of extreme points of $\Pi \cap \mathbf{H}^\nu$. We shall write $f_i(\Pi)$ for the number of i-dimensional faces of Π.

Theorem 5.1 : *For all ν, there is a constant $c(\nu) > 0$ such that if $\Pi \subseteq \mathbf{H}^\nu_C$ is a polytope with n vertices, then $\mathrm{vol}_\nu \Pi \leq nc(\nu)$.*

Before beginning the proof, we make some general observations. We shall assume that all polytopes have non-empty interior.

Suppose $\Sigma \subseteq \mathbf{H}^\nu_C$ is a ν-simplex (i.e. $f_0(\Sigma) = \nu + 1$). Then, it's not hard to see that the volume of Σ is bounded in terms of the dimension, ν. In fact it's known [HM] that $\mathrm{vol}_\nu \Sigma$ is maximised precisely when Σ is a regular ideal simplex, Σ^ν_0. Such a simplex Σ^ν_0 is unique up to isometry.

Now suppose that $\Pi \subseteq \mathbf{H}^\nu_C$ is a polytope with $f_0(\Pi) = n$, and with non-empty interior, $\mathrm{int}\, \Pi$. By subdividing, we can assume that all the codimension-1 faces of Π are simplices. By choosing an arbitrary point $v_0 \in \mathrm{int}\, \Pi$, and coning on v_0, we obtain a subdivision of Π

into $f_{\nu-1}(\Pi)$ simplices of dimension ν. Obviously, $f_{\nu-1}(\Pi) \leq \binom{n}{\nu}$ and so this immediately gives us an upper bound for $\mathrm{vol}_\nu \Pi$ which is polynomial in n. In fact, the solution of the Upper Bound Conjecture (see [MS]) gives a sharp upper bound for $f_{\nu-1}(\Pi)$ which is $O(n^{[\nu/2]})$ where $[\nu/2]$ is the integer part of $\nu/2$. Thus for $\nu \leq 3$, we get a linear bound. (This also follows directly from Euler's formula.) The 3-dimensional case is discussed in [SlTT]. In higher dimensions, we need to do some more geometry.

Suppose $\Sigma \subseteq \mathbf{H}^\nu_C$ is a ν-simplex. Let $E(\Sigma)$ be the set of edges of Σ, i.e. closed 1-dimensional faces. Suppose $x \in \Sigma$ is an interior point of some $e \in E(\Sigma)$. Let $\Omega(\Sigma, x)$ be the set of unit normal vectors to e based at x which point into the interior of Σ. Let $\Theta(\Sigma, e)$ be the $(\nu - 2)$-dimensional spherical Lebesgue measure of $\Omega(\Sigma, x)$. This is the "solid angle" of Σ in e. It is independent of the choice of x. (Thus if $\nu = 3$, then $\Theta(\Sigma, e)$ is the dihedral angle.) Given $v \in \mathrm{vert}(\Sigma)$, let $E(\Sigma, v) \subseteq E(\Sigma)$ be the set of edges incident on v (so that $|E(\Sigma, v)| = \nu$). Let $\Phi(\Sigma, v) = \sum_{e\in E(\Sigma,v)} \Theta(\Sigma, e)$.

Lemma 5.2 : *Given ν, there is some $k(\nu) > 0$ such that if $\Sigma \subseteq \mathbf{H}^\nu_C$ is a ν-simplex, and $v \in \mathrm{vert}(\Sigma)$, then $\mathrm{vol}_\nu \Sigma \leq k(\nu)\Phi(\Sigma, v)$.*

Proof : Since $\bigcup E(\Sigma, v) \subseteq \Sigma$ is starlike, and $\Sigma = \mathrm{hull}(\bigcup E(\Sigma, v))$, we have some universal constant $r > 0$ such that

$$\Sigma \subseteq N(\bigcup E(\Sigma, v), r) = \bigcup_{e\in E(\Sigma,v)} N(e, r).$$

(Note that a starlike set is quasiconvex—for example, since any two points are joined by a path consisting of at most two geodesic segments.)

