ISSN 1364-0380 (on line) 1465-3060 (printed) **521**

Geometry & Topology
Volume 5 (2001) 521–550
Published: 21 May 2001

Metric characterizations of spherical and Euclidean buildings

Ruth Charney
Alexander Lytchak

Mathematics Department, Ohio State University
231 W 18th Ave, Columbus, OH 43210, USA

and

Mathematisches Institut der Universität Bonn
Wegelerstraße 10, D-53115 Bonn, Germany

Email: charney@math.ohio-state.edu and lytchak@math.uni-bonn.de

Abstract

A building is a simplicial complex with a covering by Coxeter complexes (called apartments) satisfying certain combinatorial conditions. A building whose apartments are spherical (respectively Euclidean) Coxeter complexes has a natural piecewise spherical (respectively Euclidean) metric with nice geometric properties. We show that spherical and Euclidean buildings are completely characterized by some simple, geometric properties.

AMS Classification numbers Primary: 20E42

Secondary: 20F65

Keywords: Buildings, CAT(0) spaces, spherical buildings, Euclidean buildings, metric characterisation

Proposed: Walter Neumann | Received: 23 November 2000
Seconded: Jean-Pierre Otal, Steve Ferry | Revised: 11 May 2001

0 Introduction

In recent years, much attention has been given to curvature properties of piecewise Euclidean and piecewise spherical complexes. A notion of curvature bounded above for such complexes was introduced by Alexandrov in the 1950's and further developed in the 1980's by Gromov. The curvature bound is defined as a condition on the shape of triangles (they must be sufficiently "thin") and is known as a CAT–inequality (Comparison inequality of Alexandrov–Toponogov). Some particularly nice examples of spaces satisfying CAT–inequalities are spherical and Euclidean buildings which come equipped with a natural piecewise spherical or Euclidean metric.

Buildings also satisfy other nice metric properties. A spherical building X, for example, is easily seen to have diameter π, as does the link of any simplex in X. It is natural to ask whether Euclidean and spherical buildings are characterized by their metric properties. In this paper, we give several metric characterizations of buildings. For example we prove

Theorem *Let X be a connected, piecewise spherical (respectively Euclidean) complex of dimension $n \geq 2$ satisfying*

(1) *X is CAT(1) (respectively CAT(0)).*

(2) *Every $(n-1)$–cell is contained in at least two n–cells.*

(3) *Links of dimension ≥ 1 are connected.*

(4) *Links of dimension 1 have diameter π.*

Then X is isometric to a spherical building (respectively a metric Euclidean building).

Metric Euclidean buildings are products of irreducible Euclidean buildings, cones on spherical buildings, trees, and nonsingular Euclidean spaces (see Section 5).

Another metric property of buildings is that for every local geodesic γ, the set of directions in which γ can be geodesically continued is non-empty and discrete. We call this the "discrete extension property".

Theorem *Let X be a connected, piecewise spherical (respectively Euclidean) complex of dimension $n \geq 2$ satisfying*

(1) *X is CAT(1) (respectively CAT(0)).*

(2) *X has the discrete extension property.*

Then X is isometric to a spherical building (respectively metric Euclidean building).

Werner Ballmann and Michael Brin, studying the question of rank rigidity for piecewise Euclidean complexes of nonpositive curvature, have obtained related results in [2] and [3], including a metric characterization of spherical and Euclidean buildings of dimension 2. Bruce Kleiner has also described (unpublished) a metric characterization of Euclidean buildings under the assumption that every geodesic is contained in an n–flat.

The first two sections of the paper contain background about buildings and geodesic metric spaces. The key problem in identifying a building is the construction of enough apartments. Sections 3, 4, and 5 are devoted to this task. Section 6 considers the 1–dimensional case, and Section 7 combines these results to arrive at the main theorems.

The first author would like to thank Bruce Kleiner for helpful conversations. The second author is indebted to Werner Ballmann for suggesting the question and for his continued support during the development of this paper. The first author was partially supported by NSF grant DMS-9803374.

1 Buildings

In this section we review some definitions and terminology. For more details about buildings, see [6] and [5].

Let S be a finite set. A *Coxeter matrix* on S is a symmetric function $m\colon S\times S \to \{1,2,\dots,\infty\}$ such that $m(s,s) = 1$ and $m(s,t) \geq 2$ for $s \neq t$. The *Coxeter group* associated to m is the group W given by the presentation

$$W = \langle\ S \mid (st)^{m(s,t)} = 1,\ s,t \in S\ \rangle.$$

The pair (W,S) is called a Coxeter system. If $T \subset S$, then the subgroup W_T generated by T is the Coxeter group associated to $m|_{T\times T}$. The Coxeter system (W,S) is *irreducible* if there is no non-trivial partition $S = S_1 \amalg S_2$ such that $W = W_{S_1} \times W_{S_2}$.

The Coxeter group W can be realized as a discrete group of linear transformations of an n–dimensional vector space V, with the generators $s \in S$ acting as reflections across the walls of a simplicial cone. This action preserves a bilinear form B on V represented by the matrix $B(s,t) = -\cos\left(\frac{\pi}{m(s,t)}\right)$ (where $\frac{\pi}{\infty}$ is taken to be 0). W is finite if and only if this form is positive definite. In this case, W acts as a group of orthogonal transformations and the action restricts to the unit sphere $S(V)$ in V. Hence W is called a *spherical Coxeter group*.

If B is positive semi-definite (but not definite) and W is irreducible, then the action of W on V induces an action on an (n-1)-dimensional affine space

$\mathbb{R}^{n-1} = V/V^{\perp}$ with the generators acting as affine reflections across the walls of a simplex. In this case, W is called an *irreducible Euclidean Coxeter group.* A key fact about irreducible Euclidean Coxeter groups is that for any proper subgroup $T \subset S$, W_T is a spherical Coxeter group. More generally, we call W a *Euclidean Coxeter group* if it is a direct product of irreducible Euclidean Coxeter groups.

To any Coxeter group W, one can associate a simplicial complex Σ_W called the *Coxeter complex* for W. In the case of spherical and Euclidean Coxeter groups, the Coxeter complex has a simple, geometric description. Let $M = \mathbb{S}^{n-1} = S(V)$ if W is spherical or $M = \mathbb{R}^{n-1} = V/V^{\perp}$ if W is irreducible Euclidean. For each element $r \in W$ which acts as a reflection on M (namely, r is a generator or conjugate of a generator), r fixes some hyperplane, called a *wall* of M. The walls divide M into simplices. The resulting simplicial complex, Σ_W, is the Coxeter complex for W. If W is a product of irreducible Euclidean Coxeter groups, then Σ_W is the product of the corresponding Coxeter complexes. The top dimensional simplices (or cells) of Σ_W are called *chambers.* W acts freely transitively on the set of chambers of Σ_W and the stabilizer of any lower dimensional cell σ is conjugate to W_T for some $T \subset S$.

There are several equivalent definitions of buildings. The most convenient for our purposes is the following (see [6]).

Definition 1.1 A *building* is a simplicial complex X together with a collection of subcomplexes $\mathcal{A}$, called *apartments*, satisfying

(1) each apartment is isomorphic to a Coxeter complex,

(2) any two simplices of X are contained in a common apartment,

(3) if two apartments A_1, A_2 share a chamber, then there is an isomorphism $A_1 \to A_2$ which fixes $A_1 \cap A_2$ pointwise.

If, in addition, every codimension 1 simplex is contained in at least three chambers, then X is a *thick* building. It follows from conditions (2) and (3) that all of the apartments are isomorphic to the same Coxeter complex Σ_W. We say that a building X is *spherical* (respectively *Euclidean*) if W is spherical (respectively Euclidean).

Although the collection of apartments $\mathcal{A}$ is not, in general unique, there is a unique maximal set of apartments. We will always assume $\mathcal{A}$ to be maximal.

If X_1, X_2 are spherical buildings with associated Coxeter groups W_1, W_2, then the join $X_1 * X_2$ is a spherical building with Coxeter group $W_1 \times W_2$. In particular, the suspension $\Sigma X_2 = \mathbb{S}^0 * X_2$ is a building with Coxeter group

$\mathbb{Z}/2 \times W_2$. (Note that for $k > 0$, the simplicial structure on $\mathbb{S}^k * X_2$ depends on a choice of identification of $\mathbb{S}^k$ with the (k+1)-fold join $\mathbb{S}^0 * \cdots * \mathbb{S}^0$. Despite this slight ambiguity, we will consider the k–fold suspension of a building to be a building.) Conversely, if X is a spherical building whose Coxeter group W splits as a product $W_1 \times W_2$, then X can be decomposed as the join of a building for W_1 and a building for W_2 (see [13], Theorem 3.10). Similarly, any Euclidean building splits as a product of irreducible Euclidean buildings.

2 Metrics

A metric space (X, d) is a *geodesic metric space* if for any two points $x, y \in X$, there is an isometric embedding γ: $[0, a] \to X$ with $\gamma(0) = x$ and $\gamma(a) = y$. Such a path is called a *geodesic segment* or simply a *geodesic* from x to y. An isometric embedding of $\mathbb{R}$ into X is also called a *geodesic*, and an isometric embedding of $[0, \infty)$ is called a *ray*.

A *piecewise Euclidean* (respectively *piecewise spherical*) complex is a polyhedral cell complex X together with a metric d such that each cell of X is isometric to a convex polyhedral cell in $\mathbb{R}^n$ (respectively $\mathbb{S}^n$) for some n, and

$$d(x, y) = \inf \{ \text{ length}(\gamma) \mid \gamma \text{ is a path from } x \text{ to } y\}$$

for any $x, y \in X$. We will also assume that the metric d is a complete, geodesic metric. In particular, the infimum $d(x, y)$ is realized by the length of some path γ.

