Affine buildings for dihedral groups

Arkady Berenstein and Michael Kapovich

January 11, 2011

Abstract

We construct 2-dimensonal thick nondiscrete affine buildings associated with an arbitrary finite dihedral group.

1 Introduction

In his foundational work on buildings, J. Tits [T1, T3] proved that all thick irreducible spherical and affine buildings of dimension ≥ 2 and ≥ 3 respectively, (including nondiscrete ones) are of algebraic origin and, in particular, have Weyl groups coming from root systems (see also [W]). He then showed [T2] how to use the *free construction* to prove existence of thick spherical buildings of rank 2 modeled on arbitrary finite Coxeter groups (see Definition 4.1). Such buildings are necessarily infinite except for dihedral groups of the order $2, 4, 6, 8, 12, 16$. M. Ronan [R1] constructed a variety of irreducible 2-dimensional discrete affine buildings not coming from algebraic groups. The corresponding finite Coxeter groups are necessarily of crystallographic type, i.e., they are the dihedral groups $D_m, m = 3, 4, 6$. However, examples of nondiscrete 2-dimensional affine buildings corresponding to the rest of finite dihedral groups remained elusive. The main result of this paper is a construction of such buildings.

Consider the dihedral group $W = D_m$ of order $2m, m \geq 2$. Take a countable dense subgroup Λ of translations of the apartment $A = \mathbb{R}^2$, invariant under the action of W by conjugations, so that for the group $W_{af} = \Lambda \rtimes W$ (which is, in fact, a Coxeter group), every vertex of the Coxeter complex (A, W_{af}) is special. Such Coxeter complexes will be called *special*. For instance, when m is a prime number, the complex (A, W_{af}) is special no matter what Λ is. The main goal of this paper is to prove

Theorem 1.1. *For every special Coxeter complex (A, W_{af}) with nondiscrete countable group W_{af}, there exists a 2-dimensional thick Euclidean building X modeled on (A, W_{af}).*

After proving this theorem we found out that for $m = 8$, existence of thick Euclidean buildings (for some countable W_{af}) was also proven in a recent work of P.Hitzelberger, L.Kramer and R.Weiss [HKW] by a totally different method: Their buildings are embedded in buildings of the type F_4.

It is quite likely that Theorem 1.1 holds for semidirect products W_{af} of W with arbitrary (nondiscrete) groups of translations $\Lambda \subset \mathbb{R}^2$, or, at least countable ones, where the issue is to eliminate the assumption that (A, W_{af}) is special. Note that the only place in the paper where the latter assumption is used, is Part 2 of the proof of Proposition 6.5.

The strategy for proving Theorem 1.1 is the same as in Tits' free construction of rank 2 spherical buildings, see [T2, FS], i.e., we run a *"Ponzi-scheme"*:

We inductively construct a sequence of CAT(0) spaces

$$X_1 \to X_2 \to \ldots \to X_n \to X_{n+1} \to \ldots$$

where each X_n is *finite*, i.e., is a finite union of apartments. (The embeddings $X_n \to X_{n+1}$ are *not* isometric.) Each X_n is modeled on (A, W_{af}) but fails, rather badly, the rest of axioms of a thick building, as "most" pairs of points are not contained in a common apartment. Given points (say, vertices) $x_1, x_2 \in X_n$, we obtain X_{n+1} by attaching an apartment to X_n along a closed convex subcomplex $C \subset X_n$ containing x_1, x_2. Then $x_1, x_2 \in X_{n+1}$ belong to a common apartment. This, of course, produces more points in X_{n+1} which do not belong to a common apartment, as well as more walls which are not *thick*. However, these will be taken care of at some later stages of the construction. The hardest part of the proof is to establish the following *combinatorial convexity theorem*, proven in sections 7 and 8:

Theorem 1.2. *Let X be a CAT(0) metric space, which is a finite space modeled on (A, W_{af}). Then, given any pair of points $x_1, x_2 \in X$ and germs of Weyl chambers $\sigma_i \subset X$ with the tips at x_i, there exists a* planar *convex subcomplex $C \subset X$ containing x_i and $\sigma_i, i = 1, 2$. The same applies to each pair of* super-antipodal *chambers σ_1, σ_2 at infinity, i.e., pair of chambers which admit regular points connected by a geodesic in X.*

Planarity of C means existence of a *weakly isometric embedding* $C \hookrightarrow A$. (We introduce D-isometry and weak isometry in Sections 3 and 4 as certain generalizations of the ordinary isometry between metric spaces.) The embedding $C \hookrightarrow A$ allows one to attach A to X along C so that $X \cup_C A$ is again a CAT(0) space which is finite, modeled on (A, W_{af}). The convex subcomplexes C used in our paper are (extended) *corridors* obtained by taking a finite union of *tiles* in X. The appearance of corridors in this context is natural since, in the case when X is a Euclidean building, every pair of vertices $x_1, x_2 \in X$ admits a combinatorial convex hull, called a *Weyl hull*, which is a tile with the vertices x_1, x_2.

Since each X_n will have only countably many corridors, we enumerate them and form the sequence (X_n) so that for every convex corridor C in X_k there exists $n \geq k$ such that X_{n+1} is formed by attaching an apartment along a convex corridor $C' \subset X_n$ containing C.

What goes wrong in higher dimensions.

If an appropriate form of the combinatorial convexity theorem were true, our proof would extend in higher-dimensional case as well. However, our convexity theorem (Theorem 1.2) fails in dimensions 3 and higher: We do not even have an adequate notion of a weak isometry in this case. The reason is that our concept of a weak isometry is motivated by the Tits' free construction of 1-dimensional spherical buildings, which does not extend in higher dimensions. Thus, one is lead to the question what goes wrong in the case of 2-dimensional spherical buildings: There we have the concept of a weak isometric embedding to S^2 (essentially the same that we use in the present paper). Moreover, an analogue of our convexity theorem (Theorem 1.2) holds for 2-dimensional CAT(1) complexes modeled on (S^2, W), where W is a finite reflection group. (Instead of corridors one would use unions of minimal galleries connecting facets whose combinatorial distance is less than twice the length of w_0, the longest element of W.) The problem, however, is that if Y is a 2-dimensional CAT(1) complex, $C \subset Y$ is a connected subcomplex, $\iota : C \to S^2$ is a weakly isometric embedding, then the new complex $Y' = Y \cup_C S^2$ (obtained by gluing Y and S^2 via ι) is a simply-connected locally CAT(1) complex, which need not be a CAT(1) complex. (This situation is impossible for CAT(0) spaces due to the Gromov-Cartan-Hadamard Theorem.) Nevertheless, if one is content to work with simply-connected locally CAT(1) complexes, the rest of the proof of Theorem 1.1 goes through; in particular, an analogue of the convexity theorem extends in this setting (provided that one uses connected locally convex subcomplexes $C \subset Y$ in such spaces). The end result is a space Y_∞ (a metric simplicial complex) which satisfies all the axioms of a thick building except that it is merely a simply-connected locally CAT(1) space (rather than a CAT(1) space). In particular, intersections of apartments are locally convex subcomplexes in S^2. For instance, such complexes Y_∞ would exist for the rank 3 Weyl group $W = H_3$ (see [T2]), for which a thick spherical building cannot exist. We do not know what such spaces are good for. For instance, if we were to use the above construction in order to form 3-dimensional affine buildings, the intermediate spaces X_n would not be locally CAT(0) (although they would satisfy a local CAT(0) axiom away from the vertices). Thus, apartments in such spaces need not be convex, so the direct limit of X_n's may not even have a natural structure of a metric space.

Organization of the paper. In sections 2, 3, 4 we review standard concepts related to buildings: Coxeter complexes, CAT(0) and CAT(1) metric spaces and spaces modeled on Coxeter complexes. We review the building axioms in section 4. The only nonstandard material in these sections is the one of *D-isometries* between metric spaces and the corresponding concept of *weak isometries* in the context of spaces modeled on (A, W_{af}). In section 5 we define the main technical tool of this paper — *corridors*, which are finite unions of *tiles*, parallelograms in A with the angles $\frac{\pi}{m}$ and $\pi - \frac{\pi}{m}$. In section 6 we define two orders on corridors and prove existence of maximal elements with respect to these orders. The *maximal corridors* are proven to be convex in section 7. In the same section we prove a weaker version of Theorem 1.2, i.e., we prove that every two points in X belong to a corridor (and, therefore, to a convex corridor). In section 8 we introduce the notion of *extended corridors*, which are noncompact analogues of corridors and finish the proof of Theorem 1.2. In section 9 we define *enlargement* of a space X modeled on (A, W_{af}), which is the operation of attaching to X an apartment along an extended convex corridor in X. We then show that the image of a convex corridor under enlargement is again a corridor (possibly nonconvex). In section 10 we put it all together and prove our main result, Theorem 1.1, by constructing an increasing sequence (X_n) of spaces modeled on (A, W_{af}), whose union is an affine building. In section 11 we make some speculative remarks on possible generalizations of our results, in particular, constructing 2-dimensional buildings with highly transitive automorphism groups.

Acknowledgments. Our collaboration on this project started at the AIM workshop "Buildings and Combinatorial Representation Theory" in 2007 and we are grateful to AIM for this opportunity. The first author was supported by the NSF grant DMS-08-00247. The second author was supported by the NSF grants DMS-05-54349, DMS-09-05802. We are grateful to Linus Kramer and Petra Hitzelberger for the reference to their work and to the referee for useful remarks.

Contents

2 Coxeter complexes

Let A, the *apartment*, be either the Euclidean space E^n or the unit $n-1$-sphere $S^{n-1} \subset E^n$ (we will be primarily interested in the case of Euclidean plane and the circle). If $A = S^{n-1}$, a *Coxeter group* acting on A is a finite group W generated by isometric reflections. If $A = E^n$, a *Coxeter group* acting on A is a group W_{af} generated by isometric reflections in hyperplanes in A, so that the linear part of W_{af} is a Coxeter group acting on S^{n-1}. Thus, $W_{af} = \Lambda \rtimes W$, where Λ is a certain (countable or uncountable) group of translations in E^n.

Definition 2.1. A spherical or Euclidean[1] *Coxeter complex* is a pair (A, G), of the form (S^{n-1}, W) or (E^n, W_{af}).

[1] Also called *affine*.

A *wall* in the Coxeter complex (A, G) is the fixed-point set of a reflection in G. A *regular point* in a Coxeter complex is a point which does not belong to any wall. A *singular point* is a point which is not regular. A *half-apartment* in A is a closed half-space bounded by a wall. A *cell* in is a convex polyhedron which is the intersection of finitely-many half-apartments.

Remark 2.2. Note that in the spherical case, there is a natural cell complex in S^{n-1} associated with W, where the *cells* are those cells (in the above sense) which cannot be separated by a wall. However, the affine case, when W_{af} is nondiscrete, there will be no cell natural complex attached to W_{af}.

Chambers in (S^{n-1}, W) are the fundamental domains for the action $W \curvearrowright S^{n-1}$, i.e., the closures of the connected components of the complement to the union of walls.

A *dilation* of an affine Coxeter complex is a map of the form $x \mapsto \lambda x + v$, where v is a vector and $\lambda > 0$. Thus, we regard translations as a (limiting form of) dilations.

A *vertex* in (A, G) is a (component of, in the spherical case) the 0-dimensional intersection of walls. We will consider almost exclusively only those Coxeter complexes which have at least one vertex; such complexes are called *essential*. Equivalently, these are spherical complexes where the group G does not have a global fixed point and those Euclidean Coxeter complexes where W does not have a fixed point in S^{n-1}.

A *special point* in (E^n, W_{af}) is a point whose stabilizer in W_{af} is isomorphic to W. We will always choose coordinates in $A = E^n$ so that the origin is special. A Coxeter complex where every vertex is special is called a *special Coxeter complex*. In general, the stabilizer of a point $x \in A$ is naturally isomorphic a Coxeter subgroup of W.

Consider now the case of $n = 2$, when $W = D_m$, the dihedral group of order $2m$. Under the identification of the vertices of the regular $2m$-gon with the $2m$-th roots of unity, the group Λ always contains (or equals to) the ring $\Lambda_0 = \mathbb{Z}[\zeta]$, where ζ is a primitive $2m$-th root of unity if m is odd and a primitive m-th root of unity if m is even. It is well–known that Λ_0 is discrete in $\mathbb{C} = \mathbb{R}^2$ if and only if $m = 1, 2, 3, 4, 6$.

If m is prime, then every proper Coxeter subgroup of W is isomorphic to $\mathbb{Z}/2$ and, hence, every vertex of (E^2, W_{af}) is special. If m is not prime and Λ is the smallest nontrivial group of translations normalized by W, the complex (A, W_{af}) is non-special. Nevertheless,

Lemma 2.3. *For every (countable) $W_{af} = \Lambda \rtimes W$, there exists a (countable) group Λ' containing Λ and normalized by W, so that the Coxeter complex $(A, \Lambda' \rtimes W)$ is special.*

Proof. Set $\Lambda^1 := \Lambda$ and $W^1_{af} := W_{af}$ and $C^1 := (A, W^1_{af})$. We continue inductively: Given a Coxeter complex $C^i := (A, W^i_{af})$ with (countable) Coxeter group, we define Λ^{i+1} as the group of translations generated by all vertices of C^i (since we have chosen the origin in A, this makes sense). Then, every vertex of C^i is special in $C^{i+1} = (A, W^{i+1}_{af})$, $W^{i+1}_{af} = \Lambda^{i+1} \rtimes W$. By taking

$$\Lambda' := \bigcup_i \Lambda^i, \quad W'_{af} := \bigcup_i W^i_{af},$$

we obtain the required special Coxeter complex (A, W'_{af}). □

Given a point $x \in A = E^n$, we then let $W_{af,x}$ denote the stabilizer of x in W_{af}. Then the *type*, *type*(x), of the point x is the W-conjugacy class of the *linear part* of $W_{af,x}$ (i.e., the image under the canonical projection $W_{af} \to W = W_{af}/(A \cap W_{af})$). In particular, for special Coxeter complexes, all vertices have the same type. Note that this concept of type is weaker[2] than the standard (in the discrete case) notion of type given by the coset $W_{af} \cdot x$ of x in A/W_{af} (equivalently, the intersection of the orbit $W_{af} \cdot x$ with the fundamental alcove). However, in the case of nondiscrete Coxeter groups, the latter concept does not seem to be of much use since A/W_{af} cannot be identified with an alcove or any reasonable (commutative) geometric object.

In the spherical case, the notion of *type* is given by the projection

$$\theta : S^{n-1} \to S^{n-1}/W = \Delta_{sph},$$

where the quotient is the spherical Weyl chamber, a fundamental domain for $W \curvearrowright S^{n-1}$.

We now assume that $n = 2$, $A = E^2$ and $W = D_m$, the dihedral group of order $2m$. Then there will be at most $div(m) + 2$ types of points in (A, W_{af}), where $div(m)$ is the number of divisors of m. In what follows, we choose the coordinates in A so that the positive Weyl chamber Δ of W is bisected by the x-axis.

A *sub-wall* in A is a codimension 1 cell in A, i.e., a nondegenerate geodesic segment contained in a wall and bounded by two vertices.

3 Metric concepts

In this section we review basics of CAT(0) and CAT(1) spaces and introduce the concept of a *D-isometry* between metric spaces.

Notation 3.1. *For a subset Y in a metric space X and $r \geq 0$, we let $B_r(Y)$ denote the* open r-neighborhood *of Y in X, i.e.,*

$$B_r(Y) := \{x \in X : \exists y \in Y, d(x, y) < r\}.$$

For instance, if $Y = \{y\}$ is a single point, then $B_r(Y) = B_r(y)$ is the open r-ball centered at y.

A metric space X is called *geodesic* if every two points in X are connected by a (globally distance-minimizing) geodesic. Most metric spaces considered in this paper will be geodesic; the only occasional exceptions will be links of vertices in certain metric cell complexes.

