Minuscule Elements of Weyl Groups

John R. Stembridge

Department of Mathematics
University of Michigan
Ann Arbor, Michigan 48109–1109

August 23, 1999
Minor Revisions May 16, 2000

Contents

Partially supported by NSF Grant DMS–9700787.

0. Introduction

In some unpublished work dating back to the 1980's, Dale Peterson has defined and studied what he calls "λ-minuscule" elements of (symmetrizable Kac-Moody) Weyl groups. (The precise definition is given in Section 2 below.) The terminology presumably derives from the fact that if λ is the highest weight of a minuscule representation of a simple Lie algebra, then for every μ in the orbit of λ, the shortest element of the Weyl group such that $w\lambda = \mu$ is λ-minuscule.

Associated with any λ-minuscule element w is a partially ordered set (the "heap") whose vertices are labeled by nodes of the Dynkin diagram; the linear extensions of the heap encode the reduced expressions for w. In type A, the heaps are Young diagrams, and the reduced expressions are in one-to-one correspondence with standard Young tableaux. Similarly, plane partitions can be viewed as order-preserving assignments of integers to the vertices of a Young diagram, so there is an analogous notion of "λ-minuscule partition."

Starting with [**P1–3**], Proctor has begun the development of a combinatorial theory for simply-laced λ-minuscule elements, including classification theorems and a generalization of *jeu de taquin*. In a forthcoming paper, Peterson and Proctor give an explicit hooklength formula for the generating function for λ-minuscule partitions, generalizing well-known results for both shifted and unshifted plane partitions.

This paper has two main objectives.

First, it has been clear from the beginning of Proctor's work in [**P1**] that λ-minuscule elements are "fully commutative" in the sense of [**St1**], or "commutative" in the sense of [**F**], a property characterized by the non-existence of certain subwords in the set of reduced expressions. (For the definition, see Section 2.) Here we clarify more directly the exact nature of the relationship, providing reduced-word characterizations of minuscule elements, as well as order-theoretic characterizations of their heaps.

We should explain that the "wave" posets of [**P1**] are the simply-laced cases of the heaps we consider here, although they are defined in a completely different way. Similarly, the "colored d-complete" posets of [**P1–2**] provide an order-theoretic characterization of wave posets that, although different in appearance, is equivalent to the simply-laced case of the heap characterization we provide in Section 3.

Our second objective is to extend Proctor's classification of (dominant) λ-minuscule elements (or equivalently, their heaps) from the simply-laced case to any symmetrizable Kac-Moody Weyl group. There is a natural way to decompose heaps of dominant minuscule elements into irreducible components. In the simply-laced case, Proctor has shown that the irreducible cases can be grouped into 15 families, 14 of which are infinite [**P2**]. Here we show that in the multiply-laced cases, there are two more infinite families (see

Theorem 4.2). It is noteworthy that the members of these two families are isomorphic, as *unlabeled* posets, to simply-laced heaps. Thus every abstract poset that occurs as a dominant λ-minuscule heap can be obtained from a simply-laced (Kac-Moody) Weyl group, and is therefore "d-complete" in the sense of Proctor. It is conceivable that one could prove that every multiply-laced dominant minuscule heap is d-complete by some direct argument, bypassing the need for a classification. However, there may be applications of the theory in which the labeling of the posets plays a role, and hence increased significance for the multiply-laced cases.

In the final section of the paper, we show (Theorem 5.5) that the heap of any (dominant) λ-minuscule element w can also be obtained by restricting the standard partial ordering of the positive co-roots to those co-roots that are "inverted" by w (i.e., $\alpha^\vee > 0$ and $w\alpha^\vee < 0$). A key ingredient of the proof is a reduced-word and heap characterization of the elements w having no triple of inversions of the form $\alpha^\vee, \beta^\vee, \alpha^\vee + \beta^\vee$ (Theorem 5.3), generalizing the simply-laced result of [**FS**].

In a sequel to this paper [**St2**], we will provide an application and extension of this theory to the combinatorics of reduced expressions for reflections in finite Weyl groups.

1. Preliminaries

We begin by choosing a Cartan matrix $A = [a_{ij}]_{1 \leqslant i,j \leqslant n}$ for a symmetrizable Kac-Moody Lie algebra [**K**]. Thus A is an integer matrix satisfying

(1) $a_{ij} \leqslant 0$ for $i \neq j$; $a_{ii} = 2$,

(2) $a_{ij} = 0$ if and only if $a_{ji} = 0$,

and (by virtue of being symmetrizable), there exists a symmetric bilinear form $\langle\ ,\ \rangle$ on $\mathbf{R}^n$ and a basis $\alpha_1, \dots, \alpha_n$ of $\mathbf{R}^n$ such that

(3) $\langle \alpha_i, \alpha_i \rangle > 0$ and $a_{ij} = \langle \alpha_i, \alpha_j^\vee \rangle$ $(1 \leqslant i, j \leqslant n)$,

where $\alpha^\vee := 2\alpha / \langle \alpha, \alpha \rangle$.

It may happen that the bilinear form is degenerate, in which case we may embed $\mathbf{R}^n$ in some larger space $\mathbf{R}^N$ with the bilinear form extended in some non-degenerate way.