Fix, for the moment, some $e \in E(\Sigma, v)$, and x in the interior of e. Any unit vector $\xi \in \Omega(\Sigma, x)$, together with e, determines a 2-plane σ which intersects Σ in a hyperbolic triangle. Given $u > 0$, let $l(\xi, u)$ be the length of the arc $\sigma \cap \Sigma \cap N(e, u)$. We may obtain the tolal volume of $\Sigma \cap N(e, r)$ by integrating the quantity $l(\xi, u) \sinh^{\nu-2} u$ first in u from 0 to r, and then with respect to spherical Lebesgue measure, as ξ varies over $\Omega(\Sigma, x)$. Now, we may bound $\int_0^r l(\xi, u) \sinh^{\nu-2} u du$ independently of ξ as follows. Note that $l(\xi, u) \leq L(u)$, where $L(u)$ is the length of the boundary of the u-neighbourhood of an edge in a hyperbolic ideal triangle Σ^2_0. Thus $\int_0^\infty L(u)du = \mathrm{vol}_2 \Sigma^2_0 = \pi < \infty$, and so $k(\nu) = \int_0^r L(u) \sinh^{\nu-2} u du$ is finite. We deduce that

$$\mathrm{vol}_\nu(\Sigma \cap N(e, r)) \leq k(\nu)\Theta(\Sigma, e).$$

Finally, summing over all $e \in E(\Sigma, v)$, we obtain

$$\mathrm{vol}_\nu \Sigma \leq k(\nu)\Phi(\Sigma, v).$$

$\diamondsuit$

Proof of Theorem 5.1 : Let Π be a polytope with n vertices and non-empty interior. We subdivide Π in into a set $\mathcal{S}$ of ν-simplices, by coning over an arbitrary $v_0 \in \operatorname{int} \Pi$, as described above. In this triangulation, there are precisely n edges (1-cells) incident on v_0. If e is such an edge, then

$$\sum_{\Sigma \in \mathcal{S}(e)} \Theta(\Sigma, e) = \Delta(\nu - 2),$$

where $\mathcal{S}(e) \subseteq \mathcal{S}$ is the subset of those simplices which have e as an edge. Summing over all edges incident on v_0, we obtain

$$\sum_{\Sigma \in \mathcal{S}} \Phi(\Sigma, v_0) = n\Delta(\nu - 2).$$

Applying Lemma 5.2, we obtain

$$\mathrm{vol}_\nu \Pi \leq nc(\nu)$$

where $c(\nu) = k(\nu)\Delta(\nu - 2)$. ♢

Certainly, we cannot do better than a linear bound. I don't know what is the best multiplicative constant in dimensions greater than 3. In dimension 3, the best such constant is twice the volume of a regular ideal 3-simplex ($2\mathrm{vol}_3 \Sigma_0^3 = 2 \times 1{\cdot}01494\ldots$). In other words, the maximal volume of a polytope with n vertices, divided by $2n\mathrm{vol}_3 \Sigma_0^3$, tends to 1 as n tends to ∞.

Note that, in dimension $\nu = 2$, the same method of subdivision works with variable curvature, since convex hulls are always polygonal. Here, the lower curvature bound is irrelevant, and we obtain a best multiplicative constant of $\mathrm{vol}_2 \Sigma_0^2 = \pi$.

References.

[A] M.T.Anderson, *The Dirichlet problem at infinity for manifolds of negative curvature* : J. Diff. Geom. **18** (1983) 701–721.

[BaGS] W.Ballmann, M.Gromov, V.Schroeder, *Manifolds of non-positive curvature* : Progress in Maths. 61, Birkäuser (1985).

[Be] M.Bestvina, *Degenerations of the hyperbolic space* : Duke Math. J. **56** (1988) 143–161.

[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 with variable negative curvature* : preprint, I.H.E.S. (1990).

[Bu] H.Busemann, *The geometry of geodesics* : Pure and Applied Maths. 6, Academic Press (1955).

[CE] J.Cheeger, D.G.Ebin, *Comparison theorems in riemannian geometry* : North Holland (1975).

[F] W.J.Floyd, *Group completions and limit sets of Kleinian groups* : Invent. Math. **57** (1980) 205–218.

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

[Gr] M.Gromov, *Hyperbolic groups,* in "Essays in group theory" : ed. S.M.Gersten, M.S.R.I. Publications No. 8 (1988).

[HM] U.Haagerup, H.J.Munkholm, *Simplices of maximal volume in hyperbolic n-space* : Acta. Math. **147** (1982) 1–11.

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

[MS] P.McMullen, G.C.Shephard, *Convex polytopes and the upper bound conjecture* : L.M. S. Lecture Notes 3, Cambridge University Press (1971).

[P] F.Paulin, *Dégénérescences algébrique de représentations hyperboliques* : preprint, E.N.S. Lyon (1990).

[SlTT] D.D.Sleator, R.E.Tarjan, W.P.Thurston, *Rotation distance, triangulations, and hyperbolic geometry* : preprint.

[Sp] M.Spivak, *A comprehensive introduction to differential geometry* : Publish or Perish (1979).