Notation 3.2. *For a pair of points x, y in a metric space X we let $\overline{xy}$ denote a closed geodesic segment (if it exists) in X connecting x and y. As, most of the time, we will deal with spaces where every pair of points is connected by the unique geodesic, this is a reasonable notation. We let xy denote the* open geodesic segment, *the* interior of $\overline{xy}$:

$$xy = \overline{xy} \setminus \{x, y\}.$$

Geodesics will be always parameterized by their arc-length.

[2]Unless W_{af} is transitive on A.

***D*-isometries.** Let X, X' be metric spaces and $D > 0$ be a real number. We say that a map $f : X \to X'$ is a *D-isometry* if it satisfies the following:

1. $\forall x, y \in X, d(x, y) < D \Rightarrow d(f(x), f(y)) = d(x, y)$.
2. $\forall x, y \in X, d(x, y) \geq D \Rightarrow d(f(x), f(y)) \geq D$.

The motivation for this definition comes from considering maps where the diameter D of the target is less than the diameter of the domain, but we still would like to have maps $X \to X'$ which are as close to being isometries as possible. We will use this concept, however, even if $\mathrm{diam}(X') > D$.

The following properties of D-isometries are immediate:

- Each D-isometry is injective.
- For each D-isometry f, its inverse $f^{-1} : f(X) \to X$ is also a D-isometry.
- Composition of D-isometries is again a D-isometry.

CAT(0) and CAT(1) spaces. We refer the reader for instance to [Ba, BH] for the detailed treatment of this material.

Let X be a geodesic metric space. The space X is called CAT(0), if the geodesic triangles in X are *not thicker* than triangles in the Euclidean plane E^2. One can state this in terms of quadruples of points in X. Let $x, y, z \in X$ (vertices of a triangle $\Delta(x, y, z)$). Let $m \in X$ be such that

$$d(x, m) + d(m, y) = d(x, y)$$

(i.e., m will be a point on the side $\overline{xy}$ of the triangle). Find points $\bar{x}, \bar{y}, \bar{z}$ in E^2 so that

$$d(x, y) = d(\bar{x}, \bar{y}), \quad d(y, z) = d(\bar{y}, \bar{z}), \quad d(z, x) = d(\bar{z}, \bar{x}),$$

these are vertices of a *comparison triangle* $\Delta(\bar{x}, \bar{y}, \bar{z})$ in E^2. Next, find a point $\bar{m} \in E^2$ so that

$$d(x, m) = d(\bar{x}, \bar{m}), \quad d(m, y) = d(\bar{m}, \bar{y}),$$

i.e., $\bar{m}$ belongs to the side of $\Delta(\bar{x}, \bar{y}, \bar{z})$ which corresponds to the side $\overline{xy}$ of $\Delta(x, y, z)$. Then we require the CAT(0) inequality:

$$d(z, m) \leq d(\bar{z}, \bar{m}).$$

Remark 3.3. Note that we do not assume X to be complete: Most buildings constructed in this paper will be incomplete as metric spaces.

Similarly, one defines CAT(1) spaces Y, except:

1. Triangles of the perimeter $\leq 2\pi$ in Y are compared to triangles on the unit 2-sphere.
2. Y need not be geodesic, however, if $p, q \in Y$ cannot be connected by a geodesic, we require the distance between these points be $\geq \pi$.

We will think of the distances in CAT(1) spaces as *angles* and, in many cases, denote these distances $\angle(xy)$.

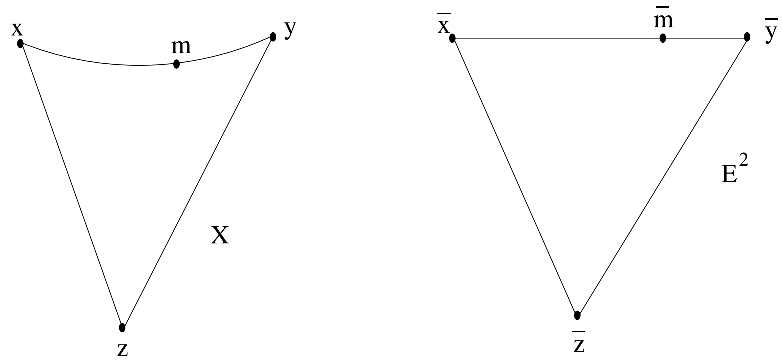

Figure 1: *Comparison triangle.*

The following characterization of 1-dimensional CAT(1) spaces will be important:

A 1-dimensional metric space (a metric graph) is a CAT(1) space if and only if the length of the shortest embedded circle in X is $\geq 2\pi$.

If Γ is a metric graph, where each edge is given the length π/m, then the CAT(1) condition is equivalent to the assumption that girth of Γ is $\geq 2m$.

In a CAT(0) space, any pair of points is connected by a unique geodesic, while in CAT(1) spaces this is true for points within distance $< \pi$. A subset Z in a CAT(0) space X is called *convex* if for every two points $p, q \in Z$ the geodesic $\overline{pq} \subset X$ is contained in Z. A connected closed subset Z in a complete CAT(0) space X is convex if and only if it is locally convex, i.e., every point $z \in Z$ possesses a neighborhood in Z which is convex in X.

Spaces of directions. Let X be a CAT(0) space. The *space of directions* $\Sigma_x(X)$ at a point $x \in X$ is defined as the set of germs of geodesics emanating from x. The distances in $\Sigma_x(X)$ are given by the *angles* between the geodesics.

If X is CAT(0), then $\Sigma_x(X)$ is CAT(1) for every $x \in X$. Conversely, suppose that X is simply-connected, has structure of a regular cell complex. Equip each cell with a metric isometric to a convex polyhedron in the Euclidean space; this determines a path-metric on X. Suppose, in addition, that $\Sigma_x(X)$ is CAT(1) for every $x \in X$. Then X is CAT(0), see [Ba, Corollary 4.3 and Theorem 4.5] and [BH].

Example 3.4. *Suppose that X is a 2-dimensional locally-finite Euclidean metric cell complex, i.e., X is obtained by gluing isometrically convex polygons in E^2 where we have only finitely many isometry classes of polygons. Equip X with the induced path-metric. Then for each $x \in X$, the space of directions $\Sigma_x(X)$ is just the link Γ of x in X. To describe the metric on $\Sigma_x(X)$ note that a small neighborhood of x in X is covered by flat Euclidean sectors S with the tip at x. The edges of Γ correspond to maximal flat sectors S (not containing any edges or vertices in its interior). The angles of these maximal sectors at x are the lengths of the corresponding edges of Γ. Then, X is CAT(0) iff X is simply-connected and for every vertex $x \in X$ its link contains no embedded cycles of length $< 2\pi$.*

4 Geometry of spaces modeled on (A, W_{af})

In this section we introduce spherical and Euclidean buildings and, more generally, spaces X modeled on Coxeter complexes. Our discussion follows [KL] and [P]. We then review basic properties of such spaces, e.g., structure of the space of directions. We also define and discuss one of the key notions of this paper, *weak isometries* between subcomplexes in spaces modeled on Coxeter complexes.

Spaces modeled on Coxeter complexes.

Definition 4.1. A space *modeled on the Coxeter complex* $(A = E^n, W_{af})$ is a metric space X together with an atlas (covering X) where charts are isometric embeddings $A \to X$ and the transition maps are restrictions of the elements of W_{af}. The maps $A \to X$ and their images are called *apartments* in X. Note that (unlike in the definition of an atlas in a manifold) we do not require the apartments to be open in X. The same definition applies to spaces modeled on the spherical Coxeter complex (S^{n-1}, W), only now the charts are maps of S^{n-1}.

Remark 4.2. This definition generalizes to the case where we replace W_{af} with an arbitrary group G of affine transformations of E^n and assume that $A \subset E^n$ is a G–invariant affine subspace. In this generality, one has to remove, of course, the assumption that X is a metric space and the charts are isometric embeddings. Instead, one has to assume that for every two charts $\phi, \psi : A \to X$,

$$\psi^{-1}\phi(A)$$

is a closed convex subset of A. Interesting examples of this setup are provided by *tropical geometry*, where $G = \mathbb{R}^n \rtimes SL(n, \mathbb{Z})$.

Therefore, all W_{af}-invariant (resp. W-invariant) notions defined in A, extend to X. In particular, we will talk about vertices, walls, sub-walls, cells, etc. Note that the intersection of finitely many cells in X is again a cell.

We also will have a well-defined notion of *type* of a point $x \in X$: It is given by computing the type of x in an apartment $A \subset X$ containing x. We retain the notation $\mathit{type}(x), \theta(x)$ in the affine, resp. spherical, case.

Example 4.3. If $W = D_m$ and Y is modeled on (S^1, W), then Y has the natural structure of a bipartite metric graph, where every edge has the length π/m.

Accordingly, if X is a CAT(0) space modeled on $(A, W_{af} = \Lambda \rtimes W)$, then for every $x \in X$, $Y = \Sigma_x(X)$ is a CAT(1) metric space modeled on (S^1, W'), where $W' = \mathit{type}(x)$, a subgroup of W.

A space X modeled on a Coxeter complex is called *locally finite* if every point of X has a neighborhood covered by finitely many apartments.

A space X modeled on (A, W_{af}) is said to be *finite* if it is covered by finitely many apartments. (This, of course, does not mean finiteness of X as a set or finiteness of its diameter.)

Notation 4.4. *We let $\mathcal{C}(A, W_{af})$ denote the collection of all finite spaces X modeled on (A, W_{af}). When (A, W_{af}) is fixed, we will use the notation $\mathcal{C}$ for this collection. Similarly, we define $\mathcal{C}_{lf}$ the collection of all locally finite spaces X modeled on (A, W_{af}).*

A *finite subcomplex* Y in $X \in \mathcal{C}_{lf}$ is a finite union of cells in X. Note that such Y need not be modeled on (A, W_{af}).

Convention 4.5. *In what follows, all spaces modeled on an affine Coxeter complexes are assumed to be CAT(0) and all metric spaces modeled on spherical Coxeter complexes are assumed to be CAT(1).*

For a space Y modeled on a spherical Coxeter complex (S^{n-1}, W), two elements $\xi, \eta \in Y$ are called *antipodal* if $\angle(\xi, \eta) = \pi$. Elements of Y are *super-antipodal* if $\angle(\xi, \eta) \geq \pi$. Accordingly, chambers $\sigma_1, \sigma_2 \subset Y$ are antipodal (rep. super-antipodal) if they contain antipodal (resp. super-antipodal) regular elements. For instance, if Y consists of a single apartment, then chambers $\sigma_1, \sigma_2 \subset Y$ are antipodal if and only if the element $w \in W$ sending σ_1 to σ_2 is the longest element of W, i.e.,

$$\sigma_2 = -\sigma_1 = \{-\xi : \xi \in \sigma_1\}.$$

Let X be modeled on a 2-dimensional complex (A, W_{af}). A *thick sub-wall* in X is a sub-wall γ so that for every $x \in \gamma$, $\Sigma_x(X)$ is not a single apartment. In other words, x has no neighborhood which is locally isometric to A. Of course, not every sub-wall is thick. If γ is a wall which is thick as a sub-wall, then γ is a *thick wall.*

Similarly, we say that $x \in X$ is a *thick vertex* of X if either x locally separates X or x belongs to at least three thick sub-walls e_1, e_2, e_3, so that

$$e_i \cap e_j = \{x\}, i \neq j.$$

Equivalently, the graph $\Sigma_x(X)$ is either disconnected or contains at least 3 vertices of valence ≥ 3.

Example 4.6. 1. Take two apartments A_1, A_2 and glue them at a vertex v. Then x is a thick vertex of the resulting space.

2. Take apartments A_1, A_2, A_3, A_4 and glue them along walls L_1, L_2 as follows:

$$X_1 := A_1 \cup_{L_1} A_2, X_2 := A_3 \cup_{L_2} A_4.$$

Now, take an isomorphism $\phi : A_1 \to A_2$, so that $\phi(L_1)$ crosses (transversally) L_2. Next, glue X_1, X_2 via ϕ to form a space X. Then the vertex $v \in X$ corresponding to $\phi(L_1) \cap L_2$ is a thick vertex in X.

A *Euclidean (or, affine) building* modeled on (A, W_{af}) is a CAT(0) space X modeled on (A, W_{af}) which satisfies the following two extra conditions:

Axiom 1. ("Connectedness") Every two points $x_1, x_2 \in X$ are contained in a common apartment.

Axiom 2. ("Angle rigidity") For every $x \in X$, the space of directions $Y = \Sigma_x(X)$ satisfies the following:

$$\forall \xi, \eta \in Y, \quad \angle(\xi, \eta) \in W \cdot \angle(\theta(\xi), \theta(\eta)).$$

Here $\theta : Y \to \Delta_{sph}$ is the *type projection.*

For instance, one can see that every (non-flat) nonpositively curved symmetric space satisfies Axiom 1 but not Axiom 2.

A *spherical building* is defined similarly, except that it is required to be modeled on a spherical Coxeter complex, be CAT(1) and the Axiom 2 is not needed in this case (it follows from *discreteness* of the Coxeter complex). It was proven in [CL] (and in an unpublished paper by B. Kleiner) that a piecewise-Euclidean CAT(0) cell complex is a Euclidean building if and only if all its links are spherical buildings.

A building X is called *thick* if every wall in X is the intersection of (at least) three apartments.

Below is an alternative set of axioms for Euclidean buildings due to Ann Parreau [P]. Replace Axioms 1 and 2 with:

Let X be a CAT(0) space modeled on (A, W_{af}) so that:

Axiom 1′. Let $x_1, x_2 \in X$ and $\sigma_i \subset \Sigma_{x_i}(X)$ be spherical chambers, where we endow the spherical buildings $\Sigma_{x_i}(X)$ with the $(S^1, type(x_i))$ structure. Then there exists an apartment $A \subset X$ so that $x_i \in A, \sigma_i \subset \Sigma_{x_i}(A)$, $i = 1, 2$.

Axiom 2′. For every $x \in X$ the following holds. Let $\Delta_i \subset X, i = 1, 2$ be antipodal *Weyl sectors* with the tip x (i.e., $\Sigma_x \Delta_i$ are antipodal in $\Sigma_x(X)$). Then there exists an apartment $A \subset X$ containing Δ_1, Δ_2.

Here a *Weyl sector* in X is a subset Δ' of an apartment $A \subset X$, so that Δ' is a parallel translate of a Weyl chamber in A. Note that if the tip x of Δ' is a special vertex, then Δ' is a Weyl chamber in X.

Remark 4.7. Parreau's numbering of axioms is different from ours. She also does not require the CAT(0) condition, but has other axioms requiring intersections of pairs of apartments to be closed and convex. These conditions follow from the CAT(0) property.

We will be using Axioms 1′ and 2′ since verifying angle rigidity would require proving a slightly weaker form of Axiom 1′ anyhow. We refer the reader to Tits' original paper [T3] and to Appendix 3 of Ronan's book [R2] for alternative axiomatization of nondiscrete Euclidean buildings.

Morphisms.

A *morphism* between two spaces X, X' modeled on (A, G), where $A = E^n, G = W_{af}$ or $A = S^{n-1}, G = W$, is defined as a map $f : X \to X'$ which, being written in "local coordinates given by apartments" is the restriction of elements of G.

More generally, given two spaces X, X' modeled on (A, G), and subsets $Y \subset X, Y' \subset X'$, a *morphism* $Y \to Y'$ is a map $f : Y \to Y'$, which, in local coordinates, appears as the restriction of an element of G. In other words, for every point $y \in Y, y' = f(y)$, and for every pair of apartments $\phi : A \to X, \psi : A \to X'$ (with images containing y, y' respectively) the (partially defined) map

$$\psi^{-1} \circ f \circ \phi : \phi^{-1}(Y) \subset A \to \psi(Y') \subset A$$

is the restriction of an element of G.