If X is a piecewise spherical or Euclidean complex and x is a point in X, then the set of unit tangent vectors to X at x is called the *link* of x and denoted $\mathrm{lk}\,(x, X)$. It comes equipped with the structure of a piecewise spherical complex, since the link of x in a single n–cell is isometric to a polyhedral cell in $\mathbb{S}^{n-1}$. If σ is a k–cell in X, we define $\mathrm{lk}\,(\sigma, X)$ to be the set of unit tangent vectors orthogonal to σ at any point x in the relative interior of σ. This set also has a natural piecewise spherical structure and we can identify

$$\mathrm{lk}\,(x, X) = \mathrm{lk}\,(x, \sigma) * \mathrm{lk}\,(\sigma, X) = \mathbb{S}^{k-1} * \mathrm{lk}\,(\sigma, X)$$

where the joins $*$ are orthogonal joins in the sense of [7]. (See [4] or the appendix of [7] for a discussion of joins of piecewise spherical complexes.)

In some cases, we may wish to consider spaces which do not have a globally defined cell structure. For this, we introduce the notion of a locally spherical space of dimension n. The definition is inductive. A locally spherical space of dimension 0 is a nonempty disjoint union of points. A *locally spherical (respectively Euclidean) space of dimension* n, $n > 0$, is a complete geodesic metric space (X,d) for which every point x has a neighborhood isometric to a spherical (respectively Euclidean) cone on a locally spherical space L_x of dimension $n-1$. We call such a neighborhood a *conelike neighborhood* of x. Clearly, $L_x = \operatorname{lk}(x,X)$. A piecewise spherical (respectively Euclidean) complex is a locally spherical (respectively Euclidean) space of dimension n if and only if every cell is contained in an n–dimensional cell.

The basis for our metric characterization of buildings will be the CAT–inequalities defined by Gromov in [9]. Let (X,d) be a complete, geodesic metric space and let T be a geodesic triangle in X. A Euclidean comparison triangle for T is a triangle T' in $\mathbb{R}^2$ with the same side lengths as T. We say X is a *CAT(0) space* if every geodesic triangle T is "thin" relative to its comparison triangle T'. That is, given any points $x,y \in T$, the distance from x to y in X is less than or equal to the distance in $\mathbb{R}^2$ between the corresponding points $x',y' \in T'$. We define a *CAT(1) space* similarly by comparing geodesic triangles T in X with spherical comparison triangles T' in $\mathbb{S}^2$. In this case, however, we only require the thinness condition to hold for triangles T of perimeter $\leq \pi$ (since no comparison triangle exists with perimeter $> \pi$).

In the next two theorems we collect some facts about CAT(0) and CAT(1) spaces. These are due to Gromov, Ballmann, Bridson and others. A good source of proofs is [4] or [1].

Theorem 2.1 *Let X be a piecewise (or locally) Euclidean geodesic metric space.*

(1) *X is locally CAT(0) if and only if $\operatorname{lk}(\sigma,X)$ is CAT(1) for every cell σ.*

(2) *X is CAT(0) if and only if it is locally CAT(0) and simply connected.*

(3) *If X is CAT(0), then any two points in X are connected by a unique geodesic and any path which is locally geodesic, is geodesic.*

Theorem 2.2 *Let X be a piecewise (or locally) spherical geodesic metric space.*

(1) *X is locally CAT(1) if and only if $\operatorname{lk}(\sigma,X)$ is CAT(1) for every cell σ.*

(2) *X is CAT(1) if and only if it is locally CAT(1), any two points of distance $< \pi$ are connected by a unique geodesic, and these geodesics vary continuously with their endpoints.*

(3) *If X is CAT(1), then any path of length $\leq \pi$ which is locally geodesic, is geodesic.*

A Euclidean (respectively spherical) building X of dimension n comes equipped with a natural piecewise Euclidean (respectively piecewise spherical) metric in which each apartment is isometric to $\mathbb{R}^n$ (respectively $\mathbb{S}^n$) with the Coxeter group W acting by isometries. In the spherical case, there is a unique such metric. In the Euclidean case, this metric is determined only up to scalar multiple on each irreducible factor. *In this paper, for a piecewise Euclidean (respectively spherical) complex (X, d), the statement that X is a Euclidean (respectively spherical) building will mean that, the cell structure on X satisfies the conditions of Definitions 1.1 and that the metric on X is the natural building metric.*

Define the diameter of X to be

$$diam(X) = \sup\{d(x,y)|x, y \in X\}.$$

The natural metric on a building satisfies a number of nice properties which are described in the proposition below.

Proposition 2.3 *Let X be a Euclidean (respectively spherical) building of dimension n with the natural metric. Then:*

(1) *X is CAT(0) (respectively X is CAT(1) and $daim(X) = \pi$).*

(2) *For any simplex σ of codimension ≥ 2, $lk(\sigma, X)$ is a spherical building. In particular, $lk(\sigma, X)$ is CAT(1) and $diam(lk(\sigma, X)) = \pi$.*

(3) *A subspace $A \subset X$ is an apartment if and only if the intrinsic metric on A is isometric to $\mathbb{R}^n$ (respectively $\mathbb{S}^n$). Moreover, the inclusion $A \hookrightarrow X$ is an isometric embedding.*

Proof (1) and (2) are well known. (1) follows from the fact that every geodesic in X is contained in an apartment (see [8] or [6]). (2) follows from the fact that the isotropy group of a simplex σ in Σ_W is a spherical Coxeter group W_σ. The action of W_σ on the sphere $\mathrm{lk}(\sigma, \Sigma_W)$ gives a natural identification of $\mathrm{lk}(\sigma, \Sigma_W)$ with the Coxeter complex for W_σ. These constitute the apartments of $\mathrm{lk}(\sigma, X)$.

For (3), let $M = \mathbb{R}^n$ if X is Euclidean and $M = \mathbb{S}^n$ if X is spherical. Consider the collection of subspaces

$$\mathcal{A} = \{A \subset X | A \text{ is isometric to } M\}.$$

By definition of the metric on X, $\mathcal{A}$ contains all the apartments of X. Since we are assuming the set of apartments to be maximal, it suffices to show that $\mathcal{A}$ satisfies the conditions in Definition 1.1 for a system of apartments.

Observe first that any subspace A isometric to M is necessarily a subcomplex of X since its intersection with any n–simplex σ must be both open and closed in σ. By induction on the dimension of X, we may assume that for any simplex $\sigma \subset A$, $\mathrm{lk}\,(\sigma, A) \hookrightarrow \mathrm{lk}\,(\sigma, X)$ is an isometric embedding. It follows that the embedding $A \hookrightarrow X$ preserves local geodesics. If X is Euclidean, then it is CAT(0), hence local geodesics are geodesics. If X is spherical, then it is CAT(1), hence local geodesics of length $\leq \pi$ are geodesics. In either case, we conclude that $A \hookrightarrow X$ maps geodesics to geodesics, so it is an isometric embedding.

Let $A \in \mathcal{A}$ and let $\sigma \subset A$ be an n–simplex. Fix an isometry α_0 of σ with the fundamental chamber σ_0 of a Coxeter complex Σ_W. This isometry extends uniquely to an isometry $\alpha\colon A \to \Sigma_W$. Since every n–simplex in A is isometric to σ_0, α is also a simplicial isomorphism. Thus, A is a Coxeter complex. Moreover, if $A' \in \mathcal{A}$ also contains σ, and $\alpha'\colon A' \to \Sigma_W$ is an isometry extending α_0, then $\alpha^{-1} \circ \alpha'\colon A \to A'$ is an isometry fixing σ and hence fixing all of $A \cap A'$. Finally, since $\mathcal{A}$ contains a system of apartments, any two simplices of X are contained in some $A \in \mathcal{A}$. Thus, $\mathcal{A}$ satisfies the conditions for a system of apartments. □

3 Spherical buildings

In this section we prove a partial converse to Proposition 2.3. It will form the inductive step to one of the main theorems in Section 7.

Theorem 3.1 *Suppose X is a connected, piecewise spherical cell complex of dimension $n \geq 2$ satisfying*

(1) *X is CAT(1),*

(2) *$\mathrm{lk}\,(x, X)$ is a spherical building for every vertex $x \in X$.*

Then X is a spherical building.

Before embarking on the proof, we make several observations about the hypotheses. First, If σ is any cell in X, then $L_\sigma = \mathrm{lk}\,(\sigma, X)$ is a spherical building. For if v is any vertex of σ, then

$$L_\sigma = \mathrm{lk}\,(\mathrm{lk}\,(v, \sigma), \mathrm{lk}\,(v, X))$$

and $\mathrm{lk}\,(v, X)$ is a spherical building by hypothesis. Thus, it follows from Proposition 2.3(2) that L_σ is a spherical building. Moreover, if x is any point in X, not necessarily a vertex, then $\mathrm{lk}\,(x, X)$ is also a spherical building. For if x lies in the relative interior of a k–cell σ, then $\mathrm{lk}\,(x, X) = \mathbb{S}^{k-1} * L_\sigma = \Sigma^k(L_\sigma)$.

By Proposition 2.3, there is an obvious candidate for a system of apartments for X, namely

$$\mathcal{A} = \{A \subset X | A \text{ is isometric to } \mathbb{S}^n\}.$$

As in the proof of Proposition 2.3(3), it is easy to show that any such subspace A is a subcomplex of X and the inclusion $A \hookrightarrow X$ is an isometric embedding.