The basis vectors α_i form the simple roots of a (crystallographic) root system $\Phi \subset \mathbf{R}^n$, and the corresponding simple reflections; viz.,

$$s_i : \lambda \mapsto \lambda - \langle \lambda, \alpha_i^\vee \rangle \alpha_i,$$

generate a Coxeter group W (the Weyl group) satisfying the relations $s_i^2 = (s_i s_j)^{m_{ij}} = 1$, where (for $i \neq j$) $m_{ij} = 2, 3, 4, 6$ or ∞, according to whether $a_{ij} a_{ji} = 0, 1, 2, 3$ or $\geqslant 4$.

The Cartan matrix is represented faithfully by the Dynkin diagram, a graph with vertex set $[n] := \{1, \ldots, n\}$, edges between pairs of vertices i, j with $a_{ij} < 0$, and various decorations on the edges to record the values of a_{ij} and a_{ji}. The main conventions that concern us here are that a simple bond is used when $a_{ij} = a_{ji} = -1$, and an oriented double bond directed from i to j is used when $a_{ij} = -2$ and $a_{ji} = -1$. If every edge is a simple bond, then A, Φ, and W are said to be *simply-laced*; otherwise, they are *multiply-laced*.

Suppose that nodes i and j are adjacent in the Dynkin diagram. We say that i, α_i, and s_i are *short* relative to j, α_j, and s_j if $a_{ij} = -1$. If Φ is finite, this is equivalent to having $\langle \alpha_i, \alpha_i \rangle \leqslant \langle \alpha_j, \alpha_j \rangle$; in particular, either i is short relative to j, or vice-versa. If Φ is simply-laced, then i is short relative to j, *and* vice-versa.

Recall that every root $\alpha \in \Phi$ is either positive or negative, according to whether the coordinates of α with respect to the simple roots α_i are nonnegative or nonpositive. We let Φ^+ and Φ^- denote the sets of positive and negative roots.

The co-root system $\Phi^\vee := \{\alpha^\vee : \alpha \in \Phi\}$ is also crystallographic, with simple roots $\alpha_1^\vee, \ldots, \alpha_n^\vee$, Cartan matrix $[a_{ji}]$ (the transpose of A), and the same Weyl group.

The weight lattice may be defined as

$$\Lambda = \{\lambda \in \mathbf{R}^N : \langle \lambda, \alpha_i^\vee \rangle \in \mathbf{Z},\ 1 \leqslant i \leqslant n\}.$$

The members of Λ are called *(integral) weights*, and if $\langle \lambda, \alpha_i^\vee \rangle \geqslant 0$ for all i, then λ is said to be *dominant*. To be careful, we should note that if $N > n$, then Λ fails to be discrete in $\mathbf{R}^N$ and hence is not strictly a lattice. To remedy this, one should view Λ as a lattice in $\mathbf{R}^N/Z$, where $Z = \{\delta \in \mathbf{R}^N : \langle \delta, \alpha_i^\vee \rangle = 0,\ 1 \leqslant i \leqslant n\}$.

There is a commonly used partial ordering of Λ defined by

$$\lambda > \mu \quad \text{if} \quad \lambda - \mu \in \mathbf{N}\alpha_1 + \cdots + \mathbf{N}\alpha_n,$$

where $\mathbf{N}$ denotes the nonnegative integers. We call this the *standard ordering* of Λ. Here we will be concerned primarily with the restriction of the standard ordering to the root system Φ, and the analogous ordering of $\Phi^\vee$.

2. Word Characterizations

Let $\lambda \in \Lambda$ be an integral weight. Following Peterson, we define $w \in W$ to be λ-*minuscule* if there is a reduced expression $w = s_{i_1} \cdots s_{i_l}$ such that

$$\langle s_{i_{k+1}} s_{i_{k+2}} \cdots s_{i_l} \lambda\, ,\, \alpha_{i_k}^\vee \rangle = 1 \quad (1 \leqslant k \leqslant l), \tag{2.1}$$

or equivalently,

$$s_{i_k} s_{i_{k+1}} \cdots s_{i_l} \lambda = \lambda - \alpha_{i_k} - \alpha_{i_{k+1}} - \cdots - \alpha_{i_l} \quad (1 \leqslant k \leqslant l).$$

We say that w is *minuscule* if it is λ-minuscule for some λ; similarly, w is *dominant minuscule* if it is λ-minuscule for some dominant λ.

In [**St1**], we defined $w \in W$ to be *fully commutative* if for every pair of non-commuting generators s_i and s_j, there is no reduced expression for w containing a subword of length m of the form $s_is_js_is_j\cdots$, where m denotes the order of s_is_j in W.

Proposition 2.1. *If w is λ-minuscule, then every reduced expression for w satisfies (2.1). Furthermore, w is fully commutative, and there is no reduced expression for w containing the subword $s_js_is_j$, unless $\langle\alpha_i,\alpha_j^\vee\rangle = -2$ and s_is_j has order $\geqslant 4$.*

Proof. Let $w = s_{i_1}\cdots s_{i_l}$ be a reduced expression satisfying (2.1), and suppose that it contains a subword of length m of the form $\cdots s_is_js_is_j$, where m denotes the order of s_is_j. Setting $\mu = s_{i_{k+1}}\cdots s_{i_l}\lambda$ (where k denotes the last position occupied by the subword), it follows that $\langle\mu,\alpha_j^\vee\rangle = 1$, $s_j\mu = \mu-\alpha_j$, and $\langle\mu-\alpha_j,\alpha_i^\vee\rangle = 1$.

