Minkowskian subspaces of non-positively curved metric spaces

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

0. Introduction.

In his article [9], Busemann developed a theory of non-positive curvature for path-metric spaces, based on a simple axiom of convexity of the distance function. This work has received renewed attention recently, particularly in the light of new ideas developed by Gromov in his paper on hyperbolic groups [12]. In this paper Gromov introduces a notion of (strict) negative curvature applicable equally to finitely generated groups and path-metric spaces. A number of results relate this "hyperbolicity" criterion to that of non-positive curvature. To describe non-positive curvature, Gromov adopts a comparison axiom he calls CAT(0). This axiom has its origins in earlier work of Aleksandrov and Toponogov. Much of the theory of Hadamard manifolds, as developed in [3], for example can be adapted to this context—see [6, 8]. The CAT(0) axiom is stronger than Busemann's convexity condition, as is exemplified by the fact that all Banach spaces are Busemann spaces, whereas a Banach space is CAT(0) if and only if it is a Hilbert space. Despite this many of the results about CAT(0) spaces generalise to this simpler and more general context. The main focus of the paper will be a proof of the following result:

Theorem 1 : *Suppose X is a Busemann space which admits a discrete cocompact isometric action of a group Γ. Then either Γ is hyperbolic (in the sense of Gromov), or else X contains a totally geodesic embedding of a minkowskian plane.*

Here, of course, a "Busemann space" is a path-metric space satisfying Busemann's axiom, as described in Section 1. A "minkowskian n-space" is an n-dimensional Banach space (Section 2).

This generalises a result announced by Gromov [12], where "Busemann space" is replaced by "CAT(0) space" and "minkowskian plane" by "euclidean plane". An exposition of some of Gromov's work, including a proof of this assertion for CAT(0) spaces was given by Heber, [14]. More recently, a detailed account of Gromov's statement has been given by Bridson [7]. The argument presented here is somewhat different, and was arrived at independently.

Note that the Cayley graph of Γ is quasiisometric to X. Since hyperbolicity is a quasiisometry invariant, saying that Γ is hyperbolic is equivalent to saying that X is hyperbolic.

Another, much simpler result, which comes directly out of the methods of this paper is:

Theorem 2 : *Suppose X is a Busemann space which admits a discrete cocompact isometric action of a group Γ. If Γ contains a free abelian group, G, of rank n, then X has a totally geodesic n-dimensional minkowskian subspace, invariant under the action of G, and on which G acts by translation.*

In fact we can weaken the hypotheses, as we describe in Section 2. (We only really need the larger group Γ in order that minimal translation distances be attained.) In the context of non-positively curved riemannian manifolds, this result is classical—see [15]. For a proof for CAT(0) spaces, see [8]. Again we get a euclidean space in these cases.

Note that an obvious question remains unanswered by the above two results; namely, with the hypotheses of Theorem 1, if Γ is not hyperbolic, then must it contain a free abelian group of rank 2? This question remains open even for CAT(0) spaces. By analogy with the existence of forced non-periodic tilings of the plane (see [13] and the references therein) one might conjecture that it be false in general, though no explicit counterexample has yet been produced. However an affirmative answer can be given in certain contexts: for smooth riemannian 3-manifolds [11]; for analytic riemannian manifolds of any dimension [4]; and, more recently [16], for cubed 3-manifolds in the sense of Aitchison and Rubinstein [1]. (In each case this is under the assumption of non-positive curvature). More generally, one can ask when the existence of an n-dimensional space implies the existence of a free abelian group of rank n. This is true for analytic riemannian manifolds (again by [4]), and for codimension-1 planes in the smooth riemannian category [18].

The structure of this paper is as follows. In Section 1, we introduce Busemann spaces. In Section 2, we discuss the the geometry of minkowskian spaces, and give a proof of Theorem 2. In Section 3, we give a proof of Theorem 1.

Most of the material for this paper was worked out while visiting the University of Melbourne with the support of an Australian Research Council Fellowship. I am particularly indebted to Iain Aitchison and Hyam Rubinstein for inviting me to Melbourne, and for introducing me to these questions in the context of cubed manifolds. I should also thank Craig Hodgson for helpful conversations.

1. Busemann Spaces.

Let (X, d) be a metric space. A *geodesic* $\alpha : I \longrightarrow X$ is a map of a closed interval $I \subseteq \mathbf{R}$ into X such that $d(\alpha(t), \alpha(u)) = |t - u|$ for all $t, u \in I$. We shall assume that (X, d) is a *length space*, i.e. that any two points of $x, y \in X$ are joined by a geodesic $\beta : [0, d(x, y)] \longrightarrow X$ with $\beta(0) = x$ and $\beta(d(x, y)) = y$. A *linearly reparameterised geodesic* is a map of the form $[t \mapsto \alpha(\lambda t)]$, where α is a geodesic, and $\lambda \in \mathbf{R}$.

