CANONICAL JOIN REPRESENTATIONS AND COMBINATORICS OF CONE TYPES IN COXETER GROUPS

YEEKA YAU

ABSTRACT. In this article, we establish some new combinatorial properties of elements in Coxeter groups. Firstly, we generalise Reading's theorem on the canonical join representations of elements in finite Coxeter groups to all finitely generated Coxeter groups.

Secondly, we show that for any element x in a Coxeter group W and root β in its inversion set $\Phi(x)$, the set of elements $y \in W$ satisfying $\Phi(x) \cap \Phi(y) = \{\beta\}$ is convex in the weak order and admits a unique minimal representative. This is strongly connected to determining the cone type of elements of W and leads to efficient computational methods to determine whether arbitrary elements of W have the same cone type.

1. INTRODUCTION

In Reading's seminal work on the Lattice Theory of the Poset of Regions of hyperplane arrangements [7, Chapters 9 & 10], important combinatorial properties of hyperplane arrangements are developed. An application of the theory gives some very nice results on the combinatorics of (mainly finite) Coxeter groups. Of particular interest in this article is Reading's theorem which gives an explicit description of the *canonical join representation* for each element in a finite Coxeter group.

Let us briefly fix some definitions and terminology in order to precisely state this result (further definitions and examples are given in Section 2). Let $(L, \leq)$ be a poset. For $U \subseteq L$, the element $x = \bigvee U$ is the *join* of U if x is the least upper bound of U with respect to $\leq$. A *join representation* for $x \in L$ is an identity of the form $x = \bigvee U$ for some $U \subseteq L$. A join representation $x = \bigvee U$ is *irredundant* if there does not exist a proper subset $U' \subset U$ such that $x = \bigvee U'$. For subsets U, V of L, denote $U << V$ if for every $u \in U$ there exists $v \in V$ such that $u \leq v$. The relation $U << V$ is called a *join-refinement* since if $U << V$ then $\bigvee U \leq \bigvee V$. We have the following key definition.

Definition 0.1. Let L be a poset and $x \in L$. An expression $x = \bigvee U$ is a *canonical join representation* if U is irredundant and if any other join representation $x = \bigvee V$ has $U << V$.

Let (W, S) be a finitely generated Coxeter system and let Φ be an associated root system. For $w \in W$, let $\ell(w)$ be the *length* of w and let $\Phi(w)$ be the inversion set of w. A basis of $\Phi(w)$ is the set $\Phi^1(w) = \{\beta \in \Phi(w) \mid \ell(s_\beta w) = \ell(w) - 1\}$, called the *short inversions* of w. The *right weak order* on W is defined by $x \preceq y$ if and only if $\ell(y) = \ell(x) + \ell(x^{-1}y)$. Equivalently, in terms of reduced words, $x \preceq y$ if and only if there is a reduced word for y which begins with a reduced word for x (we also call x a *prefix* of y). The *right descent set* of w is $D_R(w) = \{s \in S \mid ws \preceq w\}$ and $\Phi^R(w) = \{-w\alpha_s \mid s \in D_R(w)\}$ is the corresponding set of *right-descent roots*. Note that $\Phi^R(w) \subseteq \Phi^1(w)$.

An application of Reading's theory to finite Coxeter groups is the following.

Theorem 0.1. [7, Theorem 10-3.9] Suppose W is a finite Coxeter group and $w \in W$. For each $\beta \in \Phi^R(w)$ there is a unique minimal length element $j_\beta \in \{v \mid v \preceq w, \beta \in \Phi(v)\}$. The canonical join representation of w is $w = \bigvee\{j_\beta \mid \beta \in \Phi^R(w)\}$.

Let us give some geometric intuition for Theorem 0.1.

Example 1.1. Consider the Coxeter complex for the finite Coxeter group of type A_2, the dihedral group of order 6. The inversion set of the element sts, are the roots $\Phi(sts) = \{\alpha_t, \alpha_s, \alpha_s + \alpha_t\}$ which bijectively correspond to the walls (or equivalently hyperplanes) H_t, H_s, H_{sts} separating sts from the identity e (see Section 2 for a more detailed description of the concepts of walls, hyperplanes and roots). The right descent roots of sts, $\Phi^R(sts)$ are the roots $\{\alpha_s, \alpha_t\}$ corresponding to the walls H_s and H_t. The element t (shaded in red) is the unique minimal length element of $\{v \mid v \preceq sts, \alpha_t \in \Phi(v)\}$ (the prefixes of sts separated from e by H_t). Similarly s (shaded in blue) is the unique minimal length element of $\{v \mid v \preceq sts, \alpha_s \in \Phi(v)\}$. It can be seen that $sts = t \bigvee s$.
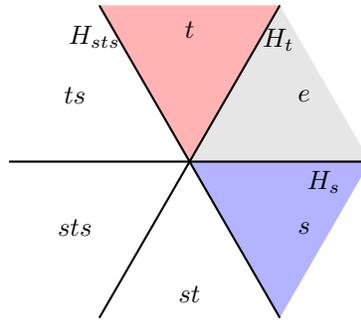
FIGURE 1. Example of Reading's canonical join representation theorem from the Coxeter group of type A_2.

Reading's proof of Theorem 0.1 relies on geometric properties of hyperplane arrangements associated with finite Coxeter groups (in particular *tightness* of the Coxeter arrangement, see [7, Section 10-2.1]); properties which generally fail in the infinite setting.

In this article, we give a uniform proof of this result; generalising Theorem 0.1 to arbitrary Coxeter groups.

Theorem 1. *Let (W, S) be a finitely generated Coxeter group and $w \in W$. For each $\beta \in \Phi^R(w)$ there is a unique minimal length element $j_\beta \in \{v \mid v \preceq w, \beta \in \Phi(v)\}$. The canonical join representation of w is $w = \bigvee \{j_\beta \mid \beta \in \Phi^R(w)\}$.*

Example 1.2. Consider the Coxeter complex for the infinite Coxeter group of type $\widetilde{A}_2$. The inversion set $\Phi(w)$ bijectively corresponds to the set of walls separating w from the identity (shaded in grey). The short inversions $\Phi^1(w)$ consists of the solid blue walls (H_α and H_β) together with the dashed blue walls and the horizontal black wall (note that the elements obtained by reflecting w in those walls has length $\ell(w) - 1$). The right-descent roots $\Phi^R(w)$ are the walls H_α and H_β. The elements x and y are the unique minimal length elements of the sets $\{v \mid v \preceq w, \alpha \in \Phi(v)\}$ and $\{v \mid v \preceq w, \beta \in \Phi(v)\}$ respectively and $w = x \bigvee y$.
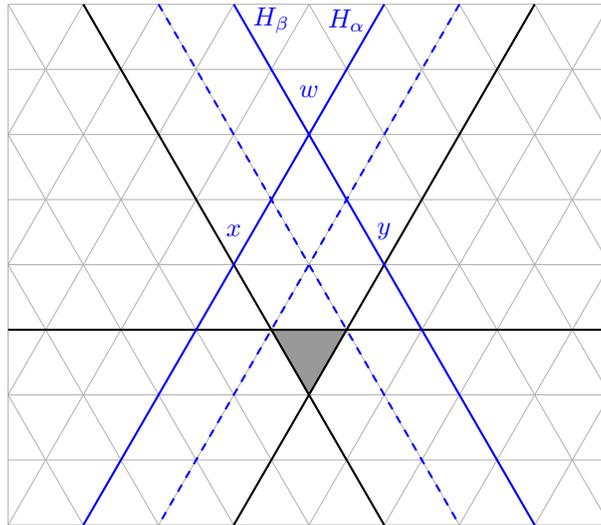

FIGURE 2. Illustration of Theorem 1

To prove Theorem 1 we first prove a slightly stronger statement than the first part of the theorem; we show that for each $\beta \in \Phi^1(w)$ there is a unique minimal length element $j_\beta \in \{v \mid v \preceq w, \beta \in \Phi(v)\}$. We also utilise the technology of the *short inversion* poset recently developed by Dyer, Hohlweg, Fishel and Marks [6] in their study of low-elements and Shi-arrangements (see Section 2.1.5 for a detailed definition and example).

The second part of this article studies some combinatorics related to the *cone type* of an element $x \in W$. The second main result of this paper is the following (see Theorem 4.1).

Theorem 2. *Let $x \in W$ and $\beta \in \Phi(x)$. If there exists $y \in W$ such that*

$$\Phi(x) \cap \Phi(y) = \{\beta\},$$

then there exists a unique element $y_\beta \in W$ of minimal length with this property.

Theorem 2 has some important consequences for understanding the structural properties of cone types in Coxeter groups. Let us now fix some additional terminology and notation and explain this view.

For $w \in W$ we denote

$$T(w) = \{v \in W \mid \ell(wv) = \ell(w) + \ell(v)\}$$

to be the *cone type* of w. Intuitively, the cone type of w consists of all elements $v \in W$ such that the concatenation of reduced words for w and v is a reduced word for $w \cdot v$. The set of cone types of W is $\mathbb{T} = \{T(w) \mid w \in W\}$. A consequence of Brink and Howlett's seminal work in [3] is that the set $\mathbb{T}$ is finite and by the Myhill-Nerode theorem, the number of cone types of W is the size (number of states) of the minimal automaton recognising the language of reduced words of (W, S). Fundamentally, it helps to visualise a cone type in terms of the inverses of elements in the following sense:

Example 1.3. Consider the Coxeter group of type $\widetilde{G}_2$. The identity is the alcove shaded in dark grey. Let x be the element as indicated below. The cone type of x^{-1}, $T(x^{-1})$ is the region shaded in grey below; and can be visualised as the intersection of the half-spaces (the side containing the identity) associated to the walls separating x from the identity e.
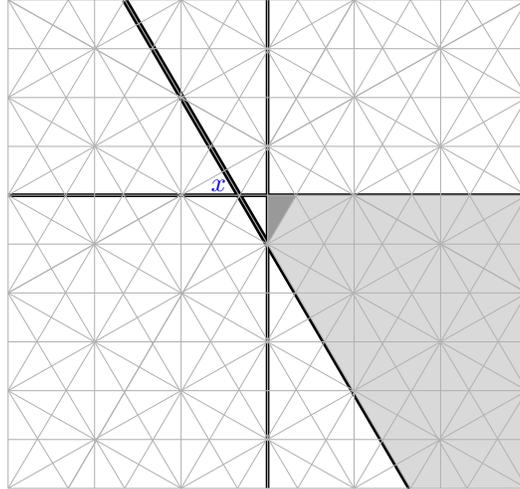

FIGURE 3. Example of a cone type in the Coxeter group of type $\widetilde{G}_2$.

Note that the cone type in Figure 3 is bounded by three walls (or hyperplanes) and that one of the walls separating x from e is redundant for determining $T(x^{-1})$. The roots corresponding to those *boundary* walls are the most precise set of roots which determine the cone type [12, Theorem 2.8]. These *boundary roots* of $T := T(x^{-1})$ are characterised in the following way

$$\partial T = \{\beta \in \Phi^+ \mid \exists\, y \in W \text{ with } \Phi(x) \cap \Phi(y) = \{\beta\}\}$$

Let us introduce a few more definitions to illuminate the structure of cone types and illustrate the significance of Theorem 2.

Definition 1.4. Let $x \in W$ and $\beta \in \Phi^+$. We say an element $y \in W$ is a *witness of β with respect to x* if

$$\Phi(x) \cap \Phi(y) = \{\beta\}$$

We denote $\partial T(x^{-1})_\beta$ to be the set of *witnesses of β with respect to x*. More precisely,

$$\partial T(x^{-1})_\beta = \{y \in W \mid \Phi(x) \cap \Phi(y) = \{\beta\}\}$$

For a cone type T, the *cone type part* $Q(T)$ are the elements $w \in W$ with $T(w^{-1}) = T$.

