



Lam, T. and Thomas, A. (2015) Infinite reduced words and the Tits boundary of a Coxeter group. *International Mathematics Research Notices*, 2015(17), pp. 7690-7733. (doi:[10.1093/imrn/rnu182](https://doi.org/10.1093/imrn/rnu182))

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/99033/>

Deposited on: 05 November 2014

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

INFINITE REDUCED WORDS AND THE TITS BOUNDARY OF A COXETER GROUP

THOMAS LAM AND ANNE THOMAS

ABSTRACT. Let (W, S) be a finite rank Coxeter system with W infinite. We prove that the limit weak order on the blocks of infinite reduced words of W is encoded by the topology of the Tits boundary $\partial_T X$ of the Davis complex X of W . We consider many special cases, including W word hyperbolic, and X with isolated flats. We establish that when W is word hyperbolic, the limit weak order is the disjoint union of weak orders of finite Coxeter groups. We also establish, for each boundary point ξ , a natural order-preserving correspondence between infinite reduced words which “point towards” ξ , and elements of the reflection subgroup of W which fixes ξ .

1. INTRODUCTION

Let (W, S) be a finite rank Coxeter system with W infinite. In this paper we compare the *limit weak order* on the *infinite reduced words* of W with the topology of the Tits boundary $\partial_T X$ of the Davis complex X of W .

1.1. Infinite reduced words and limit weak order. Let $S = \{s_i \mid i \in I\}$ denote the simple generators of W , and $\{\alpha_i \mid i \in I\}$ denote the simple roots. An infinite reduced word $\mathbf{i} = i_1 i_2 i_3 \cdots$ is an infinite word with letters in I such that each initial finite subword $i_1 i_2 \cdots i_k$ is a reduced word for W . The inversion set $\text{Inv}(\mathbf{i})$ of an infinite reduced word \mathbf{i} is the set of positive roots

$$\text{Inv}(\mathbf{i}) = \{s_{i_1} s_{i_2} \cdots s_{i_{k-1}} \alpha_{i_k} \mid k = 1, 2, \dots\}.$$

In Section 2 we review the notion of a braid limit $\mathbf{i} \rightarrow \mathbf{j}$ of two infinite reduced words, which roughly says that \mathbf{j} can be obtained from \mathbf{i} by an infinite sequence of braid moves. Braid limits give rise to an equivalence relation on infinite reduced words, and a partial order on the equivalence classes, called the limit weak order, and denoted \mathcal{W} . Equivalently, we have $\mathbf{i} \leq \mathbf{j}$ in the limit weak order if and only if $\text{Inv}(\mathbf{i}) \subseteq \text{Inv}(\mathbf{j})$. The limit weak order generalises the usual weak order of a Coxeter group (see for example [1]) to infinite reduced words.

The poset \mathcal{W} is in general quite complicated, and can, for example, contain infinite intervals. Following [15], we decompose \mathcal{W} into *blocks*, allowing us to study \mathcal{W} by studying two more manageable problems. We say that two infinite reduced words \mathbf{i} and \mathbf{j} are *in the same block* if $\text{Inv}(\mathbf{i})$ and $\text{Inv}(\mathbf{j})$ differ by finitely many elements. The limit weak order descends to a partial order on blocks.

As we shall see, on the one hand, the limit weak order on blocks can be given a purely topological description in terms of the Tits boundary of the Davis complex. On the other hand, the limit weak order restricted to a block $B(\mathbf{i})$ is the usual weak order of a (usually infinite) Coxeter group.

Let us remark on the history of infinite reduced words and the limit weak order. Cellini and Papi [7] and later Ito [14] studied certain subsets of positive roots of an affine Weyl group W called biconvex sets. These biconvex sets contain as a special case the inversion sets of infinite

Date: September 18, 2014.

T.L. was supported by NSF grants DMS-0901111 and DMS-1160726, and by a Sloan Fellowship. A.T. was supported by ARC grant DP110100440 and by an Australian Postdoctoral Fellowship.

reduced words. In particular, the inversion sets of infinite reduced words of affine Weyl groups are completely classified in these works. Cellini and Papi were motivated by the study of initial sections of Dyer’s total reflection orders, while Ito was motivated by the study of convex bases of quantized universal enveloping algebras $U_q(\mathfrak{g})$. Lam and Pylyavskyy [15] introduced the limit weak order on infinite reduced words in the context of the theory of total positivity of loop groups. They described this partial order for affine Weyl groups (explained later in the introduction). Dyer [?] studied a generalisation of weak order to biconvex sets of an arbitrary Coxeter group, and formulated a number of conjectures on the lattice structure of this generalisation.

Our present work may have applications in each of these contexts, for example to the yet-to-be-developed theory of total positivity of Kac–Moody groups, or, as pointed out to us by an anonymous referee, to the conjectures in [?].

1.2. Davis complex and its Tits boundary. The Davis complex X is a proper, complete, CAT(0) metric space on which the Coxeter group W acts properly discontinuously and cocompactly by isometries [8]. The visual boundary ∂X , consisting of equivalence classes of geodesic rays in X , can be equipped with the Tits metric, giving a metric space $\partial_T X$ which we call the Tits boundary of X (see [2]). We use the Davis complex X rather than the classical Tits cone because X has a cocompact W -action and its collection of walls is locally finite, and because the CAT(0) structure on the Davis complex allows us to define the Tits boundary and to use many results from CAT(0) geometry. The Davis complex was introduced in order to study geometric and topological properties of infinite non-affine Coxeter groups. The Tits boundary of a CAT(0) space generalises the Tits boundary of a Euclidean building or symmetric space of non-compact type, which is a spherical building. We recall background on the Davis complex and its Tits boundary in Section 3.

There is a bijection between reflections $r \in W$ and walls $M(r) \subset X$. Each infinite reduced word \mathbf{i} naturally gives rise to a path $\gamma : [0, \infty) \rightarrow X$, starting in the identity chamber of X , so that the walls crossed by γ can be identified with the inversion set $\text{Inv}(\mathbf{i})$ of \mathbf{i} (see also Section 2).

The (possibly empty) boundary of a wall M is a subset ∂M of the boundary $\partial_T X$. We show in Section 4.1 that these boundaries of walls induce an arrangement in $\partial_T X$ with many properties similar to a hyperplane arrangement; in particular, each boundary ∂M separates $\partial_T X$ into “half-spaces” ∂M^+ and ∂M^- . We define an equivalence relation on $\partial_T X$ essentially via the facial structure of this arrangement. For $\xi \in \partial_T X$, we let $C(\xi)$ denote the equivalence class of ξ . An important technical tool in Section 4.1 is the Parallel Wall Theorem of Brink and Howlett [3].

1.3. Limit weak order on blocks and the Tits boundary. In Section 4.2, we associate to each infinite reduced word \mathbf{i} a nonempty subset $\partial_T X(\mathbf{i}) \subset \partial_T X$ as follows. The set $\partial_T X(\mathbf{i})$ consists of points $\xi \in \partial_T X$ such that there is a geodesic ray $[p, \xi) \subset X$ pointing in the direction of ξ , with the property that $\text{Inv}(\mathbf{i})$ is the disjoint union of the walls crossed by $[p, \xi)$, and the walls separating the identity from p . Our first main theorem is the following:

Theorem 1. *The subsets $\partial_T X(\mathbf{i}) \subset \partial_T X$ have the following properties:*

- (1) *For each \mathbf{i} and \mathbf{j} we have $\partial_T X(\mathbf{i}) = \partial_T X(\mathbf{j})$ or $\partial_T X(\mathbf{i}) \cap \partial_T X(\mathbf{j}) = \emptyset$. We have $\partial_T X(\mathbf{i}) = \partial_T X(\mathbf{j})$ if and only if \mathbf{i} and \mathbf{j} are in the same block.*
- (2) *Each $\partial_T X(\mathbf{i})$ is an equivalence class $C(\xi)$, and each equivalence class $C(\xi)$ is of the form $\partial_T X(\mathbf{i})$ for some infinite reduced word \mathbf{i} . Thus the $\partial_T X(\mathbf{i})$ form a partition of $\partial_T X$.*
- (3) *Each $\partial_T X(\mathbf{i})$ is a path-connected, totally geodesic subset of $\partial_T X$.*
- (4) *The closure of $\partial_T X(\mathbf{i})$ in $\partial_T X$ is the following union:*

$$\overline{\partial_T X(\mathbf{i})} = \bigcup_{\mathbf{j} \leq \mathbf{i}} \partial_T X(\mathbf{j})$$

where \leq denotes the limit weak order.

Thus *the topology of the Tits boundary encodes the limit weak order on the blocks of infinite reduced words*. We prove parts (1) and (2) of Theorem 1 in Section 4.2, using results from Section 4.1, and parts (3) and (4) of Theorem 1 are established in Sections 4.3 and 4.4 respectively.

A natural direction to explore would be the question: *to what extent does the limit weak order on blocks (or the limit weak order itself) determine the homotopy or homeomorphism type of the Tits boundary?*

We thank Jean Lécureux for the observation that some of the combinatorial objects we consider in Theorem 1 appear in a different guise in his work with Caprace [6]. There is a bijection between equivalence classes of infinite reduced words and elements of the minimal combinatorial compactification $\mathcal{C}_1(X)$, and blocks of infinite reduced words are essentially the same as the spaces X^ξ defined in Section 5.1 of [6].

When W is right-angled, the Davis complex X is naturally a CAT(0) cube complex, which coincides with the CAT(0) cube complex constructed for an arbitrary Coxeter system by Niblo and Reeves (see Lemma 6 of [16]). At least in the case that W is right-angled, there seems to be a close relationship between blocks of infinite reduced words and the simplices in the simplicial boundary of X . The simplicial boundary of a CAT(0) cube complex was introduced by Hagen in [12].

1.4. Limit weak order restricted to a block. In Section 5 we associate to each $\xi \in \partial_T X$ the subgroup $W(\xi)$ of W generated by the set $R(\xi)$ of reflections in walls which have ξ in their boundary. The subgroup $W(\xi)$ is itself a Coxeter group. We prove that the set of reflections in $W(\xi)$ is equal to $R(\xi)$, and then apply results of Deodhar [10] and Dyer [11] to establish our second main result.

Theorem 2. *Suppose $C(\xi) = \partial_T X(\mathbf{i})$. Then the block $B(\mathbf{i})$ is in bijection with the set of elements of $W(\xi)$, and the limit weak order on the block $B(\mathbf{i})$ corresponds naturally to the weak order on $W(\xi)$.*

This is proved as Theorem 32.

1.5. Examples. Section 6 discusses many special cases and examples. We briefly describe the most important here.

In Section 6.1 we consider the case that (W, S) is an irreducible affine Coxeter group. Let W_{fin} be the corresponding finite Weyl group. Then $\partial_T X$ is isometric to a sphere with dimension equal to the rank of W , and the partition of $\partial_T X$ by the $C(\xi)$ or $\partial_T X(\mathbf{i})$ is essentially the (spherical) braid arrangement of W_{fin} . Furthermore, for each piece $C = C(\xi)$ the set $\{\mathbf{i} \mid \partial_T X(\mathbf{i}) = C\}$ of equivalence classes of infinite reduced words is either a singleton, or can be identified with a possibly reducible affine Coxeter group of lower rank. These statements follow from the works of Cellini–Papi [7], Ito [14] and Lam–Pylyavskyy [15].

Another interesting case is when W is word hyperbolic, as discussed in Section 6.5. Such Coxeter groups were characterised by Moussong (see [8]). We can characterise word hyperbolic Coxeter groups in several new ways, and the restriction of Theorem 1 to this case is particularly elegant:

Theorem 3. *The following are equivalent:*

- (1) *The group W is word hyperbolic.*
- (2) *For any infinite reduced word \mathbf{i} , the set $\partial_T X(\mathbf{i})$ consists of a single point $\xi = \xi(\mathbf{i})$.*
- (3) *For any $\xi \in \partial_T X$, the group $W(\xi)$ is finite.*

Moreover if W is word hyperbolic then two infinite reduced words \mathbf{i} and \mathbf{j} are comparable only if \mathbf{i} and \mathbf{j} are in the same block, and the limit weak order restricted to a block $B(\mathbf{i})$ is isomorphic to the weak order of the (possibly trivial) finite Coxeter group $W(\xi(\mathbf{i}))$.

Thus the limit weak order of a word hyperbolic W is the disjoint union of weak orders of finite Coxeter groups. (Jean Lécureux points out that the finiteness of $W(\xi(\mathbf{i}))$ in Theorem 3 can also be deduced from the proof of Theorem 5.13 in [6].) From the point of view of braid relations and infinite reduced words, these results seem far from obvious, as we illustrate in Example 1 below.

Example 1. Consider the Coxeter group

$$W = \langle s_1, s_2, s_3, s_4, s_5 \mid s_i^2 = (s_i s_{i+1})^2 = 1 \rangle.$$

The generators s_i may be viewed as reflections in the sides of a regular right-angled hyperbolic pentagon and the Davis complex X is the induced tessellation of the hyperbolic plane by such pentagons, depicted in Figure 1. (This figure uses the Poincaré disk model of the hyperbolic plane. All of the pentagons in the tessellation are isometric, although they appear to become arbitrarily small.) In this example, the Tits boundary $\partial_T X$ is the boundary of the hyperbolic plane, that is, a circle, and the distance between any two points in $\partial_T X$ is infinite. Thus in particular, $\partial_T X$ has the discrete topology. Each wall is a hyperbolic geodesic connecting two distinct boundary points. The boundaries of walls are dense in $\partial_T X$, by minimality of the action of W on the boundary. Furthermore, each point of $\partial_T X$ is in the boundary of at most one wall. For otherwise, we would have two disjoint walls that become arbitrarily close, which is impossible since each chamber of X is a regular pentagon with fixed side length.

Suppose \mathbf{i} is an infinite reduced word with $\xi = \xi(\mathbf{i})$. If ξ is in the boundary of the wall $M = M(r)$ then the block $B(\mathbf{i})$ consists of two equivalence classes of infinite reduced words, corresponding to the two elements of the finite Coxeter group $W(\xi(\mathbf{i})) = \langle r \rangle$. One of these equivalence classes consists of infinite reduced words which do not cross M but which “follow M out to ξ ”, and this equivalence class corresponds to the trivial element of $W(\xi(\mathbf{i}))$. The other equivalence class in the block $B(\mathbf{i})$ consists of infinite reduced words which cross M and then immediately “follow M out to ξ ”. If ξ is not in the boundary of any wall, then $B(\mathbf{i})$ contains a single equivalence class of infinite reduced words and $W(\xi(\mathbf{i}))$ is trivial.