Definition : We say that a length space (X, d) is a *Busemann space* if, for any two linearly reparameterised geodesics, $\alpha, \beta : I \longrightarrow X$, the map $[t \mapsto d(\alpha(t), \beta(t)]$ is convex in t.

It follows that a geodesic between any x and y in X must be unique. We write it as $[x \to y]$. We write its image in X as $[x, y]$.

Remarks:

(1) There are various equivalent ways of expressing the Busemann condition. Note that it amounts to saying that $d(\alpha(t),\beta(t)) \leq td(\alpha(0),\beta(0)) + (1-t)d(\alpha(1),\beta(1))$ for any linearly reparameterised geodesics, $\alpha,\beta : [0,1] \longrightarrow X$ and $t \in [0,1]$. It's not hard to see that we can take $\alpha(0) = \beta(0)$. Indeed we also restrict to the case where $t = 1/2$ provided we assume geodesics vary continuously in their endpoints (for example if X is complete and locally compact). This latter corresponds to Busemann's original formulation in terms of midpoints. (Busemann originally worked with an additional hypothesis of prolongation of geodesics, though we shall have no need of this.)

(2) There is also a notion of a locally Busemann space, where X is covered by open sets having the Busemann property. This local property expresses the idea of non-positive curvature. There is a "Cartan-Hadamard" theorem for such spaces [2, 17]: if (X,d) is locally Busemann, then it is (globally) Busemann if and only if it is simply connected. In fact a Busemann space is contractible.

(3) If (X,d) is a Busemann space, then the property of a path being geodesic is a local one: suppose $\alpha : I \longrightarrow X$ is a path, and I has an open cover $\mathcal{U}$ such that if $U \in \mathcal{U}$, and $x,y \in U$ then $d(\alpha(t),\alpha(u)) = |t-u|$, then it follows that α is geodesic.

As mentioned in the introduction, all CAT(0) spaces are Busemann, so this provides a large class of examples, including Hadamard manifolds, and simply connected, non-positively curved polyhedral complexes [6]. Other examples include minkowskian spaces, described in Section 2.

Definition : Suppose (X,d) is a Busemann space. We say that two bi-infinite geodesics $\alpha,\beta : \mathbf{R} \longrightarrow X$ are *parallel* if $d(\alpha(t),\beta(t))$ is bounded.

Since $[t \mapsto d(\alpha(t),\beta(t))]$ is convex, it follows that it must be constant. This defines an equivalence relation on the set of bi-infinite geodesics.

We have the following "Strip Lemma":

Lemma 1.1 (Strip Lemma) : *Suppose $\alpha,\alpha' : \mathbf{R} \longrightarrow X$ are parallel geodesics, so that $d(\alpha(t),\alpha'(t)) = r$, say. For $t \in \mathbf{R}$, let $\beta_t : [0,r] \longrightarrow X$ be the geodesic joining $\alpha(t)$ to $\alpha'(t)$. Define $\beta : \mathbf{R} \times [0,r] \longrightarrow X$ by $\beta(t,u) = \beta_t(u)$. Then:*

(1) For each $u \in [0,r]$, the map $\alpha_u = [t \mapsto \beta(t,u)]$ is a bi-infinite geodesic parallel to α and α'.

(2) For each $t_0 \in \mathbf{R}$, and $\lambda \in \mathbf{R}$, the map $[u \mapsto \beta(t_0 + \lambda u, u)]$ is a linearly reparameterised geodesic.

Proof :

(1) Fix $u \in [0,r]$, and suppose $t_0 < t_1 \in \mathbf{R}$. Let $\mu = d(\alpha_u(t_0),\alpha_u(t_1))/(t_1 - t_0)$. Thus, $|\mu - 1| \leq 2r/(t_1 - t_0)$. (In fact, by convexity, $\mu \leq 1$). Let $\gamma : [0,\mu(t_1 - t_0)] \longrightarrow X$ be the geodesic joining $\alpha_u(t_0)$ to $\alpha_u(t_1)$. Given $t \in [t_0,t_1]$, let $x = \gamma(\mu(t - t_0))$ (so that x cuts γ in the same ratio as t cuts $[t_0,t_1]$). By convexity applied to the geodesics γ and $\alpha|[t_0,t_1]$,

we see that $d(x, \alpha(t)) \le u$. Similarly, $d(x, \alpha'(t)) \le r - u$. But $d(\alpha(t), \alpha'(t)) = r$, and so $d(x, \alpha(t)) = u$ and $d(x, \alpha'(t)) = r - u$. Thus $x = \beta_t(u)$, and so $\alpha_u(t) = \beta_t(u) = \gamma(\mu(t - t_0))$. This is true for all $t \in [t_0, t_1]$, and so $\alpha_u|[t_0, t_1]$ is a geodesic scaled by a factor μ.