Remark 1.1. For each cone type T and $x, y \in Q(T)$ it is not generally true that $\partial T(x^{-1})_\beta = \partial T(y^{-1})_\beta$. Continuing the example from Figure 3 consider the following
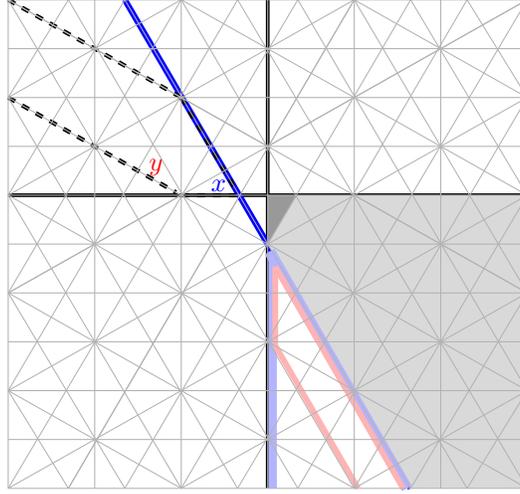

FIGURE 4. Let x and y be the elements illustrated above. One can see that $T(x^{-1}) = T(y^{-1})$, so x and y are in the same cone type part $Q(T)$; which is the region bounded by the dashed lines, and the horizontal black and blue hyperplanes (the region looks like a downward right pointing arrow). The cone type $T := T(x^{-1}) = T(y^{-1})$ is represented by the gray shaded region. Let $\beta \in \partial T$ correspond to the blue hyperplane. The region bounded by the red highlighted lines is $\partial T(y^{-1})_\beta$ and the region bounded by the blue highlighted lines is $\partial T(x^{-1})_\beta$.

Since for elements x, y in the same cone type part $Q(T)$ we do not necessarily have $\partial T(x^{-1})_\beta = \partial T(y^{-1})_\beta$ it is natural to make the following definition at the level of the cone type.

Definition 1.5. Let $T \in \mathbb{T}$ be a cone type and $\beta \in \partial T$. The *witnesses of β with respect to T* is

$$\partial T_\beta = \{y \in W \mid \Phi(x) \cap \Phi(y) = \{\beta\} \text{ for all } x \in Q(T)\}$$
$$= \bigcap_{x \in Q(T)} \partial T(x^{-1})_\beta$$

In [12] we showed that for each boundary root $\beta \in \partial T$ and $x \in Q(T)$ there is *a* witness $y \in W$ such that $\Phi^R(y) = \{\beta\}$ (see Proposition 2.4). However, it was not known whether such an element y is unique, let alone that it is the minimal length representative of $\partial T(x^{-1})_\beta$ and in fact ∂T_β.

As a consequence of Theorem 2 we obtain the following, which by the uniqueness of the element y, is a stronger statement than [12, Theorem 2.6].

Corollary 1. *For each cone type T and $\beta \in \partial T$ there is a unique minimal length element $y_\beta \in W$ such that*

$$\Phi(x) \cap \Phi(y_\beta) = \{\beta\}$$

for all $x \in Q(T)$. Moreover, $\Phi^R(y) = \{\beta\}$ and $y_\beta \preceq y$ for all $y \in \partial T(x^{-1})_\beta$ for each $x \in Q(T)$. Thus $y_\beta \preceq y$ for all $y \in \partial T_\beta$.

We give a special name for the property enjoyed by the element y_β in relation to the sets $\partial T(x^{-1})_\beta$ and ∂T_β. For a subset $X \subseteq W$, we say X is *gated* if there is a unique minimal length element $x \in X$ such that $x \preceq y$ for all $y \in X$. We say the subset X is *convex* if for all $x, y \in X$ and all reduced expressions $x^{-1}y = s_1 \cdots s_n$, the elements $xs_1 \cdots s_j$ with $0 \leq j \leq n$ is in X. In terms of the Cayley graph X^1 of (W, S) this means that all elements on geodesics between x and y are in X.

Theorem 2 and Corollary 1 reveals that each cone type T is associated with a number of convex, gated subsets of W, including its cone type part $Q(T)$ (see Theorem 2.2), the sets $\partial T(x^{-1})_\beta$ for $x \in Q(T)$ and $\beta \in \partial T$ (as well as the cone type T itself with the identity as the gate).

We illustrate in Figure 5 an example in the case of the Coxeter group of type $\widetilde{G}_2$. The figure shows a cone type T, its cone type part $Q(T)$ and the sets $\partial T(x^{-1})_\beta$ for each $\beta \in \partial T$ for the minimal element x of $Q(T)$.
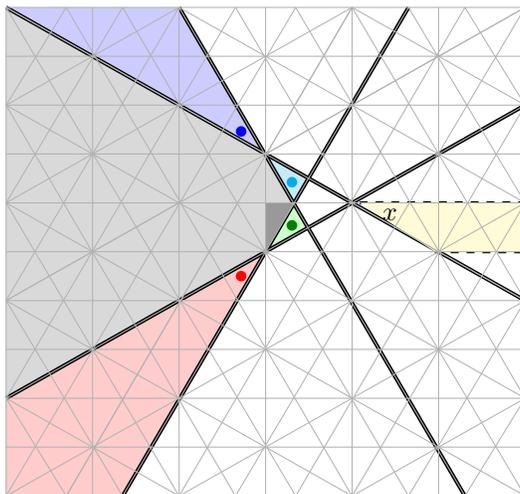

FIGURE 5. Let the element x be indicated as above. The grey shaded region is the cone type $T := T(x^{-1})$ (the darker grey alcove represents the identity) with boundary roots ∂T indicated by the thickened black walls. The yellow shaded region is the cone type part $Q(T)$. The remaining coloured shaded regions are the sets $\partial T(x^{-1})_\beta$ for each $\beta \in \partial T$ and the corresponding coloured dots are the gates of those regions.

Let us note a few more properties of the unique minimal length witnesses of boundary roots. One may notice that by symmetry, $y \in \partial T(x^{-1})_\beta$ if and only if $x \in \partial T(y^{-1})_\beta$. Furthermore, the gates of the sets $Q(T)$ for $T \in \mathbb{T}$ (we denote these elements by Γ) are precisely the elements $y \in W$ such that $\Phi^R(y) \subseteq \partial T(y^{-1})$ (see Theorem 2.3). It thus follows by Corollary 1 that the unique minimal witnesses of boundary roots are gates of cone type parts. In particular, since these minimal witnesses have a single right-descent, they are what we call *tight gates* denoted Γ^0. More precisely

$$\Gamma^0 = \{x \in \Gamma \mid |\Phi^R(x)| = 1\}$$

Theorem 3. *Let Γ be the gates of the cone type parts $\{Q(T) \mid T \in \mathbb{T}\}$. Then*

$$\Gamma^0 = \{g \in \Gamma \mid |\Phi^R(g)| = 1\}$$

is the set of join-irreducible elements of the partially ordered set $(\Gamma, \preceq)$.

A consequence of Theorem 3 is that the set of tight gates Γ^0 completely determines the set of cone types $\mathbb{T}$ of W (see Corollary 5.2). In addition, we obtain the following (see Proposition 4.1).

Corollary 2. *The set*

$$\Gamma^0 = \{g \in G \mid |\Phi^R(g)| = 1\}$$

is closed under suffix.

Corollary 2 then leads to a very efficient algorithm to compute the set Γ^0 without having to compute the entire set Γ (effectively all the cone types) first. A consequence of this result is that to determine whether two elements x, y have the same cone type, one is only required to compute Γ^0, a set strictly smaller than Γ (with only a few exceptions). Our computations using Sagemath ([4]) show that this set is in many cases much smaller than Γ as the rank of W grows (see Algorithm 6.1 and Figure 10).

2. PRELIMINARIES

To begin, we review some essential properties of Coxeter systems, the standard geometric representation, the Coxeter complex, Cayley graph and cone types. For general Coxeter group theory the primary references are [11], [2] and [14]. For more on cone types see [12].

2.1. **Coxeter Groups.** Let (W, S) be a Coxeter system with $|S| < \infty$. The *length* of $w \in W$ is

$$\ell(w) = \min\{n \geq 0 \mid w = s_1 \ldots s_n \text{ with } s_1, \ldots, s_n \in S\},$$

and an expression $w = s_1 \ldots s_n$ with $n = \ell(w)$ is called a *reduced expression* or *reduced word* for w. The *left descent set* of w is $D_L(w) = \{s \in S \mid \ell(sw) = \ell(w) - 1\}$, similarly the *right descent set* is $D_R(w) = \{s \in S \mid \ell(ws) = \ell(w) - 1\}$. It is well known that the set

$$R = \{wsw^{-1} \mid w \in W, s \in S\}$$

is the set of *reflections* of W (see [2, Chapter 1]).

2.1.1. *Standard Geometric Representation and root system of W.* Let V be an $\mathbb{R}$-vector space with basis $\{\alpha_s \mid s \in S\}$. Define a symmetric bilinear form on V by linearly extending $\langle \alpha_s, \alpha_t \rangle = -\cos(\pi/m(s,t))$. The Coxeter group W acts on V by the rule $sv = v - 2\langle v, \alpha_s \rangle \alpha_s$ for $s \in S$ and $v \in V$, and the root system of W is $\Phi = \{w\alpha_s \mid w \in W, s \in S\}$. The elements of Φ are called *roots*, and the *simple roots* are $\Delta = \{\alpha_s \mid s \in S\}$.

Each root $\alpha \in \Phi$ can be written as $\alpha = \sum_{s \in S} c_s \alpha_s$ with either $c_s \geq 0$ for all $s \in S$, or $c_s \leq 0$ for all $s \in S$. In the first case α is called *positive* (written $\alpha > 0$), and in the second case α is called *negative* (written $\alpha < 0$). Denote Φ^+ (resp. Φ^-) to be the positive roots (resp. negative roots) of Φ. By definition of Φ, each root β can be written $\beta = w\alpha_s$ for some $s \in S$. The reflection $s_\beta = wsw^{-1}$ gives rise to the hyperplane in V

$$H_\beta = \{\lambda \in V \mid s_\beta \lambda = \lambda\}$$

and this hyperplane separates W into two sets (called *half spaces*)

$$H_\beta^- = \{w \in W \mid \ell(s_\beta w) < \ell(w)\} \text{ and } H_\beta^+ = \{w \in W \mid \ell(s_\beta w) > \ell(w)\}$$

Half spaces provide an important characterisation of convexity in W.

Lemma 2.1. [1, Proposition 3.94] *A subset $A \subseteq W$ is convex if and only if it is an intersection of half spaces.*

For $A \subset \Phi^+$, $cone(A)$ is the set of non-negative linear combinations of roots in A and $cone_\Phi(A) = cone(A) \cap \Phi^+$. The *(left) inversion set* of $w \in W$ is

$$\Phi(w) = \{\alpha \in \Phi^+ \mid w^{-1}(\alpha) < 0\}.$$

Geometrically, the inversion set of w corresponds to the roots β for which w is contained in the half space H_β^- (or equivalently, separated from the identity by the hyperplane H_β). A well known fact is that $x \preceq y$ if and only if $\Phi(x) \subseteq \Phi(y)$ ([10, Proposition 2.8(3)]). An important subset of $\Phi(w)$ is the set of *short inversions* $\Phi^1(w)$ (also called the *base* of $\Phi(w)$).

Definition 2.2. [10, Proposition 4.6] Let $w \in W$. The *short inversions* of $\Phi(w)$ is the set

$$\Phi^1(w) = \{\beta \in \Phi(w) \mid \ell(s_\beta w) = \ell(w) - 1\}$$

The set $\Phi^1(w)$ determines the inversion set $\Phi(w)$ in the following way (see [5, Lemma 1.7] and [8, Corollary 2.13]).

Proposition 2.1. *For $w \in W$ we have $\Phi(w) = cone_\Phi(\Phi^1(w))$ and if $A \subset \Phi^+$ is such that $\Phi(w) = cone_\Phi(A)$ then $\Phi^1(w) \subseteq A$.*

As noted earlier, another important subset of $\Phi(w)$ is the following.

Definition 2.3. Let $w \in W$. The set of *right-descent roots* of w (also called *final roots*) is the set

$$\Phi^R(w) = \{-w\alpha_s \mid s \in D_R(w)\}$$

A root $\beta \in \Phi^R(w)$ if and only if $s_\beta w = ws$ for some $s \in D_R(w)$. In terms of the Coxeter complex $\Sigma(W, S)$ or Cayley graph X^1 the right descent roots correspond to the set of *final* or *last* walls crossed for geodesics from the identity e to w (see Section 2.1.2 and Section 2.1.3). The *left-descent roots* are $\Phi^L(w) = \{\alpha_s \mid s \in D_L(w)\}$ and correspond to *initial* walls crossed for geodesics from the identity e to w.

We collect some further useful facts regarding joins in the right weak order which are used in this article.