It follows that, if a geodesic $\gamma \subset X$ is contained in Y, then its image under a morphism $Y \to Y'$ is a piecewise-geodesic path. Therefore, every morphism $Y \to Y'$ induces a morphism (again denoted f to simplify the notation)

$$f : \Sigma_y Y \to \Sigma_{y'} Y', \quad y \in Y, \quad y' = f(y) \in Y'.$$

Weak isometries. Let Z be a subcomplex in a metric space U modeled on a spherical Coxeter complex (S^{n-1}, W). We always endow Z with the restriction of the path-metric on U (not the induced path-metric!). Let V denote the vertex set of the complex (S^{n-1}, W). Set

$$D := \max\{d(x, w(x)) : x \in V, w \in W\}.$$

Suppose now that $n = 2$ and $W = D_m$. We set $\delta = \pi - \pi/m$. Then, clearly,

$$D = \begin{cases} \pi & \text{if } m \text{ is even} \\ \pi - \pi/m = \delta & \text{if } m \text{ is odd} \end{cases}$$

More generally, for arbitrary finite Coxeter groups W we have: $D \leq \pi$ with strict inequality if and only if the Coxeter number of W is odd. This is the case when W is a product of symmetric groups S_m with odd m and dihedral groups D_m with odd m. For instance, if $W = S_m$ with odd m, then $D = \arccos(-(m-1)/(m+1))$.

We then say that a *weak isometry* between subcomplexes $Z \subset U, Z' \subset U'$ is a surjective morphism $Z \to Z'$ which is a D-isometry. It is easy to see that in the case $W = D_m$ and $n = 2$, the latter condition is equivalent to being a δ-isometry.

We now define weak isometries between subsets of spaces modeled on affine Coxeter complexes. Let X, X' be modeled on the affine Coxeter complex (A, W_{af}).

1. We first consider the easier case when $Y \subset X, Y' \subset X'$ are subcomplexes. Then a *weak isometry* $Y \to Y'$ is a morphism $Y \to Y'$ which sends cells isometrically to cells and induces a weak isometry

$$\Sigma_y(Y) \to \Sigma_{y'}(Y')$$

for every $y \in Y, y' = f(y) \in Y'$. This is the concept akin to the notion of an isometric embedding of Riemannian manifolds (rather than isometry between metric spaces).

2. Now, suppose that $Y \subset X, Y' \subset X'$ are *locally conical subsets*, i.e., each point of Y, Y' admits a neighborhood in Y, Y' which is a subcomplex in X, X'. Therefore, for every $y \in Y, y' \in Y'$,

$$\Sigma_y(Y) \subset \Sigma_y(X), \quad \Sigma_{y'}(Y') \subset \Sigma_{y'}(X')$$

are subcomplexes.

Then a *weak isometry* $Y \to Y'$ is a surjective morphism $f : Y \to Y'$ inducing a weak isometry

$$f : \Sigma_y(Y) \to \Sigma_{f(y)}(Y'),$$

for every $y \in Y$.

We will use the name *weakly isometric embedding* for the map $f : Y \to X'$, when we do not want to specify the image of f.

Local structure of a locally finite space X modeled on (A, W_{af}).

Definition 4.8. (Conical sets, cf. [CL].) Let $Z \subset X$ be a subset and $o \in Z$. Then we say that Z is *conical* with respect to o if:

1. For every $x \in U$, $\overline{xo} \subset Z$ (i.e., Z is star-like with respect to o).

2. Geodesics starting at o do not branch in Z: If $\gamma_i = \overline{ox_i}$, $i = 1, 2$ are geodesics of the same length in C, and the germs of γ_1, γ_2 agree near o, then $x_1 = x_2$.

Let Y be a metric graph. We define the *cone over* Y, denoted $Cone(Y)$, as follows. For each edge e of Y of length α, let $Cone(e)$ be the (infinite) sector in $\mathbb{R}^2$ with the tip o (at the origin) and the angle α. If v is a vertex of e, the cone $Cone(v)$ is the ray bounding $Cone(e)$. We now glue all the sectors $Cone(e)$ via isometries

$$Cone(v) \subset Cone(e) \to Cone(v) \subset Cone(e')$$

whenever the edges $e, e' \subset Y$ share the vertex v. The result is $Cone(Y)$. There is a more general definition of a metric cone over a CAT(1) space Y, obtained by taking an appropriate metric on $Y \times [0, \infty)/Y \times \{0\}$, see [KL, N], but we will not need it.

Now, if Z is a conical subset in X (modeled on (E^2, W_{af})), then at every point $o \in Z$, the set Z is isometric to a star-like subset of the Euclidean cone

$$Cone(Y), \quad Y = \Sigma_o(X).$$

Since X is locally finite, each point $x \in X$ has a convex neighborhood which is conical with respect to x. In particular, every subcomplex in X is locally conical. The converse, of course, is false. Namely, take a geodesic Z in X: Then Z is locally conical but, typically, is not a subcomplex.

The following lemma describing the geometry of X near sub-walls generalizes the above observation. We first need

Definition 4.9. Let X_0 be a building which contains at most one thick wall L, i.e., X_0 is obtained by attaching several half-apartments along L. Then, for each closed subsegment $\gamma \subset L$ and $\epsilon > 0$ we consider $B_\epsilon(\gamma) \subset X_0$. We will refer to $B_\epsilon(\gamma)$ as an *open book with the binding* γ.

Lemma 4.10. *Let X be locally finite, modeled on (A, W_{af}). Let $\gamma \subset X$ be a sub-wall which does not cross (transversally) any thick sub-walls. (We allow γ to be a sub-wall.) Then there exists $\epsilon > 0$ so that $B_\epsilon(\gamma)$ is conical with respect to every $o \in \gamma$. Moreover, $B_\epsilon(\gamma)$ is isomorphic to an* open book with the binding γ.

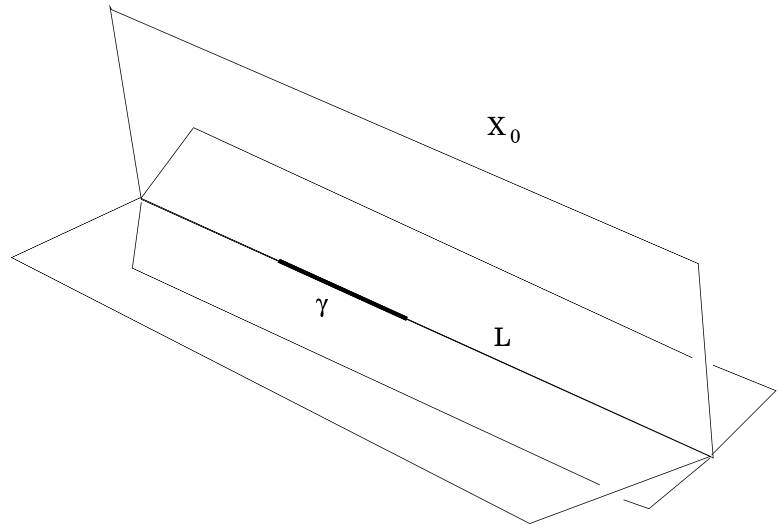

Figure 2: *Open book with the binding L.*

Proof. 1. Since X is CAT(0), the neighborhood $C = B_\epsilon(\gamma)$ is convex. In particular, it is star-like with respect to every point $o \in \gamma$. If two geodesics in C "branch", then branching occurs at a point of γ, since elsewhere, C is locally isometric to $A = \mathbb{R}^2$. If a geodesic $\beta = \overline{ox} \subset C$ intersects γ at a point different from o, then it is entirely contained in γ. The non-branching of geodesics follows.

2. Pick a point $o \in \gamma$ which is not an end-point. We identify X_0 with $Cone(\Sigma_o(X))$. We have the map

$$h : C \to X_0$$

which sends $x \in C \setminus \{o\}$ to the pair (ξ, r), where ξ is the direction of $\overline{ox}$ in $\Sigma_o(X)$ and $r = d(o, x)$. We leave it to the reader to verify that h is an isometry of C to an open book in X_0. □

Developing maps. Let $U \subset X$ (which is modeled on 2-dimensional (A, W_{af})) be a connected subset which is homeomorphic to an open 2-dimensional manifold. We also assume that each point $x \in U$ has a neighborhood which is contained in a single apartment. We thus obtain an *open* atlas $\{(U_\alpha, \phi_\alpha)\}$ on U, whose charts

$$\phi_\alpha : U_\alpha \to \phi_\alpha(U_\alpha) \subset A$$

are isomorphisms. Thus, the transition maps for this atlas are restrictions of the elements of W_{af}. This gives us a geometric structure on U modeled on (A, W_{af}) and, therefore, we obtain a multi-valued developing map $f : U \to A$ which is, locally, a morphism (see e.g. [G]).

Lemma 4.11. *Suppose that U is convex in X. Then f is single-valued and injective.*

Proof. Since U is convex, it is simply-connected. Hence, f is single-valued. In order to check injectivity, note that, being a local isomorphism, f is a local isometry. Let $x, y \in U$

be distinct points. Since U is convex, the (unique) geodesic $\gamma = \overline{xy} \subset X$ is contained in U. Since f is a local isomorphism, it sends local geodesics to local geodesics. Since a local geodesic in A is a global geodesic, $f(\gamma) \subset A$ is a straight-line segment with the end-points $f(x), f(y)$ and the length equal to that one of γ. Therefore, $f(x) \neq f(y)$. □

Definition 4.12. A subset $P \subset X$ is called *planar* if there exists a weakly isometric embedding $f : C \to A$.

5 Corridors and tiled corridors

In this section we will assume that the Coxeter complex (A, W_{af}) is *special*, the group of translations of W_{af} is dense in $\mathbb{R}^2$, and spaces X modeled on this complex are locally finite. Below we define *tiles* and *corridors*, which are the main tools for proving a combinatorial convexity theorem (Theorem 7.1), which, in turn, is the key for constructing the building X as in Theorem 1.1. Tile are analogues of (non-existing) alcoves, while corridors are analogues of (non-existing) galleries, both of which are a important combinatorial tool for studying discrete affine buildings.

We assume from now on that the coordinates on $\mathbb{R}^2$ and the chamber Δ are chosen so that Δ is invariant under the reflection in the x-axis and contains the positive sub-ray in the x-axis. This convention will be useful for interpreting pictures (except for the figures 4, 5 where we draw bridges horizontally), i.e., Δ is rotated by $-\frac{\pi}{2m}$.

A *model tile* in a Coxeter complex (A, W_{af}) is a parallelogram D whose edges are sub-walls parallel to the walls bounding Δ. In particular, the angles of D are of the form $\pi/m, \pi - \pi/m$. We will regard a single sub-wall, with the slope $\pm\tan(\pi/2m)$, as a (degenerate) tile. We let $v_- = v_-(D), v_+ = v_+(D)$ denote the left and right acute vertices of D. (I.e., the D has the angle π/m at these vertices and v_- has the smaller and the vertex v_+ has the larger x-coordinate.) Then D is uniquely determined by the pair (v_-, v_+).

We now define *model corridors* in A. We will frequently omit the adjective "model" to simplify the terminology. Every model corridor C is determined a pair of piecewise-linear functions $u_\pm : [a, b] \subset \mathbb{R} \to \mathbb{R}$ satisfying the following:

1. $u_+(a) = u_-(a), u_+(b) = u_-(b)$. The points $v_- = (a, u_\pm(a)), v_+ = (b, u_\pm(b))$ are (distinct) vertices of A: They will be called the *extreme vertices* (or *extremes*) of the corridor C.

2. $u_-(x) \leq u_+(x)$ for all $x \in [a, b]$.

3. The slopes of the graphs of $u_\pm$ are of the form $\pm\tan(\pi/2m)$.

The function u_+ is called *upper roof function* and u_- is called *lower roof function.* The slope of u_+ at v_- is $\tan(\pi/2m)$ and at v_+ it is $-\tan(\pi/2m)$, for u_- it is the other way around.

Remark 5.1. In Section 8, we will generalize the above concepts to allow tiles of infinite diameter and roof functions defined on infinite intervals.

Definition 5.2. A *(model) corridor* C is the closed region between the graphs of the functions u_+, u_-. A corridor C is called *simple* if either $u_-(x) < u_+(x)$ on the open interval (a, b) or $u_+ \equiv u_-$ on $[a, b]$. In other words, simple corridors are the ones which are homeomorphic to a closed disk or isometric to a segment.

A corridor is *degenerate* if it is a sub-wall, i.e., is a single segment. Thus, every degenerate corridor is simple.

Figure 3 provides examples of corridors.

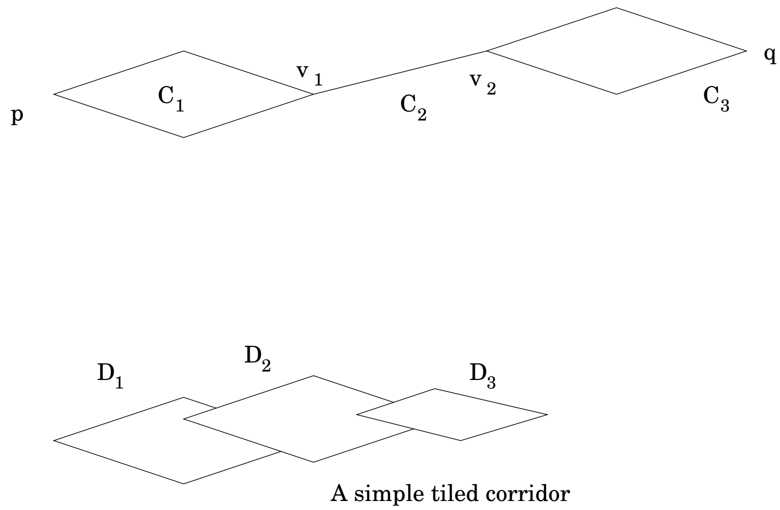

Figure 3: *Corridors.*

We note that each corridor is a union of finitely many simple corridors, every two of which intersect in at most one point, an extreme vertex of both corridors. These intersection points will be called *cut-vertices* of C, as they separate the corridor C. Away from these points, a corridor C is a manifold with boundary (of dimension 1 or 2). This decomposition of C into maximal simple sub-corridors C_i is unique; the sub-corridors C_i are called the *components* of C.

We let $\partial_+ C, \partial_- C$ denote the graphs of u_+, u_-. Then $\partial C = \partial_+ C \cup \partial_- C$.

We note that each corridor with the extremes v_-, v_+ is contained in the tile with the vertices v_-, v_+. It follows immediately from the definition that each edge of the polygon C is a sub-wall and each vertex is a vertex of (A, W_{af}).

Despite the fact that a corridor is not convex (unless it is a single tile), it satisfies the following weak convexity property:

Let L be a line in A with the slope $\pm \tan(\pi/2m)$. Then $L \cap C$ is either empty or is a segment.

Although we will not need this, one can characterize simple nondegenerate corridors in A geometrically (without using the roof functions) as follows:

1. C is a subcomplex in (A, W_{af}).

2. C has only finitely many edges and is homeomorphic to the closed 2-disk.

3. All edges of C have slopes $\pm \tan(\pi/2m)$.

4. C does not have any two consecutive (interior) angles each $> \pi$.

The following lemma is elementary and is left to the reader.

Lemma 5.3. *Suppose we are given $\epsilon > 0$ and a continuous function $u : [a_1, a_2] \to \mathbb{R}$, so that $(a_i, u(a_i))$ are vertices of the Coxeter complex (A, W_{af}). Then there are upper and lower roof functions $u_\pm : [a_1, a_2] \to \mathbb{R}$ so that:*

1. $u_\pm(a_i) = u(a_i), i = 1, 2$.

2. $\|u - u_\pm\| < \epsilon$.