Now let $t_0 \to -\infty$ and $t_1 \to \infty$. Since the scale factor is constant, we see that it must equal 1.

(2) Let $t_1 = t_0 + \lambda r$, and $h = d(\alpha(t_0), \alpha'(t_1))$. Let $\gamma : [0, h] \longrightarrow X$ be the geodesic joining $\alpha(t_0)$ to $\alpha'(t_1)$. Suppose $u \in [0, r]$, and let $t = t_0 + \lambda u$. Let $x = \gamma(hu/r)$. By convexity applied to the geodesics γ and $\alpha|[t_0, t_1]$, we see that $d(x, \alpha(t)) \le u$. Similarly, $d(x, \alpha'(t)) \le r - u$. As in part (1), we see that $\beta_t(u) = x = \beta(t_0 + \lambda u, u)$, and so $[u \mapsto \beta(t_0 + \lambda u, u)]$ is a linear reparameterisation of the geodesic γ. $\diamondsuit$

In the case where α and α' have disjoint images, the map β is injective, and $\beta(\mathbf{R} \times [0, r])$ is a totally geodesic subspace. In fact, applying the arguments of Proposition 2.2, one can show that it is isometric to a strip in a minkowskian plane. Thus, in the case of a CAT(0) space, we get a euclidean strip (c.f. [3]).

2. Minkowskian geometry.

Let V be a vector space over $\mathbf{R}$. A *pseudonorm* on V is a map $\phi : V \longrightarrow [0, \infty)$ such that for all $x, y \in V$ and $\lambda \in \mathbf{R}$, we have $\phi(\lambda x) = |\lambda|\phi(x)$ and $\phi(x + y) \le \phi(x) + \phi(y)$. If $\phi(x) = 0$ if and only if $x = 0$, then ϕ is a *norm*. Such a norm is derived from an inner-product if and only if $\phi(x + y)^2 + \phi(x - y)^2 = 2(\phi(x)^2 + \phi(y)^2)$ for all $x, y \in V$.

Given a norm, ϕ, we define a metric, d, on V by $d(x, y) = \phi(x - y)$. With respect to this metric, affine lines are linearly reparametrised geodesics. If the unit ball in V is strictly convex, then these are the only geodesics, and thus (V, d) is a Busemann space. We make the following observation:

Proposition 2.1 : *If V is a normed vector space with metric d, then the following are equivalent:*

(1) V is an inner-product space.

(2) (V, d) is CAT(0).

(3) Every 2-dimensional subspace of V is isometric to the euclidean plane.

(4) Every finite-dimensional subspace of V is isometric to a euclidean space. $\diamondsuit$

We shall be particularly interested in the case where $V \equiv \mathbf{R}^n$ is finite dimensional. In this case, $(\mathbf{R}^n, d)$ is referred to as a *minkowskian space*. There are many ways one can characterise minkowskian spaces by metric properties [10]. (For example: if d is a metric on $\mathbf{R}^n$ which induces the usual topology, and for which the affine midpoint of any two points is a midpoint for the metric d, then $(\mathbf{R}^n, d)$ is minkowskian.)

We describe one way in which minkowskian spaces arise in Busemanm spaces. Let $\{e_1, \ldots, e_n\}$ be the standard basis of $\mathbf{R}^n$, i.e., $e_i = (0, \ldots, 0, 1, 0, \ldots, 0)$.

Definition : An *n-dimensional grid* in X is a map $f : \mathbf{R}^n \longrightarrow X$ such that for all $x \in \mathbf{R}^n$, $\lambda \in \mathbf{R}$, and $i \in \{1, 2, \ldots, n\}$, we have $d(f(x), f(x + \lambda e_i)) = |\lambda|$. (In other words, $[t \mapsto f(x + te_i)] : \mathbf{R} \longrightarrow X$ is a bi-infinite geodesic.) We say that f is *non-degenerate* if it is injective.