If $m=2$ (i.e., s_i and s_j commute), then $\langle\alpha_j,\alpha_i\rangle = 0$, so $\langle\mu,\alpha_i^\vee\rangle = 1$ and $s_i\mu = \mu-\alpha_i$. It follows that if we replace the subword s_is_j with s_js_i, then the new reduced expression so obtained also satisfies (2.1).

If $m \geqslant 3$, then we also have $\langle\mu-\alpha_i-\alpha_j,\alpha_j^\vee\rangle = \langle s_is_j\mu,\alpha_j^\vee\rangle = 1$. Hence

$$\langle\alpha_i+\alpha_j,\alpha_j^\vee\rangle = \langle\mu,\alpha_j^\vee\rangle - \langle\mu-\alpha_i-\alpha_j,\alpha_j^\vee\rangle = 0,$$

and therefore $\langle\alpha_i,\alpha_j^\vee\rangle = -2$. This eliminates $m=3$, since in that case, α_i and α_j generate a root system of type A_2, whence $\langle\alpha_i,\alpha_j^\vee\rangle = -1$. However if $m \geqslant 4$, then $s_is_js_i$ also occurs as a subword, so the same reasoning implies $\langle\alpha_j,\alpha_i^\vee\rangle = -2$. Thus α_i and α_j generate an (infinite) affine root system of type $A_1^{(1)}$ and s_is_j does not have finite order, a contradiction.

The above argument shows that there are no opportunities to apply braid moves (i.e., replacing $\cdots s_is_js_is_j$ with $\cdots s_js_is_js_i$) except when $m=2$ and the two generators commute. However, one knows that any reduced expression can be obtained from any other by a sequence of braid moves (e.g., [**B**, §IV.1.5]), and we have seen that (2.1) is preserved by commutation relations, so (2.1) holds for every reduced expression. □

Remark 2.2. (a) The above result shows that $s_is_js_is_j$ cannot occur in a reduced expression for a minuscule element unless s_is_j has infinite order.

(b) The fact that dominant minuscule elements are fully commutative has been noted previously by Proctor in the simply-laced case [**P1**]. Indeed, Proctor shows that if w is dominant minuscule, then the subinterval of the weak ordering of W from 1 to w is a distributive lattice. On the other hand, from [**St1**], one knows that in any Coxeter group, the weak ordering from 1 to w is a distributive lattice if and only if w is fully commutative.

Proposition 2.3. *If $w = s_{i_1} \cdots s_{i_l} \in W$ is a reduced expression, then w is minuscule if and only if between every pair of occurrences of a generator s_i (with no other occurrences of s_i in between), there are exactly*

(i) *two terms that do not commute with s_i, and both are short relative to s_i, or*
(ii) *one term that does not commute with s_i, say s_j, and $\langle \alpha_j, \alpha_i^\vee \rangle = -2$.*

Proof. Define $\gamma_k = \alpha_{i_k} + \cdots + \alpha_{i_l}$, so that w is λ-minuscule if and only if

$$\langle \lambda - \gamma_{k+1}, \alpha_{i_k}^\vee \rangle = 1 \quad (1 \leqslant k \leqslant l). \tag{2.2}$$

Given that w is λ-minuscule, it follows that if s_i occurs in positions p and q of the reduced expression (where $p < q$), then $\langle \lambda - \gamma_{p+1}, \alpha_i^\vee \rangle = \langle \lambda - \gamma_{q+1}, \alpha_i^\vee \rangle = 1$, so we have

$$0 = \langle \gamma_{p+1} - \gamma_{q+1}, \alpha_i^\vee \rangle = \langle \alpha_{i_{p+1}}, \alpha_i^\vee \rangle + \cdots + \langle \alpha_{i_{q-1}}, \alpha_i^\vee \rangle + 2.$$

Since $a_{ji} = \langle \alpha_j, \alpha_i^\vee \rangle \leqslant 0$ for $j \neq i$, it follows that either two -1's or one -2 appear in this sum of Cartan integers, as claimed.

Conversely, for each generator s_j appearing in the reduced expression, choose a position $p = p(j)$ such that $i_p = j$. Since $\langle\ ,\ \rangle$ is non-degenerate on R^N, there is an integral weight λ such that

$$\langle \lambda, \alpha_j^\vee \rangle = 1 + \langle \gamma_{p(j)+1}, \alpha_j^\vee \rangle \tag{2.3}$$

for all such j. Conditions (i) and (ii) imply that $\langle \gamma_{p(i)+1}, \alpha_i^\vee \rangle = \langle \gamma_{q+1}, \alpha_i^\vee \rangle$ for all indices q such that $i_q = i$, so (2.2) is satisfied and w is λ-minuscule. $\square$

Remark 2.4. One can see from the above argument that if w is λ-minuscule, then λ is essentially unique. More precisely, the values $\langle \lambda, \alpha_i^\vee \rangle$ are uniquely determined for all i such that s_i appears in a reduced expression for w.