Lemma 5.4. *Let C_s ($s \in \mathbb{R}_+$) be an increasing family of corridors with common extreme points v_-, v_+: $C_s \subset C_t \iff s \le t$. Assume also that the numbers of edges in the corridors C_s are uniformly bounded from above. Then*

$$C = \overline{\bigcup_s C_s}$$

is a corridor in $(A, \overline{W_{af}})$ with the extreme points v_-, v_+.

Proof. Let $u_{\pm,s}$ be the upper/lower roof function for the corridor C_s. Since the corridors C_s are increasing, we obtain:

$$u_{+,s} \le u_{+,t}, \quad u_{-,s} \ge u_{-,t}, \quad \forall s \le t.$$

Moreover, all the roof functions in question are uniformly bounded (from above and below) by the roof functions of the tile with the vertices v_-, v_+. It is then immediate that there exist limits

$$u_\pm = \lim_{s\to\infty} u_{\pm,s}$$

and these limits are roof functions with respect to $(A, \overline{W_{af}})$. □

It is clear from the definition that every point in a corridor C belongs to the interior (relative to C) of a tile $D \subset C$. By compactness, we can therefore cover the corridor C by finitely many tiles.

Definition 5.5. A *tiled corridor* in A is a finite set of tiles $N = \{D_1, ..., D_k\}$ whose union $|N|$ is a corridor C in A, so that N is minimal with this property, i.e., any proper subset $N' \subset N$ will have strictly smaller union. When convenient, we will abuse the notation and conflate N and $C = |N|$. We set $\#N := k$.

We now extend our discussion to corridors in X, a space modeled on (A, W_{af}).

Definition 5.6. A *tile* in X is a subcomplex isomorphic to a (model) tile in A.

We now define corridors in X. They will satisfy two axioms C1 and C2, the first of which is easy to state, the second axiom will take more work.

Axiom C1. Every corridor in X is a subcomplex $C \subset X$ weakly isometric to a model corridor in A.

In particular, if v is a cut-vertex of C, it disconnects C into two sub-corridors C_1, C_2, so that $\Sigma_v(C_i)$ is a chamber or a vertex in $\Sigma_v(X)$, and the distance between these is at least $\pi - \pi/m$.

Using the embedding $f : C \to A$, we will be identifying corridors C in X with corridors in A. Suppose $C \subset X$ is a corridor with extreme vertices v_-, v_+. Since $f(C)$ is contained in a tile with the extreme vertices $f(v_-), f(v_+)$, we obtain

Lemma 5.7. *$C \subset B_d(\overline{v_- v_+})$, where $d = d(v_-, v_+)$.*

We now introduce the second corridor axiom. To motivate the definition, consider a model corridor $C = D_0 \cup D_1 \cup D_2 \subset A$ as in Figure 4, where D_0 is a degenerate tile. Let $\sigma \in W_{af}$ be the reflection in the wall containing D_0. Then the union

$$C' = D_0 \cup D_1 \cup D_2' \subset A, \quad \text{where} \quad D_2' = \sigma(D_2),$$

is not a model corridor in A, but satisfies Axiom C1: The isomorphism $f : C' \to C$ is given by the map which fixes $D_1 \cup D_0$ pointwise and sends D_2' to D_2 via σ. For a variety of reasons, we would like to exclude C' (and other similar "twisted corridors") from being corridors in spaces X modeled on (A, W_{af}). The most obvious reason (although not the critical one) is that we would like every corridor in A to be the image of a model corridor under some $w \in W_{af}$. The key reason for excluding such C' will become clear in the proof of Lemma 6.5.

One can tell apart a "twisted corridor" $C' \subset A$ from a model corridor by observing that there exists a *trapezoid T connecting* the tiles D_1, D_2' as in Figure 4, while a model corridor cannot admit such trapezoids since its existence would mean that boundary slopes of a model corridor would take at least 3 different values. Equivalently, there is a nontrivial continuous family of segments parallel to D_0 with end-points contained in C'. (These segments foliate the trapezoid T.)

Definition 5.8. A *trapezoid* in A is a nondegenerate quadrilateral bounded by sub-walls and having the angles $\pi/m, \pi/m, \pi - \pi/m, \pi - \pi/m$.

We will say that the set of polygons D_1, D_0, D_2', T is a *bridge* in A. By abusing the notation, we will refer to the union of polygons in a bridge as a bridge as well.

We define a *bridge* in a space X (modeled on (A, W_{af})) as a set of polygons whose union is isometric to a bridge in A.

We are now ready to state the second corridor axiom:

Axiom C2. A corridor in C cannot admit a "bridge", i.e., C cannot contain a degenerate sub-corridor D_0 which, together with tiles $D_1, D_2' \subset C$, and a trapezoid $T \subset X$ form a bridge. Equivalently, if D_0 is a degenerate irreducible component of C, there is no nontrivial family of segments in X parallel to D_0 whose end-points are contained in C.

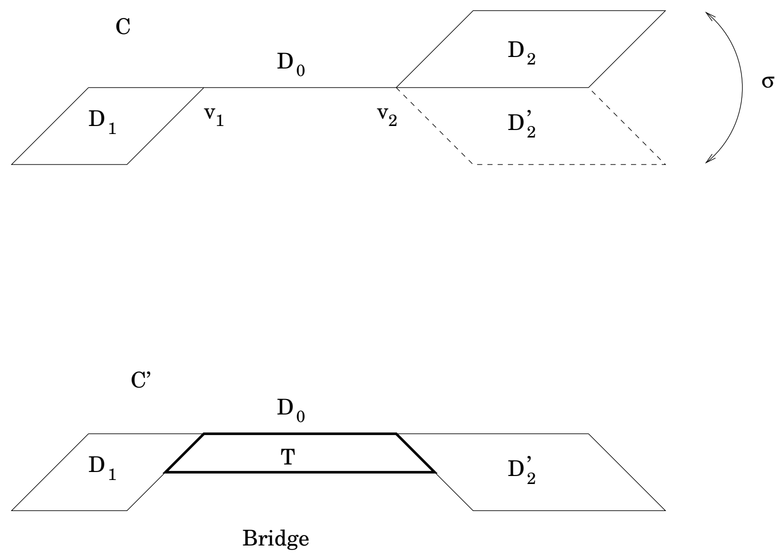

Figure 4: *Corridor, "twisted corridor" and a bridge. Here* $D_0 = \overline{v_1 v_2}$.

Given a corridor $C \subset X$ and its weak isometry $f : C \to C' = f(C) \subset A$, we declare the *boundary* ∂C of C, to be the preimage $f^{-1}(\partial C')$. It is clear that the boundary is independent of the embedding f. We then similarly extend to corridors in X all the concepts defined for model corridors.

Lemma 5.9. *Suppose that* $C \subset X$ *is a corridor and a morphism* $\iota : X \to X'$ *is an isometric embedding. Then* $\iota(C)$ *is again a corridor in* X'.

Proof. Let $f : C \to C' \subset A$ be a weakly isometric embedding. Then $\iota \circ f^{-1} : C' \to \iota(C)$ is again a weakly isometric embedding. This verifies Axiom C1.

Suppose that $\iota(C)$ admits a bridge $D_0 \cup D_1 \cup D_2 \cup T$. Then T is foliated by parallel geodesic segments whose endpoints belong to $D_1 \cup D_2 \subset \iota(X)$. Since $\iota(X)$ is convex in X', it follows that each of these geodesics is contained in $\iota(X)$. In particular, $C \subset X$ admits a bridge. Contradiction. □

Remark 5.10. It important to note that, unless X is a building, if $C \subset X$ is a convex corridor, its image $f(X) \subset A$ is not necessarily convex. This follows from the fact that a weakly isometric embedding is not in general an isometric embedding. The reader can construct examples of convex corridors with nonconvex images $f(C)$ using the example of a space X described in Example 4.6, Part 2.

6 Two orders on tiled corridors

In this section, $X \in \mathcal{C}_{lf}$ (as in Notation 4.4) is a space modeled on (A, W_{af}). We define two orders on the corridors in X. The main result of this section is Proposition 6.15, which shows existence of a maximal tiled corridor N containing the given tiled corridor N_0 (more precisely, $N_0 \ll N$). We also prove Lemma 6.3, showing that maximal corridors are convex in X.

Definition 6.1. Let $N = \{D_i, i \in I\}$ and $N' = \{D'_j, j \in J\}$ be tiled corridors in X with $C = |N|, C' = |N'|$. We say that $N \ll N'$ if:

1. C and C' have the same extreme vertices.
2. For every $D_i \in N$ there exists $D'_j \in N'$ so that $D_i \subset D'_j$.
3. $f_C = f_{C'}|_C$, where $f_C, f_{C'}$ are weak isometries to model corridors in A.

We say that $N < N'$ if $N \ll N'$ and the cardinality of J is at most the cardinality of I, i.e., $\#N' \leq \#N$.

It is easy to see that both $<$ and $\ll$ are (partial) orders.

Remark 6.2. On can replace the order $\ll$ with the inclusion relation between the corridors C, but the proofs will become somewhat more complicated.

A tiled corridor N in X is *weakly maximal* if it is maximal with respect to the order $<$. A tiled corridor N in X will be called *maximal* if it is maximal with respect to the order $\ll$.

Lemma 6.3. *If N is a maximal tiled corridor (with respect to the order $\ll$) in X, then $C = |N|$ is convex in X.*

Proof. Suppose that C is not convex in X. Then, since X is CAT(0), there exists a vertex v of C so that C fails to be locally convex at v. Since all the angles of C are of the form $\pi/m, \pi + \pi/m$, it follows that C has the angle $\pi + \pi/m$ at v. This means that $\Sigma_v(C)$ is not convex in $\Sigma_v(X)$. In other words, if η, ξ denote the tangent directions to ∂C at v, then $\angle(\eta, \xi) < \pi + \pi/m$. Recall that $\Sigma_v(X)$ is a bipartite graph with the edges of the length π/m. Since ξ_1, ξ_2 have different colors, it follows that $\angle(\xi_1, \xi_2) \leq \pi - \pi/m$. Since $\Sigma_v(X)$ has girth $\geq m$, it follows that $\angle(\xi_1, \xi_2) = \pi - \pi/m$. Therefore, there exists a small flat sector $S \subset X$ with the vertex v and $\Sigma_v(S) = \overline{\xi_1\xi_2} \subset \Sigma_v(X)$. Thus, by density of the walls of A, we can find a small tile $D \subset S$ with the vertex v and the angle $\pi - \pi/m$ at v. The union $C' = C \cup D$ is, therefore, again a corridor and $N' = N \cup \{D\}$ is a tiled corridor so that $N \ll N'$. Since $|N'|$ is strictly larger than C, it follows that N is not maximal. Contradiction. ☐

Remark 6.4. In general, weakly maximal corridors (i.e., maximal with respect to the order $<$) need not be convex.

Let X be modeled on (A, W_{af}). We will use the notation $\overline{X}$ for the same metric space X, but regarded as a space modeled on $(A, \overline{W_{af}})$, where $\overline{W_{af}} = \mathbb{R}^2 \rtimes W$ is the closure of W_{af}. Then

Proposition 6.5. *For every tiled corridor N_0 in X there exists a tiled corridor N such that N is weakly maximal in X and $N_0 < N$.*

Proof. Let $N_0 = \{D_{0,0}, ..., D_{k,0}\}$.

1. We will first prove existence of a weakly maximal corridor in $\overline{X}$; in the second part of the proof we will verify that this corridor is actually a corridor in X. By Zorn's lemma, it suffices to consider a maximal totally ordered (with respect to the order $<$) family of tiled corridors

$$\{N_s, s \in J \subset \mathbb{R}_+\},$$

and to show that this family contains a maximal element. Without loss of generality we can assume all tiled corridors have the same cardinality. Then each tiled corridor N_s is of the form

$$\{D_{0,s}, ..., D_{k,s}\},$$

where $D_{i,s} \subset D_{i,t}, s \le t$.

Set $C_s := |N_s|$ and let $f_s : C_s \to A$ be the isomorphism onto a model corridor $C'_s \subset A$, so that

$$f_t|_{C_s} = f_s, \quad s \le t.$$

Then we obtain an increasing family of (model) tiled corridors $N'_s = f(N_s)$.

By Lemma 5.4, the closure C' of the union $\cup_s C'_s$ is again a corridor in A. Defining a tiled corridor

$$N' := \{D'_0, ..., D'_k\}, \quad \text{where} \quad D'_i := \overline{\cup_s D'_{i,s}}, \quad i = 0, ..., k.$$

we get $C' = |N'|$. Similarly, set

$$N := \{D_0, ..., D_k\}, \quad \text{where} \quad D_i := \overline{\bigcup_s D_{i,s}}, \quad i = 0, ..., k, \quad C := |N|.$$

Remark 6.6. We may have to omit some of the tiles in N and N' in order to obtain minimal tiled corridors N, N' with $|N| = C, |N'| = C'$. In order to simplify the notation we will ignore this issue.

Then the maps f_s^{-1} yield a map $h : C' \to C, h(D'_i) = D_i, i = 0, ..., k$. We claim that C is a corridor in $\overline{X}$.

First, let us show that h is 1-1. Let $\gamma \subset X$ denote the geodesic connecting the extreme vertices of C; it is the image under h of a piecewise-linear path $\gamma' \subset C'$ connecting the extreme vertices of C'. Since γ' is contained in all the corridors C'_s, it follows that $h : \gamma' \to \gamma$ is 1-1. For every point $x' \in C'$ let $\alpha(x') \subset C'$ denote the shortest path from x' to γ'. Then $\alpha(x')$ is a Euclidean geodesic (possibly constant) perpendicular to γ'. The map h is clearly a local isometry on every simple subcorridor in C', therefore, it sends each $\alpha(x')$ to a geodesic perpendicular to γ. Suppose $x = h(x'_1) = h(x'_2)$; let $h(\alpha(x'_i)) := \overline{y_i x}$, $y_i \in \gamma$, $i = 1, 2$. Then $h(\alpha(x'_1)) \cup h(\alpha(x'_2)) \cup \overline{y_1 y_2}$ is a triangle in X with two right angles. Since X is CAT(0), such triangle is necessarily degenerate and, hence, $x'_1 = x'_2$. Thus, h is 1-1. Let $f : C \to C'$ denote its inverse.

We then verify that f is a weak isometry. It is clear that the restriction of f to every simple sub-corridor is an isometry[3] to its image, as the limit of isometries, is an isometry. In particular, the restriction of f to each simple sub-corridor is a weak isometry.

[3] with respect to the induced path-metric

Therefore, we have to consider the behavior of f at the cut-vertices. If v is a cut-vertex of C then it is a cut-point of every $C_s, s \in \mathbb{R}_+$. If a cut-vertex v of C was separating nondegenerate sub-corridors for some C_s, the same is true for all $t \geq s$. Moreover, in this case

$$\Sigma_v(C_t) = \Sigma_v(C), \quad \Sigma_{v'}(C'_t) = \Sigma_{v'}(C'), \quad v' = f(v), t \geq s,$$

and

$$f_t : \Sigma_v(C_t) \to \Sigma_{v'}(C'_t)$$

is independent of t for $t \geq s$. Thus, f is a weak isometry at v in this case.

Suppose now that v belongs to a degenerate component, say, $D_{0,s}$, of C_s for all s. Note that the point v need not be an end-point of this geodesic segment.

Remark 6.7. This is the part of the proof where the corridor Axiom C2 will be used.

We will consider only the case when $\Sigma_v(C_s)$ is *not* of the form

$$\{\xi_1, \xi_2\}, \quad \text{where} \ \ \angle(\xi_1, \xi_2) = \delta = \pi - \pi/m$$

for sufficiently large s. The remaining case will be similar and is left to the reader.