The key problem in the proof of Theorem 3.1 is to construct enough of these subcomplexes. The idea is as follows. For any pair of antipodal points (two points are *antipodal* if they have distance π) and any apartment A_x in $\mathrm{lk}\,(x, X)$, we construct an apartment A in X by propagating geodesics from x to y in every direction in A_x.

We begin with a key technical lemma. Some additional notation will be needed for the proof. If $x \in X$ and γ is a geodesic emanating from x, let $\gamma_x \in \mathrm{lk}\,(x, X)$ denote the tangent vector to γ at x. Let $\mathrm{st}\,(x)$ denote the closed star of x, that is, $\mathrm{st}\,(x)$ is the union of the closed simplices containing x. (In the locally spherical context, $\mathrm{st}\,(x)$ will denote a conelike polyhedral neighborhood of x.)

For σ a spherical $(n-1)$–cell, the spherical suspension $\Sigma(\sigma)$, viewed as a subspace of $\mathbb{S}^n$, is called a *spherical sector*. When $n = 2$, it is also called a *spherical lune*.

We prove the next lemma under slightly more general hypotheses for use in the next section. In particular, we do not assume that X is globally CAT(1).

Lemma 3.2 *Suppose X is a locally spherical space of dimension $n \geq 2$ such that the link of every point in X is isometric to a building. Let γ be a local geodesic of length π from x to y and let A_x be an apartment in L_x containing γ_x. Then there is a neighborhood $N_x \subset A_x$ of γ_x and a unique locally isometric map F of the spherical sector $\Sigma(N_x)$ into X such that*

(1) *for any $v \in N_x$, the restriction of F to $\Sigma(v)$ $(= \mathbb{S}^0 * \{v\})$ is a local geodesic from x to y with tangent vector v, and*

(2) *the restriction of F to $\Sigma(\gamma_x)$ is precisely γ.*

Proof Divide γ into segments $\gamma_1, \gamma_2, \ldots, \gamma_k$ with endpoints $x = x_0, x_1, \ldots, x_k$ such that each γ_i lies in $\mathrm{st}\,(x_{i-1})$. Let N_x be an ϵ–ball in A_x centered at γ_x and let S be the spherical sector $S = \Sigma(N_x)$.

For each vector $v \in N_x$, there is a unique geodesic segment γ_1^v in $\mathrm{st}\,(x)$ from x to $\partial(\mathrm{st}\,(x))$ whose tangent at x is v. Let B^1 be the subspace of $\mathrm{st}\,(x)$ consisting of the union of these geodesic segments. Identifying γ_1^v with an initial segment of $\Sigma(v)$ gives an isometry F_1 of a polyhedral subspace of S onto B_1.

Next, consider the $(n-1)$–dimensional building $L_1 = \mathrm{lk}\,(x_1, X)$. The tangent vectors to γ_1 and γ_2 at x_1 form a pair of antipodal points a_1, a_2 in L_1 (since the concatenation $\gamma_1 \cdot \gamma_2$ is geodesic), and $\mathrm{lk}\,(x_1, B_1)$ is a neighborhood of a_1 in L_1 isometric to a spherical $(n-1)$–cell. The union of geodesics in L_1 from a_1 to a_2 with an initial segment lying in this spherical cell forms an apartment $A_1 \subset L_1$. The geodesic segments emanating from x_1 in directions A_1 form a spherical n–cell C in $\mathrm{st}\,(x_1)$. Shrinking the original ϵ–neighborhood N_x if necessary, we may assume that all of the segments γ_1^v end in C. There is then a unique locally geodesic continuation of γ_1^v across C, ending in $\partial(\mathrm{st}\,(x_1))$. Call this new segment γ_2^v. Let B_2 be the union of the local geodesics $\gamma_1^v \cdot \gamma_2^v$, $v \in N_x$. Then F_1 extends in an obvious manner to a local isometry F_2 from a polyhedral subspace of S onto B_2.

We repeat this process at each x_i until we get geodesics $\gamma^v = \gamma_1^v \cdot \gamma_2^v \cdots \gamma_k^v$ for every $v \in N_x$ and a local isometry F from S onto $B_k = \bigcup \gamma^v$ as required. □

Returning to the hypothesis of Theorem 3.1, we can now construct apartments in X.

Lemma 3.3 *Let X be as in Theorem 3.1. Suppose $x, y \in X$ are antipodal points and A_x is an apartment in $L_x = lk\,(x, X)$. Then the following hold.*

1. *For every $v \in A_x$, there exists a unique geodesic γ^v from x to y whose tangent vector at x is v.*
2. *The union of all such γ^v, $v \in A_x$, is isometric to $\mathbb{S}^n$.*

Proof First note that since X is CAT(1), any local geodesic of length $\leq \pi$ is a geodesic. Moreover, if two geodesics from x to y, γ and γ', have the same tangent vectors $\gamma_x = \gamma'_x = v$, then they must agree inside $\mathrm{st}\,(x)$, hence they must agree everywhere. (Otherwise, we get a geodesic digon of length $< 2\pi$.) Thus, geodesics from x to y are uniquely determined by their tangents at x.

Consider the set $C = \{v \in A_x | \gamma^v \text{ exists}\}$. By Lemma 3.2, this set is open in A_x. We claim that it is also closed. To see this, first note that if $v_1, v_2 \in C$ are points of distance α in A_x, then for any $t \in [0, \pi]$, $d(\gamma^{v_1}(t), \gamma^{v_2}(t)) \leq \alpha$. This can be seen by comparing the digon formed by $\gamma^{v_1}, \gamma^{v_2}$ in X with a digon γ'_1, γ'_2 of angle α connecting a pair of antipodal points x', y' in $\mathbb{S}^2$. Inside $\mathrm{st}\,(x)$, these two digons are isometric. That is, for sufficiently small ϵ, the distance from $z_1 = \gamma^{v_1}(\epsilon)$ to $z_2 = \gamma^{v_2}(\epsilon)$ in X is the same as the distance between the corresponding points z_1', z_2' in $\mathbb{S}^2$. Thus, z_1', z_2', y' is a spherical comparison triangle for z_1, z_2, y. It follows from the CAT(1) condition that the distance between $\gamma^{v_1}(t)$ and $\gamma^{v_2}(t)$ in X is less than or equal to the corresponding distance in $\mathbb{S}^2$ for all t. In particular, if (v_i) is a sequence of points in A_x converging to v with $v_i \in C$, then (γ^{v_i}) converges uniformly to a path γ from x to y with $\gamma_x = v$. This path has length π since each γ^{v_i} has length π. Hence γ is geodesic and $v \in C$.

Since C is both open and closed, it is either empty or all of A_x. Since X is a geodesic metric space, there must exist at least one geodesic γ from x to y. If v is any point in L_x, then there exists an apartment A_x containing both γ_x and v. For this apartment, C is nonempty, hence $v \in C$. Thus, there is a geodesic γ^v with tangent vector v as desired. This proves (1).

By Lemma 3.2, we know that the map $F\colon \Sigma(A_x) \to X$ taking $\Sigma(v)$ to γ^v is locally isometric. Since local geodesics of length $\leq \pi$ are geodesic in X, this map is an isometry onto its image. This proves (2). ◻

The spheres constructed in Lemma 3.3 give us a large number of apartments. It is now easy to show that X is a building.

Lemma 3.4 *X has diameter π.*

Proof Since every point in $\mathrm{lk}\,(x, X)$ has an antipodal point, geodesics in X are locally extendible. Since X is CAT(1), any local geodesic of length π is a geodesic. Thus, $\mathrm{diam}(X) \geq \pi$. Suppose there exists a geodesic $\gamma\colon [0, d] \to X$ with $d > \pi$. Let $y = \gamma(0)$ and let $x = \gamma(\pi)$. Let $v \in \mathrm{lk}\,(x, X)$ be the outgoing tangent vector to γ (ie, the tangent vector to $\gamma|_{[\pi,d]}$). By Lemma 3.3, there is a geodesic α from x to y with tangent vector equal to v. But this means that α and $\gamma|_{[\pi,d]}$ agree in a neighborhood of x. This is clearly impossible since the distance from y decreases along α and increases along $\gamma|_{[\pi,d]}$. ◻

Lemma 3.5 *If $A \in \mathcal{A}$, then A is a Coxeter complex Σ_W or a suspension $\mathbb{S}^k * \Sigma_W$.*

Proof Let $A \in \mathcal{A}$. We first show that the $(n-1)$–skeleton of A, A^{n-1}, is a union of geodesic $(n-1)$–spheres which is closed under reflection across each such sphere.

Suppose x is a point in the relative interior of a $(k-1)$–simplex $\sigma \subset A^{n-1}$. Then $\mathrm{lk}\,(\sigma, A)$ is isometric to $\mathbb{S}^{n-k}$, hence it is an apartment in the $(n-k)$–dimensional building $\mathrm{lk}\,(\sigma, X)$. The $(n-k-1)$–skeleton of this apartment is a union of geodesic $(n-k-1)$–spheres closed under reflection. Taking the join with $\mathbb{S}^{k-2} = \mathrm{lk}\,(x,\sigma)$, we see that $\mathrm{lk}\,(x, A^{n-1})$ is a union of geodesic $(n-2)$–spheres in $\mathrm{lk}\,(x, A)$ closed under reflections. It follows that, in a conelike neighborhood of x in A, the $(n-1)$–skeleton consists of a union of geodesic $(n-1)$–disks, $D_1, D_2, \ldots, D_k$. In particular, any geodesic through x which enters this neighborhood through A^{n-1} must also leave through A^{n-1}. Since this is true at every point $x \in A$, we conclude that any geodesic in A containing a non-trivial segment in A^{n-1}, lies entirely in A^{n-1}.