Proposition 2.2 : *Suppose $f : \mathbf{R}^n \longrightarrow X$ is a non-degenerate grid, and that ρ is the pull-back metric on $\mathbf{R}^n$, then $(\mathbf{R}^n, \rho)$ is minkowskian. (From which it follows that ρ is a path-metric, and so $f(\mathbf{R}^n)$ is totally geodesic.)*

Proof : We claim that ρ is translation-invariant, so that $\rho(x, y) = \phi(x - y)$ where $\phi(z) = \rho(z, 0)$. First note that distances remain bounded under translation: if $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ lie in $\mathbf{R}^n$, then $\rho(x, y) \le \sum_{i=1}^n |x_i - y_i|$ (changing one coordinate at a time). It follows that, for each $i \in \{1, 2, \ldots, n\}$ the bi-infinite geodesics $[t \mapsto f(x + te_i)]$ and $[t \mapsto f(y + te_i)]$ are parallel. Thus $\rho(x + te_i, y + te_i) = \rho(x, y)$ for all t, i. Since the e_i form a basis of $\mathbf{R}^n$, it follows that $\rho(x, y) = \phi(x - y)$ as claimed.

Next we claim that if $x \in \mathbf{R}^n$, then the map $[t \mapsto f(tx)] : \mathbf{R} \longrightarrow X$ is a linearly reparameterised geodesic. We prove this by induction on the nested subspaces $\mathbf{R} \subseteq \mathbf{R}^2 \subseteq \cdots \subseteq \mathbf{R}^{n-1} \subseteq \mathbf{R}^n$. So suppose that $x = y + \lambda e_n$ with $y \in \mathbf{R}^{n-1}$ and $\lambda \in \mathbf{R}$. and $\lambda \in \mathbf{R}$. By induction, we suppose $[t \mapsto f(ty)]$ is geodesic. Since ρ is translation-invariant, $[t \mapsto f(ty + \mu e_n)]$ is geodesic for all $\mu \in \mathbf{R}$. Given, $u_0, u_1 \in \mathbf{R}$, consider the map $\beta = [(t, u) \mapsto f(yu + te_n)] : \mathbf{R} \times [u_0, u_1] \longrightarrow X$. Now, $\alpha = [t \mapsto \beta(t, u_0)]$ and $\alpha' = [t \mapsto \beta(t, u_1)]$ are parallel geodesics, and so applying the Strip Lemma (1.1), we see that $[t \mapsto f(tx) = f(ty + t\lambda e_n)]$ for $t \in [u_0, u_1]$ is a linearly reparameterised geodesic. Since, u_0 and u_1 are arbitrary, this is true as t varies in $\mathbf{R}$, so the claim follows.

It now follows immediately that $\phi(\lambda x) = |\lambda| \phi(x)$ for all $\lambda \in \mathbf{R}$ and $x \in \mathbf{R}^n$. Also, since ρ is translation invariant, the map $[t \mapsto tx + (1 - t)y]$ is a linearly reparameterised geodesic for all $x, y \in \mathbf{R}^n$. By convexity, we see that $\phi(x + y) = 2\phi(\frac{1}{2}(x + y)) \le \phi(x) + \phi(y)$. Thus ϕ is a norm on $\mathbf{R}^n$. $\diamondsuit$

Thus a non-degenerate grid is essentially the same as an embedded totally geodesic minkowskian subspace (parameterised so that all the standard basis elements have unit length). In general, a (possibly degenerate) grid will factor through such a subspace:

Proposition 2.3 : *If $f : \mathbf{R}^n \longrightarrow X$ is a grid, then $f = g \circ T$, where $T : \mathbf{R}^n \longrightarrow \mathbf{R}^m$ is a surjective linear map, and $g : \mathbf{R}^m \longrightarrow X$ is non-degenerate grid.*

Proof : In this case, the same argument as Proposition 2.2 gives us a pseudonorm ψ on $\mathbf{R}^n$ so that $d(f(x), f(y)) = \psi(x - y)$. Now $V = \{x \in \mathbf{R}^n \mid \psi(x) = 0\}$ is a subspace of $\mathbf{R}^n$, and ψ projects to a norm ϕ on $\mathbf{R}^n / V$ so that $\phi(V + x) = \psi(x)$ for all $x \in \mathbf{R}^n$. We define $g : \mathbf{R}^n / V \longrightarrow X$ by $g(V + x) = f(x)$. We now identify $\mathbf{R}^n / V$ with $\mathbf{R}^m$ in such a way that $\phi(e_i) = 1$ for all $i \in \{1, 2, \ldots, m\}$. $\diamondsuit$

In particular, we see that a degenerate 2-dimensional grid has the form $[(t, u) \mapsto \alpha(t \pm u)]$ where $\alpha : \mathbf{R} \longrightarrow X$ is a bi-infinite geodesic.

It turns out that it is enough that f be defined on a lattice in $\mathbf{R}^n$. If f is a "discrete grid" then it extends uniquely to a grid in the sence already defined.