If v is not a vertex of $D_{0,s}$ we let $\xi_i \in \Sigma_v X$ $(i = 1, 2)$ denote the antipodal directions tangent to $D_{0,s}$. Otherwise, we define $\xi_i \in \Sigma_v X$ so that one of them (say, ξ_1) is tangent to $D_{0,s}$ and the other is a fixed direction which belongs to $\Sigma_v(C_s)$ (for all s) and is super-antipodal to ξ_1:

$$\angle(\xi_1, \xi_2) \geq \pi.$$

By the construction of C, the germ $\Sigma_v(C)$ at v consists of two components σ_i containing ξ_i $(i = 1, 2)$. Each σ_i is either a chamber or equals ξ_i. Therefore, the distance between σ_1, σ_2 in $\Sigma_v X$ is at least $\pi - \frac{2\pi}{m}$. If the distance is $\geq \delta = \pi - \pi/m$ then f is indeed a weak isometry at v and we are done. Otherwise, the distance equals $\pi - \frac{2\pi}{m}$ and both σ_i are chambers $\overline{\xi_i \eta_i}$: $\sigma_i = \Sigma_v(D_i), i = 1, 2$. Moreover, we necessarily have $\angle(\xi_1, \xi_2) = \pi$. In particular, v separates nondegenerate sub-corridors in C.

Since the distance between σ_1, σ_2 is $\pi - \frac{2\pi}{m}$, there exists a flat triangle $T \subset X$ with the tip v and the angle $\pi - \frac{2\pi}{m}$ at v, so that

$$\Sigma_v(T) = \{\eta_1, \eta_2\}$$

Combining the germ of this flat triangle with the germs (at v) of $D_i, i = 1, 2$, we obtain a bridge in X containing $D_{i,t}, i = 0, 1, 2$ for sufficiently large t. This contradicts Axiom C2 for the corridors C_t. See Figure 5.

The same argument excludes existence of a bridge for C: one can regard the triangle T in the above argument as a degenerate trapezoid.

Since the family of tiled corridors $\{N_s\}$ was chosen to be maximal, it shows that $N = N_s$ for all sufficiently large s and N is a weakly maximal tiled corridor in $\overline{X}$.

2. Note that even if all the corridors C_s were corridors in X, then C might be a corridor only in $\overline{X}$, as the walls of (A, W_{af}) are dense among the walls of $(A, \overline{W_{af}})$. Suppose that one of the edges e of C is not contained in a wall of X. Since its image under f has the slope $\pm \tan(\pi/2m)$, it follows that e cannot contain any vertices of X.

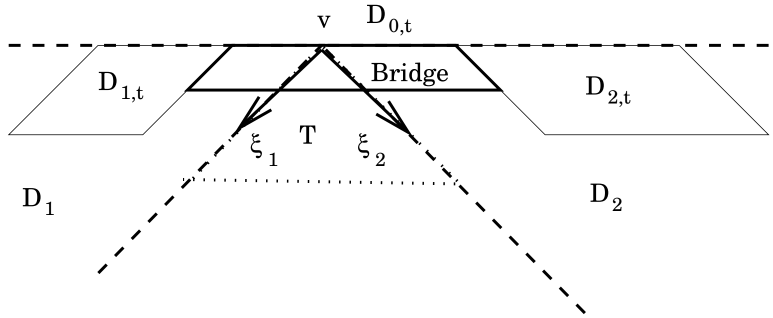

Figure 5: *Building a bridge between $D_{1,t}$ and $D_{2,t}$.*

Remark 6.8. This is the only place of the paper where we *truly need* the fact that (A, W_{af}) is special: All other arguments could be modified to handle the general case.

In particular, e contains no extreme points of C and no thick vertices of X. Therefore, all the tiles of N which overlap with the edge e can be slightly expanded in the direction of this edge to form a new corridor $\hat{C} = |\hat{N}|$ which is strictly larger than C (it will have an edge parallel to e). Thus, $N < \hat{N}, N \neq \hat{N}$ contradicting maximality of N. □

Our next goal is to prove existence of maximal tiled corridors with respect to the order $\ll$. In view of the previous proposition, it suffices to restrict our discussion to weakly maximal tiled corridors.

Let $D, D' \subset A$ be model tiles; we denote the (closed) edges of these tiles $e_1, ..., e_4$ and $e'_1, ..., e'_4$ respectively, so that the edges e_i, e'_i correspond to each other under the dilation mapping D to D'. Given a subset $E \subset A$, we define

$$I = I(D, D', E) := \#\{i \in \{1, ..., 4\} | \exists a \in E \cap e_i \cap e'_i\},$$

which is *the number of sides of the tiles D, D' "sharing" a point from E.* The following lemma is elementary:

Lemma 6.9. *Let D, D' be tiles in A and $E \subset A$.*

1. *If $I(D, D', E) \geq 4$, then $D = D'$.*
2. *If $I(D, D', E) \geq 3$, then either $D \subset D'$ or $D' \subset D$.*

In view of this lemma, we will say that a tile $D \in N$ that has a thick vertex of X (or an extreme point of $|N|$) on every edge is *rigid* and a tile that has such a point on all but one edges is *semi-rigid*. In particular, as X is locally finite, a compact subset of X contains only finitely many rigid tiles.

Lemma 6.10. *Let D be a tile in a weakly maximal corridor N, $|N| = C$, and e is an edge of D which contains neither thick vertices, nor extreme points of C. Then $e \subset \partial C$.*

Proof. If not, we can expand D slightly to the direction of e to a new tile D' (which shares with D three other sides), see Figure 6. Then $N \setminus \{D\} \cup \{D'\}$ is still a tiled corridor in X (which has the same extreme points as N) contradicting weak maximality of N. □

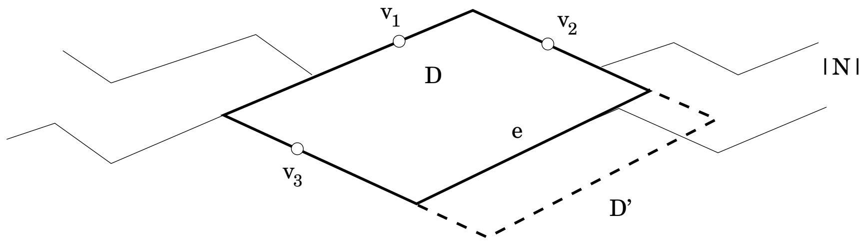

Figure 6: *The edge e is not contained in ∂C and, hence, the tile D can be expanded to a tile D'. The points v_1, v_2, v_3 are the thick vertices on $\partial D \cap \partial C$.*

Suppose that N is a weakly maximal tiled corridor in X, $C = |N|$. Let E denote the union of the set of extreme points $\{v_-, v_+\}$ of N with the set of thick vertices contained in $B_d(\overline{v_- v_+})$, where $d = d(v_-, v_+)$. Let t denote the (finite) cardinality of E. Set $E_C := E \cap C$.

Lemma 6.11. *1. If $D, D' \in N$ are semi-rigid tiles so that $I(D, D', E_C) \geq 3$, then $D = D'$.*

2. Suppose that $D, D' \in N$ are not semi-rigid tiles. If $I(D, D', E_C) \geq 2$, then $D = D'$.

Proof. 1. Suppose $D \neq D'$. By Lemma 6.9, $D \subset D'$ or $D' \subset D$. Therefore, we can eliminate one of the tiles D, D' from the tiled corridor N (without changing $|N|$), contradicting the definition of a tiled corridor.

2. Suppose that, say, D, D' are not semi-rigid and share two elements a_1, a_2 of E, both of which belong to $\partial_+ C$. Then, by Lemma 6.10, $\partial_- D, \partial_- D'$ are both contained in $\partial_- C$. Suppose that $D \neq D'$. Definition of a tiled corridor implies that D cannot contain D' and vice versa. But this means that, $\partial_- D$ has nonempty intersection with the interior of D', see Figure 7. Contradiction.

The case when one of the shared points is on $\partial_+ C$ and the other on $\partial_- C$ is similar and is left to the reader. □

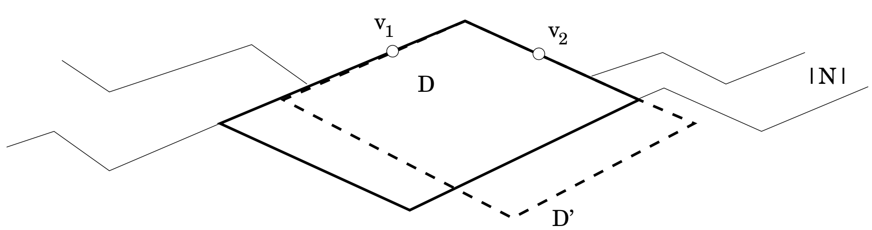

Figure 7: *$\partial_- D$ is not contained in ∂C. The points v_1, v_2 are the thick vertices on $\partial_+ D$.*

The following three corollaries are then immediate:

Corollary 6.12. *Each tile in N is uniquely determined by 2 or by 3 elements of E.*

Corollary 6.13. *Let N be a weakly maximal tiled corridor with the extreme points v_-, v_+. Then number of tiles in N is at most t^3.*

Corollary 6.14. *Given a tiled corridor N in X, there exists only finitely many weakly maximal tiled corridors N' in X so that $N \ll N'$.*

We then obtain:

Proposition 6.15. *For every tiled corridor N_0 there exists a maximal corridor N so that $N_0 \ll N$.*

Proof. By Lemma 6.5, the set of weakly maximal tiled corridors N_i so that $N_0 \ll N_i$ is nonempty. This set is finite according to Corollary 6.14. Therefore, it contains a maximal element. □

7 Combinatorial convexity

We continue with the notation of the previous section. Our next objective is to build convex corridors in X containing the given pair of points. The existence of such corridors is a generalization of the fact that the union of all minimal galleries (in a discrete Euclidean building) connecting any given pair of alcoves, is a convex subset of an apartment.

Theorem 7.1. *Let X be locally finite and modeled on a special Coxeter complex (A, W_{af}), whose group of translations is countable and dense in $\mathbb{R}^2$. Then:*

1. For any two points $x_1, x_2 \in X$, there exists a corridor C in X so that $x_1, x_2 \in C$. Moreover, if x_i is not a thick vertex and $\overline{xy}$ is not a sub-wall near x_i, then we can choose C so that $x_i \notin \partial C$.

2. For every corridor C in X, there exists a convex corridor C' so that $C \subset C'$.

The proof of this theorem will occupy the rest of the section.

Remark 7.2. 1. Since X contains only countably many walls, it contains only countably many (convex) subcomplexes.

2. The above theorem is clear in the case when X is a building, as any two points belong to an apartment. In this case, we can use $C = C'$, a single tile.

3. If X is not a building, the convex corridor C' is, in general, not a tile. Accordingly, the weakly isometric embedding $f : C' \to A$ will not send C' to a convex subset of A. See Remark 5.10.

4. The theorem is probably also true for arbitrary CAT(0) locally Euclidean complexes (not necessarily 2-dimensional ones) provided that we allow cells in C to be arbitrary convex Euclidean polyhedra. However, requiring C in Theorem 7.1 to be planar, imposes serious restrictions on the allowed subcomplexes.

5. In the main application of the above theorem, X will be finite, i.e., covered by finitely many apartments.

We prove this theorem by:

i. Constructing a corridor C containing x_1 and x_2. We will subdivide the geodesic $\gamma = \overline{x_1 x_2}$ into subsegments γ_i not passing through thick vertices. We then construct for each γ_i a simple (possibly degenerate) corridor C_i, so that $\gamma_i \subset C_i$ and the end-points of γ_i are the extreme vertices of C_i. The corridor C will be the union of the sub-corridors

C_i. The most difficult part of the proof is constructing a *planar neighborhood* U_i of γ_i. The corridor C_i will be contained in this neighborhood.

ii. Given a tiled corridor N with $|N| = C$, we will take a maximal tiled corridor N' so that $N \ll N'$. In particular, $C \subset C' = |N'|$. We then verify convexity of C'.

1. Constructing simple corridors. We first assume that $\gamma = \overline{x_1 x_2}$ is a geodesic with distinct end-points, which contains no thick vertices in its interior and is not contained in a wall.

Remark 7.3. If γ is contained in a wall $L \subset X$, then $L \subset A$ is a planar set and we obtain an isometric embedding

$$\gamma \hookrightarrow L \hookrightarrow A.$$

Any sub-wall $e \subset L$ containing γ, is therefore a degenerate convex corridor in X containing γ.

Planar neighborhoods. At first, we assume, in addition, that neither x_1, nor x_2 is a thick vertex and that γ does not cross a thick wall. Then a *planar neighborhood* of $\gamma \subset X$ is a connected closed planar subset $\overline{U}$ of X whose interior U (in X) contains γ.

Lemma 7.4. *Under the above assumptions, γ has a convex planar neighborhood $\overline{U}$. In particular, an embedding $U \hookrightarrow A$ will be distance-preserving.*

Proof. Let ϵ be the minimal distance from γ to the union of thick sub-walls of X. Our assumptions imply that $\epsilon > 0$. We then let U be the open ϵ-neighborhood $B_\epsilon(\gamma)$ and $\overline{U}$ be its closure in X. Then both U and $\overline{U}$ are convex in X, since X is a CAT(0) space. Moreover, by the choice of ϵ,

$$\overline{B_\epsilon(\gamma)} = \bigcup_{x \in \gamma} \overline{B_\epsilon(x)}$$

where each $\overline{B_\epsilon(x)}$ is planar. Therefore, by considering the developing map of B_ϵ, and using convexity of this set we obtain a morphism $f : B_\epsilon(\gamma) \to A$. Convexity of U implies that f is an isometric embedding. Therefore, it extends to an isometric embedding $\overline{U} \to A$. □

Relative planar neighborhoods of γ. We will have to generalize the concept of a planar neighborhood to take into account the cases when one of the x_i's is a thick vertex, or when γ crosses a thick sub-wall: In both cases a planar neighborhood of γ does not exist, so we will have to modify the definition.

Definition 7.5. A *relative* planar neighborhood of γ is a closed connected subset $\overline{U} \subset X$ (the closure of some $U \subset X$) so that the following holds:

1. $x_1, x_2 \in \overline{U}, \gamma \setminus \{x_1, x_2\} \subset U$.
2. If x_i is not a thick vertex, then U contains x_i.
3. If x_i is a thick vertex, then (for sufficiently small $\epsilon > 0$) $S_i = \overline{U} \cap \overline{B_\epsilon(x_i)}$ is conical with respect to x_i and $\Sigma_x(S_i)$ is a chamber in $\Sigma_{x_i}(X)$.
4. There exists an injective isometric morphism $f : U \to A$ onto an open subset of A. (Note that we do not require U itself to be open in X.)

Figure 8 gives an example of relative neighborhood. We observe that if γ is a geodesic satisfying the assumptions of Lemma 7.4, then the notions of a planar neighborhood and of a neighborhood coincide.

Lemma 7.6. *Every geodesic $\gamma = \overline{x_1 x_2} \subset X$ (which has no thick vertices in its interior) has a convex relative planar neighborhood.*

Proof. First of all, without loss of generality we may assume that for both $i = 1, 2$, either x_i is a thick vertex or it is not contained in a wall (otherwise, we can extend γ slightly in the direction of x_i).

Let $\epsilon > 0$ be such that $B_\epsilon(\gamma)$ intersects only those thick sub-walls which cross γ, possibly at its endpoint(s). As before, $B_\epsilon(\gamma)$ is convex, however, it is no longer planar (unless we are in the situation of Lemma 7.4). Our goal is to decrease this neighborhood of γ to make it planar and to satisfy the rest of the properties of a relative neighborhood.