Now let D_i be one of the geodesic disks at x as above. The geodesic segments in D_i emanating from x extend to form a geodesic $(n-1)$–sphere H_i in A which, by the discussion above, lies entirely in A^{n-1}. We call H_i a "wall" through x. Reflection of A across H_i fixes x and permutes the disks $D_1, \ldots, D_k$, hence it permutes the walls through x. Moreover, if H' is a wall through some other point $x' \in A^{n-1}$, then H' must intersect H_i (since they are two geodesic $(n-1)$–spheres in a n–sphere). Say $z \in H_i \cap H'$. Applying the argument above with x replaced by z shows that reflection across H_i takes H' to some other wall through z. Thus, it preserves A^{n-1}.

Let W be the group generated by reflection across the walls of A. By the previous lemma, W acts on A as a group of simplicial isomorphisms. Since A is a finite complex, W is a finite reflection group, or in other words, a spherical Coxeter group. Let A^W be the fixed set of W which consists of the intersection of all the walls. Then A^W is a geodesic k–sphere for some k, and A decomposes as a join, $A = A^W * \Sigma_W$. □

It follows from Lemma 3.5, that cells in X are simplices or suspensions of simplices.

Lemma 3.6 *Any two cells σ_1, σ_2 in X are contained in some $A \in \mathcal{A}$.*

Proof Since any cell is contained in an n–cell, it suffices to prove the lemma for two n–cells. Let x_1, x_2 be points in the interior of σ_1, σ_2, respectively. Let γ be a geodesic from x_1 to x_2 and continue γ to a geodesic of length π. Let y be the endpoint of γ antipodal to x_1 and note that $\mathrm{lk}\,(x_1, X) = \mathrm{lk}\,(x_1, \sigma_1) \cong \mathbb{S}^{n-1}$. It follows from Lemma 3.3 that the union of all geodesics from x_1 to y forms a subspace $A \in \mathcal{A}$. Since A is a subcomplex and contains both x_1 and x_2, it must contain σ_1 and σ_2. □

Lemma 3.7 *If $A_1, A_2 \in \mathcal{A}$ share a common chamber σ, then there is a simplicial isomorphism ϕ: $A_1 \to A_2$ fixing $A_1 \cap A_2$ pointwise. Moreover, $A_1^W = A_2^W$.*

Proof Let x be a point in the relative interior of σ. Let p denote the north pole of $\mathbb{S}^n$ and fix an isometry $\theta\colon \mathrm{lk}\,(x,\sigma) \to \mathbb{S}^{n-1} = \mathrm{lk}\,(p, \mathbb{S}^n)$. Then there is a unique isometry $\phi_i\colon A_i \to \mathbb{S}^n$ with $\phi_i(x) = p$ and the induced map on $\mathrm{lk}\,(x, A_i) = \mathrm{lk}\,(x,\sigma)$ equal to θ. Let $\phi = \phi_2^{-1} \circ \phi_1$. Since the cell structure of A_i is completely determined by reflection in the walls of σ, the isometry ϕ is also a simplicial isomorphism. For any point $y \in A_1 \cap A_2$ not antipodal to x, there is a unique geodesic γ from x to y which necessarily lies in $A_1 \cap A_2$. Since ϕ_1 and ϕ_2 agree on the tangent vector γ_x, they agree on all of γ.

The last statement of the lemma follows from the fact that $A_i^W = \sigma^W$. □

It follows from Lemma 3.7, that X itself decomposes as a join of A^W and a spherical building with Coxeter group W. This completes the proof of Theorem 3.1.

If we are not given an a priori cell structure, we can work in the setting of locally spherical spaces and use the singular set (ie, the branch set) of X to define a cell structure. In this setting we get the following theorem.

Theorem 3.8 *Suppose X is a locally spherical space of dimension $n \geq 2$ satisfying*

(1) *X is CAT(1),*

(2) *$lk\,(x, X)$ is isometric to a building for every point $x \in X$.*

Then X is isometric to a spherical building. The cell structure determined by the singular set is that of a thick, spherical building or a suspension of a thick, spherical building.

4 More on spherical buildings

In contrast to the locally Euclidean case, a locally spherical space which is simply connected and locally CAT(1) need not be globally CAT(1). However, as we now show, under the stronger hypothesis that links are isometric to buildings, a simply connected locally spherical space of dimension ≥ 3 is CAT(1), and hence is also isometric to a spherical building.

Theorem 4.1 *Suppose X is a locally spherical space of dimension $n \geq 3$ satisfying*

(1) *X is simply connected,*

(2) *$lk\,(x, X)$ is isometric to a building for every point $x \in X$.*

Then X is CAT(1), hence it is isometric to a building.

The hypothesis that $n \geq 3$ is essential here. In the 1980's there was much interest in the the relation between incidence geometries and buildings (see for example [14] and [10]). In [14], Tits proves a theorem analogous to Theorem 4.1 for incidence geometries, with the same dimension hypothesis. A counterexample in dimension $n = 2$ is given by Neumaier in [12]. It is a finite incidence geometry of type C_3, with a transitive action of A_7 (the alternating group on 7 letters). The flag complex associated to Neumaier's A_7–incidence geometry is a 2–dimensional simplicial complex, all of whose links are buildings, but which cannot be covered by a building. One can metrize Neumaier's example by assigning every 2–simplex the metric of a spherical triangle with angles $\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{4}$. Passing to the universal cover gives a counterexample to Theorem 4.1 in dimension $n = 2$.

Before proving the theorem, we will need some preliminary lemmas. We begin with an easy consequence of Lemma 3.2.

Lemma 4.2 *Let X be as in Theorem 4.1 and let γ be a local geodesic in X of length π from x to y. Then there is a unique locally isometric extension $F_\gamma\colon \Sigma(L_x) \to X$ of γ (in the sense of Lemma 3.2, (1) and (2)).*

Proof Suppose A is an apartment in L_x containing the tangent vector $v = \gamma_x$. Then by Lemma 3.2, there is a neighborhood U of v in A and a unique locally isometric map $F_U\colon \Sigma(U) \to X$ whose restriction to $\Sigma(v)$ is γ. Using the maps F_U, we can extend γ uniquely along any geodesic in L_x beginning at v. Since L_x is simply connected for $n \geq 3$, these extensions are compatible. □

Suppose γ_1 and γ_2 are two local geodesics of length π from x to y and let $v_i = (\gamma_i)_x$. Then it follows from the construction of F_{γ_i} that the following are equivalent.

(1) $F_{\gamma_1} = F_{\gamma_2}$.

(2) γ_2 is the restriction of F_{γ_1} to $\Sigma(v_2)$ (and vice versa).

(3) There exists a locally isometric map of a spherical lune into X with sides γ_1 and γ_2.

We say that a geodesic $\eta\colon [0,a] \to X$ from x is *nonbranching* if any other geodesic $\eta'\colon [0,a] \to X$ from x with $\eta_x = \eta'_x$ is equal to η. (Or in other words, η has unique continuation at every point in its interior.) In particular, if η is contained in a cone-like neighborhood of x, then it is non-branching.

The following is an immediate consequence of Lemma 4.2.

Lemma 4.3 *Suppose γ and η are geodesics of length $\leq \pi$ starting at x and assume η is nonbranching. Then there is a locally isometric map of a spherical triangle into X (possibly a geodesic or a spherical lune) which restricts on two sides to γ and η.*

The local isometry in the corollary above is essentially unique. More precisely, we have the following.

Lemma 4.4 *Let T_1 and T_2 be spherical triangles and $\Theta_1\colon T_1 \to X$ and $\Theta_2\colon T_2 \to X$ be local isometries. If Θ_1 and Θ_2 agree along two edges of the triangle, then one of the following holds.*

(1) *$T_1 = T_2$ (ie, they are isometric) and $\Theta_1 = \Theta_2$,*

(2) *T_1 and T_2 are hemispheres and Θ_1, Θ_2 agree along their entire boundary.*

(3) *T_1 and T_2 are spherical lunes and the two edges along which they agree form one entire side of the lune.*

Proof By hypothesis, Θ_1 and Θ_2 restrict along two edges to local geodesics γ and η emanating from some point x. The angle between these two edges is the distance in L_x between γ_x and η_x. Suppose this angle less than π. Then clearly $T_1 = T_2$. Since X is locally CAT(1), the subspace of T_1 on which $\Theta_1 = \Theta_2$ must be locally convex, and hence must be all of T_1.

If the angle between the edges is exactly π, then T_i is either a geodesic (if $\text{length}\,(\gamma) + \text{length}\,(\eta) < \pi$), a spherical lune (if $\text{length}\,(\gamma) + \text{length}\,(\eta) = \pi$), or a hemisphere (if $\text{length}\,(\gamma) + \text{length}\,(\eta) > \pi$). In the last case, we may assume without loss of generality that $\text{length}\,(\gamma) = \pi$. then Θ_1 and Θ_2 are both restrictions of $F_\gamma\colon \Sigma L_x \to X$. In particular, they agree along $(\Sigma\gamma_x) \cup (\Sigma\eta_x)$ which forms the boundary of T_i. □

Lemma 4.5 *The diameter of X is at most π.*

Proof The proof is the same as that of Lemma 3.4 (using Lemma 4.2 in place of Lemma 3.3). □

Proof of Theorem 4.1 Fix a point x in X. Define an equivalence relation on the set of geodesics of length π starting at x by

$$\gamma_1 \sim \gamma_2 \iff F_{\gamma_1} = F_{\gamma_2}$$

To prove Theorem 4.1, we define a covering space $f\colon \tilde{X} \to X$ as follows. As a set, $\tilde{X}$ is defined as the quotient

$$\tilde{X} = \{\gamma \mid \gamma \text{ is a local geodesic of length } \leq \pi \text{ with } \gamma(0) = x\}/\sim$$

Note that only local geodesics of length π can be identified in $\tilde{X}$.