Proposition 2.2. [10, Proposition 2.8] *If $X \subset W$ is bounded (in the right weak order), then*

$$\Phi(\bigvee X) = cone_\Phi(\bigcup_{x \in X} \Phi(x))$$

Corollary 2.1. [10, Corollary 4.7 (2)] *If $x \vee y$ exists and $\ell(s_\beta(x \vee y)) = \ell(x \vee y) - 1$, then either $\ell(s_\beta x) = \ell(x) - 1$ or $\ell(s_\beta y) = \ell(y) - 1$.*

2.1.2. *The Coxeter complex.* The Coxeter complex $\Sigma(W, S)$ is a useful combinatorial object for providing geometric intuition for many of the concepts discussed and illustrated in the figures of this article. While the formal construction of $\Sigma(W, S)$ is not required for any result in this work, it will be helpful to provide a brief description. The reader is directed to [1] for further details.

For each $w \in W$, let C_w be a combinatorial simplex with vertices indexed by S. For each $w \in W$ and $s \in S$ identify the faces of the simplices C_w and C_{ws} via their matching vertices in $S \setminus \{s\}$. The resulting simplicial complex $\Sigma(W, S)$ is known as the *Coxeter complex* of (W, S). The Coxeter complex has maximal simplices C_w, $w \in W$ called the *alcoves* of the complex which we identify with the elements of W. The Coxeter group W then acts on the alcoves simply transitively whereby $wC_v = C_{w \cdot v}$ for all $w, v \in W$.

For $\beta \in \Phi^+$ the set of simplices

$$H_\beta = \{\sigma \in \Sigma(W, S) \mid s_\beta(\sigma) = \sigma\}$$

is a *wall* of the Coxeter complex. Note that no alcoves are contained in a wall, and thus one may consider the wall H_β separating the half-spaces H_β^+ and H_β^- (the notion of half-spaces here coincide with that of half-spaces discussed in Section 2.1.1 and walls play the same role as hyperplanes). In particular, for $w \in W$, its inversion set $\Phi(w)$ bijectively corresponds to the set of walls separating the alcove C_w from C_e and the right descent roots $\Phi^R(w)$ are the walls incident to the alcove C_w which separate C_w from C_e.

2.1.3. *The Cayley graph of* (W, S). The Cayley graph of (W, S) is the dual graph of the Coxeter complex, recording adjacency of alcoves across walls labelled by simple reflections. Thus, geometric features of the Coxeter complex translate naturally into combinatorial properties of paths and convexity in the Cayley graph.

More precisely, denote X^1 to be the Cayley graph of W with vertex set $X^0 = W$ and edge-set $\{(w, ws) \mid w \in W, s \in S\}$ where each edge is of length 1. For vertices $x, y \in X^1$ a *geodesic* between x and y is a minimal length path between x and y in X^1. In the setting of X^1, for $w \in W$, $\ell(w)$ is the length of geodesics from the identity e to w.

The left action of W on X^0 induces an action of W on X^1. Under the identification of the Cayley graph with the dual graph of the Coxeter complex, walls admit a consistent interpretation in both settings. For a reflection $r \in R$, the *wall* H_r is defined as the fixed point set of r acting on the Cayley graph X^1. Equivalently, an edge (w, ws) crosses H_r if and only if $r = wsw^{-1}$, that is, if the corresponding adjacent alcoves are separated by the reflecting wall associated to r in the Coxeter complex (see Figure 6).

Under the standard geometric representation, each root $\beta \in \Phi$ determines a reflecting hyperplane $H_\beta \subset V$, whose induced wall H_{s_β} in the Coxeter complex and fixed-point set in the Cayley graph define the same separation of W into half-spaces. Throughout this article, most arguments are phrased in terms of roots and inversion sets; however, in the proofs of Proposition 3.2 and Theorem 4.1 we make essential use of the Cayley graph description, which may be visualised geometrically via the Coxeter complex.
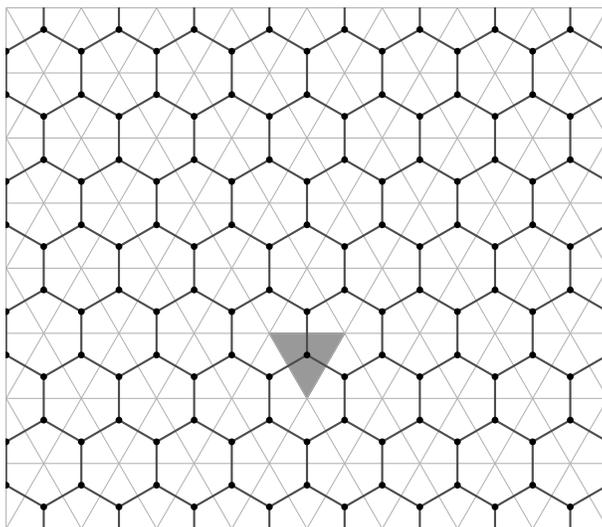

FIGURE 6. Coxeter complex for the Coxeter group of type $\widetilde{A}_2$ with Cayley graph overlaid.

Two distinct walls H_r and H_q *intersect* if H_r is not contained in a half-space for H_q (this relation is symmetric), or equivalently, $\langle r, q \rangle$ is finite. We say r, q are *sharp-angled* if r and q do not commute and there is $w \in W$ such that wrw^{-1} and wqw^{-1} are in S. Then there is a component of $X^1 \setminus (H_r \cup H_q)$ whose intersection with X^0 is a *geometric fundamental domain* for the action of $\langle r, q \rangle$ on X^0.

For $J \subseteq R$ and $w \in X^0$ the $J(w)$-*residue* is the subgraph of X^1 induced by the action of the subgroup generated by J on w. In particular, note that for $w \in W$ if $\alpha, \beta \in \Phi^R(w)$ then $w^{-1}s_\alpha w = s$ and $w^{-1}s_\beta w = t$ are in S and $\langle s_\alpha, s_\beta \rangle$ is a finite dihedral reflection subgroup with $2m(s,t)$ elements. Then $w\Phi^+_{\langle s,t \rangle}$ is the set of $m(s,t)$ roots separating the $2m(s,t)$ elements of the residue $\langle s_\alpha, s_\beta \rangle(w) = wW_{\langle s,t \rangle}$. We will use this rank–two residue structure in path–modification arguments in the proofs of Theorem 4.1 and Proposition 3.2.

In those proofs, we also consider an order on paths in X^1 using the lexicographic order on tuples with entries in $\mathbb{N}$. For tuples $\mathbf{a} = (a_1, \ldots, a_m)$ and $\mathbf{b} = (b_1, \ldots, b_n)$ we have $\mathbf{a} \leq \mathbf{b}$ if and only if one of the following holds:

(i) There exists $j \leq \min(m, n)$ such that $a_i = b_i$ for all $i < j$ and $a_j < b_j$, or;

(ii) $a_i = b_i$ for all $i \leq m$ and $m < n$.

For a path $\pi = (p_0, \ldots, p_n)$ in X^1 define the ordered tuple $\Pi = (n_L, \ldots, n_1)$ where n_j is the number of elements $p_i \in \pi$ with $\ell(p_i) = j$ and $L = max_{i=0}^n \ell(p_i)$. Then for paths π, π' we have $\pi \leq \pi'$ if and only if $\Pi \leq \Pi'$.

2.1.4. *Cone Types and Boundary roots of W.* The definition of cone types, cone type parts and boundary roots were introduced in Section 1. We now recall additional structure and properties of cone types in Coxeter groups relevant to the results in this article.

The *cone type arrangement* or *cone type partition* $\mathscr{T}$ is the partition of W into its cone type parts.

$$Q(T) = \{x \in W \mid T(x^{-1}) = T\}$$

for each $T \in \mathbb{T}$ (see Figure 7 for an example). Each part of this partition can be considered as a *state* of the minimal finite state automaton recognising the language of reduced words of (W, S) (this fact isn't necessarily required, but is discussed in some depth in [12] for interested readers). The following theorem is an interesting characterisation of $\mathscr{T}$ and reveals that the parts $\{Q(T) \mid T \in \mathbb{T}\}$ of $\mathscr{T}$ can realised as intersections of the set of cone types $\mathbb{T}$.

Theorem 2.1. [16, Theorem 3.4.3] *For $x, y \in W$ define $x \sim y$ if and only if, for every cone type $T \in \mathbb{T}$, one has*

$$x \in T \text{ if and only if } y \in T$$

Then the partition of W by the $\sim$ equivalence classes is $\mathscr{T}$.

As mentioned in Section 1, each part $Q(T)$ contains a unique minimal length representative g_T, called the *gate* of $Q(T)$ which is a prefix of all elements in $Q(T)$ (these are the grey shaded alcoves in the example of Figure 7), this was proved in [12, Theorem 1] and [13, Theorem 1.3]. These elements are also known in the literature as the *gates of $\mathscr{T}$* and are denoted by Γ. Moreover, we showed in [13, Corollary 1.5] that Γ is the smallest *Garside shadow* in W; the smallest subset of W containing S and is closed under join and suffix in the right weak order.

The notion of Garside shadows will only be briefly discussed again in Section 6; we encourage the reader to see [9], [10] and [15] for more on these fascinating structures.

Theorem 2.2. [12, Theorem 1] *For each cone type T there is a unique element $m_T \in W$ of minimal length such that $T(m_T) = T$. Moreover, if $w \in W$ with $T(w) = T$ then m_T is a suffix of w.*

Remark 2.1. Note that in Theorem 2.2, $m_T = g_T^{-1}$ where g_T is the gate of the part $Q(T)$.

The *tight gates* of $\mathscr{T}$ is the set

$$\Gamma^0 = \{g \in \Gamma \mid |\Phi^R(g)| = 1\}$$

The following theorem gives a characterisation of a gate x in terms of $\Phi^R(x)$.

Theorem 2.3. [12, Theorem 4.33] *Let $x \in W$. Then $x \in \Gamma$ if and only if for each $\beta \in \Phi^R(x)$ there exists $w \in W$ with $\Phi(x) \cap \Phi(w) = \{\beta\}$.*

We list some further elementary results regarding cone types and their boundary roots which will be used throughout. We omit their proofs here.

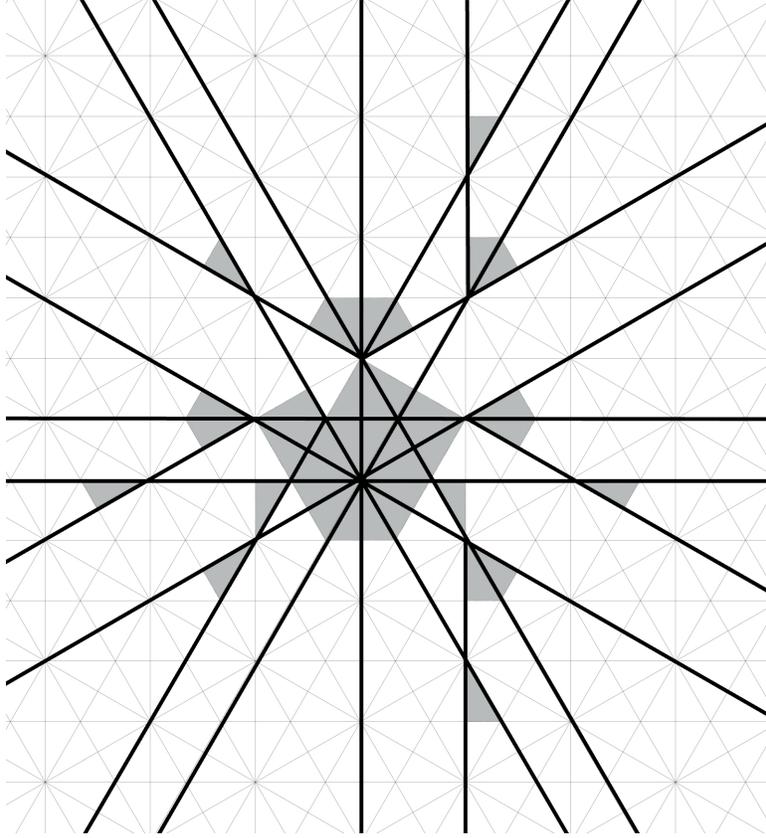**Lemma 2.4.** [12, Lemma 1.15] *If $x \preceq y$ then $T(y^{-1}) \subseteq T(x^{-1})$.*

FIGURE 7. The cone type partition $\mathscr{T}$ for the Coxeter group of type $\widetilde{G}_2$. Each part $Q(T)$ contains the elements x such that $T(x^{-1}) = T$. The shaded in alcove of each part is the gate of its part $Q(T)$.