Proposition 2.5. *If $w = s_{i_1} \cdots s_{i_l} \in W$ is a reduced expression, then w is dominant minuscule if and only if the conditions of Proposition 2.3 are satisfied, and the last occurrence of each generator s_i is followed by at most one generator that does not commute with s_i, and this generator is short relative to s_i.*

Proof. Continue the notation from the proof of Proposition 2.3. If w is λ-minuscule and λ is dominant, then (2.2) implies

$$1 + \langle \alpha_{i_{p+1}}, \alpha_i^\vee \rangle + \cdots + \langle \alpha_{i_l}, \alpha_i^\vee \rangle = 1 + \langle \gamma_{p+1}, \alpha_i^\vee \rangle = \langle \lambda, \alpha_i^\vee \rangle \geqslant 0$$

for any p such that $i_p = i$. In particular, if p is the index of the last occurrence of i, then the Cartan integers appearing in the above sum are $\leqslant 0$, so at most one of them is -1, and the remainder are 0. Conversely, if the stated conditions are satisfied, then the same calculation shows that $\langle \gamma_{p+1}, \alpha_i^\vee \rangle \geqslant -1$ if p is the position of the last occurrence of s_i. It follows that there is a dominant weight λ satisfying (2.3). □

The simply-laced case of the following is analogous to Lemma 6.8 of [**P1**].

Corollary 2.6. *If w is dominant minuscule, then the subdiagram corresponding to the simple reflections that appear in a reduced expression for w must be acyclic.*

Proof. Choose a reduced expression for w, and assume towards a contradiction that some subset of the generators that appear in it indexes a circuit in the Dynkin diagram. Among the last occurrences of each generator in this subset, the leftmost one, say s_i, is followed by at least two generators that do not commute with s_i. In that case, Proposition 2.5 implies that w cannot be dominant minuscule. □

Remark 2.7. (a) Suppose that w is dominant minuscule, and for simplicity, assume that every generator appears in a reduced expression for w. (If not, we may pass to a suitable parabolic subgroup and root subsystem.) If s_i and s_j are a pair of noncommuting generators, then the last occurrence of s_j must be followed by s_i or vice-versa, whence Proposition 2.5 implies that s_i must be short relative to s_j, or vice-versa. Thus for the study of dominant minuscule elements, we may restrict our attention to Cartan matrices satisfying $a_{ij} = -1$ or $a_{ji} = -1$ whenever $a_{ij} < 0$.

(b) If a reduced expression for a dominant minuscule element is cut into two subwords, it is clear from the definition that both subwords are minuscule, and the right subword must be dominant. On the other hand, even with an acyclic diagram, there may exist minuscule elements that cannot be obtained as initial segments of dominant minuscule reduced expressions. For example, in D_4 (with node 3 having degree 3), the minuscule element $w = s_3 s_1 s_2 s_4$ has no dominant completion.

3. Heaps of Minuscule Elements

Since a minuscule element $w \in W$ is fully commutative, there is a partial ordering naturally associated with w whose elements are labeled by nodes of the Dynkin diagram, the so-called *heap* [**St1**]. More precisely, the heap of a (not necessarily reduced) W-expression $s_{i_1} \cdots s_{i_l}$ is the triple $P = ([\,l\,], \preccurlyeq, \mathbf{i})$, where $\preccurlyeq$ is the partial ordering of $[\,l\,]$ obtained by taking the transitive closure of the relations

$$p \prec q \quad \text{whenever} \quad p < q \text{ and either } s_{i_p} s_{i_q} \neq s_{i_q} s_{i_p} \text{ or } i_p = i_q,$$

and $\mathbf{i} = (i_1, \ldots, i_l)$ is the labeling that records the fact that vertex p has label i_p.

Any word that can be obtained from $s_{i_1} \cdots s_{i_l}$ by transposing commuting generators has a heap that is isomorphic to P as a labeled poset. Conversely, the label sequence of any linear extension of P to a total order corresponds to a word in the same commuting equivalence class (Proposition 2.2 of [**St1**]).

Now suppose that $s_{i_1} \cdots s_{i_l}$ is a reduced expression for some fully commutative element w. Since any other reduced expression for w can be obtained via a sequence of commutation relations, it follows that the heaps of all reduced expressions for a fully commutative element are isomorphic. Thus we may refer to "the heap of w" without ambiguity. Note also that in this situation, the linear extensions of the heap are in one-to-one correspondence with the reduced expressions for w.

The labeled posets in Figures 1, 2, and 4–6 are examples of heaps of fully commutative elements. In some of these examples, certain covering relations are displayed as oriented double bonds as a reminder of the corresponding entries of the Cartan matrix.

Now consider an arbitrary poset P whose vertices are labeled by nodes of the Dynkin diagram of A. The following is a list of properties that P may or may not possess.

(H1) All covering pairs in P have labels that are equal or adjacent in the Dynkin diagram, and incomparable pairs have distinct, non-adjacent labels.

(H2) Every open subinterval of P between two elements labeled i (with no other elements labeled i in between), has either (i) exactly two elements whose labels are adjacent to i, and both labels are short relative to i, or (ii) exactly one element, and the label of this element, say j, satisfies $a_{ji} = -2$.

(H3) An element that is maximal in P among all elements labeled i is covered by at most one element, and this element is maximal among all elements of some label that is short relative to i.

(H4) The labels that occur in P index an acyclic subset of the Dynkin diagram.