The topology on $\tilde{X}$ is defined as follows. Let $B_r(y)$ denote the ball of radius r in X centered at y. Given a local geodesic γ from x to y and a real number r such that $B_r(y)$ is conelike, define

$$\begin{aligned} B_r(\gamma) = \{\eta \in \tilde{X} \mid \exists \text{ a locally isometric map of a spherical triangle} \\ \text{into } X \text{ which restricts on two sides to } \gamma \text{ and } \eta \\ \text{and whose third side lies in } B_r(y)\}. \end{aligned}$$

If γ has length π, then these locally isometric maps are all restrictions of F_γ. In particular, $B_r(\gamma)$ depends only on the class of γ in $\tilde{X}$. These sets form a basis for the topology on $\tilde{X}$. They also define a metric (locally) on $\tilde{X}$. Namely, the distance between γ and $\eta \in B_r(\gamma)$ is the length of the third side of the triangle.

Define $f\colon \tilde{X} \to X$ to be the map taking γ to its endpoint. By Lemma 4.3, f restricts to an isometry of $B_r(\gamma)$ onto $B_r(y)$. Letting γ run over all local geodesics from x to y, we claim that these balls make up the entire inverse image of $B_r(y)$. For suppose $\eta \in \tilde{X}$ with $f(\eta) = z \in B_r(y)$. Since z lies in a conelike neighborhood of y, the geodesic δ from z to y is nonbranching. It follows from Lemma 4.3, applied to η and δ, that there exists a local isometry of a spherical triangle into X which restricts on two side to η and δ. The restriction to the third side, γ, is a local geodesic from x to y such that $\eta \in B_r(\gamma)$.

It remains to show that for distinct $\gamma_i \in \tilde{X}$, the balls $B_r(\gamma_i)$ are disjoint. Suppose $\delta \in B_r(\gamma_1) \cap B_r(\gamma_2)$. Then there is a local isometry Θ_i of a spherical triangle with sides γ_i and δ, for $i = 1, 2$. Since $B_r(y)$ is conelike, there is a unique geodesic η from y to the endpoint z of δ. Thus Θ_1 and Θ_2 agree along two edges, δ and η. It follows from Lemma 4.4 that $\gamma_1 = \gamma_2$ in $\tilde{X}$.

This proves that $f\colon \tilde{X} \to X$ is a covering map. By hypothesis, X is simply connected, and it is easy to verify that $\tilde{X}$ is connected, thus f is injective. In particular, for any $y \in X$ of distance less than π from x, there is a unique local geodesic from x to y. Moreover, it follows from Lemma 4.3 that this geodesic varies continuously with the endpoint y. Since x was chosen arbitrarily, this applies to all x and y. By Theorem 2.2(2), we conclude that X is CAT(1). □

5 Euclidean buildings

Theorem 5.1 *Suppose X is a connected, locally Euclidean complex satisfying*

(1) *X is CAT(0),*

(2) *for every point $x \in X$, $L_x = lk(x, X)$ is isometric to a spherical building.*

Then X decomposes as an orthogonal product $X = \mathbb{R}^l \times X_1 \times \cdots \times X_k$, where $l \geq 0$, and each X_i is one of the following,

(1) *a thick, irreducible Euclidean building,*

(2) *the Euclidean cone on a thick, irreducible spherical building,*

(3) *a tree.*

Remark The reader may object that a tree is a 1–dimensional irreducible Euclidean building whose apartments are Coxeter complexes for the infinite dihedral group. However, the standard building metric on a 1–dimensional Euclidean building would assign the same length to every edge of the tree. Since this need not be the case in our situation, we list these factors separately.

In [11], Kleiner and Leeb introduce a more general notion of a Euclidean building, which we will call a "metric Euclidean building", and prove an analogous product decomposition theorem (Prop. 4.9.2) for these buildings. We review their definition in the context of locally Euclidean spaces. (Kleiner and Leeb work in a more general setting.)

Call a group W of affine transformations of $\mathbb{R}^n$ an *affine Weyl group* if W is generated by reflections and the induced group of isometries on the sphere at infinity is finite. Affine Weyl groups include Euclidean and spherical Coxeter groups, as well as nondiscrete groups generated by reflections across parallel walls.

Let $\mathcal{A}$ be a collection of isometric embeddings of $\mathbb{R}^n$ into a locally Euclidean space X of dimension n. We call $\mathcal{A}$ an *atlas* for X and the images of the embeddings are called *apartments*.

Definition 5.2 Suppose X is a CAT(0), locally Euclidean space of dimension n. Then X is a *metric Euclidean Building* if there is an atlas $\mathcal{A}$ and an affine Weyl group W such that

(1) Every geodesic segment, ray, and line is contained in an apartment.

(2) $\mathcal{A}$ is closed under precomposition with W.

(3) If two apartments $\phi_2(\mathbb{R}^n), \phi_2(\mathbb{R}^n)$ intersect, then $\phi_1^{-1} \circ \phi_2$ is the restriction of some element of W.

(In the context of locally Euclidean spaces, this definition agrees with that of Kleiner and Leeb since their first two axioms hold automatically for locally Euclidean spaces.) It is an immediate consequence of Theorem 5.1 that the space X is a metric Euclidean building.

Corollary 5.3 *Let X be as in Theorem 5.1. Then X is a metric Euclidean building.*

Conversely, it is easy to see that a metric Euclidean building satisfies the hypotheses of Theorem 5.1. Thus, for locally Euclidean spaces, Theorem 5.1 also provides another proof of Kleiner and Leeb's product decomposition theorem.

The proof of Theorem 5.1 will occupy the remainder of this section. As in the spherical case, the key is to find lots of apartments. By Proposition 2.3, the apartments in X are isometrically embedded copies of $\mathbb{R}^n$, known as n–*flats*. The crucial step to constructing n–flats is to find *flat strips* (isometrically embedded copies of $\mathbb{R} \times [0, a]$ for some $a > 0$).

Definition 5.4 Let X be a $CAT(0)$–space. We will call the triangle Δ in X a *Euclidean (or flat) triangle*, if its convex hull is isometric to a triangle in $\mathbb{R}^n$.

For a triangle $\Delta(x, y, z)$ we denote the segment from x to y by xy, etc. The angle between xy and xz is defined as the distance in L_x between the tangent vectors to xy and xz and it is denoted by $\angle_x(xy, xz)$. The following lemma follows immediately from Proposition 3.13 of [1].

Lemma 5.5 *Let X be a $CAT(0)$ space, $\Delta = \Delta(a, b, c)$ a triangle in X. Let d be a point between a and b. Suppose the triangles $\Delta(a, d, c)$ and $\Delta(b, d, c)$ are Euclidean. If in addition $\angle_d(da, dc) + \angle_d(db, dc) = \pi$, then the original triangle $\Delta(a, b, c)$ is Euclidean.*

The condition on the angles is automatically satisfied if the geodesic ab is non-branching, for example if b lies in a cone-like neighborhood of a.

If L is a locally spherical space and $r > 0$, let $C_r(L)$ denote the Euclidean cone on L of radius r (ie, the geodesics emanating from the cone point have length r). The following lemma is an analogue of Lemma 3.2. The proof is essentially the same and the details are left to the reader.

Lemma 5.6 *Suppose X is a locally Euclidean space of dimension $n \geq 2$ such that the link of every point in X is isometric to a spherical building. Let γ be a locally geodesic ray from x and let A_x be an apartment in L_x containing γ_x. Then for any $r > 0$, there is a neighborhood $N_x \subset A_x$ of γ_x and a unique locally isometric map Θ of the Euclidean cone $C_r(N_x)$ into X such that*

(1) *for any $v \in N_x$, the restriction of Θ to $C_r(v)$ is a local geodesic with tangent vector v, and*

(2) *the restriction of Θ to $C_r(\gamma_x)$ is precisely $\gamma|_{[0,r]}$.*

Suppose γ is as in Lemma 5.6, and α is a geodesic in L_x originating at γ_x. Then Lemma 5.6 implies that there is a unique extension of $\gamma|_{[0,r]}$ to a locally isometric map of $C_r(\alpha)$ into X. If X is CAT(0), the map is an isometric embedding. This enables us to construct Euclidean triangles in X since the image of any triangle in $C_r(\alpha)$ is Euclidean.

From now on, we assume that X satisfies the hypotheses of Theorem 5.1.