The following result formally states the geometric interpretation of cone types as intersection of half-spaces (as illustrated in Example 1.3) and the next relates the join of a set of elements with their cone types.

Theorem 2.4. [12, Corollary 2.9] *Let T be a cone type. If $T = T(x^{-1})$ Then*

$$T = \bigcap_{\Phi(x)} H_\alpha^+ = \bigcap_{\Phi^1(x)} H_\alpha^+ = \bigcap_{\partial T(x^{-1})} H_\alpha^+$$

Proposition 2.3. [12, Proposition 1.19] *If $X \subseteq W$ is bounded with $y = \bigvee X$ then $T(y^{-1}) = \bigcap_{x \in X} T(x^{-1})$.*

We end this section by recalling two results describing boundary roots of cone types and the existence of minimal witnesses. These results provide the foundational input for Theorem 2 and Corollary 1, where we strengthen existence statements to uniqueness, canonical minimality and *gatedness*.

Theorem 2.5. [12, Theorem 2.6] *Let T be a cone type. If $T = T(x^{-1})$ then $\beta \in \partial T$ if and only if there exists $w \in W$ with*

$$\Phi(x) \cap \Phi(w) = \{\beta\}$$

Moreover, if $\beta \in \partial T$ then there exists $w \in W$, independent of x, such that $\Phi(x) \cap \Phi(w) = \{\beta\}$ whenever $T = T(x^{-1})$.

Proposition 2.4. [12, Proposition 4.34] *Let $x \in W$ and $\beta \in \Phi^+$. Suppose there exists $w \in W$ such that $\Phi(x) \cap \Phi(w) = \{\beta\}$, and let w be of minimal length subject to this property. Then $\Phi^R(w) = \{\beta\}$ and w is a gate.*

2.1.5. *Maximal dihedral reflection subgroups and the short inversion poset.* Finally, we recall the notion of maximal dihedral reflection subgroups in Coxeter groups as this is required for the definition of the short inversion poset (introduced in [6]) which we utilise in the proof of Theorem 3.1. For a more in-depth discussion

see [10, Section 2.5]. A dihedral reflection subgroup W' of W is a subgroup of rank 2. Every dihedral reflection subgroup is contained in a *maximal* dihedral reflection subgroup (note that dihedral subgroups may be infinite).

For $\alpha, \beta \in \Phi^+$, we denote $\mathcal{M}_{\alpha,\beta}$ to be the *maximal* dihedral reflection subgroup whose associated root system $\Phi_{\mathcal{M}_{\alpha,\beta}}$ contains α and β. The simple roots of $\Phi_{\mathcal{M}_{\alpha,\beta}}$ is denoted $\Delta_{\mathcal{M}_{\alpha,\beta}}$.

We can now define the short inversion poset. Let $w \in W$. For $\alpha, \beta \in \Phi^1(w)$ define $\alpha \dot{\prec} \beta$ if $\beta \notin \Delta_{\mathcal{M}_{\alpha,\beta}}$. This is equivalent to $\alpha \in \Delta_{\mathcal{M}_{\alpha,\beta}}$ and $\beta \notin \Delta_{\mathcal{M}_{\alpha,\beta}}$ (See [6, Proposition 3.2]). Then define the relation $\preceq_w$ as the transitive and reflexive closure of $\alpha \dot{\prec} \beta$. This is a partial order on $\Phi^1(w)$.

Theorem 2.6. [6, Theorem 3.17] *Let $w \in W$. For the poset $(\Phi^1(w), \preceq_w)$, the minimal elements are the left-descent roots in $\Phi^L(w)$ and the maximal elements are the right-descent roots in $\Phi^R(w)$. In other words, for any $\beta \in \Phi^1(w)$ there is $\alpha \in \Phi^L(w)$ and $\gamma \in \Phi^R(w)$ such that $\alpha \preceq_w \beta \preceq_w \gamma$.*

For our purposes, from Theorem 2.6 we utilise the fact that for $w \in W$ and any $\beta \in \Phi^1(w)$, there is a *final* root $\gamma \in \Phi^R(w)$ such that $\beta \preceq_w \gamma$. In particular, the following fact (supplied by the next proposition) is critical for the proof of Theorem 1: If $\beta \in \Phi^1(w)$ then there is $\gamma \in \Phi^R(w)$ such that for any reduced word of w, the wall H_β is crossed before H_γ in the geodesic path in X^1 from the identity e to w (see Example 2.5).

Proposition 2.5. [6, Proposition 3.13] *The relation $\preceq_w$ is a partial order on $\Phi^1(w)$. Moreover, for any reduced word $w = s_1 s_2 \ldots s_k$ consider the following total order $\leq$ on $\Phi(w)$:*

$$\alpha_{s_1} < s_1 \alpha_{s_2} < \ldots < s_1 \ldots s_{k-1} \alpha_{s_k}$$

Then $\alpha \preceq_w \beta$ implies $\alpha \leq \beta$ for any $\alpha, \beta \in \Phi^1(w)$.

Example 2.5. Consider the same element from Example 1.2 in the Coxeter group of type $\widetilde{A}_2$. Let the roots $\alpha, \gamma \in \Phi^1(w)$ correspond to the red and blue coloured walls respectively. Together they generate an infinite dihedral reflection subgroup $W'_{\alpha,\gamma}$ and we have $W'_{\alpha,\gamma} \subset \mathcal{M}_{\alpha,\gamma}$ with simple roots $\Delta_{\mathcal{M}_{\alpha,\beta}} = \{\alpha, \beta\}$. Thus $\alpha \preceq_w \gamma$ and any geodesic path (reduced word) from e (alcove shaded in grey) to w must first cross H_α.

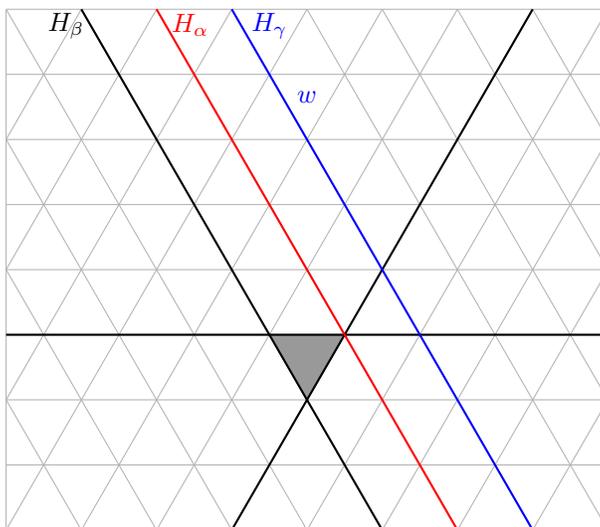

FIGURE 8. Illustration of Proposition 2.5

3. CANONICAL JOIN-REPRESENTATIONS

This section is dedicated to proving Theorem 1. We first give a few preliminary results. The following result is an obvious, but rather useful fact.

Lemma 3.1. *Let $x \in W$. If $\alpha \in \Phi(x)$ then there is $y \preceq x$ with $\Phi^R(y) = \{\alpha\}$*

Proof. Let x_0 be a prefix of x with $\alpha \in \Phi^R(x_0)$. If there is $\beta \in \Phi^R(x_0)$ with $\beta \neq \alpha$ then consider the prefix $x_1 := s_\beta x_0$ of x. Repeating this process, since each step replaces x_i by a shorter prefix $s_\beta x_i$, the process terminates and there must be some prefix x_k of x with $\Phi^R(x_k) = \{\alpha\}$.   □

Proposition 3.1. *Let $x \in W$. Then*

$$x = \bigvee \phi^1(x)$$

where $\phi^1(x) = \{z \in W \mid z \preceq x \text{ with } \Phi^R(z) = \{\alpha\} \text{ and } \alpha \in \Phi^1(x)\}$.

Proof. Clearly x is an upper bound of $\phi^1(x)$. So by [2, Theorem 3.2.1] the join of $\phi^1(x)$ exists. Let $y = \bigvee \phi^1(x)$. Then $y \preceq x$ with $\Phi(y) \subseteq \Phi(x)$. By Proposition 2.2 we have

$$\Phi(y) = cone_\Phi \Big(\bigcup_{z \in \phi^1(x)} \Phi(z) \Big) = cone_\Phi \Big(\bigcup_{z \in \phi^1(x)} cone_\Phi \Phi^1(z) \Big) = cone_\Phi \Big(\bigcup_{z \in \phi^1(x)} \Phi^1(z) \Big)$$

By definition of $\phi^1(x)$, if $\alpha \in \Phi^1(x)$ then $\alpha \in \Phi^R(z) \subseteq \Phi^1(z)$ for some $z \in \phi^1(x)$, so $\Phi^1(x) \subseteq \bigcup_{z \in \phi^1(x)} \Phi^1(z)$. Therefore,

$$\Phi(x) = cone_\Phi \big(\Phi^1(x) \big) \subseteq cone_\Phi \Big(\bigcup_{z \in X} \Phi^1(z) \Big) = \Phi(y)$$

and $x \preceq y$. Hence $x = y$. $\qquad\square$

We note that nothing in Proposition 3.1 suggests that for each $\alpha \in \Phi^1(x)$ there cannot be more than one prefix z of x with $\Phi^R(z) = \{\alpha\}$. The following proposition shows that there is a bijection between the set $\phi^1(x) = \{z \in W \mid z \preceq x \text{ with } \Phi^R(z) = \{\alpha\} \text{ and } \alpha \in \Phi^1(x)\}$ and $\Phi^1(x)$.

Proposition 3.2. *Let $x \in W$. For each $\beta \in \Phi^1(x)$ there is a unique minimal length $z \preceq x$ such that $\Phi^R(z) = \{\beta\}$. Thus the map $\Phi_R : \phi^1(x) \to \Phi^1(x)$ defined by $z \mapsto \Phi^R(z)$ is a bijection.*

Proof. Let $\beta \in \Phi^1(x)$ and denote $Z_\beta = \{z \in W \mid z \preceq x \text{ and } \beta \in \Phi(z)\}$. We show that for each $z_0, z_n \in Z_\beta$ there is $z \in Z_\beta$ such that $z_0 \succeq z \preceq z_n$. Let $\pi = (z_0, z_1, \ldots, z_n)$ be a geodesic edge path between z_0 and z_n. Note that in terms of half-spaces

$$Z_\beta = \Big(\bigcap_{\Phi^+ \setminus \Phi(x)} H_\alpha^+ \Big) \cap H_\beta^-$$

and hence by Lemma 2.1 Z_β is convex. Therefore, each $z_i \in \pi$ is also in Z_β. Thus H_β does not separate any z_j from z_j in π and each z_i is a prefix of x.

Our aim is to modify π to obtain another embedded edge path between z_0 and z_n so that there is no $z_{i-1} \prec z_i \succ z_{i+1}$. Now suppose there is $z_{i-1} \prec z_i \succ z_{i+1}$. Let H_r, H_q be the walls separating z_{i-1} from z_i and z_i from z_{i+1} respectively and let α_r, α_q be the corresponding roots. Now z_{i-1}, z_{i+1} are elements of the *residue* $R = \langle r, q \rangle (z_i)$ and $\alpha_r, \alpha_q \in \Phi^R(z_i)$. Hence by [14, Theorem 2.9] we have $y \preceq z_i$ for all $y \in R$ and as a consequence we have that $y \preceq x$ for all $y \in R$. If r and q do not commute then r, q are sharp angled and the identity e is in a geometric fundamental domain for the finite dihedral reflection subgroup $W_{\langle r,q \rangle}$ generated by the reflections r, q. We claim that all the elements of the residue $R = \langle r, q \rangle (z_i)$ lie in Z_β.

We need to show that $\beta \notin \Phi^+_{\langle r,q \rangle}$. If r and q commute then we are done since there is no other root in $\Phi^+_{\langle r,q \rangle}$, otherwise if $\beta \in \Phi^+_{\langle r,q \rangle}$ then since $\beta = k\alpha_r + j\alpha_q$ for some $k, j \in \mathbb{R}_{>0}$ and $\alpha_r, \alpha_q \in \Phi(z_i) \subseteq \Phi(x)$ this implies $\beta \in cone_\Phi(\Phi(x) \setminus \{\beta\})$, contradicting Proposition 2.1.