Suppose that x_i is a thick vertex. Consider the germ of γ at x_i. Since γ is not contained in a wall, it follows that this germ determines a regular point $\gamma'(x_i)$ of $\Sigma_{x_i}(X)$. Therefore, $\gamma'(x_i)$ is contained in a unique chamber $\sigma_i \subset \Sigma_{x_i}(X)$. Hence, there exists a closed conical (with respect to x_i) subset

$$S_i \subset \overline{B_\epsilon(x_i)},$$

so that

$$\Sigma_{x_i}(S_i) = \sigma_i.$$

Thus S_i is isomorphic to an appropriate subset of

$$Cone(\sigma_i).$$

In particular, S_i is planar. We then remove from $\overline{B_\epsilon(\gamma)}$ the subset

$$\overline{B_\epsilon(x_i)} \setminus S_i.$$

If x_i is not a thick vertex, we do no not modify $B_\epsilon(\gamma)$ near x_i. This ensures that the resulting convex subset $U' \subset B_\epsilon(\gamma)$ satisfies Parts 1, 2, 3 of the definition. We also observe that U' were planar if the interior of γ did not cross any thick sub-walls.

We now deal with the thick sub-walls which γ might cross. Let $L_1, ..., L_k$ denote the thick sub-walls of X crossed by the interior of γ at the points $z_1, ..., z_k$. Then U' is locally planar except at the points of the geodesics $e_i = L_i \cap U'$: Near those points, U' is not even a manifold, but looks like an "open book" with the binding e_i. Our goal is to "tear off" all but two pages of this open book (for each $i = 1, ..., k$), i.e., the "pages" which are not tangent to γ.

Consider the link $\Sigma_{z_i}(X)$. Since z_i is not a thick vertex, this link is a building: This graph has exactly 2 thick vertices, which are antipodal points representing the tangent directions to e_i at z_i. The tangent directions to γ at z_i are represented by two *regular* antipodal points in $\Sigma_{z_i}(X)$. Therefore, there exists a unique apartment $\sigma \subset \Sigma_{z_i}(X)$ containing these regular antipodal points. We now remove from U' the union of those open geodesic segments $z_i u$ which are *not* tangent to σ. (The closure of the removed set is conical with respect to

z_i.) We do this for each $z_i, i = 1, ..., k$. See Figure 8. The resulting set U is still convex in X. Now, however, it is a topological 2-manifold. We construct the injective morphism $f : U \to A$ as in the proof of Lemma 7.4. The fact that the image of f is open follows from the invariance of domain theorem. □

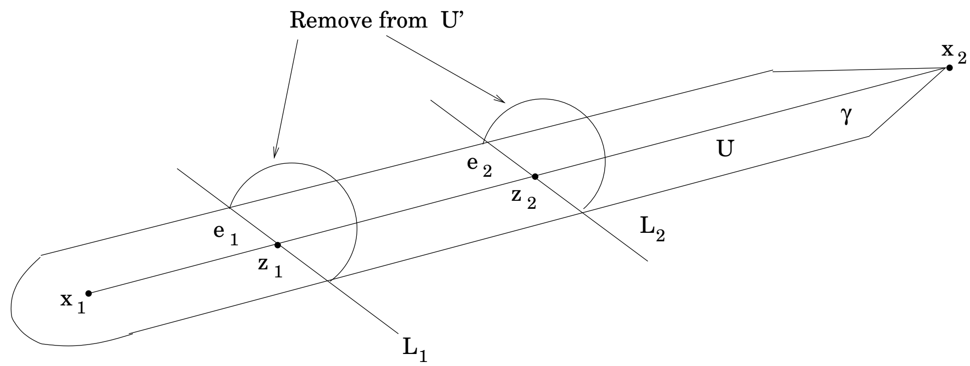

Figure 8: *Constructing a planar neighborhood.*

Corollary 7.7. *For any pair of points $x_1, x_2 \in X$ so that the open geodesic segment $x_1 x_2$ is not a sub-wall and contains no thick vertices, there exists a simple corridor $C \subset X$ containing x_1, x_2. If x_i is a vertex then we can choose C that has x_i as one of its two extremes. If x_i is not a thick vertex, we can also choose C so that $x_i \notin \partial C$.*

Proof. Let $\overline{U}$ be a convex (relative) planar neighborhood of the geodesic $\gamma = \overline{x_1 x_2}$, given by the above lemma.

1. We begin by considering the case when both x_i are thick vertices; thus, the germ of $\overline{U}$ near x_i is a sector S_i (with the angle π/m at x_i). Since $\overline{U}$ is planar, it suffices to consider the case when it is contained in A. Moreover, we can assume that it is embedded in A so that the slopes of the boundary of S_1 have the form $\pm \tan(\pi/2m)$. Then, the same will be true for the slopes of ∂S_2. The geodesic segment $\gamma \subset A$ is the graph of a linear function $u : [a_1, a_2] \to \mathbb{R}$, with

$$x_i = (a_i, u(a_i)), \quad i = 1, 2.$$

Now, approximate u (from above and from below) by upper and lower roof functions $u_\pm$ using Lemma 5.3. Note that the graphs of the germs of $u_\pm$ near a_i are the boundary segments of S_i, $i = 1, 2$. By taking a sufficiently fine approximation, we obtain that the graphs of $u_\pm$ are contained in $\overline{U}$. Therefore, the region bounded by these graphs is the required simple corridor.

2. Suppose that neither x_1 nor x_2 is a thick vertex and, hence, $x_1, x_2 \in U$. If the slope of $\gamma \subset A$ represents a vertex in (S^1, W), we can as well assume that this slope equals $\pm \tan(\pi/2m)$. In this case, using convexity of U, it is easy to construct a single ("nearly degenerate") tile $D \subset U$ containing x_1, x_2. We therefore assume that the slope of γ is in the open interval $(-\tan(\pi/2m), \tan(\pi/2m))$. Pick disjoint congruent tiles $D_i \subset U$ containing $x_i, i = 1, 2$. Let y_i denote the extreme vertices of the smallest tile $D \subset A$ containing $D_1 \cup D_2$. Then $y_i \in D_i, i = 1, 2$.

Remark 7.8. If x_i is a vertex, we can, if necessary, choose D_i so that $y_i = x_i$.

We assume that $a_1 < a_2$ are the projections of y_1, y_2 to the horizontal axis. The convex hull

$$H = Hull(D_1 \cup D_2) \subset U \subset A$$

is a hexagon, with four edges $e_i^{\pm} \subset \partial_{\pm} D_i$ incident to y_i, $i = 1, 2$ and the remaining two parallel edges $d^{\pm}$, so that the boundary of H is the union of two arcs

$$e_1^+ \cup d^+ \cup e_2^+, \quad e_1^- \cup d^- \cup e_2^-.$$

These arcs are the graphs of functions $h^{\pm} : [a_1, a_2] \to \mathbb{R}$. Approximating $h^{\pm}$ by the respectively upper and lower roof functions $u_{\pm}$ we obtain a simple model corridor $C \subset U$ containing x_1, x_2. Since $C \subset A$ is a simple model corridor, as a subcomplex in X is automatically satisfies Axiom C2.

3. The construction in the case when one of the points x_1, x_2 is a vertex and the other is not, is a combination of the two arguments in Case 1 and 2, and is left to the reader. □

2. Constructing general corridors. We now drop the assumption that the open geodesic x_1x_2 does not contain any thick vertices.

Proposition 7.9. *For every pair of points $x_1, x_2 \in X$, there is a corridor C containing x_1, x_2. If x_i is not a thick vertex and the germ of the geodesic $\gamma = \overline{x_1x_2}$ at x_i is a regular direction in $\Sigma_{x_i}(X)$, we can assume that x_i does not belong to ∂C.*

Proof. We subdivide the geodesic γ into maximal subsegments $\gamma_1 \cup ... \cup \gamma_k$, so that the interior of each $\gamma_i = \overline{y_iy_{i+1}}$ does not contain any thick vertices. According to Corollary 7.7, there exist simple (possibly degenerate) corridors $C_1, ..., C_k$ so that each C_i contains $\overline{y_iy_{i+1}}$. Moreover, for $1 < i < k$, the vertices y_i, y_{i+1} are the extremes of C_i. If C_1 or C_k is non-degenerate and the corresponding point x_i in this corridor is not a thick vertex, by Corollary 7.7, we can assume that x_i does not belong to the boundary of this corridor.

Let C denote the union $C_1 \cup ... \cup C_k$. We claim that C is the required corridor.

For each i we have an isomorphic (and isometric) embedding $f_i : C_i \to A$, whose image is a simple model corridor $C_i' \subset A$. We combine these maps to obtain a weakly isometric embedding $f : C \to A$ as follows: We choose f_1 arbitrarily. We then normalize each f_i (by a translation in W_{af}), so that

$$y_i' := f_i(y_i) = f_{i+1}(y_i)$$

and y_i' is the extreme rightmost point of C_i' and extreme leftmost point of C_{i+1}'. Since each C_i' is a model corridor, this ensures that the germs of C_i', C_{i+1}' at y_i' are antipodal in $\Sigma_{y_i'}(A)$, provided that both are chambers. In any case,

$$\delta = \pi - \pi/m \leq \angle(\Sigma_{y_i'}(C_i'), \Sigma_{y_{i+1}'}(C_{i+1}')) \leq \pi$$

with equality to δ, except for the case when both C_i', C_{i+1}' are degenerate and their union is a geodesic segment in A. In the latter case the angle in A equals π and the angle in $\Sigma_{y_i}(X)$ is $\geq \pi$, since γ is geodesic.

We thus obtain a morphism $f : C \to A$ whose restriction to each C_i is f_i. For each i, the convex hull of C'_i is a tile $D'_i \subset A$ with the extreme vertices y'_i, y'_{i+1}. By the construction, for $i \neq j$, $D'_i \cap D'_j$ is at most one point (the common extreme vertex). Therefore, f is injective.

We next need to verify the weak isometry at the vertices y_i. Suppose that both C_i, C_{i+1} are nondegenerate corridors. Then their germs at y_i are chambers $\sigma_i, \sigma_{i+1} \subset \Sigma_{y_i}(X)$. Since γ is a geodesic, its angle at y_i is $\geq \pi$. Therefore, the chambers σ_i, σ_{i+1} are *super-antipodal*: The angular distance between these chambers in $\Sigma_{y_i}(X)$ is $\geq \delta = \pi - \frac{\pi}{m}$. On the other hand, the germs of C'_i at $y'_i = f(y_i)$ are antipodal in $\Sigma_{y'_i}(A) = S^1$. Therefore, the map

$$f : \Sigma_{y_i} C \subset \Sigma_{y_i} X \to \Sigma_{y'_i}(A)$$

is a δ-isometry. Thus, $f : C \to A$ is a weakly isometric embedding.

Suppose that C_i is degenerate, but C_{i+1} is not. The germ σ_i represents the germ of the geodesic γ_i at y_i and σ_{i+1} is a chamber containing the germ of the geodesic γ_{i+1} at y_i. Therefore, since σ_{i+1} contains a point antipodal to σ_i, the angular distance in $\Sigma_{y_i}(X)$ between the vertex σ_i and the chamber σ_{i+1} is $\geq \delta$. The angular distance between their images $f(\sigma_i), f(\sigma_{i+1})$ in S^1 is exactly δ. Therefore, f is a weak isometry at $\Sigma_{y_i}(C)$.

The argument when both C_i, C_{i+1} are degenerate follows from our discussion of the angles between C_i, C_{i+1} and C'_i, C'_{i+1}.

Thus, C satisfies Axiom C1 of corridors. We need to check Axiom C2. Suppose that C admits a bridge

$$D_0 \cup D_1 \cup \cup D_2 \cup T,$$

where $D_0 = \overline{v_1 v_2}$ is degenerate. Then, by the construction of the corridor C, one of the v_1, v_2 (say, v_1) is a vertex y_i and $D_0 \subset C_{i+1}$ (one can actually assume that $D_0 = C_{i+1}$). Therefore, the germ of the geodesic γ at v_1 is contained in $D_1 \cup D_0$. Since $T \cup D_1$ is a planar region, this means that the angle of γ at $v_1 = y_i$ is less than π, contradicting the fact that γ is a geodesic. □

Corollary 7.10. *Every corridor $C \subset X$ is contained in a convex corridor $C' \subset X$.*

Proof. Let $C = |N|$. Let N' be a maximal corridor so that $N \ll N'$ (see Proposition 6.15). Then $C \subset C' = |N'|$. The corridor C' is convex by Lemma 6.3. □

Remark 7.11. Although we do not need this, given N so that $C = |N|$, we can find a canonical convex corridor C' containing C. Namely, the set of weakly maximal tiled corridors N_i so that $N \ll N_i$ is finite according to Corollary 6.14. The number of maximal corridors among the corridors N_i is finite as well. Let C' denote the intersection

$$C' = \bigcap_i |N_i|$$

taken over all maximal corridors N_i. One can verify that C' is indeed a corridor. Convexity of C' follows from the fact that it is intersection of convex sets.

We can now finish the proof of Theorem 7.1.

Proof. The first assertion of Theorem follows from Proposition 7.9. The second assertion is Corollary 7.10. □

8 Generalizations of the convexity theorem

In this section we prove two generalizations of Theorem 7.1. We first extend this theorem to allow corridors of infinite diameter (Theorem 8.2). This theorem will be used in order to verify Axiom $2'$ of buildings. We conclude the section by proving the existence of a convex corridor containing not only given pair of points x_1, x_2 but also chambers $\sigma_i \in \Sigma_{x_i}(X)$ (Theorem 8.3). This theorem will be used to check Axiom $1'$ of buildings.

We start by discussing extended corridors (corridors of infinite diameter).

We first have to generalize the notion of a *tile*. We define *extended (infinite) tiles*, which are infinite sectors bounded by walls with the slopes $\pm\tan(\pi/2m)$ and having the vertex angle π/m or $\pi - \pi/m$. We will also allow the whole apartment as a single tile. Such infinite tiles appear as limits (in the Chabauty topology on the space of closed subsets of A) of sequences of the ordinary tiles. Accordingly, we generalize the upper and lower roof functions to allow functions defined on $\mathbb{R}$. We still require such piecewise-linear functions to have only finitely many break-points: Every roof function admits a partition of $\mathbb{R}$ in finitely many subintervals, where the function is linear. Then, a *extended model corridor* is the closed region bounded from above and below by graphs of upper and lower roof functions; we also declare A to be a corridor (consisting of a single infinite tile). To simplify the terminology, we will retain the name *corridor* for extended corridors.

It is clear that each corridor $C \subset A$ is a union of finitely many tiles, two of which are infinite, with the vertex angles π/m. Therefore, for every corridor C, there exists a extended tiled corridor N so that $|N| = C$: The corridor N still has only finitely many tiles, but some of them are allowed to be infinite.

Thus the "extremes" of the extended corridors C are points at infinity of C. Namely, if $\gamma \subset A$ is any complete geodesic contained in C, we declare the points at infinity of γ to be the extremes of C. Therefore, C now has (typically) more than 2 extreme points.

Let $X \in \mathcal{C}$, a space modeled on (A, W_{af}) and covered by finitely many apartments. (It is no longer enough to assume local finiteness.) We define (infinite) corridors in X by repeating the definition for the ordinary corridors.

Proposition 8.1. *Every complete geodesic $\gamma \subset X$ is contained in a corridor.*

Proof. The arguments remain pretty much the same as in the proof of Proposition 7.9, except that we start by observing that, for sufficiently large t, the subrays $\gamma(-\infty, -t), \gamma(t, \infty)$ in γ do not cross any walls in X. Therefore, there are infinite tiles $D_1, D_2 \subset X$ containing these subrays.