For example, as shown in Figure 1, we have a braid limit

$$\mathbf{i} = 125252525 \cdots \rightarrow 25252525 \cdots = \mathbf{i}'$$

of two inequivalent infinite reduced words \mathbf{i} and \mathbf{i}' such that $\xi = \xi(\mathbf{i}) = \xi(\mathbf{i}')$ lies on the boundary of the wall $M = M(s_1)$ (see Section 2 for notation). Thus $B(\mathbf{i}) = B(\mathbf{i}') = \{\mathbf{i}, \mathbf{i}'\}$ contains exactly two equivalence classes. On the other hand the infinite reduced word

$$\mathbf{j} = 1352413524 \cdots$$

is such that $\xi(\mathbf{j})$ is not in the boundary of any wall. To see this it is enough to check that there does not exist an infinite reduced word \mathbf{j}' , not braid equivalent to \mathbf{j} , satisfying either $\mathbf{j} \rightarrow \mathbf{j}'$ or $\mathbf{j}' \rightarrow \mathbf{j}$. The fact that $\mathbf{j} \rightarrow \mathbf{j}'$ is impossible is easy: each initial finite subword of \mathbf{j} is the *unique* reduced word of the corresponding element of W . So suppose $\mathbf{j}' = j'_1 j'_2 \cdots \rightarrow \mathbf{j}$. Then the letter 1 must occur in \mathbf{j}' in such a way that it can be moved to the front via the Coxeter relations of W . Thus \mathbf{j}' starts like one of $1 \cdots$, $21 \cdots$ or $51 \cdots$, or starts with a string of alternating 2s and 5s followed by a 1. The latter three cases are impossible, since s_5 and s_2 do not commute, and s_5 and s_3 do not commute. Thus \mathbf{j}' starts with $1 \cdots$. Continuing this argument we see that in fact $\mathbf{j}' = \mathbf{j} = 1352413524 \cdots$.

More generally, one can deduce the following result: an infinite reduced word $\mathbf{i} = i_1 i_2 \cdots$ satisfies $\xi(\mathbf{i}) \in \partial M$ for some wall M if and only if there exists $N > 0$ such that $i_N i_{N+1} i_{N+2} i_{N+3} \cdots = a(a+2)a(a+2) \cdots$ for some $a \in \mathbb{Z}/5\mathbb{Z}$.

Another case where the topology of the Tits boundary $\partial_T X$ is known reasonably explicitly is the case where X has *isolated flats*, which we discuss in Section 6.6. Hruska and Kleiner [13] have shown that the Tits boundary $\partial_T X$ of such spaces is a disjoint union of isolated points and standard Euclidean spheres, while Caprace [4] has classified Coxeter groups W where X has isolated flats.




FIGURE 1. The Davis complex for Example 1

Using Caprace's work, we show that the partition of each Euclidean sphere in $\partial_T X$ into equivalence classes $C(\xi)$ is the arrangement induced by an affine Coxeter group.

The examples we consider in Section 6 suggest the following questions: *does the Tits boundary of a Coxeter group always have the homotopy type of a disjoint union of a wedge of spheres? When is the limit weak order (or limit weak order on blocks) a disjoint union of posets P that are graded, or Eulerian, or Cohen–Macaulay, or shellable?* See for example [1] for a discussion of these poset properties in the Coxeter setting.

Throughout this paper, we make use of the standard theory of Coxeter groups, as found in for instance [?], and results about CAT(0) spaces from [2].

2. INFINITE REDUCED WORDS AND THE LIMIT WEAK ORDER

This section has brief background on infinite reduced words in Coxeter groups, and the limit weak order.

Let (W, S) be a Coxeter system of finite rank, where $S = \{s_i \mid i \in I\}$. Thus S is finite and W has presentation

$$W = \langle \{s_i\}_{i \in I} \mid (s_i s_j)^{m_{ij}} = 1 \rangle$$

where $m_{ii} = 1$ for all $i \in I$, and for all distinct $i, j \in I$, $m_{ij} = m_{ji}$ and $m_{ij} \in \{2, 3, 4, \dots\} \cup \{\infty\}$. We shall assume that W is an infinite group.

An *infinite reduced word* in W is a sequence

$$\mathbf{i} = i_1 i_2 i_3 \cdots$$

where each $i_k \in I$, and for each $n \geq 1$, the finite initial subword $i_1 i_2 \cdots i_n$ is a reduced word, or equivalently, the product $s_{i_1} s_{i_2} \cdots s_{i_n}$ is a reduced expression.

Let Φ^+ denote the set of positive roots of W . For each infinite reduced word $\mathbf{i} = i_1 i_2 i_3 \cdots$, we define an *inversion set* $\text{Inv}(\mathbf{i}) \subset \Phi^+$ by

$$\text{Inv}(\mathbf{i}) = \{s_{i_1} s_{i_2} \cdots s_{i_{k-1}} \alpha_{i_k} \mid k = 1, 2, \dots\}$$

where α_{i_k} denotes the simple root indexed by i_k . The fact that all the roots in $\text{Inv}(\mathbf{i})$ are positive is a consequence of the reducedness condition.

There are bijections between the set Φ^+ of positive roots, the set R of reflections r of W and the set \mathcal{M} of walls $M = M(r)$ in the Davis complex of Section 3. We shall often identify $\text{Inv}(\mathbf{i})$ with a set of reflections, or a set of walls, without further comment. Thus $\text{Inv}(\mathbf{i})$ also denotes the set of walls crossed by \mathbf{i} in the Davis complex.

Let \mathbf{j} and \mathbf{i} be infinite reduced words for W . We shall say that \mathbf{j} is a *braid limit* of \mathbf{i} if it can be obtained from \mathbf{i} by a possibly infinite sequence of braid moves. More precisely, we require that one has $\mathbf{i} = \mathbf{j}_0, \mathbf{j}_1, \mathbf{j}_2, \dots$ such that $\lim_{k \rightarrow \infty} \mathbf{j}_k = \mathbf{j}$ and for each k , there exists ℓ such that the finite initial subword $(\mathbf{j}_k)_1 (\mathbf{j}_k)_2 (\mathbf{j}_k)_3 \cdots (\mathbf{j}_k)_\ell$ of \mathbf{j}_k is a reduced word for the same element of W as the finite initial subword $(\mathbf{j}_{k+1})_1 (\mathbf{j}_{k+1})_2 (\mathbf{j}_{k+1})_3 \cdots (\mathbf{j}_{k+1})_\ell$ of \mathbf{j}_{k+1} , and we have $(\mathbf{j}_k)_r = (\mathbf{j}_{k+1})_r$ for $r > \ell$. Here, the limit $\lim_{k \rightarrow \infty} \mathbf{j}_k = \mathbf{j}$ of words is taken coordinate-wise: $j_r = \lim_{k \rightarrow \infty} (\mathbf{j}_k)_r$. We write $\mathbf{i} \rightarrow \mathbf{j}$ to mean that there is a braid limit from \mathbf{i} to \mathbf{j} . We say that \mathbf{i} and \mathbf{j} are (limit) *braid equivalent* if $\mathbf{i} \rightarrow \mathbf{j}$ and $\mathbf{j} \rightarrow \mathbf{i}$, and write $[\mathbf{i}]$ for the equivalence class of \mathbf{i} .

It is well known, see for example [1, Proposition 3.1.3], that for two Coxeter group elements $v, w \in W$, we have $v \leq w$ in weak order if and only if their inversion sets satisfy $\text{Inv}(v) \subseteq \text{Inv}(w)$. In an analogous fashion, there are two equivalent ways to define the limit weak order on the set of braid equivalence classes of infinite reduced words. The relation between the two possible definitions is given by the following result established in [15, Section 4]. See also [14], which uses different language.

Lemma 4. *Two infinite reduced words \mathbf{i} and \mathbf{j} are braid equivalent if and only if $\text{Inv}(\mathbf{i}) = \text{Inv}(\mathbf{j})$. We have $\text{Inv}(\mathbf{i}) \subset \text{Inv}(\mathbf{j})$ if and only if there is a braid limit from \mathbf{j} to \mathbf{i} .*

Define a partial preorder on infinite reduced words by $\mathbf{i} \leq \mathbf{j}$ if and only if $\mathbf{j} \rightarrow \mathbf{i}$. Note that by Lemma 4, we have $\mathbf{i} \leq \mathbf{j}$ if and only if $\text{Inv}(\mathbf{i}) \subset \text{Inv}(\mathbf{j})$. This preorder descends to a partial order on braid equivalence classes of infinite reduced words, called the *limit weak order* in [15]. We let (\mathcal{W}, \leq) denote the set of equivalence classes of infinite reduced words in W , equipped with the limit weak order.

As in [15], we shall divide \mathcal{W} into *blocks*. As explained in the introduction, this allows us to study the rather complicated poset \mathcal{W} by studying two simpler classes of posets. We say that two infinite reduced words \mathbf{i} and \mathbf{j} are *in the same block* if $\text{Inv}(\mathbf{i})$ and $\text{Inv}(\mathbf{j})$ differ by finitely many roots. We write $B(\mathbf{i}) \subset \mathcal{W}$ to denote the set of equivalence classes of infinite reduced words in the same block as \mathbf{i} . The partial order \leq on \mathcal{W} induces a partial order on blocks: $B(\mathbf{i}) \leq B(\mathbf{j})$ if there exist $[\mathbf{i}'] \in B(\mathbf{i})$ and $[\mathbf{j}'] \in B(\mathbf{j})$ so that $\mathbf{i}' \leq \mathbf{j}'$. It does not seem easy to give a simple criterion for two infinite reduced words \mathbf{i} and \mathbf{j} to lie in the same block, without discussing inversion sets.

3. PRELIMINARIES ON GEOMETRIC GROUP THEORY

This section recalls brief background on notions from geometric group theory, from the references [2] and [8]. In Section 3.1 we review geodesics and the CAT(0) condition. We then recall the construction and relevant properties of the Davis complex X , and discuss the ends of W , in Section 3.2, and discuss boundaries of X , including the Tits boundary $\partial_T X$, in Section 3.3.

3.1. Geodesics and the CAT(0) condition. We briefly recall definitions concerning geodesics and the CAT(0) condition, referring the reader to [2] for details.

Let (X, d) be a metric space. A *geodesic segment* from $x \in X$ to $y \in X$ is a map $\gamma : [0, l] \rightarrow X$ such that $\gamma(0) = x$, $\gamma(l) = y$, and $d(\gamma(t), \gamma(t')) = |t - t'|$ for all $t, t' \in [0, l]$. Similarly we define *geodesic rays* $\gamma : [0, \infty) \rightarrow X$ and *geodesic lines* $\gamma : (-\infty, \infty) \rightarrow X$. The metric space X is said to be *geodesic* if every pair of points $x, y \in X$ is connected by a geodesic segment.

Let X be a metric space. Let Δ be a geodesic triangle in X , and $\bar{\Delta}$ be a Euclidean comparison triangle, that is, a triangle in the Euclidean plane with the same side lengths as Δ . There is then a bijection $x \mapsto \bar{x}$ between points of Δ and points of $\bar{\Delta}$. We say that Δ is CAT(0) if for any $x, y \in \Delta$ we have $d(x, y) \leq d(\bar{x}, \bar{y})$, and that X is CAT(0) if X is a geodesic metric space all of whose geodesic triangles are CAT(0). Similarly one defines CAT(1) spaces using comparison triangles on the unit sphere in \mathbb{R}^3 .

Let $\gamma : [0, a] \rightarrow X$ and $\gamma' : [0, a'] \rightarrow X$ be two geodesic segments with the same start point $\gamma(0) = p = \gamma'(0)$. The (*Alexandrov*) *angle* between γ and γ' is defined as

$$\angle_p(\gamma, \gamma') := \limsup_{t, t' \rightarrow 0} \bar{Z}_p(\gamma(t), \gamma(t'))$$

where $\bar{Z}_p(\gamma(t), \gamma(t'))$ denotes the angle at p of a Euclidean triangle which has side lengths equal to the pairwise distances between $\{p, \gamma(t), \gamma(t')\}$.

3.2. The Davis complex and ends. Let (W, S) be a Coxeter system. In this section we briefly recall the construction and relevant properties of the Davis complex X for (W, S) , following [8]. We also discuss the ends of W and show that if W is one-ended then X has the geodesic extension property (defined below).

For each $T \subset S$, let W_T be the *special subgroup* of W generated by the elements of T . We are following the terminology of Davis [8] here; the subgroups W_T are often called (*standard*) *parabolic subgroups* in the literature. By convention, W_\emptyset is the trivial group. We say $T \subset S$ is *spherical* if W_T is finite. Denote by \mathcal{S} the set of all spherical subsets of S , partially ordered by inclusion. The poset $\mathcal{S}_{>\emptyset}$ is an abstract simplicial complex, denoted L and called the *nerve* of (W, S) . Thus the vertex set of L is S , and a nonempty set T of vertices spans a simplex σ_T in L if and only if T is spherical.

We denote by K the geometric realisation of the poset \mathcal{S} . Equivalently, K is the cone on the barycentric subdivision of the nerve L . Note that K is compact, since it is the cone on a finite simplicial complex. We call the cone point of K its *centre*. For each $s \in S$ let K_s be the union of the (closed) simplices in K which contain the vertex s but do not contain the centre. In other words, K_s is the closed star of the vertex s in the barycentric subdivision of L . Note that K_s has nonempty intersection with K_t if and only if s and t generate a finite subgroup of W . For each $x \in K$, let

$$S(x) := \{s \in S \mid x \in K_s\}.$$

Now define an equivalence relation \sim on the set $W \times K$ by $(w, x) \sim (w', x')$ if and only if $x = x'$ and $w^{-1}w' \in W_{S(x)}$. The *Davis complex* $X = X(W, S)$ for (W, S) is then the quotient space

$$X := (W \times K) / \sim.$$

The natural W -action on $W \times K$ descends to an action on X .

We identify K with the subcomplex (e, K) of X , where e is the identity element of W . Then K , as well as any of its translates by an element of W , is called a *chamber* of X . It is immediate that each chamber is compact, the set of chambers is in bijection with the set of elements of W and W acts transitively on the set of chambers. We denote by 0 the centre of the chamber $K = (e, K)$ in X .

If (W, S) is irreducible affine then the Davis complex $X = X(W, S)$ is just the barycentric subdivision of the Coxeter complex for (W, S) .

The Davis complex X may be equipped with a piecewise Euclidean metric so that it is a proper, complete CAT(0) metric space. Then in particular, given any two points $x, y \in X$, there is a unique geodesic segment from x to y , denoted $[x, y]$. Since X is a complete CAT(0) metric space, by Lemma II.5.8(2) of [2] we may define X to have the *geodesic extension property* if every geodesic segment can be extended to a geodesic line.

In the special case that W is generated by the set of reflections in codimension one faces of a compact convex hyperbolic polytope P , the Davis complex X may instead be equipped with a piecewise hyperbolic metric, so that X is isometric to the induced tessellation of hyperbolic space by copies of P . This natural metrisation of X is used in Example 1 above. Although we work with a piecewise Euclidean metric on X , our results hold for this piecewise hyperbolic metric as well.

By construction, the Coxeter group W acts properly discontinuously and cocompactly by isometries on X . Hence W with its word metric is quasi-isometric to X .

Recall that for a finitely generated group G , the *ends* of G counts the number of path components of the Cayley graph of G as larger and larger finite subgraphs are removed. By a result of Hopf, a finitely generated group G has 0, 1, 2 or infinitely many ends. The first statement in the following theorem is due to Hopf, while the latter three statements are due to Davis and are Theorems 8.7.2, 8.7.3 and 8.7.4 of [8], respectively.