We now modify π and replace the subpath (z_{i-1}, z_i, z_{i+1}) with the other embedded edge-path in the residue R from z_{i-1} to z_{i+1}. This path modification decreases the complexity of π (under lexicographic order); considering π as the tuple $\Pi = (n_L, \ldots, n_2, n_1)$, where n_j is the number of elements z_i in π with $\ell(z_i) = j$ and $L = max_{i=0}^n \ell(z_i)$.

After finitely many such modifications we obtain the desired path. The fact that the minimal element z necessarily has $\Phi^R(z) = \{\beta\}$ follows by Lemma 3.1. $\qquad\square$

The next result shows that the set $\phi^1(x)$ *join-refines* every join-representation Y of x.

Corollary 3.1. *Let $x \in W$ and suppose $x = \bigvee Y$ for $Y \subseteq W$. Then for every $z \in \phi^1(x)$ there is $y \in Y$ such that $z \preceq y$.*

Proof. Since $x = \bigvee Y$, by Proposition 2.2 we have

$$
\begin{aligned}
\Phi(x) &= cone_\Phi\big(\bigcup_{y \in Y} \Phi(y)\big) \\
&= cone_\Phi\big(\bigcup_{y \in Y} cone_\Phi(\Phi^1(y))\big) \\
&= cone_\Phi\big(\bigcup_{y \in Y} \Phi^1(y)\big)
\end{aligned}
$$

Again by Proposition 2.1 we have $\Phi^1(x) \subseteq \bigcup_{y \in Y} \Phi^1(y)$ and thus for $\alpha \in \Phi^1(x)$, we have $\alpha \in \Phi^1(y)$ for some $y \in Y$. Hence by Proposition 3.2 there is $z \preceq y \preceq x$ with $\Phi^R(z) = \{\alpha\}$. Then again by Proposition 3.2, z is the unique element in $\phi^1(x)$ such that $\Phi^R(z) = \{\alpha\}$ and the result follows. $\qquad\square$

We can now prove Theorem 1.

Theorem 3.1. *Let $w \in W$. For each $\beta \in \Phi^R(w)$ there is a unique minimal length element $j_\beta \in \{v \mid v \preceq w, \beta \in \Phi(v)\}$. The canonical join representation of w is $w = \bigvee\{j_\beta \mid \beta \in \Phi^R(w)\}$.*

Proof. The first part of the theorem follows by Proposition 3.2 since $\Phi^R(w) \subseteq \Phi^1(w)$. For the second part, we first show that $w = \bigvee\{j_\beta \mid \beta \in \Phi^R(w)\}$. First, note that $j_\beta \preceq w$ for all $\beta \in \Phi^R(w)$ and so $v := \bigvee\{j_\beta \mid \beta \in \Phi^R(w)\}$ exists with $v \preceq w$. Our aim is to show that $v = w$. For $\alpha \in \Phi^1(w)$, let j_α be the unique minimal length element of $\{v \mid v \preceq w, \alpha \in \Phi(v)\}$.

By Proposition 2.2, it suffices to show that for each $\alpha \in \Phi^1(w) \setminus \Phi^R(w)$ there is $\beta \in \Phi^R(w)$ such that $w_\alpha \preceq j_\beta$ for some $w_\alpha \preceq w$ with $\alpha \in \Phi^1(w_\alpha)$; since then we have $\Phi^1(w) \subseteq \bigcup_{\beta \in \Phi^R(w)} \Phi(j_\beta)$ and thus by Proposition 2.1 and Proposition 2.2 it would then follow that

$$
\Phi(w) = cone_\Phi(\Phi^1(w)) \subseteq cone_\Phi\big(\bigcup_{\beta \in \Phi^R(w)} \Phi(j_\beta)\big) = \Phi(v)
$$

and so $w \preceq v$.

Hence let $\alpha \in \Phi^1(w) \setminus \Phi^R(w)$. By Theorem 2.6 there is $\beta \in \Phi^R(w)$ such that $\alpha \preceq_w \beta$ and by Proposition 2.5 this implies that for any reduced word $\bar{w}$ of w we have $\alpha \leq \beta$ in the total order on $\Phi(w)$. Intuitively, this means that the wall H_α is crossed *before* the wall H_β for any reduced word $\bar{w}$ of w, when considered as a walk from the identity e to w in the Coxeter complex or Cayley graph of (W, S).

Now let $\bar{w}$ be a reduced word of w and let $\bar{w}_\beta$ be the initial path of $\bar{w}$ with $\beta \in \Phi^R(w_\beta)$, where w_β is the element represented by the word $\bar{w}_\beta$. Note that $\bar{w}_\beta$ is a reduced word since $\bar{w}$ is reduced. By Proposition 3.2, since $\beta \in \Phi^1(w_\beta)$ and $w_\beta \preceq w$, we have $j_\beta \preceq w_\beta$ and so we can modify $\bar{w}$ to obtain another reduced word of w, $\bar{\bar{w}}$ with a reduced word $\bar{j}_\beta$ the initial path of $\bar{\bar{w}}$ with $\beta \in \Phi^R(j_\beta)$.

Applying the same argument to the reduced word $\bar{\bar{w}}$, by Proposition 2.5 again, there must be an initial path of $\bar{\bar{w}}$ with reduced word $\bar{w}_\alpha$ with $\alpha \in \Phi^R(w_\alpha)$ such that $w_\alpha \preceq j_\beta \preceq w$ as required.

Now let $\phi^R(w) = \{j_\beta \mid \beta \in \Phi^R(w)\}$. We now show that $\phi^R(w)$ is irredundant. If $|\Phi^R(w)| = 1$, then w is join-irreducible and the result is clear, hence assume $|\Phi^R(w)| > 1$. Suppose for some $\alpha \in \Phi^R(w)$ we have

$$
\begin{aligned}
w &= \bigvee\{j_\beta \mid \beta \in \Phi^R(w) \setminus \{\alpha\}\} \\
&= \bigvee\{j_\beta \mid \beta \in \Phi^1(w) \setminus \{\alpha\}\} \quad \text{by Proposition 3.1}
\end{aligned}
$$

It then follows that $s_\alpha w \prec w$ and $\Phi^1(s_\alpha w) = \Phi^1(w) \setminus \{\alpha\}$ and thus

$$
w \succ s_\alpha w = \bigvee\{j_\beta \mid \beta \in \Phi^1(w) \setminus \{\alpha\}\} \quad \text{again by Proposition 3.1}
$$

a contradiction. Finally, the fact that the set $\phi^R(w)$ join-refines every join representation of w follows by Corollary 3.1 since $\phi^R(w) \subseteq \phi^1(w)$. $\qquad\square$

4. Witnesses of boundary roots

In this section, we show that for any element $x \in W$ and $\beta \in \Phi(x)$, the set of elements $y \in W$ satisfying $\Phi(x) \cap \Phi(y) = \{\beta\}$ is convex in the right weak order and admits a unique minimal representative. By Definition 1.4 these are the *witnesses* of β with respect to x and means that β is a *boundary root* of the cone type $T := T(x^{-1})$.

Furthermore, we show that this minimal witness is the same one for all $z \in Q(T)$. In addition, we show that the set of *tight gates* Γ^0 is closed under suffix. We begin with a basic characterisation of the set of witnesses for a boundary root in terms of intersections of half spaces.

Lemma 4.1. *Let* $x \in W$ *and* $T := T(x^{-1})$. *Then*

$$\partial T(x^{-1})_\beta = H_\beta^- \cap \big(\bigcap_{\alpha \in \Phi(x) \setminus \{\beta\}} H_\alpha^+ \big)$$

and the set $\partial T(x^{-1})_\beta$ *is convex.*

Proof. The proof is straightforward. If $w \in \partial T(x^{-1})_\beta$ then $\alpha \notin \Phi(w)$ for all $\alpha \neq \beta \in \Phi(x)$, so $w \in H_\alpha^+$ for all $\alpha \in \Phi(x) \setminus \{\beta\}$. Since $\beta \in \Phi(w)$ we have $w \in H_\beta^-$. The reverse inclusion is also clear. Convexity follows by Lemma 2.1 since $\partial T(x^{-1})_\beta$ is an intersection of half spaces. $\square$

Theorem 4.1. *Let* $x \in W$ *and* $\beta \in \Phi(x)$ *such that there exists* $y \in W$ *with* $\Phi(x) \cap \Phi(y) = \{\beta\}$. *Then there exists a unique minimal length element* $y_\beta \in W$ *such that*

$$\Phi(x) \cap \Phi(y_\beta) = \{\beta\}$$

Furthermore, $y_\beta \in \Gamma^0$ *and the set of witnesses* $\partial T(x^{-1})_\beta$ *is a convex, gated set in* W.

Proof. The proof is very similar to the proof of Proposition 3.2. We pass between reflections in the Cayley graph X^1 of W and their corresponding roots in Φ^+. Hence consider $\partial T(x^{-1})_\beta$ to be the intersection of the half spaces H_t^- ($t := s_\beta$) and $H_{u_i}^+$ ($u_i := s_{\alpha_i}$) in X^1 corresponding to the roots β and $\{\alpha_i \mid \alpha_i \in \Phi(x) \setminus \{\beta\}\}$ (per the characterisation of $\partial T(x^{-1})_\beta$ in Lemma 4.1).

It suffices to show that for each $y_0, y_n \in \partial T(x^{-1})_\beta$ there is $y \in \partial T(x^{-1})_\beta$ with $y_0 \succeq y \preceq y_n$. Let $\pi = (y_0, y_1, \ldots, y_n)$ be a geodesic edge path between y_0 and y_n. By Lemma 4.1, each $y_i \in \partial T(x^{-1})_\beta$ for $0 \leq i \leq n$.

We proceed as in the proof of Proposition 3.2; our aim is to modify π to obtain another embedded edge path between y_0 and y_n (not necessarily geodesic) so that there is no $y_{i-1} \prec y_i \succ y_{i+1}$. Now suppose there is $y_{i-1} \prec y_i \succ y_{i+1}$. Let H_r, H_q be the walls separating y_{i-1} from y_i and y_i from y_{i+1} respectively and let α_r, α_q be the corresponding roots. Note that $\alpha_r \neq \beta$ and $\alpha_q \neq \beta$ since by convexity $y_{i-1}, y_i, y_{i+1} \in \partial T(x^{-1})_\beta$. Also, y_{i-1}, y_{i+1} are elements of the *residue* $R = \langle r, q \rangle(y_i)$ and $\alpha_r, \alpha_q \in \Phi^R(y_i)$. If r and q do not commute then r, q are sharp angled and the identity e is in a geometric fundamental domain for the finite dihedral reflection subgroup $W_{\langle r,q \rangle}$ generated by the reflections r, q. We claim that all the elements of the residue $R = \langle r, q \rangle(y_i)$ lie in $\partial T(x^{-1})_\beta$.

Since $y_i \in \partial T(x^{-1})_\beta$ this implies that $\alpha_r, \alpha_q \notin \Phi(x) \setminus \{\beta\}$. We then need to show that $\alpha' \notin \Phi(x) \setminus \{\beta\}$ for any $\alpha' \in \Phi_{\langle r,q \rangle}^+$. Hence we can assume that r and q do not commute, otherwise there is no other root in $\Phi_{\langle r,q \rangle}^+$. If there is $\alpha' \in \Phi_{\langle r,q \rangle}^+ \cap (\Phi(x) \setminus \{\beta\})$ then since $\alpha' = k\alpha_r + j\alpha_q$ for some $k, j \in \mathbb{R}_{>0}$ this implies that either $\alpha_r \in \Phi(x) \cap \Phi(y_i)$ or $\alpha_q \in \Phi(x) \cap \Phi(y_i)$. Since $\beta \in \Phi(x) \cap \Phi(y_i)$, this then implies that $|\Phi(x) \cap \Phi(y_i)| > 1$, a contradiction.