Proposition 3.1. *A labeled poset P is isomorphic to the heap of a ...*

(a) *W-expression (not necessarily reduced) if and only if (H1) holds.*

(b) *minuscule element of W if and only if (H1) and (H2) hold.*

(c) *dominant minuscule element of W if and only if (H1)–(H4) hold.*

Proof. (a) The fact that heaps of W-expressions satisfy (H1) is clear from the definition. For the converse, proceed by induction on l, the number of elements in P. Assuming $l > 1$, we may choose a maximal element q of P. Since $P - \{q\}$ satisfies (H1), it follows by induction that $P - \{q\}$ is isomorphic to the heap of some W-expression. Consider the heap Q obtained by appending s_i to this expression, where i denotes the label of q in P. Since (H1) implies that elements with the same label are totally ordered, it follows that the (labeled) isomorphism between $P-\{q\}$ and $Q-\{l\}$ is unique; the chains of $P-\{q\}$ and $Q-\{l\}$ labeled j must correspond, for each j.

Each element covered by l in Q must have a label adjacent (or equal) to i, and hence corresponds to an element below q in P, by (H1). Furthermore, each element incomparable to l in Q must correspond to an element p that is incomparable to q in P. Otherwise, there would exist a maximal chain $p \prec \cdots \prec q' \prec q$ in P. In particular, since $q' \prec q$ is a covering relation, (H1) implies that the label of q' must be adjacent (or equal) to i. Since $P-\{q\} \cong Q-\{l\}$, it follows that the corresponding elements also form a chain in Q, a contradiction. Thus the isomorphism can be extended to P and Q.

(b) First suppose that P is the heap of some minuscule $w \in W$, and consider a subinterval of P between two elements labeled i. By choosing a suitable linear extension of P, one may obtain a reduced expression for w in which the terms corresponding to the subinterval appear as a subword. Applying Proposition 2.3, we obtain that the subinterval must satisfy (H2). In particular, the (open) subinterval has only one element when subcase (ii) applies, since a 3-element chain is the only (finite, bounded) poset that has only one element that covers or is covered by one or both endpoints.

Conversely, assume (H1) and (H2). From (a), it follows that P is the heap of some W-expression, say $s_{i_1} \cdots s_{i_l}$. Furthermore, (H2) implies that there can be no linear extension of P in which (i) two elements labeled i appear consecutively, or (ii) three elements labeled i, j, i appear consecutively, unless $a_{ji} = -2$, or (iii) four elements labeled i, j, i, j, unless $a_{ij} = a_{ji} = -2$. Since $a_{ij} = -2$ can occur only if $s_i s_j$ has order $\geqslant 4$, and $a_{ij} = a_{ji} = -2$ only if $s_i s_j$ does not have finite order, Proposition 3.3 of [**St1**] implies that P is the heap of a fully commutative element w; in particular, the expression is reduced. Applying Proposition 2.3, it follows that w is minuscule.

(c) If P is the heap of some dominant minuscule element, then (b) implies (H1)–(H2), Proposition 2.5 implies (H3), and Corollary 2.6 implies (H4). Conversely, if P satisfies (H1)–(H4), then from (b) we know that P is (isomorphic to) the heap of a reduced expression for some minuscule $w \in W$; say, $w = s_{i_1} \cdots s_{i_l}$. Assume towards a contradiction that w fails to be dominant. By Proposition 2.5 and (H3), the last occurrence of some generator s_i must be followed by at least two generators that do not commute with s_i. If the first two of these are in positions p and p' ($p < p'$), and the last s_i is in position q, then (H3) implies that p must be the unique element that covers q (whence $q \prec p \prec p'$), and p and p' must have different labels. Since q is the last element labeled i, it follows that a maximal chain in P from p to p' traces a path in the diagram that is disjoint from i but whose (distinct) endpoints are adjacent to i, contradicting (H4). $\square$

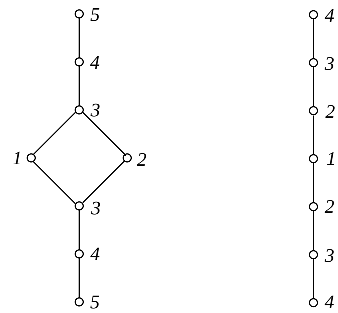

FIGURE 1: Subintervals of type D_5 and C_4.

REMARK 3.2. If P is a dominant minuscule heap, then the same is true of any order filter of P. (A subset F of P is an order filter if $p \in F$ and $p \prec q$ implies $q \in F$.) This follows easily from (H1)–(H4), or simply the fact that dominant minuscule reduced expressions are closed under deletion of initial segments (cf. Remark 2.7(b)).

In the minuscule case, an acyclic diagram tightly constrains the possibilities for subintervals bounded by elements of the same label (cf. Proposition 7.2 of [**P1**]).

PROPOSITION 3.3. *In the heap of a minuscule element satisfying (H4), every closed subinterval between two elements labeled i (with no other element labeled i in between) is isomorphic as a labeled poset to the heap of $s_k \cdots s_3 s_1 s_2 s_3 \cdots s_k$ in D_k ($k \geqslant 3$), or $s_k \cdots s_2 s_1 s_2 \cdots s_k$ in C_k ($k \geqslant 2$).*

The two subinterval types are illustrated in Figure 1.

It is important to note that we are making no claims about the entries of the Cartan matrix corresponding to the labels that appear in these subintervals beyond what can be inferred about certain entries being zero or nonzero. For example, in the (dominant) minuscule heaps of Figure 5, there are subintervals of type D_3 whose labels index subdiagrams of type B_3.

Proof. Consider a subinterval whose endpoints are labeled i.