Theorem 5. *Let W be a Coxeter group.*

- (1) *W has 0 ends if and only if W is spherical (that is, finite).*
- (2) *W is one-ended if and only if, for each $T \in S$, the punctured nerve $L - \sigma_T$ is connected.*
- (3) *W is two-ended if and only if (W, S) decomposes as the direct product $(W, S) = (W_0 \times W_1, S_0 \cup S_1)$ where W_1 is finite and W_0 is the infinite dihedral group.*
- (4) *W has infinitely many ends if and only if it is infinite, not as in (3) and there is at least one $T \in S$ so that the punctured nerve $L - \sigma_T$ is disconnected.*

The following result will be needed only in Sections 5.2 and 6.5, however we include it here since its proof uses details of the construction of the Davis complex.

Proposition 6. *If W is one-ended then X has the geodesic extension property.*

Proof. By [2, Proposition II.5.10], it suffices to show that X has no free faces. Let K be the standard chamber of X and L the nerve of W . Then by construction of the Davis complex, X has no free faces if and only if every free face of K is contained in a mirror of K . Since K is the cone on the barycentric subdivision of L , there is a free face of K not contained in a mirror if and only if there is an $s \in S$ such that the vertex $\{s\}$ is contained in a unique maximal simplex of L of dimension ≥ 1 . (Such a free face will contain the edge between $\{s\}$ and the cone point of K .)

Now suppose $s \in S$ is contained in a unique maximal simplex. Then the link of s in L is this maximal simplex with s removed, which is itself a simplex $\sigma_T \in L$. Thus L is either a simplex (in which case W has 0 ends), or $L - \sigma_T$ is disconnected. Then by Theorem 5, W cannot be one-ended. \square

3.3. Visual boundary and Tits boundary of the Davis complex. We briefly recall definitions and some results concerning these boundaries. See [2] for details.

We denote by ∂X the visual boundary of the Davis complex X . That is, ∂X is the set of equivalence classes of geodesic rays in X , where two rays $c, c' : [0, \infty) \rightarrow X$ are defined to be *equivalent* if there is a $0 < K < \infty$ so that for all $t \geq 0$, $d(c(t), c'(t)) < K$. We say that a point $\xi \in \partial X$ is *represented* by a geodesic ray c if c is in the equivalence class ξ , and we may then write $c(\infty) = \xi$. Since X is complete and CAT(0), for any $\xi \in \partial X$ and $p \in X$, there exists a unique geodesic ray from p to ξ , denoted $[p, \xi)$.

We denote by $\partial_T X$ the Tits boundary of the Davis complex X , that is, the visual boundary ∂X equipped with the length metric d_T induced by the angular metric \angle . By definition, the distance between points $\xi, \xi' \in \partial X$ in the angular metric is $\angle(\xi, \xi') = \sup_{p \in X} \angle_p(\xi, \xi')$, where $\angle_p(\xi, \xi')$ is the (Alexandrov) angle at p between the geodesic rays $[p, \xi)$ and $[p, \xi')$.

Lemma 7. *Let ξ and ξ' be distinct points in $\partial_T X$. Then at least one of the following holds:*

- (1) ξ and ξ' are connected by a geodesic arc in $\partial_T X$;
- (2) ξ and ξ' are connected by a geodesic in X .

Moreover, if ξ and ξ' are not connected by a geodesic in X and $d_T(\xi, \xi') < \pi$, then there is a unique geodesic arc in $\partial_T X$ which connects ξ and ξ' .

Proof. Since X is a proper CAT(0) space, Proposition II.9.21(1) and (2) of [2] say that if ξ and ξ' are not connected by a geodesic in X , they are connected by a geodesic arc in the boundary $\partial_T X$, of length $d_T(\xi, \xi') = \angle(\xi, \xi') \leq \pi$. By Theorem 9.20 of [2], since X is a complete CAT(0) space the boundary $\partial_T X$ is a complete CAT(1) space, so $d_T(\xi, \xi') < \pi$ implies that this geodesic arc is unique. \square

We note that it is possible for both (1) and (2) in Lemma 7 to occur. For example if W is irreducible affine of rank n , then $\partial_T X$ is the $(n - 1)$ -dimensional sphere with its usual metric, and every pair of antipodal points in $\partial_T X$ will be connected by infinitely many geodesic arcs of length π in $\partial_T X$ as well as by infinitely many (parallel) geodesics in X .

4. THE LIMIT WEAK ORDER AND THE TOPOLOGY OF THE TITS BOUNDARY

In this section we define the terms in and then prove our first main result, Theorem 1, which is stated in the introduction. In Section 4.1 we use boundaries of walls to define the equivalence class $C(\xi)$ of a point $\xi \in \partial_T X$, and show that the induced arrangement in $\partial_T X$ has properties similar to a hyperplane arrangement. Then in Section 4.2 we define the subsets $\partial_T X(\mathbf{i}) \subset \partial_T X$ and use results from Section 4.1 to prove parts (1) and (2) of Theorem 1. In particular, the sets $C(\xi)$ and $\partial_T X(\mathbf{i})$ induce the same partition of $\partial_T X$. We then investigate the metric and topological properties of this partition, establishing parts (3) and (4) of Theorem 1 in Sections 4.3 and 4.4, respectively.

4.1. The equivalence class $C(\xi)$. In this section we use boundaries of walls to define an equivalence relation \sim on $\partial_T X$, with $C(\xi)$ denoting the equivalence class of $\xi \in \partial_T X$, and we establish some properties of this partition. The separation properties of walls and their boundaries will be crucial in this work, and so we include a careful discussion of these topics before defining the relation \sim in Definition 1. We then define various inversion sets and relate them to \sim (see in particular Proposition 15). We will use many results from this section in proving our main results.

A *wall* in X is by definition the fixed set of a reflection $r \in W$. Each wall $M = \text{Fix}(r)$ is a closed, convex subcomplex of X , which determines two closed half-spaces M^+ and M^- of X such that $M = M^+ \cap M^-$ and the reflection r interchanges M^+ and M^- . The half-space M^+ is by definition that containing the point 0. The collection of walls in X is locally finite, and the maximum number of walls intersecting at any point of X is bounded by the maximum size of a subset $T \subset S$ for which W_T is finite.

We will say that two points $x, y \in X$ are on *opposite sides* of a wall M if $x \in M^+ \setminus M$ and $y \in M^- \setminus M$, or vice versa. We then say that a wall M *separates* two points $x, y \in X$ if x and y are on opposite sides of M , so in particular neither x nor y is in M . By the local structure of walls, if M separates x and y then the geodesic segment $[x, y]$ intersects M in a single point. Given $x, y \in X$, we denote by $\mathcal{M}(x, y)$ the set of walls which separate x and y .

We now consider the boundaries of walls. We define \overline{M} to be the closure of M in $\overline{X} = X \cup \partial X$, and define $\partial M = \overline{M} \cap \partial X = \overline{M} \setminus M$. Similarly we define $\overline{M}^+, \partial M^+$ and $\overline{M}^-, \partial M^-$.

Lemma 8. *Let M be a wall in X .*

- (1) *The sets ∂M and ∂M^\pm are closed in $\partial_T X$ and $\partial M^\pm \setminus \partial M$ is open in $\partial_T X$.*
- (2) *The intersection of ∂M^+ and ∂M^- is ∂M .*

Proof. For (1), by [2, Proposition 9.7(1)], it is enough to establish the claims for ∂X . But then the claims follow from the definitions.

For (2), the inclusion $\partial M \subset \partial M^+ \cap \partial M^-$ is clear. Now suppose $\xi \in \partial M^+ \cap \partial M^-$. Since M^+ and M^- are totally geodesic sets, ξ can be represented by geodesic rays c, c' where c (respectively c') stays completely inside M^+ (respectively M^-). Since c and c' are equivalent we have $d(c(t), c'(t)) < K < \infty$ so in particular both c and c' are a bounded distance from M . It follows that $\xi \in \partial M$. \square

We say that two points $\xi, \xi' \in \partial X$ are on *opposite sides* of ∂M if $\xi \in \partial M^+ \setminus \partial M$ and $\xi' \in \partial M^- \setminus \partial M$, or vice versa.

Lemma 9. *Let M be a wall. Suppose $\xi, \xi' \in \partial X \setminus \partial M$ are on opposite sides of ∂M . If we are in the situation of Lemma 7(1), then the geodesic arc in $\partial_T X$ connecting ξ and ξ' crosses ∂M . If we are in the situation of Lemma 7(2), then the geodesic in X connecting ξ and ξ' crosses M .*

Proof. In the case of Lemma 7(1), let $\gamma \subset \partial_T X$ be the geodesic arc. Then $\gamma \cap \partial M^+$ and $\gamma \cap \partial M^-$ are both closed subsets of γ , and must thus intersect. Hence $\gamma \cap \partial M^- \cap \partial M^+ \neq \emptyset$. In the case of Lemma 7(2), if the geodesic $\gamma \subset X$ does not intersect M , then it must stay completely inside either M^+ or M^- . But then ξ, ξ' will both lie in ∂M^+ or ∂M^- , contradicting the assumption. \square

If $\xi, \xi' \in \partial X \setminus \partial M$ are on opposite sides of ∂M , then using Lemma 9, we will abuse terminology and say that ξ and ξ' are on *opposite sides* of M and are *separated* by M . Two points $\xi, \xi' \in \partial X \setminus \partial M$ are defined to be on the *same side* of M if they are not on opposite sides.

We are now ready to define the relation \sim , which partitions the boundary $\partial_T X$ into equivalence classes $C(\xi)$.

Definition 1 (The set $C(\xi)$). Define an equivalence relation on $\partial_T X$ by $\xi \sim \xi'$ if and only if for each wall M , $\xi \in \partial M^\pm \Leftrightarrow \xi' \in \partial M^\pm$. Let $C(\xi) \subset \partial_T X$ denote the equivalence class of ξ .

That is, $\xi \sim \xi'$ if and only if, for every wall M such that $\xi, \xi' \in \partial X \setminus \partial M$, ξ and ξ' are on the same side of M .

To explore the properties of this partition, we first consider geodesic rays from various points in the Davis complex to various points in its boundary. The following lemma will be used repeatedly.

Lemma 10. *Let M be a wall.*

- (1) *The geodesic ray from any $p \in M$ to $\xi \in \partial M$ must be contained completely in M .*
- (2) *Let c be a geodesic ray which intersects M but is not contained in M . Then $c(\infty)$ is not in ∂M .*
- (3) *Let $\xi \in \partial M$ and $p \in X \setminus M$. Then the geodesic ray $[p, \xi)$ never intersects M .*

Proof. For (1), if $[p, \xi)$ is not contained in M then reflecting this ray in M gives another geodesic ray from p to ξ , contradicting the uniqueness of the geodesic ray. For (2), by local considerations there is a unique point of intersection of c with M , at say $c(t_0) = p_0$. The unique geodesic ray from

p_0 to $\xi = c(\infty)$ is then given by $c(t)$ for $t \geq t_0$. If $\xi \in \partial M$ this contradicts (1). For (3), if $[p, \xi]$ intersects M this contradicts (2). \square

We thus make the following definition: a wall M *separates* $p \in X$ from $\xi \in \partial X$ if p is not in M and the geodesic ray $[p, \xi]$ intersects M , hence by Lemma 10(3) the boundary point ξ is not in ∂M .

We now define several inversion sets for points in X or its boundary $\partial_T X$, and relate them to the equivalence relation \sim . First, for $\xi \in \partial X$, let $\text{Inv}(\xi)$ be the set of walls separating ξ from 0. By our definition of separation, if ξ lies in ∂M , then $M \notin \text{Inv}(\xi)$ regardless of the relative position of 0 and M . Next, define $\text{Inv}(p)$ for $p \in X$ to be the set of walls separating p from 0. Finally, define $\text{Inv}(p, \xi)$ to be the set of walls separating p from ξ , so that $\text{Inv}(\xi) = \text{Inv}(0, \xi)$. A first observation, which we use in Section 4.2 below, is the following.

Lemma 11. *Suppose ξ and ξ' are equivalent. Then for each $p \in X$, $\text{Inv}(p, \xi) = \text{Inv}(p, \xi')$.*

Proof. Suppose first that $M \in \text{Inv}(p)$, that is, that $p \in M^- \setminus M$. Then $\xi \in \partial M^+ \setminus \partial M$ (respectively, $\xi' \in \partial M^+ \setminus \partial M$) if and only if $[p, \xi]$ (respectively, $[p, \xi']$) crosses M , and $\xi \in \partial M$ (respectively, $\xi' \in \partial M$) if and only if $[p, \xi]$ (respectively, $[p, \xi']$) does not cross M . The case that $p \in M^+ \setminus M$ is similar. If $p \in M$ then neither $[p, \xi]$ nor $[p, \xi']$ crosses M . \square

The next two lemmas will be used to prove Proposition 15 below, which says that if ξ and ξ' are not equivalent then $\text{Inv}(\xi)$ and $\text{Inv}(\xi')$ differ by finitely many walls (compare the definition of blocks in Section 2 above). Lemma 13 will also be used in Section 4.4 below. We shall need the following result [3, Theorem 2.8].

Theorem 12 (Parallel Wall Theorem). *Let (W, S) be a Coxeter system. Then there is a constant $k = k(W, S)$ such that for every reflection $r \in R$ and every $x \in X$, if the distance in X from x to the wall $M(r)$ is greater than k , then there is another reflection $r' \in R$ such that $M(r')$ is disjoint from $M(r)$ and $M(r')$ separates x from $M(r)$.*

Remark 1. The statement of [3, Theorem 2.8] is in the framework of root systems. As remarked on p. 182 of [3] their Theorem 2.8 can be shown to be equivalent to a result called the Parallel Wall Theorem, stated as Theorem 1.7 but not completely proved in a preprint of Davis and Shapiro [?]. The result [?, Theorem 1.7] applies to the Cayley graph of a Coxeter system (W, S) , rather than to the Davis complex. Since the Cayley graph of (W, S) is naturally quasi-isometric to the Davis complex, it is not hard to obtain from the statement of [?, Theorem 1.7] the formulation of Theorem 12 we give above.

Lemma 13. *Suppose $\xi \in \partial X$, M is a wall and $p \in M$. If $\xi \notin \partial M$ then $\text{Inv}(p, \xi)$ contains infinitely many pairwise disjoint walls M_i that separate ξ from ∂M .*

Proof. Let c be the (unique) geodesic ray such that $c(0) = p \in M$ and $c(\infty) = \xi$. By the Parallel Wall Theorem (Theorem 12), there is a wall M_1 which separates $c(t_1)$ from M , for some sufficiently large t_1 . Now, c must intersect M_1 at some time $t'_1 \in (0, t_1)$, so in particular ξ does not lie in ∂M_1 , by Lemma 10(2). It follows that c does not stay bounded distance from M_1 , so we can repeat the argument to find M_2, M_3 and so on. \square

Lemma 14. *Let $\xi \in \partial X$ and $p \in X$. Then*

$$(\text{Inv}(0, \xi) \setminus \text{Inv}(p, \xi)) \cup (\text{Inv}(p, \xi) \setminus \text{Inv}(0, \xi))$$

is finite.