Definition : Suppose $\lambda_1, \lambda_2, \dots, \lambda_n > 0$. Let Λ be the lattice $\lambda_1 \mathbf{Z} \oplus \lambda_2 \mathbf{Z} \oplus \cdots \oplus \lambda_n \mathbf{Z} \subseteq \mathbf{R}^n$. A *discrete grid* is a map $f : \Lambda \longrightarrow X$ such that $d(f(x), f(x + m\lambda_i e_i)) = |m|\lambda_i$ for all $x \in \Lambda$, $m \in \mathbf{Z}$ and $i \in \{1, 2, \dots, n\}$.

Proposition 2.4 : *If (X, d) is a Busemann space, then a discrete grid $f : \Lambda \longrightarrow X$ extends to a grid $f : \mathbf{R}^n \longrightarrow X$.*

Proof : To simplify notation, we shall assume that $\lambda_i = 1$ for each i, so that $\Lambda = \mathbf{Z}^n$. The general argument is exactly the same.

Note first that the 1-dimensional case is trivial. A 1-dimensional discrete grid $f : \mathbf{Z} \longrightarrow X$ extends to a bi-infinite geodesic $f : \mathbf{R} \longrightarrow X$ by joining $f(m)$ to $f(m+1)$ by a geodesic segment $(f|[m, m+1]) : [m, m+1] \longrightarrow X$ for each $m \in \mathbf{Z}$.

It follows that an n-dimensional discrete grid $f : \mathbf{Z}^n \longrightarrow X$ extends to map $f : \mathbf{Z}^{n-1} \oplus \mathbf{R} \longrightarrow X$ such that $[t \mapsto f(x, t)]$ is a bi-infinite geodesic for each $x \in \mathbf{Z}^{n-1}$. Moreover any two such geodesics remain a bounded distance apart, and are thus parallel (c.f. Proposition 2.2). It follows that for each $t \in \mathbf{R}$, the map $[x \mapsto f(x, t)] : \mathbf{Z}^{n-1} \longrightarrow X$ is an $(n-1)$-dimensional discrete grid. By induction on dimension we can assume that f extends to a map $f : \mathbf{R}^n \longrightarrow X$ such that $[x \mapsto f(x, t)] : \mathbf{R}^{n-1} \longrightarrow X$ is an $(n-1)$-dimensional grid for all $t \in \mathbf{R}$.

We claim that $f : \mathbf{R}^n \longrightarrow X$ is an n-dimensional grid. For this we need to know that the map $[t \mapsto f(x, t)] : \mathbf{R} \longrightarrow X$ is bi-infinite geodesic for all $x \in \mathbf{R}^n$. But this is a consequence of the Strip Lemma (1.1), and Propositions 2.3 and 2.4, for all $x \in \mathbf{Q}^n$. The result now follows by continuity. ♢

We immediately arrive at a proof of Theorem 2:

Proof of Theorem 2 : Let $\{g_1, \dots, g_n\}$ be the standard generating set for $\mathbf{Z}^n \subseteq \Gamma$. Choose $x \in X$ so as to minimise $\sum_{i=1}^n d(x, g_i(x))$. This is attained since Γ acts cocompactly. Let $\lambda_i = d(x, g_i(x))$. Now, for all $m \in \mathbf{Z}$, we must have $d(x, (mg_i)(x)) = |m|\lambda_i$, otherwise we could replace x by the midpoint of the geodesic joining x to $(mg_i)(x)$, and contradict minimality. Let $\Lambda = \bigoplus_{i=1}^n \lambda_i \mathbf{Z} \subseteq \mathbf{R}^n$, and define $f : \Lambda \longrightarrow X$ by $f(m_1\lambda_1, \dots, m_n\lambda_n) = \left(\sum_{i=1}^n m_i g_i\right)(x)$. Since $\mathbf{Z}^n$ acts by isometry, we see that f is a $\mathbf{Z}^n$-equivariant discrete grid. By Proposition 2.4, it extends to a grid $f : \mathbf{R}^n \longrightarrow X$. By naturality, this is also $\mathbf{Z}^n$ invariant. Using Proposition 2.3, since $\mathbf{Z}^n$ acts discretely, we see easily that f must be non-degenerate. Thus, by Proposition 2.2, $f(\mathbf{R}^n)$ is the required $\mathbf{Z}^n$-invariant minkowskian subspace. ♢

As remarked, in the introduction, all we really require is that we should have an action by isometry of $\mathbf{Z}^n$ so that for each $i \in \{1, 2, \dots, n\}$, the minimal translation distance $\min\{d(x, g_i(x)) \mid x \in X\}$ is attained. From this, a simple induction argument shows that there is some $x \in X$ which attains each of these minima simultaneously, and the result follows, as above.

Note that in the case of a CAT(0) space, we get an embedded euclidean n-space. We refer to [8] for details.