Case 1: If the upper endpoint covers two or more elements then there must be exactly two such elements, and there can be no other elements in the subinterval whose labels are adjacent to i, by (H2). Since the two elements are incomparable, they must have distinct, non-adjacent labels, by (H1). Furthermore, since there must be maximal chains from these elements to the lower endpoint, these must be the *only* elements in the subinterval; otherwise, the elements covering the lower endpoint would exceed the limit of two elements in the interval with labels adjacent to i. It follows that the subinterval is of type D_3.

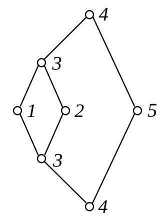

FIGURE 2: A fully commutative heap in D_5.

Case 2: If the upper endpoint covers only one element, labeled j, then the lower endpoint can only be covered by one element, also of label j. Otherwise, there would be a maximal chain from the lower endpoint to the upper endpoint whose labels trace a circuit in the Dynkin diagram, contradicting (H4). If the elements labeled j are in fact the same, then we obtain a subinterval of type C_2. Otherwise, (H2) implies that they must be the only two elements labeled j in the subinterval. By induction, the subinterval between these two elements is of type D_k or C_k, so the full subinterval is of type D_{k+1} or C_{k+1}. □

COROLLARY 3.4. *The heap of any minuscule element satisfying (H4) is ranked.*

Proof. Let q be a maximal element of a minuscule heap P satisfying (H4). By induction, there is a rank function for $P - \{q\}$. Allowing shifts of the rank function on connected components of $P - \{q\}$, we can extend the rank function to all of P unless there are two elements p and p' covered by q that are in the same connected component and have unequal rank. By following a path in the Hasse diagram of $P - \{q\}$ from p to p', we trace a path in the Dynkin diagram between two distinct nodes that are adjacent to the label i of q. Given (H4), this is possible only if the path passes through a vertex of $P - \{q\}$ labeled i. Hence there are at least two vertices in P labeled i, so by Proposition 3.3, the top two must form a subinterval of type D_3, with q at the top and p, p' the two unrelated elements in the middle. However, since p and p' both cover a fourth element, they must have the same rank in $P - \{q\}$. □

REMARK 3.5. The heaps of fully commutative elements, even those satisfying (H4), need not be ranked. An example involving D_5 is illustrated in Figure 2. Similarly, it is easy to give examples of minuscule heaps that are not ranked (but violate (H4)).

4. The Classification of Dominant Minuscule Heaps

Let P be the heap of dominant minuscule element $w \in W$. By passing to a suitable parabolic subgroup and root subsystem, we may assume that every available generator appears in a reduced expression for w. Recall that this forces the Dynkin diagram to be acyclic (Corollary 2.6).

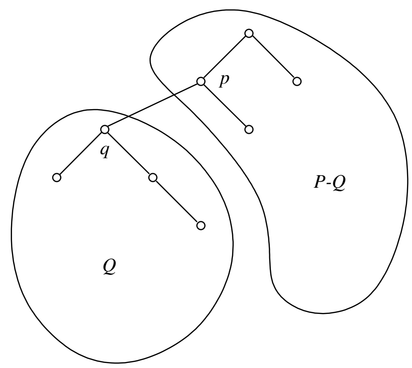

FIGURE 3: Decomposition of a reducible dominant minuscule heap.

If the Dynkin diagram is disconnected, then P is the disjoint union of the heaps of certain dominant minuscule elements belonging to the parabolic subgroups corresponding to the connected components. Conversely, the union of dominant minuscule heaps whose label sets are supported on distinct connected components is itself the heap of a dominant minuscule element. Thus we now restrict our attention to connected Dynkin diagrams.

Let T denote the set of vertices of P consisting of the top elements of each label. Property (H3) shows that every member of T is covered in P by at most one element, and this element is also a member of T. Thus T is an order filter of P and has the order structure of a forest of rooted trees. Given the hypothesis that the diagram is connected, it follows that T is in fact a single rooted tree and has a maximum element. Following Proctor, we refer to T as the *top tree* of P. However, it should be noted that in multiply-laced cases (unlike the simply-laced cases in [**P2**]), the top tree is not necessarily a maximal tree-filter of P. (Compare the two posets in Figure 1.)

We say that P is *irreducible* if the label of every vertex that is not minimal in T occurs at least twice in P. In the simply-laced case, this is equivalent to being "slant-irreducible" as defined in [**P2**].

Suppose that P is not irreducible. Thus there is some label i that is assigned only once in P, say to p, and there is some $q \in T$ covered by p. Let j denote the label of q, let Q be the labeled subposet of P consisting of all $q' \preccurlyeq q$, and let J denote the set of labels that occur in the portion of T that is $\preccurlyeq q$. (See Figure 3.) Every $q' \in Q$ has a maximal chain from q' to q; this chain cannot pass through p, the unique vertex labeled i, so the labels along the path must stay within J. Thus Q consists of all members of P whose labels are in J. There also cannot be any covering relations between members of Q and $P - Q$ other than $q \prec p$; otherwise, there would be a path in the diagram between i and j in addition to the edge between i and j. Furthermore, it follows easily from Proposition 3.1 that Q and $P - Q$ are heaps of dominant minuscule elements of W.