Proof. Assume that there are infinitely many walls $M_i \in \text{Inv}(0, \xi) \setminus \text{Inv}(p, \xi)$. Note that by definition $\xi \notin \partial M_i$, so $\xi \in \partial M_i^- \setminus \partial M_i$ for each i . Since only finitely many walls cross each point of the finite length geodesic segment $[0, p]$, and the intersection points are a discrete set, we may assume that $[0, p]$ does not cross any of the walls M_i . Thus in particular $p \in M_i^+ \setminus M_i$ for each i . But $\xi \notin \partial M_i$

and $[p, \xi]$ does not cross M_i , so this implies $\xi \in \partial M_i^+ \setminus \partial M_i$ for each i , a contradiction. The argument if there are infinitely many walls in $\text{Inv}(p, \xi) \setminus \text{Inv}(0, \xi)$ is similar. \square

Proposition 15. *Suppose ξ and ξ' are not equivalent. Then*

$$(\text{Inv}(\xi) \setminus \text{Inv}(\xi')) \cup (\text{Inv}(\xi') \setminus \text{Inv}(\xi))$$

is infinite.

Proof. Without loss of generality, we may consider the following two cases: $\xi \in \partial M^- \setminus \partial M$ and $\xi' \in \partial M^+$, and $\xi \in \partial M^+ \setminus \partial M$ and $\xi' \in \partial M$, for some wall M . In the first case, by Lemma 13, there are infinitely many walls which separate ξ from ∂M , and so there are infinitely many walls which separate ξ from ξ' . In the second case, let p be a point in M . By Lemma 13 and Lemma 14, there are infinitely many walls M_i which are in both $\text{Inv}(p, \xi)$ and $\text{Inv}(\xi) = \text{Inv}(0, \xi)$, and which separate ∂M from ξ . None of these walls M_i are in $\text{Inv}(p, \xi')$, since $[p, \xi']$ is contained in M , so again by Lemma 14 only finitely many of the M_i can be in $\text{Inv}(\xi') = \text{Inv}(0, \xi')$. Hence there are infinitely many M_i in $\text{Inv}(\xi) \setminus \text{Inv}(\xi')$. This completes the proof. \square

4.2. The sets $\partial_T X(\mathbf{i})$ and the proof of (1) and (2) of Theorem 1. In this section we begin by defining the subsets $\partial_T X(\mathbf{i}) \subset \partial_T X$ and establishing some initial properties of these sets. The main result of this section is then Proposition 17 below, which gives a new description, using the sets $\partial_T X(\mathbf{i})$, for the partition of $\partial_T X$ induced by the equivalence classes $C(\xi)$ (see Definition 1 above). Proposition 17 is exactly the statement of Theorem 1(2), and also implies the first statement in Theorem 1(1). A corollary of Proposition 17 will complete the proof of Theorem 1(1).

Definition 2 (The set $\partial_T X(\mathbf{i})$). Let \mathbf{i} be an infinite reduced word. We define $\partial_T X(\mathbf{i})$ to be the set of $\xi \in \partial_T X$ such that there exists a geodesic ray c with $c(\infty) = \xi$ such that when $c(0) = p \in X$, the walls crossed by c , together with the walls separating p from 0, are exactly the inversions of \mathbf{i} (and no wall is crossed twice).

Note that by Lemma 4, if \mathbf{i} and \mathbf{j} are equivalent infinite reduced words, then $\partial_T X(\mathbf{i}) = \partial_T X(\mathbf{j})$. The following lemma provides some justification for Definition 2, by showing that the sets $\partial_T X(\mathbf{i})$ are nonempty and cover $\partial_T X$.

Lemma 16. (1) *For each infinite reduced word \mathbf{i} , the set $\partial_T X(\mathbf{i})$ is nonempty.*
(2) *Every $\xi \in \partial_T X$ lies in some $\partial_T X(\mathbf{i})$.*

Proof. For the first part, let $0 = x_0, x_1, x_2, \dots$ be the centres of the chambers visited by an infinite reduced word \mathbf{i} . Since the Davis complex X is a proper metric space, the space $\overline{X} = X \cup \partial X$ obtained by adjoining to X its visual boundary is compact [2, p. 264]. Thus the sequence $\{x_i\}$ has a convergent subsequence $\{x_{i_n}\}$, with limit say $\xi \in \overline{X}$. In fact, $\xi \in \partial X$, since the x_i are centres of chambers and so do not become arbitrarily close.

As noted in Definition 8.5 of [2], the sequence $\{x_{i_n}\}$ in $X \subset \overline{X}$ converges to a point $\xi \in \partial X$ if and only if the geodesic segments $[x_0, x_{i_n}]$ converge (uniformly on compact subsets) to the geodesic ray $c = [x_0, \xi]$. We claim that this ray crosses the same set of walls as \mathbf{i} , hence $\xi \in \partial_T X(\mathbf{i})$ and so $\partial_T X(\mathbf{i})$ is nonempty. Suppose that a wall M is in $\text{Inv}(\mathbf{i})$. Then there is an $n > 0$ such that M separates 0 from all of $x_{i_n}, x_{i_{n+1}}, \dots$. Thus for all large enough n the geodesic segment $[x_0, x_{i_n}]$ crosses M , and so c crosses M . Now suppose c crosses a wall M . Then for all n large enough, the geodesic segment $[x_0, x_{i_n}]$ crosses M , and thus the wall M separates x_0 from x_{i_n} , hence $M \in \text{Inv}(\mathbf{i})$.

For the second part, let $\xi \in \partial_T X$ and let $c : [0, \infty) \rightarrow X$ be a geodesic ray in the direction of ξ such that $c(0)$ lies in the identity chamber. The sequence of chambers encountered by c will give rise to sequence $w^{(1)}, w^{(2)}, \dots \in W$ of elements of the Coxeter group such that for each i , $\ell(w^{(i)}) = \ell(w^{(i-1)}) + \ell((w^{(i-1)})^{-1}w^{(i)})$. In particular such a sequence arises from taking finite subsequences of a single infinite reduced word. (Note that $\ell(w^{(i)}) - \ell(w^{(i-1)}) > 1$ can occur if

multiple walls are crossed at one time.) Finally, we remark that the sequence $w^{(1)}, w^{(2)}, \dots \in W$ must be infinite by Lemma 10(2). \square

We are now ready to state and prove the main result of this section, which is exactly the statement of Theorem 1(2).

Proposition 17. *Each $\partial_T X(\mathbf{i})$ is an equivalence class $C(\xi)$, and each equivalence class $C(\xi)$ is of the form $\partial_T X(\mathbf{i})$ for some infinite reduced word \mathbf{i} . Thus the $\partial_T X(\mathbf{i})$ form a partition of $\partial_T X$.*

Proof. We first show that each $\partial_T X(\mathbf{i})$ is a union of equivalence classes $C(\xi)$. Suppose that ξ and ξ' are equivalent. By Lemma 16(2), we have $\xi \in \partial_T X(\mathbf{i})$ for some infinite reduced word \mathbf{i} . Let c be a geodesic ray with $c(\infty) = \xi$ such that when $c(0) = p$, the walls crossed by c , together with the walls separating p from 0, are exactly the inversions of \mathbf{i} (with no wall crossed twice). Now by Lemma 11, since ξ and ξ' are equivalent, the set of walls crossed by $[p, \xi]$ is the same as the set of walls crossed by $[p, \xi']$. Thus by definition of $\partial_T X(\mathbf{i})$, we have $\xi' \in \partial_T X(\mathbf{i})$.

We now show that each $\partial_T X(\mathbf{i})$ is equal to a single equivalence class $C(\xi)$. Suppose now that ξ and ξ' are *not* equivalent. Then by Proposition 15, there are infinitely many walls which separate one but not both of the two points ξ and ξ' from 0. For any $p \in X$, there are only finitely many walls separating p from 0. It follows from the definition of $\partial_T X(\mathbf{i})$ that not both of ξ and ξ' can lie in $\partial_T X(\mathbf{i})$. Thus each $\partial_T X(\mathbf{i})$ is equal to a single equivalence class $C(\xi)$.

To complete the proof, we note that by Lemma 16(2), the sets $\partial_T X(\mathbf{i})$ cover $\partial_T X$, and so as the $C(\xi)$ are equivalence classes the $\partial_T X(\mathbf{i})$ form a partition of X . \square

Since the $\partial_T X(\mathbf{i})$ partition $\partial_T X$, we have also established the statement from Theorem 1(1) that for each \mathbf{i} and \mathbf{j} , either $\partial_T X(\mathbf{i}) = \partial_T X(\mathbf{j})$ or $\partial_T X(\mathbf{i}) \cap \partial_T X(\mathbf{j}) = \emptyset$. The remainder of Theorem 1(1) is a consequence of the definition of blocks and the next result.

Corollary 18. *We have $\partial_T X(\mathbf{i}) = \partial_T X(\mathbf{j})$ if and only if $\text{Inv}(\mathbf{i}) \setminus \text{Inv}(\mathbf{j})$ and $\text{Inv}(\mathbf{j}) \setminus \text{Inv}(\mathbf{i})$ are both finite sets.*

Proof. This follows from Proposition 17 together with Lemma 14 and Proposition 15. \square

We note another corollary of Proposition 17, which will be used in Section 4.3.

Corollary 19. *Let \mathbf{i} be an infinite reduced word and suppose $\xi, \xi' \in \partial_T X(\mathbf{i})$. Then for each $p \in X$, $\text{Inv}(p, \xi) = \text{Inv}(p, \xi')$.*

Proof. This is immediate from Proposition 17 together with Lemma 11. \square

4.3. Proof of (3) of Theorem 1. We now move to describing the geometric and topological properties of our partition of $\partial_T X$ into sets $C(\xi)$ or $\partial_T X(\mathbf{i})$. In this section we establish Theorem 1(3). Since we have shown that the sets $C(\xi)$ and $\partial_T X(\mathbf{i})$ induce the same partition, Theorem 1(3) is equivalent to Proposition 24 below, which states that each $C(\xi)$ is a path-connected, totally geodesic subset of $\partial_T X$.

In order to prove Proposition 24, the first result needed is the following. This lemma will also be used in our proof of Theorem 1(4) in Section 4.4 below.

Lemma 20. *Let ξ and ξ' be distinct points in $\partial_T X$, with $\xi = c(\infty)$ and $\xi' = c'(\infty)$ where c and c' are geodesic rays based at the same point $c(0) = c'(0) = p \in X$. If ξ and ξ' are connected by a geodesic γ in X then the set of walls crossed by both c and c' is finite.*

Proof. Suppose M is a wall such that c and c' both intersect M , and M does not contain $c(0) = c'(0)$. Then M separates $c(\infty)$ and $c'(\infty)$ from $c(0) = c'(0)$. Since γ is a geodesic it can intersect M at most once, but $c(\infty)$ and $c'(\infty)$ are on the same side of M , so γ does not intersect M at all. Thus M separates γ from $c(0) = c'(0)$.

Now let α be any geodesic segment connecting $c(0) = c'(0)$ to some point on γ . Every wall M which intersects both c and c' must intersect α . But only finitely walls go through each point, and the intersection points on α are a discrete set. Since α has finite length the statement of the lemma follows. \square

The proof of Proposition 24 will also use the following two technical results.

Lemma 21. *For each $L > 0$ there is a constant $f(L)$ so that a geodesic segment of length L intersects at most $f(L)$ walls.*

Proof. A sufficiently small ball in X intersects at most as many walls as the maximum number incident at a point. A geodesic segment of length L is covered by a fixed number (depending on L) of these sufficiently small balls, which completes the proof. \square

The *angle* between a geodesic c and a wall M that it meets is by definition the infimum of the angles between c and geodesics contained in M .

Lemma 22. *There exist constants $L > 0$ and $\epsilon > 0$ such that any geodesic segment of length L must intersect some wall at angle greater than ϵ .*

Proof. Suppose otherwise, that is, suppose that for all $L > 0$ and $\epsilon > 0$ there is a geodesic segment of length L such that every wall it intersects is met at angle $\leq \epsilon/2$. By Lemma 21, this geodesic segment intersects at most $f(L)$ walls. By reflecting this geodesic segment, we may consider the images of its subsegments inside the fundamental chamber. Let this “broken geodesic” have initial point a_1 and final point a_n and let a_2, \dots, a_{n-1} be the points at which the path turns. In other words the maximal geodesic subsegments inside the fundamental chamber are $[a_k, a_{k+1}]$ for $1 \leq k < n$. Note that $n \leq f(L)$, $\sum_{k=1}^{n-1} d(a_k, a_{k+1}) = L$, and for $1 \leq k < n$ the angles between the geodesic segments $[a_k, a_{k+1}]$ and the walls containing the points a_k and a_{k+1} are both $\leq \epsilon/2$. By [2, I.1.13(2)] it follows that for $1 \leq k \leq n-2$ the angle between $[a_k, a_{k+1}]$ and $[a_{k+1}, a_{k+2}]$ is greater than $\pi - \epsilon$. To simplify notation we will denote by $\angle(a_i a_j, a_j a_k)$ the angle between the geodesic segments $[a_i, a_j]$ and $[a_j, a_k]$, for $1 \leq i, j, k \leq n$.

We claim that for every $L > 0$, there is an $\epsilon = \epsilon(L)$ such that $d(a_0, a_n) > L/\sqrt{2}$. This suffices to complete the proof, since for L large enough there are no points inside the fundamental chamber at distance $L/\sqrt{2}$ apart.

We first show that for $1 \leq k < n$, the angle between $[a_0, a_k]$ and $[a_k, a_{k+1}]$ is greater than $\pi - k\epsilon$. Suppose by induction that $\angle(a_0 a_k, a_k a_{k+1}) > \pi - k\epsilon$. In a CAT(0) space the angle sum of a triangle is at most π , so it follows that $\angle(a_0 a_{k+1}, a_{k+1} a_k) \leq k\epsilon$. Let a'_0 be a point such that the geodesic segment $[a_0, a'_0]$ contains a_{k+1} in its interior and let a'_k be a point such that the geodesic segment $[a_k, a'_k]$ contains a_{k+1} in its interior. Since $\angle(a_k a_{k+1}, a_{k+1} a_{k+2}) > \pi - \epsilon$, we have $\angle(a'_k a_{k+1}, a_{k+1} a_{k+2}) \leq \epsilon$, and since $\angle(a_0 a_{k+1}, a_{k+1} a_k) \leq k\epsilon$ we also have $\angle(a'_k a_{k+1}, a_{k+1} a'_0) \leq k\epsilon$. By [2, Proposition I.1.14] it follows that $\angle(a'_0 a_{k+1}, a_{k+1} a_{k+2})$ is less than or equal to $(k+1)\epsilon$, so $\angle(a_0 a_{k+1}, a_{k+1} a_{k+2}) > \pi - (k+1)\epsilon$ as required.

In particular, $\angle(a_0 a_k, a_k a_{k+1}) > \pi - n\epsilon$ for each k . Now using [2, Proposition II.1.7(5)], we get that for every k we have

$$(1) \quad d(a_0, a_{k+1}) > C(n\epsilon) (d(a_0, a_k) + d(a_k, a_{k+1}))$$

for some constant C depending only on $n\epsilon$. The function $C = C(\delta)$ has the property that $C \rightarrow 1$ as $\delta \rightarrow 0$. In fact, one can pick $C(\delta) = \sqrt{(1 - \cos(\pi - \delta))/2}$. Repeatedly using (1), we obtain

$$d(a_0, a_n) > C(n\epsilon)^n \sum_{k=0}^{n-1} d(a_k, a_{k+1}) = C(n\epsilon)^n L.$$

Now pick $\epsilon = \epsilon(L) > 0$ so that $1 > C(f(L)\epsilon) > (1/\sqrt{2})^{1/f(L)}$. Then with this ϵ , every geodesic segment of length L must intersect some wall at angle greater than ϵ . \square

The last result needed for the proof of Proposition 24 is the following.