3. Hyperbolicity and minkowskian subspaces.

The aim of this section is to prove the dichotomy given by Theorem 1. The idea will be, under the assumption of non-hyperbolicity, to construct a discrete grid as a limit of "approximate grids". The method of constructing these approximating grids can be described quite generally. The assumption of non-hyperbolicity will only enter to verify that the limiting grid must be non-degenerate. Propositions 2.4 and 2.2 then show that it gives rise to an embedded minkowskian plane.

Let (X, d) be a Busemann space. (For the moment, this is our only assumption.) Given $n \in \mathbf{N}$, let $I_n = \{0, 1, 2, \ldots, n\}$. Suppose $\epsilon > 0$.

Definition : An (n, ϵ)-grid is a map $f : I_n^2 \longrightarrow X$ such that for all $i, i', j, j' \in I_n$, we have

$$|i - i'| - \epsilon \leq d(f(i, j), f(i', j)) \leq |i - i'|$$

and

$$|j - j'| - \epsilon \leq d(f(i, j), f(i, j')) \leq |j - j'|.$$

An equivalent way to express this is as follows. Define a metric ρ on I_n^2 by $\rho((i, j), (i', j')) = |i - i'| + |j - j'|$. An (n, ϵ)-grid is then a distance non-increasing map $f : (I_n^2, \rho) \longrightarrow (X, d)$ satisfying $d(f(0, j), f(n, j)) \geq n - \epsilon$ and $d(f(i, 0), f(i, n)) \geq n - \epsilon$ for all $i, j \in I_n$.

We aim to construct a sequence of such grids, with $n \to \infty$ and $\epsilon \to 0$. We describe the construction in this general setting, though it will only be of interest when additional hypotheses are added.

Given $n \in \mathbf{N}$ and $\epsilon > 0$, let q be a natural number greater than $2n^2/\epsilon$. Suppose that $\alpha : [0, qn] \longrightarrow X$ is a geodesic, and that $a \in X$ is any point with $d(a, \alpha([0, qn])) \geq n$. For each $i \in [0, qn] \cap \mathbf{N}$, let $\beta_i : [0, d(a, \alpha(i))] \longrightarrow X$ be the geodesic joining $\alpha(i)$ to a. For each $p \in \{0, 1, \ldots, q-1\}$ and $i, j \in I_n$, set $f_p(i, j) = \beta_{pn+i}(j)$. This defines a map $f_p : I_n^2 \longrightarrow X$ for each such p (Figure 1).

By construction, if $0 \leq p \leq q - 1$ and $i, i', j, j' \in I_n$, then

$$d(f_p(i, j), f_p(i, j')) = |j - j'|$$

and

$$d(f_p(i, 0), f_p(i', 0)) = |i - i'|.$$

Also, by convexity applied to the geodesics β_{pn+i} and $\beta_{pn+i'}$, we see that

$$d(f_p(i, j), f_p(i', j)) \leq |i - i'|.$$

For $0 \leq p \leq q - 1$ and $j \in I_n$, let $h(p, j) = n - d(f_p(0, j), f_p(n, j)) \geq 0$. Let $J_j = \{p \mid h(p, j) > \epsilon\} \subseteq \{0, 1, \ldots, q - 1\}$. Now since

$$d(f_0(0, j), f_{q-1}(n, j)) = d(\beta_0(j), \beta_{qn}(j)) \geq qn - 2j \geq qn - 2n,$$

we see that

$$\epsilon|J_j| \leq \sum_{p=0}^{q-1} h(p,j) \leq 2n.$$

Thus

$$\left| \bigcup_{j=1}^{n} J_j \right| \leq 2n^2/\epsilon < q,$$

and so there must be some $p \in \{0, 1, \ldots, q-1\} \setminus \bigcup_{j=1}^{n} J_j$. For such a p, we have $h(p,j) \leq \epsilon$ for all $j \in I_n$. In other words, $d(f_p(0,j), f_p(n,j)) \geq n - \epsilon$, and so f_p is an (n,ϵ)-grid.

Provided that X is unbounded (as a metric space) this construction gives a sequence $f(n)$ of (n,ϵ)-grids with $n \to \infty$ and $\epsilon \to 0$.

Now suppose in addition, X admits a cocompact group action. For each $m \in \mathbf{N}$, we define $g(m) : ([-m,m] \cap \mathbf{Z})^2 \longrightarrow X$ by $g(m)(i,j) = f(2m)(m+i, m+j)$. In other words, we shift the origin to the centre of our grid. Now, a diagonal subsequence argument allows us to find a subsequence $g(m_i)$ converging on a map $g : \mathbf{Z}^2 \longrightarrow X$. It is easily verified that g must be a discrete grid. In general, we would expect the grid we arrive at in this way to be degenerate. We shall show that if (X,d) is not hyperbolic in the sense of Gromov, then, with care, we can ensure that the limiting grid is non-degenerate. By Propositions 2.2 and 2.4, it follows that g gives rise to a totally geodesic minkowskian subspace.