Conversely, suppose that P and Q heaps of dominant minuscule elements whose labels are supported on two disjoint (but connected) Dynkin diagrams, and that p is a vertex in the top tree of P whose label i occurs only once in P. Let q be the maximum element of Q and j the label of q. Again via Proposition 3.1, it easily follows that the labeled poset obtained from $P \cup Q$ by adding the covering relation $q \prec p$ is a dominant minuscule heap relative to any Dynkin diagram obtained by taking the union of the two original diagrams and adding an edge between i and j, with i short relative to j. We call this new labeled poset the *sum* of P and Q at p. In the simply-laced case, this is equivalent to the "slant sum" defined in [**P2**].

The preceding remarks reduce the classification of dominant minuscule heaps to the irreducible case; all other connected heaps can be built from sums of irreducible heaps.

Lemma 4.1. *Let P be an irreducible dominant minuscule heap with top tree T. If $q \in T$ covers two members of P, then every $p \prec q$ in T covers an element not in T.*

Proof. Let i denote the label of q. Since $p \prec q$, there is some $q' \in T$ covered by q with $p \preccurlyeq q' \prec q$. Since q is not minimal in T and P is irreducible, there must be another vertex labeled i in P. Given that q covers two elements, both must occur in some subinterval of P between two elements labeled i, whence (Proposition 3.3) this must be a subinterval of type D_3, and the (open) subinterval has *exactly* two elements, including q'. In particular, q' covers the lower endpoint of this subinterval, and this lower endpoint is not a member of T, being the second highest element labeled i. If $p = q'$, we are done. Otherwise, q' is not minimal in T so it covers a second element (a member of T). We therefore replace $q \leftarrow q'$ and proceed by induction on the length of a maximal chain from p to q. $\square$

In the Weyl group of B_n (with the Cartan matrix arranged so that $a_{i,i+1} = a_{i+1,i} = -1$ except $a_{21} = -2$), define M_n to be the heap of $w = s_1(s_2s_1)(s_3s_2s_1)\cdots(s_n\cdots s_2s_1)$. For example, M_4 is illustrated in Figure 4. It is not hard to show that w is a dominant minuscule element; for example, one can easily see that M_n satisfies (H1)–(H4).

Theorem 4.2. *If P is the heap of a dominant minuscule $w \in W$ that is irreducible, then either Φ is simply-laced, or the Dynkin diagram of Φ has the form*

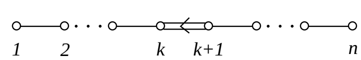

for some k $(1 \leqslant k < n)$, the maximum element of P is labeled 1, and

(i) *$k = 1$ (i.e., $\Phi \cong B_n$) and P is isomorphic to an order filter of M_n, or*
(ii) *$k > 1$ and $w = (s_{n-1}\cdots s_{k+1})(s_n\cdots s_{k+2})(s_1\cdots s_k s_{k+1} s_k\cdots s_1)$.*

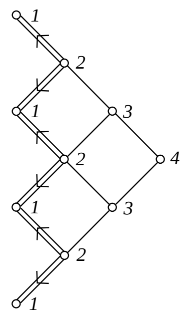

FIGURE 4: M_4.

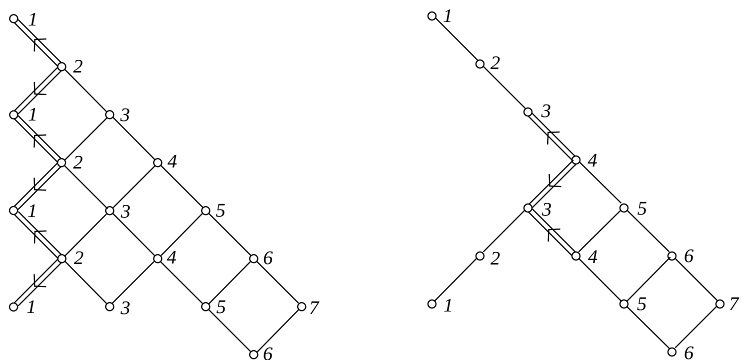

FIGURE 5: Dominant minuscule heaps.

Examples of dominant minuscule heaps of types (i) and (ii) are illustrated in Figure 5.

The simply-laced dominant minuscule heaps have been classified by Proctor in [**P2**]. Combined with the above result, this classifies all dominant minuscule heaps.

Proof. Assuming Φ is not simply-laced, there must be at least one covering pair $p \prec q$ in T whose respective labels j and i are connected by a non-simple edge of the Dynkin diagram. Since (H3) requires $a_{ij} = -1$, it must be the case that $a_{ji} \leqslant -2$. Since P is irreducible and q is not minimal in T, there must be another element labeled i in P. Hence p occurs in some subinterval of P between two elements labeled i, whence $a_{ji} = -2$ by (H2). Thus every edge of the Dynkin diagram is of type A_2 or B_2, and the B_2-edges are oriented in T so that the short end is higher.

A second consequence of (H2) is that in the above situation, p covers an element labeled i not in T. Therefore if p is not minimal in T, it covers at least two elements in P, and by Lemma 4.1, every element below p in T must cover an element not in T. It follows that the portion of T below p must be a chain. Otherwise some $p' \preccurlyeq p$ in T would cover at least three elements: two members of T and an element not in T. As a non-minimal member of T, there must be a second element of P with the same label i' as p', and hence there is a subinterval of P with endpoints labeled i' that contains three elements whose labels are adjacent to i', contradicting (H2). Similarly, no $p' \preccurlyeq p$ may cover a member of T whose label is not short relative to i'; otherwise, we would again contradict (H2). Thus, all of the covering relations below p in T must correspond to simple edges of the diagram.