Lemma 23. *Suppose $\xi, \xi' \in \partial_T X(\mathbf{i})$. Then $d_T(\xi, \xi') < \pi$.*

Proof. It follows from Lemma 20 that if $\xi' \in C(\xi)$ then ξ and ξ' are not connected by a geodesic in X . Hence by Proposition 9.21(1) and (2) of [2], $d_T(\xi, \xi') = \angle(\xi, \xi') \leq \pi$. Suppose that $d_T(\xi, \xi') = \pi$. Fix $\varepsilon > 0$. Then there is a point $p \in X$ such that $\angle_p(\xi, \xi') > \pi - \varepsilon$. Let c be the geodesic ray from p to ξ and c' the geodesic ray from p to ξ' . Then by Corollary 19, c and c' cross the same set of walls. By Lemma 22, there is a wall M that intersects c at angle greater than ε . This wall M also intersects c' , and the triangle with vertices the two intersection points and p will have total angle greater than π . This is a contradiction. \square

Proposition 24. *Each $C(\xi)$ is a path-connected, totally geodesic subset of $\partial_T X$.*

Proof. It follows from Lemma 20 that if $\xi' \in C(\xi)$ then ξ and ξ' are not connected by a geodesic in X . Hence by Lemma 7 and Lemma 23, there is a unique geodesic arc in $\partial_T X$ connecting ξ and ξ' , of length strictly less than π . We will show that each point η on the geodesic arc in $\partial_T X$ connecting ξ' and ξ in fact lies in $\partial_T X(\mathbf{i})$.

By Proposition 17, it suffices to show that $\eta \in C(\xi)$. Suppose that there is a wall M such that $\xi, \xi' \in \partial M^+$ but $\eta \in \partial M^- \setminus \partial M$. Then the arc $[\xi, \xi']$ has a (possibly equal) subarc $\gamma = [\zeta, \zeta']$ such that $\zeta, \zeta' \in \partial M$ and every point in the interior of γ is in $\partial M^- \setminus \partial M$. Let r be the reflection in the wall M . Then r induces an isometry of the Tits boundary, which fixes ζ and ζ' but does not fix γ . But then γ and $r(\gamma)$ are two distinct geodesic arcs in $\partial_T X$ connecting ζ and ζ' of length strictly less than π . This is impossible, so there is no wall M such that $\xi, \xi' \in \partial M^+$ but $\eta \in \partial M^- \setminus \partial M$. The argument is similar if $\xi, \xi' \in \partial M^-$ but $\eta \in \partial M^+ \setminus \partial M$. Therefore $\eta \in C(\xi)$, and so $C(\xi)$ is a path-connected totally geodesic subset of $\partial_T X$. \square

This completes the proof of (3) of Theorem 1.

4.4. Proof of (4) of Theorem 1. We now complete the proof of Theorem 1 by establishing (4), which describes the closure of the sets $\partial_T X(\mathbf{i})$. Since each $\partial_T X(\mathbf{i})$ is an equivalence class $C(\xi)$, we first consider the closures of the $C(\xi)$, in Proposition 25. Theorem 26 then proves Theorem 1(4).

We will need the following definitions. For each M , let

$$\epsilon_M(\xi) = \begin{cases} 0 & \text{if } \xi \in \partial M \\ + & \text{if } \xi \in \partial M^+ \setminus \partial M \\ - & \text{if } \xi \in \partial M^- \setminus \partial M. \end{cases}$$

Define $\partial \dot{M}^\pm := \partial M^\pm \setminus \partial M$ and $\partial \dot{M}^0 := \partial M$. Then

$$C(\xi) = \bigcap_M \partial \dot{M}^{\epsilon_M(\xi)}.$$

Now let

$$C'(\xi) = \bigcap_M \partial M^{\epsilon_M(\xi)}$$

where $\partial M^0 := \partial M$.

Proposition 25. *The closure $\overline{C(\xi)}$ is equal to $C'(\xi)$.*

Proof. Let $\eta' \in C'(\xi)$ and $\eta \in C(\xi)$. Suppose η' and η are in the situation of Lemma 7(1). Then as in the proof of Proposition 15, there are infinitely many walls M that separate η' from 0 which do not separate η from 0. This is impossible from the definition of $C'(\xi)$. Thus there is a geodesic arc $\gamma : [0, l] \rightarrow \partial_T X$ with $\gamma(0) = \eta'$ and $\gamma(l) = \eta$. The interior of this arc cannot intersect any walls, again by the definition of $C'(\xi)$. Thus $\gamma((0, l]) \subset C(\xi)$, and so $\gamma(0) \in \overline{\gamma((0, l])} \subset \overline{C(\xi)}$. Finally, we note that $C'(\xi)$ is closed, since it is an intersection of closed subspaces (Lemma 8(1)). \square

The following result proves Theorem 1(4), and so completes the proof of Theorem 1.

Theorem 26. *The $\partial_T X(\mathbf{i})$ form a partition of $\partial_T X$ satisfying*

$$\overline{\partial_T X(\mathbf{i})} = \bigsqcup \{\partial_T X(\mathbf{j}) \mid \mathbf{j} \text{ a set of representatives for the blocks } \leq B(\mathbf{i})\} = \bigcup_{\mathbf{j} \leq \mathbf{i}} \partial_T X(\mathbf{j}).$$

Remark 2. Note that it is not true that if $\partial_T X(\mathbf{j}') \subset \overline{\partial_T X(\mathbf{i})}$ then $\mathbf{j}' \leq \mathbf{i}$. One may have to first increase $\text{Inv}(\mathbf{i})$ by a finite set of walls before this inequality holds.

Proof. Let $\partial_T X(\mathbf{i}) = C(\xi)$ and $\xi' \in \partial_T X$. Suppose $C(\xi') \subset \overline{C(\xi)}$. It follows from Proposition 25 that $\text{Inv}(\xi') \subset \text{Inv}(\xi)$. Let $p \in X$ be such that the set of walls crossed from 0 to p , and then from p to ξ , is equal to $\text{Inv}(\mathbf{i})$ (with no wall crossed twice). Then using Lemma 10(3), one sees that the set of walls crossed from 0 to p , and then from p to ξ' , includes no wall twice, and is equal to $\text{Inv}(p) \cup \text{Inv}(\xi')$. Thus $\overline{\partial_T X(\mathbf{i})} \subset \bigcup_{\mathbf{j} \leq \mathbf{i}} \partial_T X(\mathbf{j})$.

Conversely, suppose \mathbf{j} is such that $C(\xi') = \partial_T X(\mathbf{j}) \notin \overline{C(\xi)}$. If ξ' is not in the same path-component of $\partial_T X$ as $C(\xi)$, then by Lemma 20, we know that \mathbf{j} and \mathbf{i} are incomparable. Now assume that ξ and ξ' are in the same path component of $\partial_T X$. Let $\gamma : [0, l] \rightarrow \partial_T X$ be the geodesic arc such that $\gamma(0) = \xi'$ and $\gamma(l) = \xi$. If the interior of γ intersects some wall ∂M , then by Lemma 13 there is a line of pairwise non-intersecting walls M_i separating ξ from ∂M , and hence separating ξ from ξ' . Note that the origin (or any other point) cannot lie on the same side of all these walls. Otherwise any (finite length) geodesic joining the origin to M will intersect infinitely many of these walls. It follows that $\text{Inv}(\xi) \setminus \text{Inv}(\xi')$ and $\text{Inv}(\xi') \setminus \text{Inv}(\xi)$ are both infinite, and so \mathbf{j} and \mathbf{i} are incomparable.

Finally, we consider the case that $\gamma((0, l))$ intersects no walls, but that ξ is contained in some boundary of a wall ∂M , and ξ' is not contained in ∂M . In this case, the same argument shows that $\text{Inv}(\xi') \setminus \text{Inv}(\xi)$ is infinite and so we cannot have $\mathbf{j} < \mathbf{i}$. \square

5. THE REFLECTION GROUP $W(\xi)$

In this section we study the group $W(\xi)$ generated by the reflections in the set $\mathcal{M}(\xi)$ of walls which have $\xi \in \partial_T X$ in their boundary. We prove Theorem 2 of the introduction in Section 5.1, and then in Section 5.2 relate the geometry of X to $\mathcal{M}(\xi)$. The results of Section 5.2 will be used in several of the special cases we consider in Section 6 below.

5.1. Proof of Theorem 2. We now define the group $W(\xi)$ and prove Theorem 2, which is stated in the introduction.

Given $\xi \in \partial_T X$ define $\mathcal{M}(\xi)$ to be the (possibly empty) set of walls in X which have ξ in their boundary, that is,

$$\mathcal{M}(\xi) := \{M \in \mathcal{M} \mid \xi \in \partial M\}.$$

Now for each $\xi \in \partial_T X$, define $R(\xi)$ to be the (possibly empty) set of reflections in walls in $\mathcal{M}(\xi)$. That is, $R(\xi)$ is the set of reflections in walls which have ξ in their boundary. If $R(\xi)$ is nonempty, let $W(\xi)$ be the subgroup of W generated by the elements of $R(\xi)$. If $R(\xi)$ is empty, then define $W(\xi)$ to be the trivial group.

By a theorem of Deodhar [10], proved independently by Dyer [11], the group $W(\xi)$ is itself a Coxeter group, with respect to a canonical system of generators. Dyer [11, Corollary (3.4)(i)] also established the following.

Lemma 27. *Let $w \in W(\xi)$ be arbitrary and $r \in W(\xi)$ be a reflection. Then $\ell(rw) < \ell(w)$ in $W(\xi)$ if and only if $\ell(rw) < \ell(w)$ in W .*

In order to say more about $W(\xi)$, we consider the set $R(\xi)$ and the set of reflections of $W(\xi)$.

Lemma 28. *Let $\xi \in \partial_T X$.*

- (1) *The set $R(\xi)$ is closed under conjugation, that is, if r, t lie in $R(\xi)$ then rtr^{-1} lies in $R(\xi)$.*
- (2) *Every reflection in $W(\xi)$ is conjugate via an element of $W(\xi)$ to a reflection in $R(\xi)$.*

Proof. For the first part, we have to show that if we reflect $M(t)$ in $M(r)$ then the resulting wall $M' = r.M(t)$, which is fixed by the reflection rtr^{-1} , lies in $\mathcal{M}(\xi)$. Let c be a geodesic contained in $M(t)$ such that $c(\infty) = \xi$. Let $c' = r.c$ be its reflection, which is contained in M' . Since ξ is fixed by the reflection r , we have $c'(\infty) = \xi = c(\infty)$, and so $M' \in \mathcal{M}(\xi)$.

The second part is Corollary (3.11)(ii) in Dyer [11]. □

An immediate consequence of Lemma 28 is that:

Corollary 29. *The set $R(\xi)$ is exactly the set of reflections of $W(\xi)$.*

For use in Section 6.5 we note:

Corollary 30. *The group $W(\xi)$ is finite if and only if the set of walls $\mathcal{M}(\xi)$ is finite.*

Proof. This follows from Corollary 29 and the observation that $W(\xi)$ is finite if and only if it contains finitely many reflections. □

For $w \in W$ and $\xi \in \partial_T X$, denote by $\text{Inv}_{\mathcal{M}(\xi)}(w)$ the set $\text{Inv}(w) \cap \mathcal{M}(\xi)$. The next lemma is used in the proof of Theorem 32.

Lemma 31. *Suppose $w \in W$. Then there exists $w' \in W(\xi)$ such that $\text{Inv}_{\mathcal{M}(\xi)}(w)$ can naturally be identified with $\text{Inv}_{W(\xi)}(w')$.*

Proof. Suppose first that w lies in $W(\xi)$. Then it follows from Lemma 27 that $\text{Inv}_{\mathcal{M}(\xi)}(w)$ can naturally be identified with the inversion set $\text{Inv}_{W(\xi)}(w)$, when w is considered as an element of $W(\xi)$. So take $w' = w$.

Now suppose that w is not in $W(\xi)$. We proceed by induction on $k = |\text{Inv}_{\mathcal{M}(\xi)}(w)|$. For the base case $k = 0$, take w' trivial. For $k > 0$, pick a wall $M \in \text{Inv}_{\mathcal{M}(\xi)}(w)$ so that no other walls in $\mathcal{M}(\xi)$ separate w from M . Let the reflection corresponding to M be denoted $r \in W(\xi)$. Let $v \in W$ be such that $\text{Inv}_{\mathcal{M}(\xi)}(v) = \text{Inv}_{\mathcal{M}(\xi)}(w) \setminus \{M\}$, and by induction we suppose that we have found $v' \in W(\xi)$ such that $\text{Inv}_{\mathcal{M}(\xi)}(v) = \text{Inv}_{W(\xi)}(v')$. We claim that $w' = rv'$ works, where geometrically w' is obtained from v' by reflecting in M . This follows easily from the observation that no wall in $\mathcal{M}(\xi)$ separates v' from M , which in turn follows from the same property of v . □

We now prove Theorem 32, which establishes Theorem 2 in the introduction.

Theorem 32. *Suppose $W(\xi)$ is nontrivial. Then the set of equivalence classes of infinite reduced words \mathbf{i} such that $C(\xi) = \partial_T X(\mathbf{i})$ is in bijection with the set of elements of $W(\xi)$, and the partial order on this set of equivalence classes of infinite reduced words corresponds naturally to the weak partial order on $W(\xi)$.*

Proof. Let w be an element of $W(\xi)$. The following lemma and its corollaries will allow us to construct an infinite reduced word \mathbf{i} such that $C(\xi) = \partial_T X(\mathbf{i})$ and $\text{Inv}(\mathbf{i}) = \text{Inv}_{\mathcal{M}(\xi)}(w) \sqcup \text{Inv}(\xi)$.

Lemma 33. *Let p_0 be in the interior of the chamber of X corresponding to w . Then there is a point p on the geodesic ray joining p_0 to ξ , such that every wall intersecting the geodesic segment $[0, p]$ either has ξ in its boundary, or separates 0 from ξ .*

Proof. Let M be a wall intersecting the geodesic segment $[0, p_0]$. If M does not intersect the geodesic ray $[p_0, \xi)$ and M does not separate 0 from ξ , then since neither of the geodesic rays $[p_0, \xi)$ and $[0, \xi)$ cross M , and M separates 0 from p_0 , we must have $\xi \in \partial M$, that is, $M \in \mathcal{M}(\xi)$.

Now suppose the geodesic segment $[0, p_0]$ intersects a wall M_1 such that $M_1 \notin \mathcal{M}(\xi)$ and M_1 does not separate 0 from ξ . Then by the above argument, M_1 intersects $[p_0, \xi)$. So we may define p_1 to be a point on the geodesic ray $[p_0, \xi)$ such that p_1 is on the same side of M_1 as ξ . Note that the geodesic segment $[0, p_1]$ intersects fewer walls that do not have ξ in their boundary and do not separate 0 from ξ than does the geodesic segment $[0, p_0]$. So repeating the construction, we build p_2, p_3, \dots , until we obtain $p = p_k$ with the desired property. (Note that the initial geodesic segment from 0 to p_0 intersects only finitely many walls, so this process stops.) \square

The following corollaries and their proofs continue notation from Lemma 33.