Now modify π and replace the subpath (y_{i-1}, y_i, y_{i+1}) with the other embedded edge-path in the residue R from y_{i-1} to y_{i+1}. By [14, Theorem 2.9] all the elements of R are $\preceq y_i$, since the element $\text{Proj}_R(e)$ of R is opposite to y_i. It then follows by [14, Theorem 2.15] that all elements of R lie on a geodesic edge path from $\text{Proj}_R(e)$ to y_i and hence on a geodesic edge path from e to y_i through $\text{Proj}_R(e)$. Thus this path modification decreases the complexity of π (under lexicographic order); considering π as the tuple $\Pi = (n_L, \ldots, n_2, n_1)$, where n_j is the number of elements p_i in π with $\ell(p_i) = j$ and $L = max_{i=0}^n \ell(p_i)$.

After finitely many such modifications we obtain the desired path. The fact that the minimal element y_β is necessarily a tight gate follows by Proposition 2.4 and convexity follows by Lemma 4.1. $\square$

Corollary 4.1. *For each* $T \in \mathbb{T}$ *and* $\beta \in \partial T$, *there is a unique minimal length element* $y_\beta \in W$ *such that*

$$\Phi(x) \cap \Phi(y_\beta) = \{\beta\}$$

for all $x \in Q(T)$

Proof. Let $u, v \in Q(T)$ and $\beta \in \partial T$. By Theorem 4.1 there are unique minimal length elements u', v' such that $\Phi(u) \cap \Phi(u') = \{\beta\}$ and $\Phi(v) \cap \Phi(v') = \{\beta\}$. Let g_T be the minimal element of $Q(T)$. Then $g_T \preceq u$ and $g_T \preceq v$ and $\Phi(g_T) \cap \Phi(u') = \{\beta\}$ and $\Phi(g_T) \cap \Phi(v') = \{\beta\}$. If $u' \neq v'$ then applying Theorem 4.1 to g_T, there is $z \in W$ with $z \prec u'$ and $z \prec v'$ such that $\Phi(g_T) \cap \Phi(z) = \{\beta\}$. But then also $\Phi(u) \cap \Phi(z) = \{\beta\}$ and $\Phi(v) \cap \Phi(z) = \{\beta\}$ contradicting the minimality of u' and v'. Hence we must have $u' = v'$. $\square$

As a consequence of Theorem 4.1 we also obtain the following properties for ∂T_β.

Corollary 4.2. *For each $T \in \mathbb{T}$, the following properties hold:*

(i) $\partial T_\beta \neq \emptyset$.
(ii) $\partial T_\beta = \bigcap_{y \in Q(T)} \partial T(y^{-1})_\beta$.
(iii) ∂T_β *is convex and gated.*

Proof.   (i) This directly follows from Corollary 4.1.
(ii) Follows by definition of ∂T_β.
(iii) This follows by (i) and (ii) and by the fact that each $\partial T(y^{-1})_\beta$ is convex. Moreover, by Corollary 4.1 the gate y_β is independent of $y \in Q(T)$, so the intersection remains gated with gate y_β.
$\square$

The results of this section show that for any cone type T and any boundary root $\beta \in \partial T$, there exists a unique tight gate y_β of minimal length such that $\Phi(x) \cap \Phi(y_\beta) = \{\beta\}$ for all $x \in Q(T)$. Moreover, by symmetry, there exists a unique tight gate x_β of minimal length satisfying $\Phi(x_\beta) \cap \Phi(y_\beta) = \{\beta\}$. Thus each positive root β that arises as a boundary root of some cone type canonically determines an unordered pair $\{x_\beta, y_\beta\}$ of tight gates associated to that cone type. We define the following.

Definition 4.2. A root $\beta \in \Phi^+$ is *super-elementary* if there exists $x, y \in W$ with $\Phi(x) \cap \Phi(y) = \{\beta\}$. We denote the set of super-elementary roots by $\mathcal{S}$.

By the finiteness of the set $\mathbb{T}$, the set $\mathcal{S}$ is finite for every finitely generated Coxeter group. The following example demonstrates that there are potentially multiple pairs of minimal-length tight gates x, y with $\Phi(x) \cap \Phi(y) = \{\beta\}$ and $\Phi^R(x) = \Phi^R(y) = \{\beta\}$ for a super-elementary root β.

Example 4.3. Let W be the rank 4 compact hyperbolic Coxeter group whose graph is given by
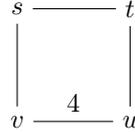
Let $\beta = \alpha_t + \sqrt{2}\alpha_u + 2\alpha_v$. Denote the elements $a = tvutv, b = uvutv, c = utvutv, d = vutv$. Then these elements are all tight gates with $\Phi^R(a) = \Phi^R(b) = \Phi^R(c) = \Phi^R(d) = \{\beta\}$ and we have

$$\Phi(a) \cap \Phi(b) = \{\beta\}$$
$$\Phi(c) \cap \Phi(d) = \{\beta\}$$

But d is not a witness of β with respect to a and b is not a witness of β with respect to c.

In subsequent work, we investigate classes of Coxeter groups in which every super-elementary root $\beta \in \mathcal{S}$ admits a unique pair of tight gates, up to order, whose inversion sets intersect precisely in $\{\beta\}$.

The following result implies that the set of tight gates Γ^0 is closed under suffix and gives more precise information about the relationship between tight gates and super-elementary roots.

Proposition 4.1. *Let $x \in \Gamma$ and $\Phi(x) \cap \Phi(y) = \{\beta\}$ and $\Phi^R(y) = \{\beta\}$ where y is of minimal length with this property. If $s \in D_L(y)$ then*

$$\Phi(sx) \cap \Phi(sy) = \{s\beta\}$$

and $sy \in \Gamma^0$.

Proof. Let $x, y \in W$ possess the properties as stated. Since $\alpha_s \in \Phi(y)$ it follows that $\alpha_s \notin \Phi(x)$. Denote $v = sx$. It then follows that $\Phi(v) = \{\alpha_s\} \sqcup s\Phi(x)$ and $\Phi(sy) = s\Phi(y) \setminus \{\alpha_s\}$. Therefore $\Phi(v) \cap \Phi(sy) = \{s\beta\}$ and it is clear that $s\beta \in \Phi^R(sy)$. Now suppose there is $\alpha \neq s\beta \in \Phi^R(sy)$. Then we have $s_\alpha sy \preceq sy$ and $\alpha_s \notin \Phi(s_\alpha sy) \cap \Phi(sy)$. Thus $ss_\alpha sy \preceq y$ with $\ell(s_\alpha sy) = \ell(sy) - 1 = \ell(y) - 2$ and $s\beta \in \Phi(s_\alpha sy)$. Then $\Phi(ss_\alpha sy) = \{\alpha_s\} \sqcup \Phi(s_\alpha sy)$ and $\beta \in \Phi(ss_\alpha sy)$. This then implies that

$$\Phi(x) \cap \Phi(ss_\alpha sy) = \{\beta\}$$

and $\ell(ss_\alpha sy) = \ell(y) - 1$ contradicting the minimality of y.
$\square$

The following two additional results demonstrate how boundary roots and the set of witnesses can be iteratively computed.

Lemma 4.4. [16, Lemma 3.3.4] *Let $w \in W$ and $s \in S$ where $s \notin D_L(w)$. Then*

$$\partial T(w^{-1}s) = \{\alpha_s\} \sqcup s\big(\{\alpha \in \partial T(w^{-1}) \mid \exists\, x \in \partial T(w^{-1})_\alpha \text{ with } s \in D_L(x)\}\big)$$

Proof. We first show the left to right inclusion. First, $\Phi(sw) \cap \Phi(s) = \{\alpha_s\}$ so $\alpha_s \in \partial T(w^{-1}s)$. If $\alpha \neq \alpha_s \in \partial T(w^{-1}s)$ then by Theorem 2.5 there is $x \in W$ such that $\Phi(sw) \cap \Phi(x) = \{\alpha\}$. We claim that $s\alpha \in \partial T(w^{-1})$ and that sx is a witness of $s\alpha$ with respect to w. Since $\alpha_s \in \Phi(sw)$ we have $\alpha_s \notin \Phi(x)$ and hence $\Phi(sx) = \{\alpha_s\} \sqcup s\Phi(x)$. Also, $\Phi(w) = s\big(\Phi(sw) \setminus \{\alpha_s\}\big)$. Therefore

$$\Phi(w) \cap \Phi(sx) = s\big(\Phi(sw) \setminus \{\alpha_s\}\big) \cap \big(\{\alpha_s\} \sqcup s\Phi(x)\big)$$
$$= s\big(\Phi(sw) \cap \Phi(x)\big)$$
$$= s(\{\alpha\})$$

Thus the claim follows. We now show the right to left inclusion. For a root $\alpha \in \partial T(w^{-1})$ where there is $x \in W$ with $\Phi(w) \cap \Phi(x) = \{\alpha\}$ and $s \in D_L(x)$ we aim to show that $\Phi(sw) \cap \Phi(sx) = \{s\alpha\}$. It then follows by Theorem 2.5 again that $s\alpha$ is a boundary root of $T(w^{-1}s)$. Since $\ell(sw) > \ell(w)$ we have that $\Phi(sw) = \{\alpha_s\} \sqcup s\Phi(w)$. If $x \in W$ is such that $s \in D_L(x)$ then $\Phi(sx) = s\big(\Phi(x) \setminus \{\alpha_s\}\big)$. Hence we have

$$\Phi(sw) \cap \Phi(sx) = \big(\{\alpha_s\} \sqcup s\Phi(w)\big) \cap s\big(\Phi(x) \setminus \{\alpha_s\}\big)$$

Since $\Phi(w) \cap \Phi(x) = \{\alpha\}$ by the formula for $\Phi(sw) \cap \Phi(sx)$ above it follows that $s\alpha \in \Phi(sw) \cap \Phi(sx)$. We now show that $\Phi(sw) \cap \Phi(sx) = \{s\alpha\}$. If $\beta \in \Phi(sw) \cap \Phi(sx)$ then $\beta \in s\Phi(w) \cap s\Phi(x) = s\big(\Phi(w) \cap \Phi(x)\big) = s(\{\alpha\})$. So $\beta = s\alpha$ as required. $\square$

The next result shows how the set $\partial T(sx^{-1})_{s\beta}$ is obtained from $\partial T(x^{-1})_\beta$ if $\beta \in \partial T(x^{-1})$ and $s\beta \in \partial T(sx^{-1})$.

Corollary 4.3. *Let $x \in W$, $s \in S$ with $s \notin D_L(x)$. Let $y = sx$. If $\beta \in \partial T(x^{-1})$ and $s\beta \in \partial T(y^{-1})$ then*

$$\partial T(y^{-1})_{s\beta} = \{w \in \partial T(x^{-1})_\beta \mid s \in D_L(w)\}$$

Proof. If $v \in \partial T(y^{-1})_{s\beta}$ then $\Phi(y) \cap \Phi(v) = \{s\beta\}$. Since $\alpha_s \in \Phi(y)$ this implies that $\alpha_s \notin \Phi(v)$ and so $\Phi(sv) = \{\alpha_s\} \sqcup s\Phi(v)$. Since $\Phi(x) = s\Phi(y) \setminus \{\alpha_s\}$ we have

$$\Phi(x) \cap \Phi(sv) = \big(s\Phi(y) \setminus \{\alpha_s\}\big) \cap \big(\{\alpha_s\} \sqcup s\Phi(v)\big)$$

Therefore, as in the proof of Lemma 4.4, we have $\Phi(x) \cap \Phi(sv) = \{\beta\}$. The reverse inclusion directly follows from the second half of the proof of Lemma 4.4. $\square$

5. The $\mathscr{T}^0$ partition and join-irreducible elements of Γ

In this section, we study additional properties of the tight gates Γ^0 of $\mathscr{T}$. We show that the partition induced by the tight gates Γ^0 produces the cone type arrangement $\mathscr{T}$. We also show that the set Γ^0 is the set of join-irreducible elements of the poset Γ under the right weak order. In particular, it follows from [13, Theorem 1.4] that the set Γ^0 are the most fundamental elements of the smallest Garside shadow of W.

Lemma 5.1. *Let $x \in W$. Then the set of prefixes of x, $P(x) = \{z \preceq x \mid z \in W\}$ is convex.*

Proof. Clearly, we have $P(x) = \bigcap_{\Phi^+ \setminus \Phi(x)} H_\alpha^+$. Hence the result follows by Lemma 2.1. $\square$

The following result shows that every gate of $\mathscr{T}$ is a join of tight gates.

Theorem 5.1. *Let $x \in \Gamma$. Then*

$$x = \bigvee \partial t(x^{-1})$$

where $\partial t(x^{-1}) = \{z \in W \mid z \preceq x \text{ with } \Phi^R(z) = \{\alpha\} \text{ and } \alpha \in \partial T(x^{-1})\}$. Hence, every gate is a join of tight gates.