The characterisation of hyperbolicity we shall use is the following. It is a slight variant of definition H5 described in [5].

Proposition 3.1 : *A path-metric space (X,d) is hyperbolic if and only if there is some $a \in X$ and constants $L > \eta > 0$ such that the following holds.*

Suppose $x, y \in X$ so that $d(a,x) = d(a,y)$ and $L \leq d(x,y) \leq 4L$. Suppose β is a path joining x to y with $d(a,\beta) \geq d(a,x)$, then $\mathrm{length}(\beta) \geq d(x,y) + \eta$.

Sketch of Proof : We may as well assume that (X,d) is a length-space. Also we are only interested here in the implication in one direction. Thus we assume that (X,d) satisfies the hypothesis with a, L and η as described. We can deduce a form of the linear isoperimetric property as follows.

Suppose that γ is a closed curve in X of length greater than $8L$. For large enough r we have that γ lies in the uniform r-ball, $N(a,r)$, about a. By continuously decreasing r, we find that at some point we will have an arc $\beta \subseteq \gamma$ with endpoints $x, y \in \partial N(a,r)$, lying outside the interior of $N(a,r)$ and satisfying $2L \leq \mathrm{length}(\beta) \leq 4L$. In particular, $d(x,y) \leq 4L$. Now either $d(x,y) \leq L$, or else $\mathrm{length}(\beta) \geq d(x,y) + \eta$. Either way, $\mathrm{length}(\beta) - d(x,y) \geq \eta$.

Let $\alpha = \gamma \setminus \beta$, and let $\gamma_1 = \alpha \cup [x,y]$ and $\delta_1 = \beta \cup [x,y]$. We have cut γ into two closed curves γ_1 and δ_1 with $\mathrm{length}(\gamma_1) \leq \mathrm{length}(\gamma) - \eta$ and $\mathrm{length}(\delta_1) \leq 8L$.

Now repeat the construction with γ_1 replacing γ, to get two sequences of closed curves $\gamma_1, \gamma_2, \ldots$ and $\delta_1, \delta_2, \ldots$. After at most $p < \mathrm{length}(\gamma)/\eta$ steps, we arrive at a curve γ_p of length at most $8L$. We have thus cut γ into $p + 1 = O(\mathrm{length}(\gamma))$ curves $\delta_1, \delta_2, \ldots, \delta_p, \gamma_p$, each of length at most $8L$. This is one form the linear isoperimetric inequality (H3 as described in [5]). We refer to [5] for more details. $\diamondsuit$

We shall also need the following observation:

Lemma 3.2 : *Suppose (X,d) is a Busemann space which admits a cocompact isometric group action. Then for all $k,h>0$, there is some $\eta>0$ such that the following holds. Suppose $x,y\in X$ with $d(x,y)\leq k$. If β is any path joining x to y with* $\mathrm{length}(\beta)\leq d(x,y)+\eta$, *then for all $z\in[x,y]$, we have $d(z,\beta)\leq h$.*

Proof : Suppose the conclusion fails. Then we can find sequences $x_i,y_i\in X$ and $z_i\in[x_i,y_i]$, and paths β_i joining x_i to y_i so that $d(x_n,y_n)\leq k$, $\mathrm{length}(\beta_i)\leq d(x_i,y_i)+1/i$ and $d(z_i,\beta_i)>h$. We can take all the β_i to be defined on the same domain and parameterised proportionally to arc length. Up to the action of the group, we can assume, after taking a subsequence, that $x_i\to x$, $y_i\to y$, and $z_i\to z\in[x,y]$. From the uniqueness of geodesics, we see that the β_i must converge to $[x,y]$, contradicting the fact that $d(z,\beta_i)$ is bounded away from 0. $\diamondsuit$

(Note that if X happens to be a CAT(0) space, then we can obtain η directly as a function if h and k without the assumption of a cocompact group action.)

We are now ready to finish the proof of our main result.

Proof of Theorem 1 : Suppose that (X,d) is a non-hyperbolic Busemann space which admits a cocompact group action. We show that X admits a non-degenerate discrete grid $g:\mathbf{Z}^2\longrightarrow X$.

Given $n>0$ and $\epsilon>0$, we construct an (n,ϵ)-grid as follows. Let q be some natural number greater than $2n^2/\epsilon$. Given $k=4qn$ and $h=1/2$, let η be the constant given by Lemma 3.2. We can assume that $\eta<qn$.