Then $\gamma \setminus (D_1 \cup D_2)$ is a finite geodesic segment and we can cover it with a tiled corridor C using Proposition 7.9. One then shrinks the tiles D_1, D_2, to D_1', D_2' so that

$$C' = D_1' \cup C \cup D_2'$$

is a corridor covering γ, see Figure 9. □

The definitions of the orders $<$ and $\ll$ on corridors generalize almost verbatim to the corridors having infinite tiles. The only modification one needs is that instead of dealing

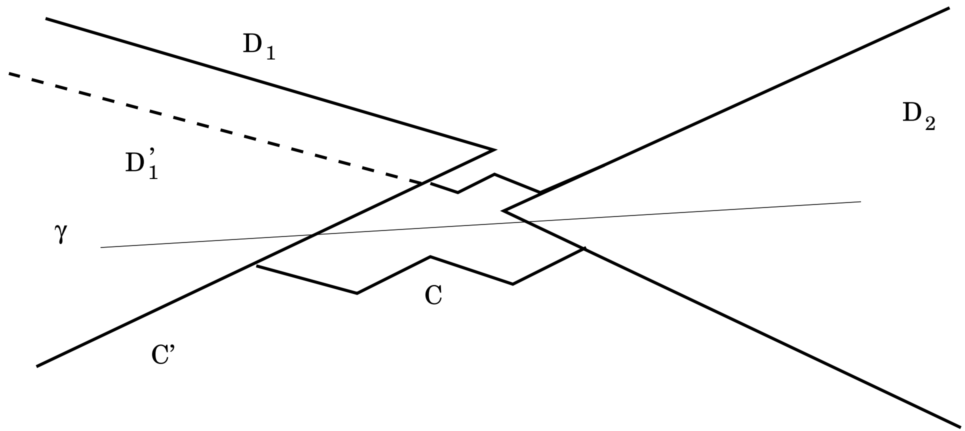

Figure 9: *Constructing an extended corridor.*

with finite extreme vertices, we require an extended corridors to contain certain points at infinity. Equivalently, one can ask for corridors covering a fixed complete geodesic in X.

One again sees that for every extended tiled corridor N, there exists an extended weakly maximal tiled corridor N' so that $N < N'$: While considering an infinite increasing sequence of tiled corridors we no longer can argue that the tiles have uniformly bounded diameter, but such unbounded sequences of tiles still converge to infinite tiles. (The limit could be the whole apartment A.) Given a weakly maximal extended corridor N we construct a maximal extended corridor N' so that $N \ll N'$ by appealing to the fact that X has only finitely many thick vertices.

The proof that every maximal extended corridor is convex is exactly the same as that one of Lemma 6.3. By combining these observations, we obtain

Theorem 8.2. *1. Every complete geodesic $\gamma \subset X$ is contained in an extended corridor C.*

2. Every extended corridor C is contained in a maximal extended corridor C'. The corridor $C' \subset X$ is convex.

3. Let $\Delta_i, i = 1, 2$ be Weyl sectors which share the tip $x \in X$ and which are super-antipodal at x. Then there exists an extended corridor $C \subset X$ so that

$$\Delta_1 \cup \Delta_2 \subset C.$$

Moreover, the corridor C contains a geodesic with the ideal points in $\partial_\infty \Delta_i, i = 1, 2$.

Proof. It remains to prove Part 3. We start with the case when x is a vertex and, hence, Δ_i are bounded by sub-walls. Then, $\Delta_i, i = 1, 2$, are extended tiles in X and, by our assumptions, $C = \Delta_1 \cup \Delta_2$ is an extended corridor. The geodesic contained in this corridor is the union of rays ρ_i which are super-antipodal at x and asymptotic to points in $\Delta_i, i = 1, 2$.

If x is not a vertex, Δ_1, Δ_2 have to be antipodal at x. Moreover, the boundary rays of Δ_i are not sub-walls. In particular, they do not contain thick vertices. Therefore, there exist (special) vertices $x_i \in \Delta_{i+1}$ close to x and Weyl sectors Δ'_i with the tips x_i, so that

$$\Delta_i \subset \Delta'_i, i = 1, 2$$

and

$$\Delta_1' \cap \Delta_2' = D'$$

is a tile. Therefore, we obtain an extended corridor $C = \Delta_1' \cup \Delta_2'$ containing $\Delta_1 \cup \Delta_2$. $\square$

We now return to the discussion of ordinary corridors. In the following theorem it suffices to assume that X is locally finite, modeled on (A, W_{af}).

Theorem 8.3. *For $x_1, x_2 \in X$ let σ_i be chambers in the buildings $\Sigma_{x_i}(X)$, which are given structures of thick buildings, $i = 1, 2$. Then there exists a convex corridor $C \subset X$ containing x_1, x_2, so that*

$$\sigma_i \subset \Sigma_{x_i}(C), i = 1, 2.$$

Proof. As before, set $\gamma = \overline{x_1 x_2}$. We start with a convex corridor C containing x_1, x_2, given by Theorem 7.1. We modify C twice: First to ensure

$$\sigma_1 \subset \Sigma_{x_1}(C)$$

and next, to ensure

$$\sigma_2 \subset \Sigma_{x_2}(C).$$

We describe only the first modification, since the second is obtained by applying the same procedure to the modified convex corridor and switching the roles of x_1, x_2. To simplify the notation, set $x := x_1$, $\sigma := \sigma_1$. We let η be the tangent direction to γ at x.

Case 1 (the generic case): x is not a thick vertex, is not contained in a thick sub-wall and the germ of γ at x is not contained in a sub-wall. In particular, $\sigma = \Sigma_x(X)$ is the apartment. Then, according to Theorem 7.1, we can assume that $x \notin \partial C$. Therefore $\sigma = \Sigma_x(X) = \Sigma_x(C)$.

Case 2: We now allow x to be contained in a thick sub-wall, but retain the rest of the restrictions. Since x is not a thick vertex, $\Sigma_x(X)$ is a building (but not a single apartment). The chamber σ is a half-apartment in $\Sigma_x(X)$ and η is a regular direction.

(a) If $\eta \in \sigma$, then $\sigma = \Sigma_x(C)$ and we are done.

(b) If $\eta \notin \sigma$, then, there exists a (necessarily regular) direction $\xi \in \sigma$ antipodal to η and we can replace γ with a slightly longer geodesic $\tilde{\gamma} = \overline{\tilde{x} x_2}$ tangent to ξ. The convex corridor $\tilde{C}$ containing $\tilde{\gamma}$ will then contain x in its interior and, therefore

$$\gamma \subset \tilde{C}, \sigma \subset \Sigma_x(\tilde{C}).$$

Hence, we are done in this case as well.

Case 3: x is a thick vertex, but η is a regular direction, i.e., the germ of γ at x is not contained in a sub-wall.

(a) Suppose that

$$\angle(\xi, \eta) \leq \pi, \quad \forall \xi \in \sigma.$$

This means that σ is contained in an arc $\tau = \overline{\eta\zeta} \subset \Sigma_x X$ of the length $\leq \pi$. Since $\Sigma_x(X)$ has extendible geodesics, we can assume that the length of τ is π and, hence, ζ is antipodal to η. We now use the same arguments as in case 2a: Extend γ to a longer geodesic $\tilde{\gamma}$ tangent to ζ. Let $\tilde{C}$ be the convex corridor containing $\tilde{\gamma}$. By convexity, $\Sigma_x(\tilde{C})$ will contain a geodesic in $\Sigma_x(X)$ connecting η, ζ. Since η is regular, this geodesic is unique and equals τ. In particular, $\sigma \subset \Sigma_x(\tilde{C})$.

(b) If (a) fails then σ contains a regular direction ξ so that

$$\angle(\xi, \eta) \geq \pi.$$

We again repeat the arguments in the case 2(b): Take a longer geodesic $\tilde{\gamma}$ tangent to ξ. Then the corresponding convex corridor $\tilde{C}$ will contain a germ of σ at x.

Case 4: η is a singular direction and the germ of γ at x is contained in a sub-wall. Then C contains a degenerate irreducible component C_1 so that $x \in C_1$. We have the same dichotomy as in Case 3:

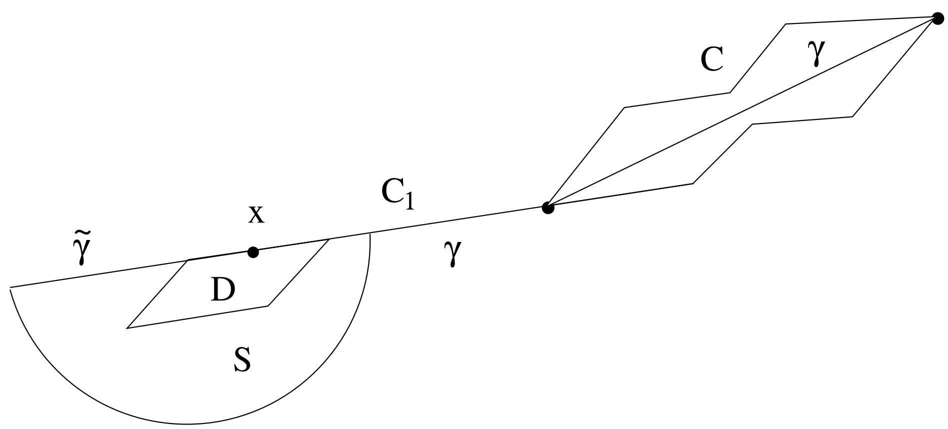

Figure 10: *Constructing the corridor* $\tilde{C} = D \cup C$.

(a) Suppose that σ is contained in an arc $\tau = \overline{\eta\zeta} \subset \Sigma_x X$ of the length π. We construct $\tilde{\gamma}$ as before. There exists a flat sector S (germ of a half-apartment) with the tip x so that $\tau = \Sigma_x(S)$. We then find a tile $D \subset S$ so that $\tau = \Sigma_x(D)$, i.e., one of the edges of D is contained in $\tilde{\gamma}$, so that $C_0 = D \cup C_1$ is a corridor. See Figure 10. Set

$$\tilde{C} := D \cup C.$$

By the construction, $\sigma \subset \tau = \Sigma_x(D) = \Sigma_x(\tilde{C})$. We claim that $\tilde{C}$ is a corridor. We obtain a weakly isometric embedding $\tilde{C} \to A$ by combining the weakly isometric embeddings of C_0 and C. If $\tilde{C}$ were to admit a bridge, there would have to be at least two vertices where $\tilde{C}$ fails to be locally convex. One of them is in C_0. Since C is convex, it does not have any nonconvex vertices. Thus, $\tilde{C}$ is a corridor.

(b) If σ contains a regular direction ξ so that $\angle(\xi\eta) \geq \pi$, we repeat the arguments in Case 3(b). $\square$

9 Enlargement along convex corridors

In this section we define the procedure of *enlargement* of spaces $X \in \mathcal{C}$, consisting of attaching an apartment A along a convex corridor C in X. We then verify that X' is again CAT(0). We also show that the embedding $X \to X' = X \cup_C A$ is a weak isometry which sends corridors to corridors and preserves convexity of apartments.

For a space $X \in \mathcal{C}$ we define $Corr(X)$, the set of all *convex extended corridors* in X. In this section we will frequently omit the adjective *extended* since almost all the discussion of corridors will be local.

Definition 9.1. Given $X \in \mathcal{C}$, we define an *enlargement of X along a (extended) convex corridor* $C \subset X$ to be the space $X' = X_C = X \cup_C A \in \mathcal{C}$ obtained by attaching to X the apartment A along $C \subset X$. The attachment is given by a weakly isometric embedding $f : C \hookrightarrow A$. We will also use the term *enlargement* for the tautological embedding $\iota : X \to X'$.

It is important to note that every wall in X is an extended corridor, therefore, enlargements include attaching apartments along walls.

Lemma 9.2. *If X' is obtained from X by enlargement, then X' is CAT(0).*

Proof. Since C is contractible, X' is simply-connected. Since X' is a locally-Euclidean cell complex, it suffices to check that the links of vertices v of X are CAT(1) spaces. Let v be a vertex of C. If $f(C)$ is locally convex at $f(v)$ then

$$\Sigma_v(X') = \Sigma_v(X) \cup_Z \Sigma_{f(v)}(A)$$

is again CAT(1), since $Z = \Sigma_v(C)$ is convex in both $Y = \Sigma_v(X)$ and $\Sigma_{f(v)}(A)$.

Consider therefore, a vertex $f(v)$ where $f(C) \subset A$ is not locally convex. Then $Y' = \Sigma_v(X')$ is obtained from $Y = \Sigma_v(X)$ by attaching either one arc α (if v is not a cut-vertex of C') or two arcs α_1, α_2 (if v is a cut-vertex); in either case, the length of α (and α_1, α_2) is greater than or equal to $\delta = \pi(1 - \frac{1}{m})$. We will consider the case of one arc since the case of two arcs is similar. Since $C \subset X$ is convex, the arc α is attached to vertices ξ_1, ξ_2 of Y, so that $\angle_Y(\xi_1, \xi_2) \geq \pi(1 + \frac{1}{m})$ (which is the inner angle of C at v). Let $\sigma \subset Y'$ be the shortest embedded circle. If σ does not contain α, then $\sigma \subset Y$ and $length(\sigma) \geq 2\pi$. If σ contains α, then $\sigma = \alpha \cup \beta$, where β is an arc in Y connecting ξ_1, ξ_2. Then, $length(\beta) \geq \pi(1 + \frac{1}{m})$ and, hence

$$length(\alpha) \geq \pi(1 - \frac{1}{m}) + \pi(1 + \frac{1}{m}) = 2\pi.$$

Thus, Y' is CAT(1) and X' is CAT(0). $\square$

We therefore obtain a category $\mathcal{E}$, whose objects are element of $\mathcal{C}$ (i.e., finite spaces modeled on (A, W_{af})) and whose morphisms, $Mor(\mathcal{E})$, are compositions of enlargements.

Lemma 9.3. *Every enlargement $\iota : X \to X' = X_C$ is a weakly isometric embedding. The map ι is 1-Lipschitz, in particular, its restriction to every apartment in X is an isometry.*

Proof. Verifying that a map is a weakly isometric embedding is a local problem. Let $v \in X$ be a vertex. Then $Y' = \Sigma_v(X')$ is obtained from $Y = \Sigma_v(X)$ by attaching at most two arcs, each having length $\geq \delta = \pi - \frac{\pi}{m}$. Therefore, the inclusion $Y \hookrightarrow Y'$ is a δ-isometric embedding. It follows that the map $X \hookrightarrow X'$ is a weak isometry.

The fact that ι does not increase distances is clear. Suppose that $\sigma \subset \Sigma_x(X) = Y$ is an embedded circle of length 2π (say, an apartment). Set $Y' := \Sigma_x(X')$. Suppose that there exist $\xi, \eta \in \sigma$ so that

$$\angle_Y(\xi, \eta) > \angle_{Y'}(\xi, \eta).$$

Then, combining the geodesic $\overline{\xi\eta} \subset Y'$ with the geodesic $\overline{\xi\eta} \subset Y$, we obtain an embedded circle of length $< 2\pi$ in Y. This contradicts the fact that X' is CAT(0) and Y' is CAT(1). Therefore, for every apartment $A' \subset X$, its image $\iota(A')$ is convex in X' and, hence, the restriction of ι to A' is an isometry. □

Note that $\iota : X \to X'$ is very seldom an isometric embedding, namely, if and only if $f(C) \subset A$ is convex, which means that C is a single (extended) tile.

Lemma 9.4. *Let $X' = X \cup_{C'} A$ be obtained from X by enlargement along a (extended) convex corridor $C' \subset X$ and $\iota : X \to X'$ be the tautological embedding. Let $C \in Corr(X)$. Then the image $\iota(C) \subset X'$ is again a corridor (not necessarily convex).*

Proof. Proof of Axiom C1 is the same as in Lemma 5.9. We need to verify Axiom C2. We will identify C and $\iota(C)$. Suppose that $C \subset X'$ does admit a bridge $D_1 \cup D_2 \cup D_0 \cup T$. Let $e_i \subset \partial D_i$ be the edges adjacent to the trapezoid.