Proof. Let $x \in \Gamma$. By Lemma 3.1, for each $\alpha \in \partial T(x^{-1})$, there is $z_\alpha \preceq x$ with $\Phi^R(z_\alpha) = \{\alpha\}$. Since $\alpha \in \partial T(x^{-1})$, by Theorem 2.5 and Proposition 2.4 there exists $w_\alpha \in \Gamma^0$ such that $\Phi(x) \cap \Phi(w_\alpha) = \{\alpha\}$. Therefore, we also have $\Phi(z_\alpha) \cap \Phi(w_\alpha) = \{\alpha\}$. By Theorem 2.3 it follows that $z_\alpha \in \Gamma^0$ and x is an upper bound of

$$\partial t(x^{-1}) := \{z \in W \mid z \preceq x \text{ with } \Phi^R(z) = \{\alpha\} \text{ and } \alpha \in \partial T(x^{-1})\} \subset \Gamma^0$$

Now let $y = \bigvee \partial t(x^{-1})$ and let $\Phi(\partial t(x^{-1})) := \bigcup_{z \in \partial t(x^{-1})} \Phi(z)$. By Proposition 2.3 and Theorem 2.4 we have

$$T(y^{-1}) = \bigcap_{z \in \partial t(x^{-1})} T(z^{-1}) = \bigcap_{\Phi(\partial t(x^{-1}))} H_\beta^+$$

Since $y \preceq x$ and $\partial T(x^{-1}) \subseteq \Phi(\partial t(x^{-1}))$ it follows by Lemma 2.4 and Theorem 2.4 again that

$$T(x^{-1}) \subseteq T(y^{-1}) = \bigcap_{\Phi(\partial t(x^{-1}))} H_\beta^+ \subseteq \bigcap_{\partial T(x^{-1})} H_\beta^+ = T(x^{-1})$$

So $T(x^{-1}) = T(y^{-1})$ and since $x \in \Gamma$, by Theorem 2.2 we must have $y = x$. $\qquad\square$

Corollary 5.1. *For $x \in W$, $\partial t(x^{-1}) \subseteq \phi^1(x)$.*

Proof. Proposition 3.2 shows that for each $\beta \in \Phi^1(x)$ there is a unique minimal length element $j_\beta \in \phi^1(x)$. Since $\partial T(x^{-1}) \subseteq \Phi^1(x)$ the result follows by Theorem 5.1 and Proposition 3.1. $\qquad\square$

Corollary 5.2. *Each cone type $T \in \mathbb{T}$ is expressible as an intersection of cone types of tight gates.*

Proof. By Theorem 5.1 it follows that every gate is a join of a set of tight gates $\partial t(x^{-1})$. Then by Proposition 2.3 we have that $T(x^{-1}) = \bigcap_{z \in \partial t(x^{-1})} T(z^{-1})$. $\qquad\square$

Corollary 5.3. *Define $x \sim y$ if the following holds: $x \in T(w^{-1})$ if and only if $y \in T(w^{-1})$ for all $w \in \Gamma^0$.*
Let $\mathscr{T}^0$ be the partition of W with parts P defined by the relation $x \sim y$. Then $\mathscr{T}^0 = \mathscr{T}$ and hence $\mathscr{T}^0$ is a regular partition.

Proof. By Theorem 2.1 it suffices to show that if x, y are in the same part of $\mathscr{T}^0$ then x, y are in the same part of $\mathscr{T}$. Suppose $x \in T(z^{-1})$ if and only if $y \in T(z^{-1})$ for all $z \in \mathscr{T}^0$. Let $x \in T(g^{-1})$ for $g \in \Gamma \setminus \Gamma^0$. By Theorem 5.1 and Proposition 2.3 we then have $x \in T(w^{-1})$ for all $w \in \partial t(g^{-1})$ where

$$\partial t(g^{-1}) = \{w \in W \mid w \preceq g \text{ with } \Phi^R(w) = \{\alpha\} \text{ and } \alpha \in \partial T(g^{-1})\}$$

Since $\partial t(g^{-1}) \subset \Gamma^0$ this implies $y \in T(w^{-1})$ for all $w \in \partial t(g^{-1})$ and hence by Theorem 5.1 and Proposition 2.3 again this implies $y \in T(g^{-1})$. Similarly, if $x \notin T(g^{-1})$ for some $g \in \Gamma \setminus \Gamma^0$ then by Theorem 5.1, there must be some $w \preceq g$ with $\Phi^R(w) = \{\alpha\}$ and $\alpha \in \partial T(g^{-1}) \cap \partial T(w^{-1})$ such that $x \notin T(w^{-1})$. Therefore, also $y \notin T(w^{-1})$. $\qquad\square$

Remark 5.1. A consequence of Corollary 5.3 is that to determine whether two elements x, y have the same cone type (or equivalently, whether x^{-1} and y^{-1} live in the same part of the cone type partition), one only needs to compute the set of tight gates Γ^0 and then check whether $x \in T(w^{-1})$ if and only if $y \in T(w^{-1})$ for all $w \in \Gamma^0$. Our computations using Sagemath ([4]) in Section 6.1 show that often, the set Γ^0 is much smaller than Γ. We illustrate the $\mathscr{T}^0$ arrangements for the rank 3 Coxeter groups of affine type in Figure 9.

Recall that for a poset P an element $x \in P$ is *join-irreducible* if x is not the minimal element in P and if $x = a \vee b$ then either $x = a$ or $x = b$. Equivalently, if $X \subseteq P$ and $x = \bigvee X$ then $x \in X$.

Theorem 5.2. *The set*

$$\Gamma^0 = \{x \in \Gamma \mid |\Phi^R(x)| = 1\}$$

is the set of join-irreducible elements of $(\Gamma, \preceq)$.

Proof. Let $x \in \Gamma^0$ and $\Phi^R(x) = \{\beta\}$. If x is not join-irreducible, then $x = \bigvee A$ for some $A \subseteq \Gamma$ with $x \notin A$. Since $\beta \in \Phi^1(x)$ by Corollary 2.1 $\beta \in \Phi^1(y)$ for some $y \in A$ and $y \prec x$. Therefore, there must be $\alpha \neq \beta \in \Phi^R(x)$, a contradiction.

Now let $x \in \Gamma$ be join-irreducible in $(\Gamma, \preceq)$ and suppose $|\Phi^R(x)| > 1$. Let $\Phi^R(x) = \{\alpha_1, \ldots \alpha_k\}$ and $x_i := s_{\alpha_i} x$. Let $\pi(x_i) \in \Gamma$ be the minimal element of the cone type part containing x_i, that is the part $Q(T(x_i^{-1}))$. We claim that $x = \bigvee \pi(x_i)$. Let $y = \bigvee \pi(x_i)$. Note that since x is an upper bound for each $\pi(x_i)$, this implies $y \preceq x$. Then since $x = \bigvee x_i$ and $T(x_i^{-1}) = T(\pi(x_i)^{-1})$ for each i, by Proposition 2.3 we have

$$T(x^{-1}) = \bigcap_{x_i} T(x_i^{-1}) = \bigcap_{x_i} T(\pi(x_i)^{-1}) = T(y^{-1})$$
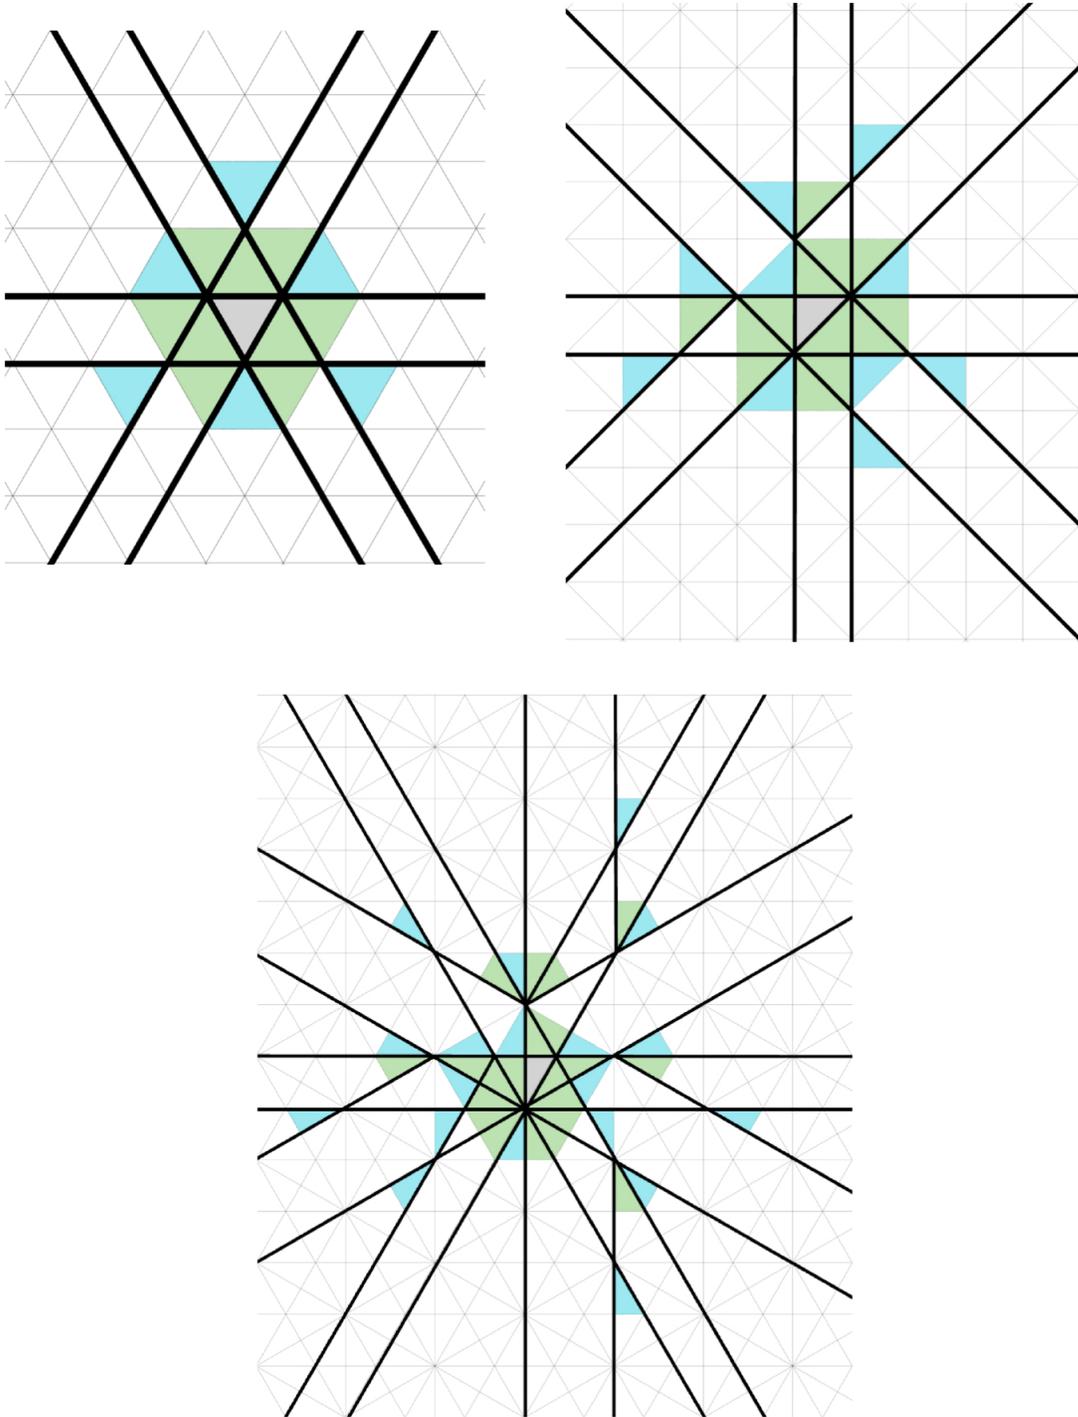
Since $x \in \Gamma$ we must have $x = y$. $\qquad\square$

FIGURE 9. The partition $\mathscr{T}^0 = \mathscr{T}$ for the rank 3 Coxeter groups of affine type ($\widetilde{A}_2$, $\widetilde{B}_2$ and $\widetilde{G}_2$). The green alcoves represent the tight gates and the cyan alcoves are the gates which are non-trivial joins of tight gates.