Lemma 5.7 *Let $\gamma\colon \mathbb{R} \to X$ be a geodesic ray with $x = \gamma(0)$. If y lies in a conelike neighborhood of x, then for any $t \in \mathbb{R}$, the triangle $\Delta(x, y, \gamma(t))$ is Euclidean.*

Proof Assume, without loss of generality, that $t > 0$. Since y lies in a conelike neighborhood of x, the geodesic η from x to y is non-branching. Choose a geodesic α in L_x from γ_x to η_x and extend $\gamma|_{[0,t]}$ to an isometric embedding Θ of $C_t(\alpha)$ into X. Since η is non-branching, it agrees with Θ on $C_t(\eta_x)$. Thus, $x, y, \gamma(t)$ span a Euclidean triangle. □

Lemma 5.8 *Let γ, x, and y be as above and let η be the geodesic from x to y. Suppose $\angle_x(\eta, \gamma^+) + \angle_x(\eta, \gamma^-) = \pi$. Then y and γ span a Euclidean strip.*

Proof By the previous lemma, for each $t \in \mathbb{R}$, the triangle $\triangle(x, y, \gamma(t))$ is Euclidean. By Lemma 5.5, every triangle of the form $\triangle(y, \gamma(t_1), \gamma(t_2))$ is Euclidean. Since any two points in the span of y and γ lie on such a triangle, the lemma follows. □

For any subspace Y of X, and any point $x \in Y$, we denote the link of x in Y by $L_x Y$.

Lemma 5.9 *Let F be an m–flat in X. Let $\eta\colon [0, r] \to X$ be a geodesic from $x \in F$ to a point y lying in a conelike neighborhood of x. Suppose that the distance in L_x from η_x to any point in $L_x F$ is $\frac{\pi}{2}$. Then y and F span a flat $R^m \times [0, r]$.*

Proof Let $Z = \mathbb{R}^m \times [0,r] = F \times [0,r]$. Let Y be the subspace of X spanned by F and y. Consider the natural map $f\colon Z \to Y$ that takes $F \times \{0\}$ via the identity map to F and takes the line segment between $(z,0)$ and $(0,r)$ to the geodesic in Y from z to y. By Lemma 5.8, the restriction of f to the strip $\gamma \times [0,r]$ is an isometry for any geodesic $\gamma\colon \mathbb{R} \to F$ through x. We must show that f is isometric on all of Z.

Any two points y_1, y_2 in Y lie on the image of a triangle $T \subset Z$ with vertices $(0,r), (z_1,0), (z_2,0)$. By the discussion above, T is a comparison triangle for its image $f(T)$ in Y. Hence, by the CAT(0) condition, the distance between y_1 and y_2 is at most the distance between the corresponding points in T. Thus, f is distance non-increasing. Moreover, if y_1 or y_2 lies on η, then they both lie in a Euclidean strip, as above, so these two distances agree.

To prove the reverse inequality, choose $r' < r$ and let $y' = \eta(r') = f((0,r'))$ (Figure 1). Consider the induced map *df* between the links $L_{(0,r')}Z \cong \mathbb{S}^m$ and $L_{y'}Y$. It suffices to prove that *df* is an isometry, for in this case, the triangle with vertices y', y_1, y_2 has the same angle and same two side lengths at y' as its comparison triangle in Z, so by the CAT(0) condition, the opposite side is at least as long as in the comparison triangle.

To see that *df* is an isometry, note that the fact that f restricts to an isometry on strips $\gamma \times [0,r]$ implies that *df* maps points of distance π in $L_{(0,r')}Z$ to points at distance π in $L_{y'}Y$. On the other hand, since f is distance non-increasing at all points, it must also be distance non-increasing on links. But these two facts contradict each other unless *df* is an isometry. □

Lemma 5.10 *Every geodesic and every flat strip in X is contained in an n–flat.*

Proof Let F be an m–flat in X with $m < n$. Let $a \in [-\infty, 0]$, $b \in [0,\infty]$ be chosen so that F is contained in a flat embedded $F \times [a,b]$ (with $F = F \times \{0\}$) and such that a, b are maximal, (that is, one cannot embed this $F \times [a,b]$ in a bigger $F \times [a_0, b_0]$).

We claim that $a = -\infty, b = \infty$. Assume the contrary. Say, for example, $b < \infty$. Let x be a point in $F \times \{b\}$. $F \times [a,b]$ determines an m–dimensional hemisphere H in L_xX containing L_xF. (If $a = b = 0$, choose any hemisphere containing L_xF.) Choose an m–sphere in L_xX containing H and let v be the point of the sphere, which has distance $\frac{\pi}{2}$ from H. Let η be a ray emanating from x in the direction of v. By the previous lemma, for small r, $\eta(r)$ and $F \times \{b\}$ span a flat $F \times [0,r]$. The choice of v, together with Lemma 5.5, guarantee that $F \times [a,b]$ and $F \times [0,r]$ fit together to form a flat strip $F \times [a, b+r]$. This contradicts the maximality of b.

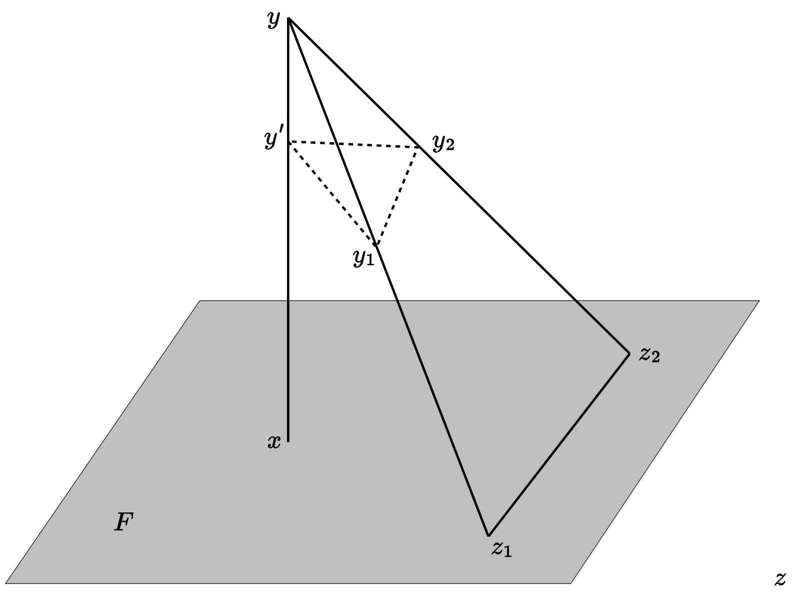

Figure 1

Thus, any m–flat, and more generally, any geodesically embedded $\mathbb{R}^m \times [a,b]$, $m < n$, can be embedded in an $(m+1)$–flat. It follows inductively, that every geodesic and every flat strip is contained in an n–flat. □

Proof of Theorem 5.1 Let F be an n–flat in X. The set Y of the singular points in F is closed and locally it is a union of hyperplanes, so Y is globally a union of hyperplanes too. We call these the singular hyperplanes. The set of singular hyperplanes is locally invariant under reflection in each of these hyperplanes, so if two singular hyperplanes H_1 and H_2 are not parallel, the reflection, H_1', of H_1 across H_2 is also singular. Moreover, if H_1 and H_2 are parallel, and there exists a singular hyperplane not orthogonal to H_1 and H_2, then a simple exercise shows that H_1' can be obtained by a series of reflections across intersecting singular hyperplanes. Thus, again, H_1' must be a singular hyperplane.

Let $Y_1, Y_2, \ldots, Y_k$ be a maximal decomposition of Y into mutually orthogonal families of singular hyperplanes. It follows from the discussion above that each Y_i is either a discrete family of parallel hyperplanes or the set of walls of an irreducible spherical or Euclidean Coxeter group (compare [6], VI.1).

Taking F_i to be the subspace of F generated by the normal vectors to the hyperplanes in Y_i, we obtain an orthogonal decomposition of F into $F = F_0 \times F_1 \times \cdots \times F_k$, where F_0 is a subspace parallel to all of the singular hyperplanes. If Y_i is a family of parallel hyperplanes, then F_i is 1–dimensional. Otherwise, Y_i gives F_i the structure of a Euclidean Coxeter complex or an (infinite) Euclidean cone on a spherical Coxeter complex.

We now choose an n–flat E such that, in the decomposition $E = E_0 \times E_1 \times \cdots \times E_k$, the dimension l of the factor E_0 is minimal, and among all flats with $\dim(E_0) = l$, we require that E minimize the number k. (As we will see below, the number k is actually independent of the choice of E.) In this flat E, we choose a point x, such that for $j = 1, \ldots, k$, the hyperplanes Y_j determine a full-dimensional, irreducible, spherical Coxeter group in the sphere L_xE_j. Intuitively, we have chosen E to be the most complicated n–flat and x to be the most complicated point in E.

The product decomposition of E gives rise to a decomposition of L_xE as an orthogonal join, $L_xE = L_xE_0 * L_xE_1 * \cdots * L_xE_k$. Since L_xE is an apartment in the spherical building L_xX, it follows that L_xX splits as a join of k irreducible buildings and an $(l-1)$–sphere, $L_xX = L_0 * L_1 * \cdots * L_k$.

Let G be another flat through x. As above, we can factor G as an orthogonal product $G = G_0 \times G_1 \times \cdots \times G_{k'}$ Since L_xG is also an apartment in L_xX, it must have the same simplicial structure as L_xE. Clearly this is possible only if $\dim(G_0) = \dim(E_0)$, $k' = k$ and (up to permutation) $L_xG_j \cong L_xE_j$ for all j. (Note that although we have shown that the simplicial structure of G_j and E_j agree in a neighborhood of x, we do not yet know that they agree globally.)

If y is any other point in X, then by Lemma 5.10, there is a flat G containing x and y, so the decomposition of G gives rise to a decomposition of L_yX as a join $L_{y,0} * L_{y,1} * \cdots * L_{y,k}$ of spherical buildings. In this case, however, the factors need not be irreducible. (Consider, for example, the case where y is a nonsingular point.)