Corollary 34. *Let M be a wall crossed by $[0, p]$ such that $\xi \in \partial M$. Then $M \in \text{Inv}_{\mathcal{M}(\xi)}(w)$. Thus $\text{Inv}(p)$ is the disjoint union of $\text{Inv}_{\mathcal{M}(\xi)}(w)$ with a subset of $\text{Inv}(\xi)$.*

Proof. Since $M \in \mathcal{M}(\xi)$ it suffices to show that $M \in \text{Inv}(w)$, that is, that M is crossed by $[0, p_0]$. But the geodesic ray $[p_0, \xi)$ does not cross M , so p_0 and p are on the same side of M . Since M separates 0 from p the result follows. \square

Corollary 35. *The set of walls crossed by $[0, p]$, together with the set of walls crossed by $[p, \xi)$, is the disjoint union of $\text{Inv}_{\mathcal{M}(\xi)}(w)$ with $\text{Inv}(\xi)$ (and no wall is crossed twice).*

Proof. By construction, no wall crosses both $[0, p]$ and $[p, \xi)$. If $M \in \text{Inv}(\xi)$ and M is not crossed by $[0, p]$, then M separates both 0 and p from ξ , and so M is crossed by $[p, \xi)$. \square

Continuing with the proof of Theorem 32, consider the sequence of chambers visited by $[0, p] \cup [p, \xi)$, where p is constructed as in Lemma 33. Then similarly to the proof of Lemma 16(2), this sequence of chambers gives rise to an infinite reduced word \mathbf{i} such that $\xi \in \partial_T X(\mathbf{i})$, hence $C(\xi) = \partial_T X(\mathbf{i})$, and $\text{Inv}(\mathbf{i}) = \text{Inv}_{\mathcal{M}(\xi)}(w) \sqcup \text{Inv}(\xi)$.

Conversely, if \mathbf{i} is an infinite reduced word such that $\xi \in \partial_T X(\mathbf{i})$, let p be a point such that the set of walls crossed by $[0, p] \cup [p, \xi)$ is equal to $\text{Inv}(\mathbf{i})$ (and no wall is crossed twice). Without loss of generality we may choose p to be in the interior of a chamber of X . Let p belong to the interior of the chamber for $w' \in W$, so that $\text{Inv}(p) = \text{Inv}(w')$. Using Lemma 31, we define $w \in W(\xi)$ to be the unique element such that $\text{Inv}_{\mathcal{M}(\xi)}(w) = \text{Inv}(w') \cap \mathcal{M}(\xi)$. We claim that $\text{Inv}(\mathbf{i})$ is the disjoint union of $\text{Inv}_{\mathcal{M}(\xi)}(w)$ with $\text{Inv}(\xi)$. To see this, let $M \in \text{Inv}(\mathbf{i})$. If M intersects $[p, \xi)$ then M separates 0 from ξ so $M \in \text{Inv}(\xi)$. If M does not intersect $[p, \xi)$ then M intersects $[0, p]$, so either $\xi \in \partial M$ or $M \in \text{Inv}(\xi)$. If M intersects $[0, p]$ and $\xi \in \partial M$ then $M \in \text{Inv}_{\mathcal{M}(\xi)}(w)$ by definition. Therefore $\text{Inv}(\mathbf{i}) \subset \text{Inv}_{\mathcal{M}(\xi)}(w) \sqcup \text{Inv}(\xi)$. The opposite inclusion is clear. Thus to each \mathbf{i} such that $\partial_T X(\mathbf{i}) = C(\xi)$, we have associated an element $w \in W(\xi)$ such that $\text{Inv}(\mathbf{i}) = \text{Inv}_{\mathcal{M}(\xi)}(w) \sqcup \text{Inv}(\xi)$.

These constructions are inverses, and the natural correspondence between the partial order on equivalence classes of infinite reduced words in $\partial_T X(\mathbf{i})$ and the weak partial order on $W(\xi)$ follows from these constructions. \square

5.2. The geometry of X and the set $\mathcal{M}(\xi)$. In this section we investigate the set $\mathcal{M}(\xi)$ and its implications for the geometry of X . The main results are Propositions 38 and 42, which both give sufficient conditions for X to contain an isometrically embedded Euclidean plane. Many results from this section will be used in Section 6 below.

We first note that:

Lemma 36. *Any collection of pairwise intersecting walls in X is finite.*

Proof. This follows from Lemma 3 of Niblo–Reeves [16]. \square

Corollary 37. *For each $\xi \in \partial_T X$, the set $\mathcal{M}(\xi)$ is infinite if and only if there are disjoint walls $M, M' \in \mathcal{M}(\xi)$.*

Proof. Since the collection of walls is locally finite in X , if $\mathcal{M}(\xi)$ is infinite the conclusion follows from Lemma 36. Conversely, suppose $M, M' \in \mathcal{M}(\xi)$ are disjoint and let r, r' be the reflections fixing M, M' respectively. Then r and r' generate an infinite dihedral group which fixes ξ . Apply this group to the wall M to obtain infinitely many walls in $\mathcal{M}(\xi)$. \square

Proposition 38. *Suppose that M, M' are disjoint walls in $\mathcal{M}(\xi)$. Then X contains an isometrically embedded Euclidean plane.*

Proof. Since M and M' are closed subsets of the complete metric space X , M and M' are complete in the induced metric. Hence by [2, Proposition II.2.4(1)], we may choose points $p \in M$ and $p' \in M'$ such that $d(p, p') = d(M, M')$. Let $c = [p, \xi)$ and $c' = [p', \xi)$ and consider the function $t \mapsto d(c(t), c'(t))$. Since the geodesic rays c and c' are equivalent and $d(p, p')$ realises the distance $d(M, M')$, this function is bounded and non-decreasing. As X is CAT(0), this function is convex. Hence the function $t \mapsto d(c(t), c'(t))$ is constant.

Now for any $t > 0$, consider the four points $p, p', q = c(t)$ and $q' = c'(t)$. Since the wall M' is a convex and complete subset of X , and $d(p, p')$ realises the distance $d(p, M')$, by [2, Proposition II.2.4] the Alexandrov angle between the geodesic segments $[p, p']$ and $[p', q']$ is $\geq \pi/2$. By considering $d(p', M)$ instead, we obtain that the Alexandrov angle between the geodesic segments $[p', p]$ and $[p, q]$ is also $\geq \pi/2$. But since $d(q, q') = d(p, p') = d(M, M')$, the same argument shows that the Alexandrov angles between $[q, q']$ and the segments $[q, p]$ and $[q', p']$ are also $\geq \pi/2$. Thus by the Flat Quadrilateral Theorem [2, II.2.11], each of these four angles is equal to $\pi/2$, and moreover the convex hull of the four points p, p', q, q' is isometric to the convex hull of a rectangle in the Euclidean plane. Therefore the convex hull of c and c' is a flat half-strip which is orthogonal to both M and M' .

Let r and r' be the reflections fixing M and M' respectively. Since M and M' are disjoint these reflections generate an infinite dihedral group. The union of the images of the flat half-strip bounded by c and c' under the action of $\langle r, r' \rangle$ is a flat half-plane in X (which is orthogonal to all of the images of M and M' under this action). Hence for all $n \geq 1$ there is an isometric embedding into X of the ball of radius n centred at the origin in the Euclidean plane. Thus by [2, Lemma 9.34], there is an isometric embedding of the Euclidean plane into X . \square

Corollary 39. *If $\mathcal{M}(\xi)$ is infinite then X contains an isometrically embedded Euclidean plane.*

Recall that by Proposition 6, X has the geodesic extension property if, for instance, W is one-ended. The next two lemmas and also Proposition 42 will be used in Section 6.5, when we consider the case that W is word hyperbolic, although they do not require hyperbolicity.

Lemma 40. *If $\mathcal{M}(\xi)$ is finite and nonempty then $\bigcap_{M \in \mathcal{M}(\xi)} M$ is nonempty. If in addition X has the geodesic extension property, then this intersection contains a geodesic line.*

Proof. Since $W(\xi)$ is finite and X is CAT(0), there is a point $p \in X$ which is fixed by all $w \in W(\xi)$. In particular, p is fixed by all reflections in $R(\xi)$, hence p is contained in every wall $M \in \mathcal{M}(\xi)$. That is, $\bigcap_{M \in \mathcal{M}(\xi)} M$ is nonempty. Now consider the geodesic ray $[p, \xi)$. By Lemma 10, this geodesic ray is contained in every $M \in \mathcal{M}(\xi)$. Let η be the other endpoint of a geodesic extension of the ray $[p, \xi)$. Then by the local properties of walls, the ray $[p, \eta)$ is also contained in every $M \in \mathcal{M}(\xi)$, so the intersection of the walls in $\mathcal{M}(\xi)$ is nonempty. If X has the geodesic extension property then by extending $[p, \xi)$ we obtain that $\bigcap_{M \in \mathcal{M}(\xi)} M$ contains a geodesic line. \square

Lemma 41. *Assume X has the geodesic extension property. If M and M' are distinct walls such that $M \cap M' \neq \emptyset$, then $\partial M \neq \partial M'$.*

Proof. Let $p \in M \cap M'$. Since M and M' are distinct walls, we may choose $q \in M \setminus M'$. Then as X has the geodesic extension property, the geodesic segment $[p, q]$ extends to a geodesic ray $[p, \xi)$. By the local properties of walls, $\xi \in \partial M$. If $\xi \in \partial M'$ as well, then by Lemma 10 the entire geodesic ray $[p, \xi)$ is contained in M' , a contradiction. \square

Proposition 42. *Assume X has the geodesic extension property. Let M and M' be disjoint walls such that $\partial M = \partial M'$ is nonempty. Then M and M' are constant distance apart and X contains an isometrically embedded Euclidean plane which is orthogonal to both M and M' .*

Proof. Let $\xi \in \partial M = \partial M'$. As in the proof of Proposition 38, we may choose points $p \in M$ and $p' \in M'$ such that $d(p, p') = d(M, M')$, and the geodesic rays $[p, \xi)$ and $[p', \xi)$ then bound a flat half-strip in X .

Now since X has the geodesic extension property, there exists a geodesic line $c : \mathbb{R} \rightarrow X$ such that $c([0, \infty)) = [p, \xi)$. By the local structure of walls, the entire image of c must be contained in M . Put $\eta = c(-\infty)$. Then $\eta \in \partial M = \partial M'$. Since $d(p, p') = d(M, M')$, the geodesic rays $[p, \eta)$ and $[p', \eta)$ also bound a flat half-strip.

Let r' be the reflection which fixes the wall M' . Then $r'c$ is a geodesic line which is contained in the wall $r'M$ and passes through the point $r'p$. The geodesic ray $[r'p, \xi)$ is constant distance from $[p', \xi)$, and the geodesic ray $[r'p, \eta)$ is constant distance from $[p', \eta)$. Thus the geodesic line $r'c$ is at uniformly bounded distance from the geodesic line c . Therefore by the Flat Strip Theorem [2, II.2.13], the convex hull of c and $r'c$ is isometric to a flat strip. Note that the width of this flat strip is $d(p, r'p) = 2d(p, p')$.

We have so far shown that there is a geodesic line c in M , and passing through p , which is at constant distance from the geodesic line $r'c$ in $r'M$.

Now let x be any point in M which is distinct from p . Then by the geodesic extension property again, the geodesic segment $[p, x]$ may be extended to a geodesic line c_x passing through p , and moreover by the local structure of walls the line c_x is contained in M . The above argument may be repeated with the endpoints of c_x to show that c_x is at constant distance $2d(p, p')$ from the geodesic line $r'c_x$ in $r'M$. In particular, x is at distance $2d(p, p')$ from $r'M$. Therefore we have that the walls M and $r'M$ are at constant distance $2d(p, p')$.

Now for any point $x' \in M'$, we have $d(x', M) = d(x', r'M)$. Thus the wall M' is exactly halfway in between M and $r'M$, and we conclude that M and M' are at constant distance $d(p, p')$.

To obtain an isometrically embedded Euclidean plane which is orthogonal to both M and M' , keep on reflecting the flat strip between lines c and $r'c$. \square

6. SPECIAL CASES

We conclude by discussing numerous special cases and examples. These include the cases that W is irreducible affine (Section 6.1), W is reducible or virtually abelian (Section 6.2), W has two ends (Section 6.3), W has infinitely many ends (Section 6.4) and W is word hyperbolic (Section 6.5), and the case that X has isolated flats (Section 6.6). In particular, in Section 6.5 we prove Theorem 3 of the introduction.

6.1. Irreducible affine Coxeter groups. Infinite reduced words for affine Weyl groups were studied by Cellini–Papi [7] and Ito [14] from a root system point of view. The limit weak partial order on equivalence classes of infinite reduced words was studied in [15]. We summarise the results here in the notation of the present paper.

Let (W, S) be an irreducible affine Coxeter group with corresponding finite Weyl group W_{fin} . Then the boundaries of walls $\partial M \subset \partial_T X$ are in bijection with the reflections of W_{fin} . The pieces

of the partition $\partial_T X = \sqcup C(\xi)$ are in bijection with the faces of the Coxeter arrangement of W_{fin} (with the origin of the Coxeter arrangement excluded). This arrangement could also be described as the (thin) spherical building at infinity.

The pieces $C(\xi)$ which are open (and thus not contained in the closure of any other $C(\xi')$) are in bijection with $w \in W_{\text{fin}}$. For any other piece $C(\xi)$, we have that $\mathcal{M}(\xi)$ is infinite but is partitioned into finitely many parallelism classes. These parallelism classes correspond to a collection of walls in the Coxeter arrangement of W_{fin} . Furthermore, $W(\xi)$ is a (possibly reducible) affine Coxeter group.

Thus the maximal elements of (\mathcal{W}, \leq) are in bijection with W_{fin} . The blocks of \mathcal{W} are in bijection with the faces of the Coxeter arrangement of W_{fin} with the origin excluded. The limit weak order on each block $B(\mathbf{i})$ is isomorphic to the weak order on some affine Coxeter group.

6.2. Reducible Coxeter groups and virtually abelian Coxeter groups. Suppose that (W, S) is a reducible Coxeter system, with $S = S_1 \sqcup S_2$ so that $W = W_1 \times W_2$ where $W_k = \langle S_k \rangle$ for $k = 1, 2$. Let $I_k = \{i \mid s_i \in S_k\}$. Assume first that W_1 and W_2 are both infinite and let \mathbf{i} be an infinite reduced word in W . Then exactly one of the following occurs:

- (1) \mathbf{i} is equivalent to an infinite reduced word consisting of a finite reduced word in W_2 followed by an infinite reduced word in W_1 ;
- (2) \mathbf{i} is equivalent to an infinite reduced word consisting of a finite reduced word in W_1 followed by an infinite reduced word in W_2 ; or
- (3) \mathbf{i} contains infinitely many elements of I_1 and infinitely many elements of I_2 .