6. Computing the sets $\mathcal{S}$ and Γ^0

The super elementary roots $\mathcal{S}$ and tight gates Γ^0 are intimately connected by definition. In this section, we utilise the results of Section 5 to give a direct and efficient method for computing Γ^0 without computing Γ first; effectively obtaining the minimal data to determine whether two elements have the same cone type.

Let us first give a bit more background to set up the algorithm and illustrate the connection between our work and some important recent developments in Coxeter group theory.

We first recall the definition of the set of *elementary* roots $\mathcal{E}$ in Coxeter groups. For roots $\alpha, \beta \in \Phi^+$ we say β *dominates* α if whenever $w\beta < 0$ implies $w\alpha < 0$ for all $w \in W$. A root is *elementary* if it does not dominate any root besides itself. Introduced by Brink and Howlett in [3], the set of elementary roots $\mathcal{E}$ were shown to be finite for any finitely generated Coxeter group [3, Theorem 2.8] and play a fundamental role in their finite state automaton recognising the language of reduced words of (W, S). They have been extensively studied (see for example [2, Chapter 4], [10], [6]) and are computed in a straightforward manner, as described in the aforementioned work. For our purposes, it is enough to know that the set $\mathcal{E}$ can be easily computed.

The *Shi arrangement* or *Shi partition* $\mathcal{S}$ is the partition of W induced by the elementary walls, and like the cone type partition $\mathcal{T}$, each part corresponds to a state of a finite state automaton recognising the language of reduced words of (W, S) (see the example of $\widetilde{A}_2$ in Figure 9, in this case $\mathcal{S} = \mathcal{T}$). This is in fact one way to visualise Brink and Howlett's automaton. Also like the cone type partition $\mathcal{T}$, each part P of $\mathcal{S}$ contains an element of minimal length, which is a gate of P (proven in [6, Theorem 1.1] and [13, Theorem 1.1]). The unique minimal length representatives of these parts are known as the *low*-elements L. Formally, we have the following definition of low-elements.

Definition 6.1. [10, Definition 3.24] An element $w \in W$ is *low* if $\Phi(w) = cone_\Phi(A)$ for some $A \subseteq \mathcal{E}$. Equivalently, w is *low* if $\Phi^1(w) \subseteq \mathcal{E}$.

For the purposes of Algorithm 6.1 we pay particular attention to the set of *tight* low elements, that is the set
$$L^0 = \{w \in W \mid |\Phi^R(w)| = 1 \text{ and } \Phi^R(w) \subseteq \mathcal{E}\}$$

Crucially for Algorithm 6.1, by [6, Theorem 4.8 (3)] we have $L^0 \subseteq L$ and since every super-elementary root is elementary ([12, Proposition 7.2]) it follows by Theorem 2.3 that $\Gamma^0 \subseteq L^0 \subseteq L$. Note also that L is finite (since $\mathcal{E}$ is finite).

In general, the set $\mathcal{S}$ can be determined in finite time by considering each root $\beta \in \mathcal{E}$ and checking whether there exists a pair of gates or low-elements $\{x, y\}$ with $\Phi(x) \cap \Phi(y) = \{\beta\}$. An application of Theorem 4.1 and Proposition 4.1 gives an alternative method which concurrently computes the set $\mathcal{S}$ and Γ^0 without first having to compute the entire set Γ or L. It only requires that the elementary roots $\mathcal{E}$ is computed first (which is quite efficient, see [2, Chapter 4.7]; we note that elementary roots are also called *small* roots in the literature). We remind the reader that Δ is the set of simple roots in the following algorithm.

Algorithm 6.1 Compute tight gates and super-elementary roots

1: Let $L' = S$, $\mathcal{S}' = \Delta$, $\Gamma' = \emptyset$
2: Initialize $L'_i = \{x \in L' \mid \ell(x) = i\}$ with $i = 1$
3: **while** $L'_i \neq \emptyset$ **do**
4:     **for** $x \in L'_i$ **do**
5:         **for** $s \notin D_L(x)$ **do**
6:             Set $y = s \cdot x$
7:             **if** $|\Phi^R(y)| = 1$ and $\{\beta\} = \Phi^R(y) \in \mathcal{E}$ **then**
8:                 Add y to L'
9:                 **if** there is $z \neq y \in L'$ with $\Phi^R(z) = \{\beta\}$ and $\Phi(z) \cap \Phi(y) = \{\beta\}$ **then**
10:                     Add β to $\mathcal{S}'$
11:                     Add y, z to Γ'
12:                 **end if**
13:             **end if**
14:         **end for**
15:     **end for**
16:     Set $i = i + 1$
17: **end while**

            The algorithm terminates in finite time and upon termination $\Gamma' = \Gamma^0$ and $\mathcal{S}' = \mathcal{S}$

Proof. By Proposition 4.1, the set of tight gates Γ^0 is closed under suffix: if $x \in \Gamma^0$ and $s \in D_L(x)$, then $sx \in \Gamma^0$ and $\ell(sx) = \ell(x) - 1$. Since Γ^0 is finite, it follows that every element $g \in \Gamma^0$ admits a length–decreasing chain contained in Γ^0

$$g = g_k \xrightarrow{s_k} g_{k-1} \xrightarrow{s_{k-1}} \cdots \xrightarrow{s_2} g_1,$$

where $g_1 \in S$, $\ell(g_i) = i$, and $g_{i-1} = s_i g_i$ with $s_i \in D_L(g_i)$. Equivalently, every tight gate can be obtained from a simple reflection by successively left-multiplying by simple reflections $s \notin D_L(x)$, increasing length by one at each step. The algorithm hence performs a breadth–first search with respect to length, starting from the simple reflections and iterating over all possible left multiplications $x \mapsto sx$ with $s \notin D_L(x)$. Consequently, every element of Γ^0 is eventually encountered and added to L'. As discussed on the previous page, since $\mathcal{S} \subseteq \mathcal{E}$, by the characterisation of Γ given in Theorem 2.3 we have

$$\Gamma^0 \subseteq L' \subseteq L^0 = \{y \in W \mid |\Phi^R(y)| = 1 \text{ and } \Phi^R(y) \subseteq \mathcal{E}\} \subseteq L$$

Since L is finite, the algorithm terminates in finite time. Finally, once an element $y \in \Gamma^0$ is encountered, Theorem 4.1 guarantees the existence of a unique element $z \in \Gamma^0$ with

$$\Phi(y) \cap \Phi(z) = \{\beta\} = \Phi^R(y) = \Phi^R(z)$$

so the algorithm records β and adds y and z to Γ'. Thus, upon termination, $\Gamma' = \Gamma^0$ and $\mathcal{S}' = \mathcal{S}$. $\square$

We also note the following consequence of Theorem 4.1 which shows that the number of non-simple tight gates is always even.

Corollary 6.1. *Let $K := |\Gamma^0 \setminus S|$. Then K is even, and $|\mathcal{S} \setminus \Delta| \leq K/2$.*

Proof. Theorem 4.1 shows that tight gates $g \in \Gamma^0 \setminus S$ are naturally paired according to their final roots $\Phi^R(g)$ in the following way: For each $\beta \in \mathcal{S} \setminus \Delta$, there is some $x \in \Gamma^0$ with $\beta \in \partial T(x^{-1}) \cap \Phi^R(x)$. Hence by Theorem 4.1 there is a unique element $y \in \Gamma^0 \setminus S$ such that $\Phi(x) \cap \Phi(y) = \{\beta\}$ and $\Phi^R(y) = \{\beta\}$. By symmetry, x is also the unique element of Γ^0 with $\Phi(x) \cap \Phi(y) = \{\beta\}$. $\square$

6.1. **Data for select Coxeter groups.** We utilise Algorithm 6.1 to compute the size of $\mathcal{S}$ and Γ^0 for some Coxeter systems in low rank. We note that as the size of Γ increases, in general, the ratio $|\Gamma^0|/|\Gamma|$ tends to decrease, making the computation of Γ^0 generally much more efficient to determine whether two elements have the same cone type.

W	$|\mathcal{E}|$	$|\mathcal{S}|$	$|L|$	$|L^0|$	$|\Gamma|$	$|\Gamma^0|$
$\widetilde{A_2}$	6	6	16	9	16	9
$\widetilde{B_2}$	8	8	25	14	24	13
$\widetilde{G_2}$	12	12	49	26	41	21
$\widetilde{A_3}$	12	12	125	28	125	28
$\widetilde{B_3}$	18	18	343	66	315	58
$\widetilde{C_3}$	18	18	343	66	317	58
$\widetilde{A_4}$	20	20	1296	75	1296	75
$\widetilde{B_4}$	32	32	6561	270	5789	227
$\widetilde{C_4}$	32	32	6561	270	5860	227
$\widetilde{D_4}$	24	24	2401	140	2400	139
$\widetilde{F_4}$	48	48	28561	1054	22428	715
$\widetilde{A_5}$	30	30	16807	186	16807	186
$\widetilde{B_5}$	50	50	161051	1030	137147	836
$\widetilde{C_5}$	50	50	161051	1030	139457	836
$\widetilde{D_5}$	40	40	59049	608	58965	596

W	$|\mathcal{E}|$	$|\mathcal{S}|$	$|L|$	$|L^0|$	$|\Gamma|$	$|\Gamma^0|$
$X_4(4)$	25	25	438	112	392	98
$X_4(5)$	32	32	516	158	462	138
Y_4	32	32	687	166	578	150
Z_4	30	30	513	142	473	132
$X_5(3)$	114	114	101412	5767	52542	4071
$X_5(4)$	83	83	25708	3128	22886	2871
$X_5(5)$	135	135	42064	6014	37956	5523
Z_5	120	120	41385	5476	39138	5391

FIGURE 10. Data for low rank affine and compact hyperbolic Coxeter groups.

References

[1] Peter Abramenko and Kenneth Brown. *Buildings: Theory and Applications*, volume 248 of *Graduate Texts in Mathematics*. Springer, 2008.

[2] Anders Björner and Francesco Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.

[3] Brigette Brink and Robert Howlett. A finiteness property and an automatic structure for coxeter groups. *Math. Ann.*, 296:179–190, 1993.

[4] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.4)*. https://www.sagemath.org, 2024.

[5] Matthew Dyer. On the weak order of coxeter groups. *Canad. J. Math.*, 71(2):299–336, 2019.

[6] Matthew Dyer, Susanna Fishel, Christophe Hohlweg, and Alice Mark. Shi arrangements and low elements in coxeter groups. *Proc. Lond. Math. Soc. (3)*, 129(2):e12624, 2024.

[7] George Grätzer and Friedrich Wehrung, editors. *Lattice Theory: Special Topics and Applications, Volume 2*. Encyclopaedia of Mathematical Sciences. Springer International Publishing, Cham, Switzerland, 2016.

[8] Christophe Hohlweg and Jean-Philippe Labbé. On inversion sets and the weak order in coxeter groups. *European J. Combin.*, 55:1 – 19, 2016.

[9] Christophe Hohlweg, Philippe Nadeau, and Nathan Williams. Automata, reduced words and garside shadows in coxeter groups. *J. Algebra*, 457:431–456, 2016.

[10] Dyer M. Hohlweg, C. Small roots, low elements, and the weak order in coxeter groups. *Advances in Math.*, 301:739–784, 2016.

[11] James E. Humphreys. *Reflection groups and Coxeter groups*, volume 29 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990.

[12] James Parkinson and Yeeka Yau. Cone types, automata, and regular partitions in coxeter groups. *Adv. Math.*, 398:108146, 2022.

[13] Piotr Przytyck and Yeeka Yau. A pair of Garside shadows. *Algebr. Comb.*, 7(6):1879–1885, 2024.

[14] Mark Ronan. *Lectures on buildings*. University of Chicago Press, Chicago, IL. Updated and revised.

[15] Fabricio Dos Santos. Garside shadows and biautomatic structures in coxeter groups, (Preprint) 2025. https://arxiv.org/abs/2505.21718.

[16] Yeeka Yau. *Automatic Structures for Coxeter Groups*. PhD Thesis, University of Sydney, 2021.

School of Mathematics and Statistics, The University of Sydney
Email address: yeeka.yau@sydney.edu.au