Now choose any $a\in X$, and apply Proposition 3.1 with $L=qn>\eta$. We find that there must exist points $x,y\in X$ with $qn\leq d(x,y)\leq 4qn$ and $d(a,x)=d(a,y)=r$, say, together with a path β joining x to y such that $d(a,\beta)\geq r$ and $\mathrm{length}(\beta)\leq d(x,y)+\eta$. Suppose $z\in[x,y]$. By convexity, we have $d(a,z)\leq r$. Also, from the choice of η, we have $d(z,\beta)\leq\frac{1}{2}$, so $d(a,z)\geq r-\frac{1}{2}$. Note that $2r\geq d(x,y)\geq L=qn$, so we can choose q large enough so that $r-\frac{1}{2}\geq n$.

Writing $\alpha:[0,d(x,y)]\longrightarrow X$ for the geodesic from x to y, we can perform the construction described above. This gives an (n,ϵ)-grid, $f=f_p:I_n^2\longrightarrow X$, with the property that for all $i\in I_n$,

$$r-\frac{1}{2}\leq d(f(i,0),a)\leq r.$$

For $i, j \in \{0, 1, \ldots, n-1\}$, let $\delta(i,j) = d(f(i,j), f(i+1, j+1))$. Now,

$$\begin{aligned} r - \frac{1}{2} &\leq d(f(i,0), a) \\ &\leq d(f(i,0), f(i,j)) + d(f(i,j), f(i+1, j+1)) + d(f(i+1, j+1), a) \\ &= j + \delta(i,j) + d(f(i+1, j), a) - 1 \\ &= \delta(i,j) + d(f(i+1, 0), a) - 1 \\ &\leq \delta(i,j) + r - 1, \end{aligned}$$

and so $d(f(i,j), f(i+i, j+i)) = \delta(i,j) \geq \frac{1}{2}$. By a similar argument, we see that $d(f(i, j+1), f(i+1, j)) \geq \frac{1}{2}$.

We may now take a sequence of such (n, ϵ)-grids, $f(n)$, with $n \to \infty$ and $\epsilon \to 0$. This gives rise to a limiting grid $g : \mathbf{Z}^2 \longrightarrow X$. From the construction, we see that $d(g(0,0), g(1,1)) \geq \frac{1}{2}$, and $d(g(0,1), g(1,0)) \geq \frac{1}{2}$. By Proposition 2.4, g extends to a grid $g : \mathbf{R}^2 \longrightarrow X$. From the form of degenerate 2-dimensional grids described after Proposition 2.3, we see that g must be non-degenerate. The theorem now follows from Proposition 2.2. $\diamondsuit$

References.

[1] I.R.Aitchison, J.H.Rubinstein, *An introduction to polyhedral metrics of non-positive curvature on 3-manifolds* : in “Geometry of low-dimensional manifolds: 2”, ed. S.K.Donaldson & C.B.Thomas, L.M.S. Lecture Notes Series 151, Cambridge University Press (1990) 127–161.

[2] S.B.Alexander, R.L.Bishop, *The Hadamard-Cartan theorem in locally convex spaces* : Enseign. Math. **36** (1990) 309–320.

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

[4] V.Bangert, V.Schroeder, *Existence of flat tori in analytic manifolds of nonpositive curvature* : Ann. Scient. Éc. Norm. Sup. **24** (1990) 605–634.

[5] 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.

[6] M.R.Bridson, *Geodesics and curvature in metric simplicial complexes* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 373–463.

[7] M.R.Bridson, *On the existence of flat planes in spaces of non-positive curvature* : preprint (1992).

[8] M.R.Bridson, A.Haefliger, *An introduction to CAT(0) spaces* : in preparation.

[9] H.Busemann, *Spaces with non-positive curvature* : Acta Math. **80** (1948) 259–310.

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

[11] S.V.Buyalo, *Euclidean planes in three-dimensional manifolds of nonpositive curvature* : Mat. Zametki **43** (1988) 103–114; English translation, Math. Notes **43** (1988) 60-66.

[12] M.Gromov, *Hyperbolic groups* : in “Essays in group theory”, (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer (1988) 75–263.

[13] B.Gruenbaum, G.C.Shephard, *Tilings and patterns* : W.H.Freeman and Co. (1987).

[14] J.Heber, *Hyperbolische geodätische Räume* : Diplomarbeit, Universität Bonn (1987).

[15] H.B.Lawson, S.T.Yau, *Compact manifolds of non-positive curvature* : J. Diff. Geom. **7** (1972) 211–228.

[16] L.Mosher, *Cubulations of 3-manifolds* : preprint, Rutgers, Newark (1993).

[17] F.Paulin, *Constructions of hyperbolic groups via hyperbolisations of polyhedra* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 313–372.

[18] V.Schroeder, *Codimension-one flats in manifolds of non-positive curvature* : Geom. Dedicata **33** (1990) 251–263.