If case (1) holds for two infinite reduced words \mathbf{i} and \mathbf{i}' , denote by w_2 and w'_2 the elements of W_2 and by \mathbf{j} and \mathbf{j}' the infinite reduced words in W_1 such that $\text{Inv}(\mathbf{i}) = \text{Inv}(w_2) \sqcup \text{Inv}(\mathbf{j})$ and $\text{Inv}(\mathbf{i}') = \text{Inv}(w'_2) \sqcup \text{Inv}(\mathbf{j}')$. Then $\mathbf{i} \leq \mathbf{i}'$ if and only if both $w_2 \leq w'_2$ in the weak partial order on W_2 and $\mathbf{j} \leq \mathbf{j}'$ in the partial order on infinite reduced words in W_1 , with strict inequality $\mathbf{i} < \mathbf{i}'$ if and only if, in addition, at least one of the containments $\text{Inv}(w_2) \subseteq \text{Inv}(w'_2)$ or $\text{Inv}(\mathbf{j}) \subseteq \text{Inv}(\mathbf{j}')$ is strict. Similarly for the partial order on pairs of infinite reduced words for which (2) holds. If (1) holds for \mathbf{i} and (2) holds for \mathbf{i}' , then \mathbf{i} and \mathbf{i}' are clearly incomparable.

If case (3) holds then for $k = 1, 2$, the infinite reduced word \mathbf{i} is strictly greater in the partial order than any infinite reduced word \mathbf{j}_k obtained by concatenating a finite initial subword of \mathbf{i} with the infinite reduced word consisting of all subsequent elements of I_k in \mathbf{i} . Note that in this case the set $\text{Inv}(\mathbf{i})$ differs from $\text{Inv}(\mathbf{j}_k)$ by infinitely many walls.

Now let us discuss the Tits boundary $\partial_T X$ of the Davis complex X of W . The Davis complex X is a product $X = X_1 \times X_2$, and thus the Tits boundary is the spherical join $\partial_T X = \partial_T X_1 * \partial_T X_2$ (see [2, Corollary II.9.11]). We shall write $\xi = \cos \theta \xi_1 + \sin \theta \xi_2$ with $\theta \in [0, \pi/2]$ for a point of $\partial_T X$, where $\xi_1 \in \partial_T X_1$ and $\xi_2 \in \partial_T X_2$. If $\xi = \xi_1 + 0$ (respectively $\xi = 0 + \xi_2$), then $C(\xi)$ lies in all the walls of W_2 (respectively all the walls of W_1), and $C(\xi) = \partial_T X(\mathbf{i})$ for an infinite reduced word \mathbf{i} of the form (1) (respectively (2)). A point of the form $\xi = \cos \theta \xi_1 + \sin \theta \xi_2$ with $\theta \in (0, \pi/2)$ has equivalence class given by

$$C(\xi) = \{\cos \theta \xi'_1 + \sin \theta \xi'_2 \mid \xi'_1 \in C(\xi_1) \text{ and } \xi'_2 \in C(\xi_2)\}$$

and we then have $C(\xi) = \partial_T X(\mathbf{i})$ for an infinite reduced word \mathbf{i} of the form (3).

If W_2 is finite (hence W_1 is infinite), then only case (1) can occur, and similarly if W_1 is finite only case (2) can occur. In particular, it follows from [8, Theorem 12.3.5] that a Coxeter group is virtually abelian if and only if it is a direct product of finite and affine Coxeter groups.

It is thus straightforward to extend the results for irreducible affine Coxeter groups described in Section 6.1 to all affine and all virtually abelian Coxeter groups.

6.3. Coxeter groups with two ends. Recall from Theorem 5(3) above that it is equivalent for W to have two ends and for W to be reducible of the form $W_0 \times W_1$ where $W_0 = \langle s, t \rangle$ is the infinite

dihedral group and W_1 is spherical. Any infinite reduced word in W is then obtained by inserting into either $ststst \cdots$ or $tststs \cdots$ the letters of a reduced word for some $w_1 \in W_1$. The boundary of X thus consists of two points, say ξ_+ and ξ_- . The only walls of X with nonempty boundary are the finitely many walls corresponding to the reflections in W_1 , and for each such wall M we have $\partial M = \{\xi_+, \xi_-\}$. In other words, $R(\xi_{\pm})$ is equal to the set of reflections in W_1 , and $W(\xi_{\pm}) = W_1$. Clearly $C(\xi_+) = \{\xi_+\}$ and $C(\xi_-) = \{\xi_-\}$. There are $2|W_1|$ equivalence classes of infinite reduced words, with $\mathbf{i} \sim \mathbf{j}$ if and only if \mathbf{i} and \mathbf{j} project to the same infinite reduced word in W_0 and \mathbf{i} and \mathbf{j} are eventually on the same side of every wall in W_1 (equivalently, the projections of \mathbf{i} and \mathbf{j} to W_1 define the same group element). The partial order on the set of infinite reduced words \mathbf{i} with $\partial_T X(\mathbf{i}) = \{\xi_+\}$ is then just the weak partial order on W_1 , and similarly for ξ_- .

6.4. Coxeter groups with infinitely many ends. Assume now that W has infinitely many ends. In particular, the case that W is virtually free and the case that W is a free product of special subgroups are contained in the case that W has infinitely many ends. In this section we will use some results about graphs of groups, a reference for which is Chapter I of [17].

By [8, Proposition 8.8.2], every Coxeter group W has a tree-of-groups decomposition (not in general unique) in which every vertex group is a spherical or a one-ended special subgroup, and every edge group is a spherical special subgroup. More precisely, the edge groups are spherical special subgroups W_T where the punctured nerve $L - \sigma_T$ is disconnected. Since W has infinitely many ends, by Theorem 5 any such tree-of-groups decomposition has at least one edge. Our analysis will depend upon the nature of the vertex and edge groups in such a decomposition.

First, the group W is virtually free if and only if all vertex groups in such a tree-of-groups decomposition of W are spherical [8, Proposition 8.8.5]. The Davis complex X is then quasi-isometric to a tree, with $\partial_T X$ homeomorphic to the set of ends of this tree. The topology on $\partial_T X$ is thus totally disconnected and so for each $\xi \in \partial_T X$, we have $C(\xi) = \{\xi\}$.

Next, the group W is a free product of spherical special subgroups if and only if all vertex groups in such a tree-of-groups decomposition of W are spherical *and* all edge groups are trivial. The Davis complex X then consists of infinitely many copies of the (finite) Davis complexes for the vertex groups, glued together at certain centres of chambers. Thus all walls are finite subcomplexes of X and so have empty boundary. It follows that for each $\xi \in \partial_T X$, there is a unique equivalence class of infinite reduced words \mathbf{i} such that $\partial_T X(\mathbf{i}) = C(\xi) = \{\xi\}$ is a single point.

We next discuss the subcase that all vertex groups are spherical and there is some nontrivial edge group. Fix a tree-of-groups decomposition of W and let r be a reflection in W . Since reflections have finite order, each reflection is contained in a conjugate of at least one vertex group. If r is not contained in any conjugates of edge groups, then the wall $M = M(r)$ is contained in the finite Davis complex for a unique (conjugate of a) vertex group, and so M has empty boundary. If r is contained in some (conjugate of an) edge group, then the wall $M = M(r)$ may or may not have empty boundary, as shown by the following examples. (We do not provide a complete characterisation of the reflections r such that the corresponding wall M has nonempty boundary.)

Examples. (1) Consider $W = \langle s, t, u \mid s^2 = t^2 = u^2 = (st)^3 = (tu)^3 = 1 \rangle$. Then W splits as the amalgamated free product $W = W_{\{s,t\}} *_{W_{\{t\}}} W_{\{t,u\}} \cong S_3 *_{C_2} S_3$ and hence has a tree-of-groups decomposition consisting of an edge with vertex groups $W_{\{s,u\}}$ and $W_{\{t,u\}}$ and edge group $W_{\{t\}}$. The reflection t is obviously contained in an edge group, but the wall $M(t)$ is finite hence has empty boundary. In fact all walls in this example have empty boundary.

(2) Let $W = W_{T_1} *_{W_T} W_{T_2}$ where T_1 and T_2 are spherical and $T = T_1 \cap T_2$ is nonempty. Let r be a reflection in W_T . Suppose that for $i = 1, 2$ there is a reflection $r_i \in W_{T_i} \setminus W_T$ which commutes with r . Then for $i = 1, 2$ and every $w \in \langle r_1, r_2 \rangle \cong D_{\infty}$, the reflection r is contained in $W_{T_i}^w$. It follows from the fact that there are infinitely many distinct such

conjugates $W_{T_i}^w$ and the construction of X that the wall $M = M(r)$ has nonempty boundary, consisting of at least two points.

Continuing the analysis of the subcase in which all vertex groups are spherical and there is a non-trivial edge group W_T , let $M = M(r)$ be a wall with nonempty boundary, where the reflection r is without loss of generality in W_T . For each $\xi \in \partial_T X$ such that $M \in \mathcal{M}(\xi)$, there is then a nontrivial partial order on the set of equivalence classes of infinite reduced words \mathbf{i} such that $\partial_T X(\mathbf{i}) = \xi$. Since X does not contain an isometrically embedded Euclidean plane, Corollary 39 implies that $\mathcal{M}(\xi)$ is finite. Hence if $M' \neq M$ is also in $\mathcal{M}(\xi)$, with r' the reflection fixing M' , then r and r' generate a finite subgroup of W . Thus $\langle r, r' \rangle$ fixes a point p in X , and we obtain a geodesic ray $[p, \xi)$ which is contained in both M and M' . It follows that r' must be in W_T as well. Applying this argument to every wall in $\mathcal{M}(\xi)$, we conclude that every reflection in $R(\xi)$ is contained in the edge group W_T . There may however be reflections in W_T which are not in $R(\xi)$.

We finally consider the case that W has infinitely many ends and is not virtually free, equivalently there is at least one one-ended vertex group in the tree-of-groups decomposition of W . Then an infinite reduced word \mathbf{i} either eventually stays within a fixed conjugate W_T^w of a one-ended vertex group W_T , or \mathbf{i} visits infinitely many distinct conjugates of vertex groups. For each $w \in W$ and each one-ended vertex group W_T , the partial order on the set of equivalence classes of infinite reduced words that eventually stay in W_T^w is the same as the partial order on the set of equivalence classes of infinite reduced words in W_T . In a similar manner to the analysis when all vertex groups are spherical, if there is some nontrivial edge group, then there may also be a nontrivial partial order on equivalence classes of certain infinite reduced words \mathbf{i} which visit infinitely many distinct conjugates of vertex groups.

6.5. Word hyperbolic Coxeter groups. In this section we consider the important special case that W is word hyperbolic. In particular, we prove Theorem 3, stated in the introduction. See [2] for the definition and basic properties of δ -hyperbolic spaces and word hyperbolic groups.

We will use the following characterisation of word hyperbolic Coxeter groups, which is due to Moussong, and which appears in [8, Corollary 12.6.3].

Theorem 43 (Moussong). *A Coxeter group (W, S) is word hyperbolic if and only if there is no subset T of S satisfying either of the two conditions:*

- (1) W_T is affine of rank at least 2.
- (2) (W_T, T) decomposes as $(W_T, T) = (W_{T'} \times W_{T''}, T' \cup T'')$ with both $W_{T'}$ and $W_{T''}$ infinite.

We now characterise the word hyperbolicity of W in terms of the sets $\mathcal{M}(\xi)$ in Proposition 44, in terms of the groups $W(\xi)$ in Corollary 45 and finally in terms of the sets $C(\xi)$ in Proposition 46.

Proposition 44. *W is word hyperbolic if and only if for all $\xi \in \partial_T X$, the set $\mathcal{M}(\xi)$ is finite.*

Proof. Suppose there is a $\xi \in \partial_T X$ such that $\mathcal{M}(\xi)$ is infinite. Then by Corollary 37 and Proposition 38, the space X contains an isometrically embedded Euclidean plane. By the Flat Plane Theorem [2, III.H.1.5], this implies X is not δ -hyperbolic, hence W is not word hyperbolic.

For the converse, let T be a subset of S satisfying either of the conditions in Theorem 43. We denote by X_T the Davis complex for (W_T, T) , which embeds isometrically in the Davis complex X for (W, S) . It suffices to prove there is a ξ in the Tits boundary of X_T such that ξ is in the boundary of infinitely many walls in X_T .

This is immediate when W_T is affine, by taking ξ in the boundary of a family of parallel walls in X_T . If $W_T = W_{T'} \times W_{T''}$ as in (2) of Theorem 43, let ξ be any point in the boundary of $X_{T''} \subset X_T$ (since $W_{T''}$ is infinite, $X_{T''}$ has nonempty boundary). Now as $W_{T'}$ is infinite, $W_{T'}$ contains infinitely many distinct reflections r_i . The r_i are also reflections in W_T , with corresponding walls say M_i in X_T . Since $W_{T''}$ commutes with $W_{T'}$, the point ξ is fixed by each of the reflections r_i , and so $\xi \in \partial M_i$ for all i . \square

Corollary 45. *W is word hyperbolic if and only if for all $\xi \in \partial_T X$, the group $W(\xi)$ is finite.*

Proof. This follows from Corollary 30 and Proposition 44. □

Proposition 46. *W is word hyperbolic if and only if for all $\xi \in \partial_T X$ the equivalence class $C(\xi)$ is reduced to $\{\xi\}$.*

Proof. Assume W is word hyperbolic. Since W is quasi-isometric to the Davis complex X (equipped with its piecewise Euclidean metric), we have that X is δ -hyperbolic for some $\delta > 0$. Suppose that there are $\xi \neq \xi' \in \partial_T X$ with $\xi' \in C(\xi)$. Fix a basepoint $p \in X$. Then by Proposition 17 and Lemma 19, the geodesic rays $c = [p, \xi)$ and $c' = [p, \xi')$ cross the same set of walls.

For each wall M crossed by c (and thus by c'), denote by γ_M the geodesic segment in M which connects the points of intersection of c and c' with M , and let Δ_M be the geodesic triangle in X with one vertex at p , with γ_M its side not containing p , and with its other two sides initial subsegments of c and c' . Then by δ -hyperbolicity of X , the geodesic segment γ_M is contained in the δ -neighbourhood of the union of c and c' .

Since $\xi \neq \xi'$, the δ -neighbourhoods of c and c' have bounded intersection. Thus the geodesic segment $\gamma_M \subset M$ passes within bounded distance of p . By the local finiteness of the collection of walls, there can be only finitely many distinct walls M such that a geodesic segment contained in this wall is within bounded distance of p . But c (and thus c') crosses infinitely many walls M , so we obtain a contradiction. Thus c and c' are equivalent, and moreover their distance from each other is bounded in terms of δ . Hence $\xi = \xi'$.

Now assume that W is not word hyperbolic, and let W_T be an affine or reducible special subgroup which is an obstruction to the hyperbolicity of W , as in Theorem 43 above. If $C(\xi) = \{\xi\}$ then the closure $\overline{C(\xi)}$ equals $C(\xi)$, so by Proposition 17 and Theorem 26 above, it suffices to find infinite reduced words \mathbf{i} and \mathbf{j} such that $\mathbf{j} < \mathbf{i}$ (strict inequality) and $\partial_T X(\mathbf{i}) \neq \partial_T X(\mathbf{j})$. The existence of suitable pairs \mathbf{i} and \mathbf{j} in such W_T follows from our discussion of the affine and reducible cases above. □

Since each set $\partial_T X(\mathbf{i})$ is an equivalence class $C(\xi)$, we have established the characterisation of word hyperbolicity of W in Theorem 3. Now if W is hyperbolic, then as each $\partial_T X(\mathbf{i})$ is a singleton, we have $\overline{\partial_T X(\mathbf{i})} = \partial_T X(\mathbf{i})$ for each infinite reduced word \mathbf{i} . Hence by Theorem 1(1) and (4), infinite reduced words \mathbf{i} and \mathbf{j} are comparable only if they are in the same block. The final claim of Theorem 3 then follows from Theorem 2. This completes the proof of Theorem 3.