Next, we prove that if the decomposition of E is not the trivial one, we can decompose X as a product. Let X_i be the union of the geodesic rays emanating from x in the direction of L_i. Then X_i is connected and by Lemma 5.7, it is locally convex, hence it is convex. To prove that $X = X_0 \times X_1 \times \cdots \times X_k$, we apply Theorem II.9.24 of [4] which states that splittings of X as a product correspond to splittings of the Tits boundary $\partial_T X$ as a join. The Tits boundary may be viewed as the set of rays emanating from x. To show that $X = X_0 \times X_1 \times \cdots \times X_k$, it suffices to show that $\partial_T X = \partial_T X_0 * \partial_T X_1 * \cdots * \partial_T X_k$. For this we must verify (cf [4] Lemma II.9.25):

(1) Every ray γ lies in the span of rays $\gamma_0, \gamma_1, \ldots, \gamma_k$ with $\gamma_i \subset X_i$.

(2) Rays γ_i, γ_j lying in distinct factors X_i, X_j have distance $\frac{\pi}{2}$ in the Tits metric.

The first of these conditions follows from the fact that any ray γ at x is contained in an n–flat G which, by the discussion above, decomposes as a product $G = G_0 \times G_1 \times \cdots \times G_k$ with $G_i \subset X_i$.

For the second condition, it is enough to show, that any geodesic γ in X_i through x and any ray η in X_j emanating from x $(i \neq j)$ span a flat halfplane. If not, then there is a largest t such that $\eta(t)$ and γ span a flat strip. By Lemma 5.8, we can assume that t is bigger than 0. By Lemma 5.10, we can embed this strip in an n–flat T. Let $z = \eta(t)$ and let γ_1 be the line in T parallel to γ through z. The product decomposition of T gives rise to a join decomposition of the building $L_zX = L_{z,0} * L_{z,1} * \cdots * L_{z,k}$. The tangent vectors to γ_1 at z lie in $L_{z,i}$ while the tangent vectors to η at z lie in $L_{z,j}$. It follows that the angle condition of Lemma 5.8 is satisfied, so that for small ϵ, $\eta(\epsilon + t)$ and γ_1 span a flat strip. This strip fits together with the strip between γ_1 and γ, contradicting the maximality of t. This proves condition (2) and we have shown that X decomposes as an orthogonal product of the X_i's.

It remains to identify the factors X_i. Since X_i is convex in X, it is CAT(0), and at any point $y \in X_i$, $L_yX_i = L_{y,i}$ is a building. Thus, each factor satisfies the hypotheses of Theorem 5.1. The factor X_0 contains no singular points, for if $y \in X_0$ is singular, then we could find an n–flat F through y with $dim F_0 < dim E_0$. In other words, $X_0 = \mathbb{R}^l$.

Assume $i > 0$. If X_i is 1–dimensional, then the CAT(0) condition implies that it is a tree. If $\dim X_i \geq 2$, then X_i contains an apartment E_i which is either an irreducible Euclidean Coxeter complex or the cone on an irreducible spherical Coxeter complex. Thus, chambers in E_i are simplices or cones on simplices, and the simplicial structure of E_i is completely determined by a single chamber. It follows that any other apartment in X_i sharing a chamber with E_i has the same simplicial structure. By Lemma 5.10, any two points (hence any two chambers) in X_i lie in a common apartment, so any two apartments are isomorphic. If E_i is a Euclidean Coxeter complex, then we conclude that X_i is an irreducible Euclidean building. In the case that E_i is the cone on a spherical Coxeter complex, it has only one vertex, namely x, which must be the cone point for every apartment. Thus, in this case, X_i is the cone on the irreducible spherical building L_i. In either case, since the simplicial structure on X_i is defined by its singular set, these buildings must be thick. This completes the proof of Theorem 5.1. □

It is reasonable to ask whether an analogous theorem holds for locally hyperbolic spaces; that is, if X is locally hyperbolic and CAT(-1), and every link in X is isometric to a spherical building, can we conclude that X is a hyperbolic building (ie, a buildings whose apartments are copies of $\mathbb{H}^n$ cellulated by the walls of a discrete, hyperbolic reflection group)? The answer is no. Although one can construct lots of embedded copies of $\mathbb{H}^n$ under these hypotheses, these "apartments" need not have the structure of a Coxeter complex. For example, let $\gamma_1, \dots \gamma_5$ be five geodesics $\gamma_i\colon \mathbb{R} \to \mathbb{H}^2$ which form a right-angled pentagon. The reflections across these lines generate a hyperbloic Coxeter group. Now let γ_6 be another geodesic line intersecting γ_1 orthogonally and disjoint from the other γ_i's. We can choose γ_6 so that the three geodesics intersecting γ_1 generate a nondiscrete group of reflections. Then $\gamma_1, \dots, \gamma_6$ cannot be the walls of a reflection group acting on $\mathbb{H}^2$. (Note that these six lines are locally, but not globally closed under reflections.) It is now possible to construct a simply connected, locally hyperbolic space X, all of whose links are spherical buildings (in fact we can take the link of every vertex to be the $K_{3,3}$ graph with edge lengths $\frac{\pi}{2}$) and such that X contains an isometrically embedded $\mathbb{H}^2$ whose singular set consists of the six lines above. This cannot be a hyperbolic building.

6 One-dimensional spherical buildings

In this section we give a metric characterization of one-dimensional spherical buildings. An equivalent characterization appears as Lemma 6.1 in [3]. We include a proof here for completeness.

Theorem 6.1 *Suppose X is a connected, one-dimensional piecewise spherical complex satisfying*

(1) *X is CAT(1) and $diam(X) = \pi$,*

(2) *every vertex of X has valence ≥ 3.*

Then X is either a thick spherical building or $X = \Sigma Y$ for some discrete set Y.

Remark We can also apply this theorem in cases where every vertex has valence ≥ 2 simply by ignoring those vertices of valence 2. In this case, however, some information about the original cell structure will be lost since vertices of valence 2 are invisible to the metric.

Assume throughout this section that X satisfies the hypotheses of Theorem 6.1. In this situation, the analogue of Lemma 3.3 is easy to prove.

Lemma 6.2 *Let $x, y \in X$ be antipodal points and $v \in lk(x, X)$ a tangent vector at x. Then there exits a unique geodesic γ from x to y such that $\gamma_x = v$.*

Proof Choose a point z in the open star of x so that the geodesic α from x to z satisfies $\alpha_x = v$. We first observe that z cannot be antipodal to y. For if w is a point on α between x and z, then the geodesic from w to y must pass through either z or x. If both z and x are antipodal to y, then $d(w,y) > \pi$, contradicting $\operatorname{diam}(X) = \pi$. Thus, $d(x,z) < \pi$ and hence there is a unique geodesic β from z to y. This geodesic cannot pass through x (since this would imply $d(z,y) > d(z,x) = \pi$), so it must leave z in the opposite direction from α. Thus, the concatenation $\alpha \cdot \beta$ is a local geodesic from x to y. Moreover, since α contains no antipodal points to y (other than x), the length of $\alpha \cdot \beta$ cannot be more than π. Thus, $\gamma = \alpha \cdot \beta$ is the desired geodesic. □

Lemma 6.3 *Suppose $x, y \in X$ are antipodal points. If x is a vertex of X, then so is y.*

Proof If x is a vertex of X, then $\operatorname{lk}(x,X)$ contains at least 3 distinct points v_1, v_2, v_3. By the previous lemma, there are geodesics $\gamma_1, \gamma_2, \gamma_3$ from x to y with $(\gamma_i)_x = v_i$. Any two of these geodesics intersect only at x and y since, otherwise, they would form a circuit of length $< 2\pi$. Thus, $\operatorname{lk}(y,X)$ has at least 3 distinct points, so y is a vertex. □

Lemma 6.4 *All edges of X have the same length, $\frac{\pi}{m}$ for some integer $m \geq 1$.*

Proof Since X is connected, if any two edges have different lengths, then we can find a pair of adjacent edges, e_1, e_2 of different lengths. Let x be their common vertex and let x_1, x_2 be the other endpoints of e_1, e_2 respectively. Say $d(x,x_1) < d(x,x_2)$. Choose a point y antipodal to x and a third edge e_3 emanating from x. By Lemma 6.2, there are geodesics $\gamma_1, \gamma_2, \gamma_3$ from x to y which begin with e_1, e_2, e_3 respectively. Any two of these geodesics, γ_i, γ_j, form a loop of length 2π. Consider the loop $\gamma_1 \cdot \gamma_3$. Let z be the point on γ_3 antipodal x_1. By Lemma 6.3, z is a vertex. Next, consider the loop $\gamma_2 \cdot \gamma_3$. Let z' be the point on γ_2 antipodal to z. By Lemma 6.3, z' is a vertex. But $d(x,z') = d(y,z) = d(x,x_1) < d(x,x_2)$. Thus, z' lies in the interior of the edge e_2. This is a contradiction. We conclude that all edges of X have the same length. Since antipodal points to vertices are also vertices, this length must be $\frac{\pi}{m}$ for some integer $m \geq 1$. □

If all edges of X have length π, then the distance between any two vertices of X is exactly π. It follows that X has only two vertices (otherwise, it would have diameter $> \pi$). Thus, X is the suspension of the discrete set Y consisting of the midpoints of the edges.

From now on, we will assume that every edge has length $\frac{\pi}{m}$ for some $m \geq 2$. In this case, every circuit of length 2π is a Coxeter complex for the dihedral group D_{2m}. We will refer to such circuits as apartments.