Since C admits a bridge, it is not convex in X'. Since C was convex in X, this means that ι is not a local isometry at the vertices v_1, v_2 of the bridge (where $D_0 = \overline{v_1 v_2}$). It follows that v_1, v_2 are vertices of $\partial C'$. By convexity, $D_0 \subset C'$ as well. Moreover, the germs of $e_i \cup D_0$ at v_i have to be contained in $\partial C'$ ($i = 1, 2$), otherwise, they would be convex in X'. Thus, $D_0 \subset \partial C'$ and, therefore, $C' \subset X'$ admits a bridge with trapezoid T. Since the convex hull of $C' \subset X'$ is contained in the apartment A (which was attached to X' along C'), we conclude that $T \subset A$. This, however, contradicts the fact that the image of C' in A is a model corridor. □

Corollary 9.5. *Let $\iota : X \to X'$ be a morphism in $\mathcal{E}$. Then for every $C \in Corr(X)$, there exists a (extended) convex corridor $C' \subset X'$ containing $\iota(C)$.*

Proof. Since ι is a composition of enlargements, we construct C' inductively starting with $C_1 = C$. If $C_i \subset X_i$ is a convex corridor and $\iota_i : X_i \to X_{i+1}$ is an enlargement, then $\iota_i(C_i) \subset X_{i+1}$ is a corridor according to Lemma 9.4. We then replace the corridor $\iota_i(C_i)$ with a convex corridor in X_{i+1} containing $\iota_i(C_i)$ (see Part 1 of Theorem 8.2). □

10 Construction of the building

We are now ready to prove Theorem 1.1. We proceed by constructing inductively an increasing sequence $X_n \in \mathcal{C}$, $n \in \mathbb{N}$, by using enlargements along (extended) convex corridors[4] $C \in Corr(X_n)$. Then the building X will be the direct limit of the sequence (X_n).

Set

$$Corr := \bigsqcup_{n \in \mathbb{N}} Corr(X_n),$$

Since $Corr(X_n)$ is countable for every $X_n \in \mathcal{C}$, we get a bijection $\omega_n : \mathbb{N} \to Corr(X_n)$. Define

$$\omega : \mathbb{N}^2 \to Corr$$

by

$$\omega(m, n) = \omega_n(m).$$

Let $\nu = (\nu_1, \nu_2) : \mathbb{N} \to \mathbb{N}^2$ denote the standard enumeration, so that both ν_i are 1-Lipschitz functions $\mathbb{N} \to \mathbb{N}$. In particular,

$$\nu_2(n) \le n, \quad \forall\, n \in \mathbb{N}. \tag{1}$$

Lastly, set $\Omega := \omega \circ \nu$.

We start with $X_1 = A$, the model apartment. However, we could as well start with any $X_1 \in \mathcal{C}$.

Then, for $n \in \mathbb{N}$ we compute the (extended) convex corridor $C \in Corr(X_k)$, $C = \Omega(n)$, where $k = \nu_2(n) \le n$. We have a morphism $\iota : X_k \to X_n$ which is a composition of enlargements. According to Corollary 9.5, there exists a convex (extended) corridor $C' \in Corr(X_n)$ containing $\iota(C) \subset X_n$. The corridor C' is not unique. This does not matter with one exception: If C is a geodesic in an apartment in X_k, $\iota(C)$ is still a geodesic in X_n; in this case we take $C' := C$.

We then set

$$X_{n+1} = X_{n,C'} = X_n \cup_{C'} A$$

be the enlargement of X_n along C'.

To make sure that X_{n+1} is well-defined, we need that $\Omega(n)$ is a corridor in X_k, for some $k \le n$. This immediately follows from the inequality (1).

Our next task is to give X structure of a building. Observe that each apartment in X_n is also an apartment in X_{n+1} (see Lemma 9.3). Therefore, we define the set of apartments in X as the union of the sets of apartments in X_n, $n \in \mathbb{N}$. Since each X_n is modeled on (A, W_{af}), we conclude that X is also modeled on (A, W_{af}): The transition maps are in W_{af} and for every pair of apartments $\phi, \psi : A \to X$

$$\phi^{-1}\psi(A) \subset A$$

is a closed and convex subset in A bounded by sub-walls.

[4] Recall that every wall in X_n is in $Corr(X_n)$.

Lemma 10.1. *Let $x_i \in X, i = 1,2$ and $\sigma_i \in \Sigma_{x_i}(X)$ be germs of Weyl chambers. Then there exists an apartment in X containing x_1, x_2 and the germs σ_1, σ_2. In particular, every two points of X belong to a common apartment.*

Proof. Take the smallest k so that $x_i \in X_k$ and $\sigma_i \subset \Sigma_{x_i}(X_k)$, $i = 1,2$. According to Theorem 8.3, there exists a convex corridor $C \subset X_k$ containing x_i and $\sigma_i, i = 1,2$. Then there exists $n \geq k$ so that $C = \Omega(n)$. Therefore, X_{n+1} is obtained by attaching an apartment A along a convex corridor $C' \subset X_n$ containing $\iota(C)$, $\iota : X_k \to X_n$. It follows that the images of x_i, σ_i in X_{n+1} will belong to an apartment A. Hence, the image of A in X will contain x_i, σ_i, $i = 1,2$. □

We can now give X structure of a CAT(0) space, so that every apartment is isometrically embedded. Even though the morphisms $\iota : X_k \to X_n$ are not isometric, nevertheless, for every $x_1, x_2 \in X_k$ there exists n so that x_1, x_2 belong to an apartment in X_n; therefore, for every $l \geq n$ we have:

$$d_{X_l}(x_1, x_2) = d_{X_n}(x_1, x_2).$$

We, hence, obtain a metric on X. Since every morphism in $\mathcal{E}$ is 1-Lipschitz, it follows that for every n, $X_n \to X$ is also 1-Lipschitz. The metric on X is a direct limit of the metrics on the CAT(0) spaces X_n; therefore X is also CAT(0) (as the CAT(0) condition is a condition on quadruples of points).

Lemma 10.2. *Let $\Delta_1, \Delta_2 \subset X$ be antipodal Weyl sectors with the common tip x. Then there exists an apartment $A \subset X$ containing $\Delta_1 \cup \Delta_2$.*

Proof. We first find k so that $\Delta_1 \cup \Delta_2 \subset X_k$. Since Δ_1, Δ_2 are antipodal in X, and $X_k \to X$ is 1-Lipschitz, it follows that Δ_1, Δ_2 are super-antipodal in X_k. Therefore, according to Theorem 8.2, there exists an extended convex corridor $C \subset X_k$ containing $\Delta_1 \cup \Delta_2$. The corridor C equals $\Omega(n)$ for some $n \geq k$. Recall that X_{n+1} is obtained from X_n by enlargement along a convex corridor containing the image $\iota(C), \iota : X_k \to X_n$. It follows that the images of Δ_1, Δ_2 in X_{n+1} (and, hence, in X) are contained in an apartment. □

We thus verified axioms A1′ and A2′ of Euclidean buildings for X.

Lemma 10.3. *X is thick.*

Proof. Let L be a wall in X. Then L is contained in an apartment $A \subset X$, which means that L is a wall in an apartment $A \subset X_k$ for some k. Since L is also a (degenerate) corridor it follows that there exists $n \geq k$ so that X_{n+1} is obtained from X_n by attaching an apartment A' along L. Therefore, the wall $L \subset X$ will be the intersection of three (actually, four) half-apartments in X and, hence, thick. □

This concludes the proof of Theorem 1.1.

Remark 10.4. One can modify the construction of the building X as follows. Besides *enlargements*, one allows isometric embeddings $X_n \to X_{n+1}$, where X_{n+1} is obtained from X_n by attaching an apartment (or another $X'_n \in \mathcal{C}$) along a convex subcomplex *isometric* to a convex subcomplex in A (resp. in X'_n).

As a corollary of Theorem 1.1, we prove

Corollary 10.5. *Let $\overline{W_{af}} = \mathbb{R}^2 \rtimes W$, the closure of W_{af} in the full group of isometries of A. Then there exists a metrically complete thick 2-dimensional Euclidean building X' modeled on $(A, \overline{W_{af}})$.*

Proof. Let X be a building as in Theorem 1.1. Pick a nonprincipal ultrafilter ω and consider the ultralimit X_ω of the constant sequence (X) with fixed base-point (see e.g. [KaL] for the definitions). The metric space X_ω is complete (as any ultralimit, see [KaL]). According to [KL], $X' = X_\omega$ is the required building: It is a thick Euclidean building modeled on $(A, \overline{W_{af}})$. □

11 Concluding remarks

One would like to generalize Theorem 1.1 in several directions:

1. Eliminate the assumption that (A, W_{af}) is special.
2. Allow for uncountable Coxeter groups (besides $\overline{W_{af}}$).
3. Construct buildings X where the automorphism group acts transitively on the set of apartments.

At the moment, we do not know how to deal with 1. However, 2 and 3 can be approached along the same lines. Namely, instead of constructing a building X as a direct limit of a totally-ordered (by inclusion) set of spaces $X_n \in \mathcal{C}$, one would like to allow direct limits using partially ordered sets. This hinges on the following problem, to which we expect a positive solution:

Problem 11.1. *Fix a Coxeter complex (A, W_{af}) on which all spaces $X \in \mathcal{C}$ are modeled. Recall that in Section 9, we defined a category $\mathcal{E}$ of finite complexes modeled on (A, W_{af}), where morphisms are given by compositions of enlargements.*

Find a category $\mathcal{D}$, so that $\mathcal{E} \subset \mathcal{D} \subset \mathcal{C}$, and that $\mathcal{D}$ admits amalgamations*: Every diagram*

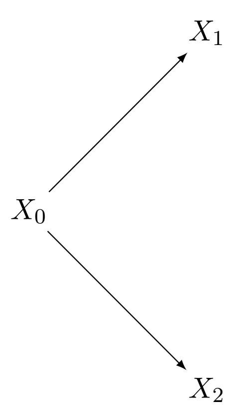

extends to a commutative diagram

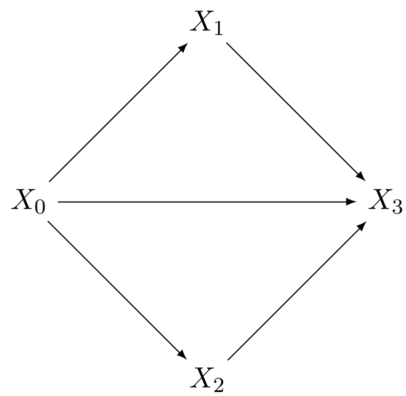

It is conceivable that $\mathcal{D} = \mathcal{E}$, i.e., that $\mathcal{E}$ is closed under amalgamations.

Defining such $\mathcal{D}$ will solve both 2 and 3: see [Te] for the explanation of how to use the amalgamation property to form buildings with "large" automorphism groups. Constructing $\mathcal{D}$ in the context of spherical buildings modeled on (S^1, W) is the main technical result of the paper by K. Tent [Te] and is nontrivial already in this case. Doing this for affine Coxeter groups appears to be much harder. Existence of the category $\mathcal{D}$ would imply that the affine buildings constructed in our paper embed in 2-dimensional affine buildings with highly transitive automorphism group.

Remark 11.2. Suppose that $X_1, X_2 \in \mathcal{C}$ are obtained from $X_0 \in \mathcal{C}$ by attaching apartments A_1, A_2 along convex subcomplexes $C_1, C_2 \subset X_0$. If C_1, C_2 were disjoint or convex in both X_1, X_2, then attaching both A_1, A_2 to X_0 would give us a CAT(0) space X_3. Imagine, however, that the corridor $C = C_1 = C_2$ is a geodesic $\overline{xy}$ (which is a union of two tiles $\overline{vx}$ and $\overline{vy}$) in X_0. Assume also that $f_i : C \to A_i$ $(i = 1, 2)$ are weakly isometric embeddings which are not isometric $v \in C$. Then, attaching both A_1, A_2 to X_0 along C via f_1, f_2, leads to a space X_3 which is not locally CAT(0) at v. To see this, consider the space of directions $Y_0 = \Sigma_v(X_0)$. This graph contains two vertices ξ, η which correspond to the germs of subsegments $\overline{vx}, \overline{vy}$ at v. Then, attaching the apartments A_i to X_0 implies attaching to Y_0 two distinct arcs a_i of the length $\pi - \pi/m$ connecting ξ, η. The resulting space $Y_3 = \Sigma_v(X_3)$ then contains a cycle $a_1 \cup a_2$ of length $2\pi(1 - 1/m) < 2\pi$ and, hence, is not CAT(1).

Of course, in this situation one should form a new CAT(0) space X_3' by identifying $A_1, A_2 \subset X_3$. However, attempting something like this in the case $C_1 \neq C_2$ (and using a partial identification of the apartments A_i), quickly leads to complications which we are currently unable to resolve.

References

[Ba] W. Ballmann, *Lectures on spaces of nonpositive curvature*, Birkhäuser 1995.

[BH] M. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren, vol. 319, Springer, 1999.

[CL] R. Charney and A. Lytchak, *Metric characterizations of spherical and Euclidean buildings*, Geom. Topol., vol. 5 (2001), p. 521–550.

[FS] M. Funk and K. Strambach, *Free constructions*, In: "Handbook of incidence geometry," 739–780, North-Holland, Amsterdam, 1995.

[G] W. Goldman, *Geometric structures on manifolds and varieties of representations*, In "Geometry of group representations (Boulder, CO, 1987)," p. 169–198, Contemp. Math., vol. 74, Amer. Math. Soc., 1988.

[HKW] P. Hitzelberger, L. Kramer and R. Weiss, *Non-discrete Euclidean buildings for the Ree and Suzuki groups,* American Journal of Mathematics, Volume 132, Number 4, August 2010, p. 1113–1152

[KaL] M. Kapovich and B. Leeb, *Asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds*, Geometric Analysis and Functional Analysis, vol. 5, no. 3 (1995), p. 582–603.

[KL] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Publ. Math. Inst. Hautes Études Sci., vol. 86 (1997) p. 115–197.

[N] I. Nikolaev, *The tangent cone of an Aleksandrov space of curvature* $\leq K$, Manuscripta Math., vol. 86 (1995), no. 2, p. 137–147.

[P] A. Parreau, *Immeubles affines: construction par les normes et étude des isométries*, In: "Crystallographic groups and their generalizations" (Kortrijk, 1999), p. 263–302, Contemp. Math., vol. 262, Amer. Math. Soc., Providence, RI, 2000.

[R1] M. Ronan, *A construction of buildings with no rank 3 residues of spherical type*, In: "Buildings and the geometry of diagrams, (Como, 1984)", p. 242–248. Lecture Notes in Math., vol. 1181, Springer, Berlin, 1986.

[R2] M. Ronan, "Lectures on buildings", Perspectives in Mathematics, vol. 7, Academic Press, 1989.

[Te] K. Tent, *Very homogeneous generalized n-gons of finite Morley rank*, J. London Math. Soc. (2), vol. 62 (2000), no. 1, p. 1–15.

[T1] J. Tits, "Buildings of spherical types and finite BN-pairs", Lecture Notes in Math., vol. 386, Springer, Berlin, 1974.

[T2] J. Tits, *Endliche Spiegelungsgruppen, die als Weylgruppen auftreten*, Invent. Math., vol. 43 (1977), no. 3, p. 283–295.

[T3] J. Tits, *Immeubles de type affine*, In "Buildings and the geometry of diagrams (Como, 1984)", p. 159–190. Lecture Notes in Math., vol. 1181, Springer, Berlin, 1986.

[W] R. Weiss, "Structure of affine buildings," Princeton University Press, Annals of Math. Studies, 2008.

Addresses:

Arkady Berenstein: Department of Mathematics, University of Oregon, Eugene, OR 97403, USA. (arkadiy@uoregon.edu)

Michael Kapovich: Department of Mathematics, University of California, Davis, CA 95616, USA. (kapovich@math.ucdavis.edu)