Example 2. Let P be a regular dodecahedron in 3-dimensional (real) hyperbolic space with all dihedral angles right angles, and let W be the Coxeter group generated by reflections in the codimension one faces of P . Then W is word hyperbolic and the Davis complex X may be equipped with a piecewise hyperbolic metric so that X is isometric to the induced tessellation of hyperbolic 3-space by copies of P . This tessellation is depicted in Figure 2.

The Tits boundary $\partial_T X$ is a 2-dimensional sphere and the distance between any two points in $\partial_T X$ is infinite. Each wall of X is a hyperplane tessellated by right-angled pentagons, and may be viewed as a copy of the Davis complex from Example 1 in the introduction. Each $\xi \in \partial_T X$ is thus contained in the boundaries of at most 2 walls of X , with these walls perpendicular to each other.

Suppose \mathbf{i} is an infinite reduced word with $\xi = \xi(\mathbf{i})$. If ξ is in the boundary of two walls $M = M(r)$ and $M' = M'(r')$ then the block $B(\mathbf{i})$ consists of four equivalence classes of infinite reduced words, corresponding to the four elements of $W(\xi) = \langle r, r' \rangle \cong C_2 \times C_2$. If ξ is in the boundary of a unique wall $M = M(r)$ then the block $B(\mathbf{i})$ consists of two equivalence classes of infinite reduced words, corresponding to the two elements of the finite Coxeter group $W(\xi(\mathbf{i})) = \langle r \rangle$. If ξ is not in the boundary of any wall, then $B(\mathbf{i})$ contains a single equivalence class of infinite reduced words and $W(\xi(\mathbf{i}))$ is trivial.




FIGURE 2. The Davis complex for Example 2

Given $\xi \in \partial_T X$, define

$$\partial\mathcal{M}(\xi) := \{\partial M \mid \xi \in \partial M\} = \{\partial M \mid M \in \mathcal{M}(\xi)\}.$$

That is, $\partial\mathcal{M}(\xi)$ is the set of boundaries of walls which contain the point ξ . It is immediate that for all $\xi \in \partial M$, $|\partial\mathcal{M}(\xi)| \leq |\mathcal{M}(\xi)|$.

Proposition 47. *Suppose W is word hyperbolic and one-ended. Then for all $\xi \in \partial_T X$, we have $|\partial\mathcal{M}(\xi)| = |\mathcal{M}(\xi)| < \infty$.*

Proof. Since $|\partial\mathcal{M}(\xi)| \leq |\mathcal{M}(\xi)| < \infty$, it suffices to show that if M, M' are distinct walls in $\mathcal{M}(\xi)$ then $\partial M \neq \partial M'$. By Lemma 40, the intersection $M \cap M'$ is nonempty. By Proposition 6, since W is one-ended X has the geodesic extension property. Thus by Lemma 41, $\partial M \neq \partial M'$. \square

We do not know how to characterise those W such that $\partial\mathcal{M}(\xi)$ is finite for all $\xi \in \partial_T X$ (noting that, by results above, if W is affine or word hyperbolic then all $\partial\mathcal{M}(\xi)$ are finite).

6.6. Isolated flats. In this section we use results of Caprace [4] to discuss some W for which X is a space with isolated flats, in the sense studied by Hruska–Kleiner in [13]. A *flat* in X is defined to be a subset which is isometric to Euclidean space of dimension ≥ 2 .

Given $T \subset S$ and a chamber wK of X , the *residue of type T* containing wK is the subcomplex of X consisting of all chambers $w'K$ such that $w^{-1}w' \in W_T$. We denote by R_T the residue of type T containing the base chamber K . The stabiliser in W of the residue R_T is W_T , and for any $w \in W$ the stabiliser of wR_T is $W_T^w = wW_Tw^{-1}$. Therefore every parabolic subgroup of W stabilises some collection of residues.

By the construction and metrisation of X , there is a canonical isometric embedding of the Davis complex for (W_T, T) , denoted X_T , into R_T . (In [4], the sets X_T and R_T appear to have been identified, although in [5] and elsewhere a residue is defined as a union of chambers of X .) Hence

for every $w \in W$, there is a canonical isometrically embedded copy of X_T in the residue wR_T . The action of W_T^w on wR_T preserves this copy of X_T .

Assume that W is not word hyperbolic. Then by Theorem 43 above, either there is some $T \subset S$ such that W_T is affine of rank ≥ 2 , or there are non-spherical $T', T'' \subset S$ such that $[T', T''] = 1$. If W_T is affine of rank ≥ 2 then X_T is isometric to Euclidean space of dimension ≥ 2 . Hence if W_T is affine of rank ≥ 2 , then every residue wR_T contains a canonical flat which is preserved by the action of the affine parabolic subgroup W_T^w .

Let \mathcal{T} be the collection of all maximal subsets $T \subset S$ such that W_T is affine (possibly reducible) and of rank ≥ 2 . We make the following additional assumptions on the set \mathcal{T} .

- (RH1) For each pair of irreducible non-spherical subsets $T', T'' \subset S$ such that $[T', T''] = 1$, there exists $T \in \mathcal{T}$ such that $T' \cup T'' \subset T$.
- (RH2) For all $T', T'' \in \mathcal{T}$ with $T' \neq T''$, the intersection $T' \cap T''$ is spherical.

Since W is not word hyperbolic, by (RH1) the set \mathcal{T} is nonempty.

As in the proof of Theorem A in [4], for each $T \in \mathcal{T}$, denote by $N(W_T)$ the normaliser of W_T in W . Now define \mathcal{R} to be the collection of all residues wR_T where w runs through a transversal for $N(W_T)$ in W . Note that we are considering representatives w for the cosets of $N(W_T)$ in W , not for the cosets of W_T in W , and so by Proposition 5.5 of [9] the set \mathcal{R} will *not* contain all residues of types $T \in \mathcal{T}$. Let \mathcal{F} be the collection of canonical flats contained in the residues in \mathcal{R} .

By Theorem A, Lemma 4.2 and Corollary D of [4], the Davis complex X then has *isolated flats*, meaning that the following conditions hold:

- (1) there is a constant $D < \infty$ such that every flat of X is in the D -neighbourhood of some element of \mathcal{F} ; and
- (2) for each $0 < \rho < \infty$, there is a constant $\kappa = \kappa(\rho) < \infty$ so that for any two distinct flats $F, F' \in \mathcal{F}$, the intersection of the ρ -neighbourhoods of F and F' has diameter $< \kappa$.

(Also, W is relatively hyperbolic with respect to its collection of maximal affine parabolic subgroups.)

By Theorem 1.2.1 of [13], since X has isolated flats the Tits boundary $\partial_T X$ is a disjoint union of isolated points and standard Euclidean spheres. Note that by condition (1) in the definition of isolated flats, for every flat F in X there is a flat $F' \in \mathcal{F}$ such that $\partial_T F \subset \partial_T F'$. By Theorem 5.2.5 of [13], the spherical components of $\partial_T X$ are precisely the Tits boundaries of the flats in \mathcal{F} . We now analyse these spherical components.

For any flat F in X , define $\mathcal{M}(F)$ to be the set of walls which separate points in F , and let $W(\mathcal{M}(F))$ be the subgroup of W generated by the set of reflections in walls in $\mathcal{M}(F)$. Then by Corollary 3.2 of [4], the group $P := \text{Pc}(W(\mathcal{M}(F)))$ is a direct product of irreducible affine Coxeter groups. (Here, $\text{Pc}(W(\mathcal{M}(F)))$ denotes the parabolic closure of the group $W(\mathcal{M}(F))$.) Moreover, by Proposition 3.3 of [4], the flat F is contained in some residue $R = wR_T$ whose stabiliser is P . Thus in particular if $F \in \mathcal{F}$ then by construction, F is contained in a unique residue $R \in \mathcal{R}$, and since the stabiliser of each residue in \mathcal{R} is a maximal affine parabolic subgroup of W we have that $P = \text{Pc}(W(\mathcal{M}(F)))$ is a maximal affine parabolic. Further, since P preserves R , the canonical flat $F \subset R$ intersects all of the walls corresponding to the reflections in P , and so $P = \text{Pc}(W(\mathcal{M}(F))) = W(\mathcal{M}(F))$.

Proposition 48. *Let R, R' be any two residues in \mathcal{R} which have stabiliser the same maximal affine parabolic P , and let F, F' be the corresponding flats in \mathcal{F} . Then:*

- (1) $F = F'$, hence $R = R'$ is the unique residue in \mathcal{R} which is stabilised by P ;
- (2) $\mathcal{M}(F) = \mathcal{M}(F')$ induce the same arrangement of boundaries of walls in $\partial_T F$; and
- (3) if M is any wall of X such that $M \notin \mathcal{M}(F)$ and ∂M intersects $\partial_T F$, then ∂M contains $\partial_T F$.

Proof. By condition (2) in the definition of isolated flats above, to show that $F = F'$ it suffices to show that for some $0 < \rho < \infty$, the intersection of the ρ -neighbourhoods of F and F' is unbounded. Using the same argument as in the last paragraph of the proof of Proposition 3.3 of [4], we may obtain R' from R by applying a sequence of reflections which commute with P . So we may without loss of generality assume that $R' = r.R$ for some reflection r which commutes with P . Fix a point $x_0 \in F$ and let $x'_0 = r.x_0 \in F'$. Choose a finite ρ so that $\rho > d(x_0, x'_0)$ and let w be an element of infinite order in P . Then for each $n \geq 1$, we have

$$\rho > d(x_0, x'_0) = d(w^n.x'_0, w^n.x_0) = d(w^n.r.x_0, t^n.x_0) = d(r.(w^n.x_0), w^n.x_0).$$

Since the set $\{w^n.x_0 \mid n \geq 1\}$ is unbounded, we conclude that $F = F'$.

Part (2) of this result is a formal consequence of part (1). For (3), suppose M is a wall such that $M \notin \mathcal{M}(F)$, ∂M intersects $\partial_T F$ but ∂M does not contain $\partial_T F$. Then ∂M must separate two points say ξ and ξ' of the sphere $\partial_T F$. Let c and c' be geodesic rays from $x_0 \in R$ such that $c(\infty) = \xi$ and $c'(\infty) = \xi'$. Since $M \notin \mathcal{M}(F)$, the rays c and c' lie on the same side of M . But then their limit points ξ and ξ' must lie on the same side of ∂M , a contradiction. \square

Thus we have that for each maximal flat F in X there is a (unique) flat $F' \in \mathcal{F}$ such that $\partial_T F = \partial_T F'$ is a sphere component of the boundary $\partial_T X$, and the arrangement of boundaries of walls in this sphere component is precisely that induced by $\mathcal{M}(F)$.

On the other hand, if $\xi \in \partial_T X$ is an isolated point then $C(\xi) = \{\xi\}$, since the $C(\xi)$ are connected, and we can say little else.

6.7. Acknowledgements. We thank the University of Sydney for travel support, Ákos Seress for a helpful suggestion and Martin Bridson for helpful conversations. We also thank Pierre-Emmanuel Caprace and Jean Lécureux for bringing to our attention the references [6], [12] and [16]. The first author thanks Pasha Pylyavskyy for collaboration which led to the line of thinking in this work. Figures 1 and 2 were produced using Jeff Weeks' free software KaleidoTile and CurvedSpaces, respectively. We thank an anonymous referee for thorough reading and many helpful suggestions.

T.L. was supported by NSF grants DMS-0901111 and DMS-1160726, and by a Sloan Fellowship. A.T. was supported by ARC grant DP110100440 and by an Australian Postdoctoral Fellowship.

REFERENCES

- [1] A. BJÖRNER AND F. BRENTI, *Combinatorics of Coxeter groups*, vol. 231 of Graduate Texts in Mathematics, Springer, New York, 2005.
- [2] M. R. BRIDSON AND A. HAEFLIGER, *Metric spaces of non-positive curvature*, vol. 319 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer-Verlag, Berlin, 1999.
- [3] B. BRINK AND R. B. HOWLETT, *A finiteness property and an automatic structure for Coxeter groups*, Math. Ann., 296 (1993), pp. 179–190.
- [4] P.-E. CAPRACE, *Buildings with isolated subspaces and relatively hyperbolic Coxeter groups*, Innov. Incidence Geom., 10 (2009), pp. 15–31.
- [5] P.-E. CAPRACE AND F. HAGLUND, *On geometric flats in the CAT(0) realization of Coxeter groups and Tits buildings*, Canad. J. Math., 61 (2009), pp. 740–761.
- [6] P.-E. CAPRACE AND J. LÉCUREUX, *Combinatorial and group-theoretic compactifications of buildings*, Ann. Inst. Fourier (Grenoble), 61 (2011), pp. 619–672.
- [7] P. CELLINI AND P. PAPI, *The structure of total reflection orders in affine root systems*, J. Algebra, 205 (1998), pp. 207–226.
- [8] M. W. DAVIS, *The geometry and topology of Coxeter groups*, vol. 32 of London Mathematical Society Monographs Series, Princeton University Press, Princeton, NJ, 2008.
- [9] V. V. DEODHAR, *On the root system of a Coxeter group*, Comm. Algebra, 10 (1982), pp. 611–630.
- [10] ———, *A note on subgroups generated by reflections in Coxeter groups*, Arch. Math. (Basel), 53 (1989), pp. 543–546.
- [11] M. DYER, *Reflection subgroups of Coxeter systems*, J. Algebra, 135 (1990), pp. 57–73.

- [12] M. F. HAGEN, *The simplicial boundary of a CAT(0) cube complex*, *Algebr. Geom. Topol.*, 13 (2013), pp. 1299–1367.
- [13] G. C. HRUSKA AND B. KLEINER, *Hadamard spaces with isolated flats*, *Geom. Topol.*, 9 (2005), pp. 1501–1538. With an appendix by the authors and Mohamad Hindawi.
- [14] K. ITO, *Parameterizations of infinite biconvex sets in affine root systems*, *Hiroshima Math. J.*, 35 (2005), pp. 425–451.
- [15] T. LAM AND P. PYLYAVSKYY, *Total positivity for loop groups II: Chevalley generators*, *Transform. Groups*, 18 (2013), pp. 179–231.
- [16] G. A. NIBLO AND L. D. REEVES, *Coxeter groups act on CAT(0) cube complexes*, *J. Group Theory*, 6 (2003), pp. 399–413.
- [17] J.-P. SERRE, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, 2074 EAST HALL, 530 CHURCH STREET, ANN ARBOR, MI 48109-1043, USA

E-mail address: `tfylam@umich.edu`

SCHOOL OF MATHEMATICS AND STATISTICS F07, UNIVERSITY OF SYDNEY NSW 2006, AUSTRALIA

Current address: School of Mathematics and Statistics, University of Glasgow, 15 University Gardens, Glasgow G12 8QW, United Kingdom

E-mail address: `anne.thomas@glasgow.ac.uk`