Lemma 6.5 *Any two simplices of X are contained in a common apartment. If A_1 and A_2 are apartments containing a common edge e, then there is a simplicial isomorphism ϕ: $A_1 \to A_2$ fixing $A_1 \cap A_2$ pointwise.*

Proof It suffices to prove the first statement for two edges e_1, e_2. Choose a point x in the interior of e_1. Suppose e_2 contains a point y antipodal to x (which, by Lemma 6.3 must lie in the interior of e_2). Then $\mathrm{lk}\,(x, e_1)$ consists of two points v_1, v_2 each of which gives rise to a geodesic γ_1, γ_2 from x to y. Together, these two geodesics form an apartment containing e_1 and e_2.

Suppose, on the other hand, that no point of e_2 is antipodal to x. Let γ_1 be a geodesic from x to a point y in the interior of e_2. Extend γ_1 to a geodesic of length π. Note that this extended geodesic (which we still denote γ_1) must contain all of e_2. Let z be the endpoint of γ_1, so z is antipodal to x. Then there exists a geodesic γ_2 from x to z which begins along e_1 in the direction opposite to γ_1. The union, $\gamma_1 \cup \gamma_2$ is the desired apartment.

The second statement is obvious since the CAT(1) hypothesis implies that $A_1 \cap A_2$ is connected. □

7 Main theorems

Combining the results of the previous sections, we arrive at our main theorems.

Theorem 7.1 *Suppose X is a connected, piecewise spherical complex of dimension $n \geq 2$ satisfying the following conditions.*

(1) *X is CAT(1).*

(2) *Every $(n-1)$–cell is contained in at least two n-cells.*

(3) *The link of every k–cell, $k \leq n-2$, is connected.*

(4) *The link of every $(n-2)$–cell has diameter π.*

Then X is isometric to a spherical building. The cell structure determined by the singular set is that of a thick spherical building or a suspension of a thick spherical building.

Proof Let $L_\sigma = \mathrm{lk}\,(\sigma, X)$. By Theorem 2.2, if X is CAT(1), then so is L_σ for every cell σ. In particular, for an $(n-2)$–cell σ, L_σ is a 1–dimensional, piecewise spherical complex which is CAT(1), diameter π, and has every vertex of valence ≥ 2. Ignoring vertices of valence 2 gives a complex satisfying Theorem 6.1. Thus, L_σ is isometric to a 1-dimensional spherical building. The theorem now follows by induction from Theorem 3.8. □

In the theorem above, we could have assumed that X was locally spherical instead of piecewise spherical if we interpret "k–cell" as meaning k–dimensional strata of the singular set. In fact, in ignoring vertices of valence 2, some information about the given cell structure on X may be lost. For example, *any* cell decomposition of the standard 2–sphere satisfies the conditions of the theorem, but need not be the cell structure of a building. To guarantee that the original cell structure is reflected in the metric, we would need to assume that X is thick.

Theorem 7.2 *Suppose X is a connected, piecewise spherical complex of dimension $n \geq 2$ satisfying the following conditions.*

(1) *X is CAT(1).*

(2) *Every $(n-1)$–cell is contained in at least three n–cells.*

(3) *The link of every k–cell, $k \leq n-2$, is connected.*

(4) *The link of every $(n-2)$–cell has diameter π.*

Then, with respect to its given cell structure, X is a spherical building.

Proof The proof is the same using Theorem 3.1 instead of Theorem 3.8. □

Analogous results hold in the piecewise Euclidean setting.

Theorem 7.3 *Suppose X is a connected, piecewise Euclidean complex of dimension $n \geq 2$ satisfying the following conditions.*

(1) *X is CAT(0).*

(2) *Every $(n-1)$–cell is contained in at least two n–cells.*

(3) *The link of every k–cell, $k \leq n-2$, is connected.*

(4) *The link of every $(n-2)$–cell has diameter π.*

Then X is a metric Euclidean building. If, in addition, X is thick, then X is a product of irreducible Euclidean buildings and trees with respect to its given cell structure.

Proof The first hypothesis implies that $\mathrm{lk}(\sigma, X)$ is CAT(1) for every σ. By Theorem 7.1, we conclude that for v a vertex of X, $\mathrm{lk}(v, X)$ is isometric to a spherical building. It follows by Theorem 5.1 and Corollary 5.3 that X is a metric Euclidean building and that it factors as a product of irreducible Euclidean buildings, cones on spherical buildings, trees, and a nonsingular Euclidean space. If X is thick, the components of the nonsingular set of X are the interiors of the n–cells. In particular, they are bounded. Thus, only the first and last type of factor can occur in this situation. □

Replacing the CAT(1) condition in Theorem 7.1 by a simply connectedness condition does not give a satisfying characterization because of the problem in dimension 2. It does, however, give an analogue of Tit's theorem about incidence geometries ([14], Theorem 1).

Theorem 7.4 *Suppose X is a piecewise spherical (respectively Euclidean) complex of dimension $n \geq 3$ satisfying*

(1) *X is simply connected.*

(2) *The link of every k–cell, $k \leq n-4$, is simply connected.*

(3) *The link of every $(n-3)$–cell is isometric to a building.*

Then X is isometric to a spherical building (respectively metric Euclidean building).

Proof In the spherical case, the theorem follows from Theorem 4.1 and induction. In the Euclidean case, links in X are spherical buildings (by the spherical case of the theorem) so X is locally CAT(0). Since X is simply connected, it is also globally CAT(0) (Theorem 2.1(2)). The theorem now follows from Corollary 5.3. □

Another interesting metric characterization involves extensions of geodesics. A local geodesic γ ending at x *extends discretely* if the set of directions in which γ can be geodesically continued through x is a non-empty discrete subset of $\mathrm{lk}\,(x, X)$, or equivalently, if the set of points $v \in \mathrm{lk}\,(x, X)$ at distance $\geq \pi$ from γ_x is non-empty and discrete. We say a geodesic metric space X has the *discrete extension property* if γ extends discretely for every local geodesic γ.

Theorem 7.5 *Suppose X is a connected, locally spherical (respectively Euclidean) space of dimension $n \geq 2$, and suppose*

(1) *X is CAT(1) (respectively CAT(0)),*

(2) *X has the discrete extension property.*

Then X is isometric to a spherical building (respectively metric Euclidean building).

Proof First note that if X has the discrete extension property, then so does $L_x = \mathrm{lk}\,(x, X)$ for every x. For if $v \in L_x$ and γ is a geodesic emanating from x in direction v, then for any point $y = \gamma(t)$ in a conelike neighborhood of x, $L_y = \mathbb{S}^0 * \mathrm{lk}\,(v, L_x)$. Thus, if every point in L_y has a non-empty discrete set of points at distance $\geq \pi$, then the same holds in $\mathrm{lk}\,(v, L_x)$. (See the appendix of [7] for details on distances in spherical suspensions.)

The discrete extension property also implies that L_x is connected. For if σ is a spherical $(n-1)$–cell in L_x and $v \in L_x$ is not in the connected component of σ, then all of σ has distance $\geq \pi$ from v.

The theorem now follows by induction on n. If $n = 2$, then the discrete extension property implies that L_x is connected, has diameter π, and every vertex has valence at least 2. Thus, by Theorem 6.1, and the remark following it, L_x is isometric to a 1-dimensional spherical building for all x. By Theorem 3.8, (respectively Corollary 5.3) X is isometric to a spherical (respectively metric Euclidean) building.

If $n > 2$, then L_x is a connected, locally spherical space satisfying conditions (1) and (2) of the theorem. By induction L_x is isometric to a spherical building for every x, and the conclusion follows from Theorem 3.8 (respectively Corollary 5.3). □

References

[1] **W Ballmann**, *Lectures on Spaces of Nonpositive Curvature*, DMV Seminar 25, Birkhäuser (1995)

[2] **W Ballmann**, **M Brin**, *Orbihedra of nonpositive curvature*, Inst. Hautes Études Sci. Publ. Math. 82 (1996) 169–209

[3] **W Ballmann**, **M Brin**, *Diameter rigidity of spherical polyhedra*, Duke Math. J. 97 (1999) 235–259

[4] **M Bridson**, **A Haefliger**, *Metric Spaces of Non-positive Curvature*, Springer–Verlag (1999)

[5] **N Bourbaki**, *Groupes et Algebrès de Lie*, Ch. 4–6, Masson, Paris (1981)

[6] **K Brown**, *Buildings*, Springer–Verlag, New York (1989)

[7] **R Charney**, **M W Davis**, *Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds*, Amer. J. Math. 115 (1993) 929–1009

[8] **M W Davis**, *Buildings are CAT(0)*, from: "Geometry and Cohomology in Group Theory", ed. by P Kropholler, G Niblo and R Stöhr, LMS Lecture Notes Series 252, Cambridge University Press (1998)

[9] **M Gromov**, *Hyperbolic groups*, from: "Essays in Group Theory", ed. by S M Gersten, MSRI Publ. 8, Springer–Verlag, New York (1987) 75–264

[10] **W M Kantor**, *Generalized polygons, SCAB's and GAB's*, from: "Buildings and the Geometry of Diagrams", LNM 1181, Springer–Verlag, New York (1986) 79–158

[11] **B Kleiner**, **B Leeb**, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Publ. IHES 86 (1997) 115–197

[12] **A Neumaier**, *Some aporadic geometries related to PG(3,2)*, Arch. Math. 42 (1984) 89–96

[13] **M Ronan**, *Lectures on Buildings*, Perspectives in Mathematics, Vol. 7, Academic Press, San Diego (1989)

[14] **J Tits**, *A local approach to buildings*, from: "The Geometric Vein, Coxeter Festschrift", ed. by C Davis, B Grüenbaum and F A Sherk, Springer–Verlag, New York (1981) 519–547