Next, we claim that every $q' \succcurlyeq q$ must cover only one element of P. Otherwise, Lemma 4.1 implies that q must cover a second element in addition to p. However since $a_{ji} = -2$, this contradicts the fact that the subinterval between q and the second highest element labeled i must contain only p, by (H2). Thus the entire top tree must be a chain.

Since we have seen that every edge of T below a B_2-edge must be simple, it follows that there is exactly one B_2-edge. Hence the Dynkin diagram has the claimed form, and (with the labels arranged to match the above figure), the label of the top element of P must be 1.

If the B_2-edge is at the top of the tree (i.e., the case $k = 1$), then we claim that P must be an order filter of M_n. To verify this claim, one needs only to check that any heap obtained by adding a minimal element to an order filter of M_n is either an order filter of M_n or violates (H1)–(H4).

In the case $k > 1$, irreducibility forces each of the labels $1, \dots, n-1$ to occur at least twice in P. Since $a_{k+1,k} = -2$, the second highest vertex labeled k must be covered by the top vertex labeled $k+1$, and no other vertex can appear in the subinterval between the top two vertices labeled k (by (H2)). In particular, the top pair of vertices labeled $k-1$ has two vertices labeled k in between, so the second k must cover the second $k-1$. Similarly, the second highest i must cover the second highest $i-1$ for $1 < i \leqslant k$, and the highest $j+1$ and second highest $j-1$ must cover the second highest j for $k < j < n$.

These relations account for the heap of the element w described in (ii). To see that there are no further possibilities, observe that one cannot add a third element of any label $< n$ (or a second element labeled n) to P without violating (H1)–(H4). □

Each of the dominant minuscule heaps described in the above theorem can be converted to simply-laced (dominant minuscule) heaps by a suitable relabeling. See Figure 6. In each case, the new Dynkin diagram is Y-shaped, with two branches of length k; in particular, the B_n-heaps are converted to D_{n+1}-heaps. This shows that the underlying unlabeled posets are "d-complete" (although reducible) in the sense of Proctor [**P1–2**].

Corollary 4.3. *Every dominant minuscule heap is isomorphic (as an unlabeled poset) to a simply-laced dominant minuscule heap; i.e., dominant minuscule heaps are d-complete.*

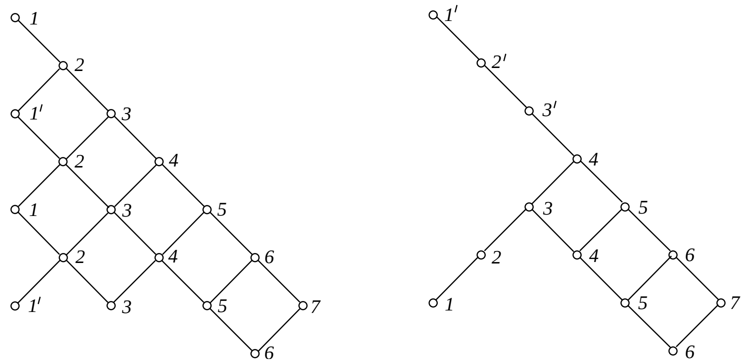

Figure 6: Simply-laced relabelings.

5. Heaps and Inversion Sets

We define an *inversion* of $w \in W$ to be a root $\gamma \in \Phi^+$ such that $w\gamma \in \Phi^-$; this generalizes in a natural way the standard notion of inversion in a permutation. We let

$$\Phi(w) := \{\gamma \in \Phi^+ : w\gamma \in \Phi^-\},$$
$$\Phi^\vee(w) := \{\gamma^\vee \in (\Phi^\vee)^+ : w\gamma^\vee \in (\Phi^\vee)^-\},$$

denote the set of inversions of w, along with the co-root analogue. The latter turns out to be more natural in some cases (e.g., Theorem 5.5).

It is well-known (e.g., Exercise 5.6.1 in [**H**]) that $\Phi(w)$ can be determined explicitly from any reduced expression $w = s_{i_1} \cdots s_{i_l}$; viz., $\Phi(w) = \{\gamma_1, \dots, \gamma_l\}$, where

$$\gamma_l = \alpha_{i_l}, \quad \gamma_{l-1} = s_{i_l}\alpha_{i_{l-1}}, \quad \dots, \quad \gamma_1 = s_{i_l} \cdots s_{i_2}\alpha_{i_1}.$$

We call $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_l)$ the *root sequence* of $s_{i_1} \cdots s_{i_l}$. Similarly, $\boldsymbol{\gamma}^\vee = (\gamma_1^\vee, \dots, \gamma_l^\vee)$ is the *co-root sequence.*

Proposition 5.1. *$w \in W$ is λ-minuscule if and only if $\langle \lambda, \gamma^\vee \rangle = 1$ for all $\gamma \in \Phi(w)$.*

Proof. Choose a reduced expression $w = s_{i_1} \cdots s_{i_l}$. Since W acts as a group of isometries relative to $\langle \ , \ \rangle$, it follows immediately from (2.1) that w is λ-minuscule if and only if $\langle \lambda, \gamma_k^\vee \rangle = 1$ $(1 \leqslant k \leqslant l)$, where $\gamma_k = s_{i_l} \cdots s_{i_{k+1}}\alpha_{i_k}$. $\square$

Given a reduced expression $w = s_{i_1} \cdots s_{i_l}$ with root sequence $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_l)$, we define the *heap ordering* of $\boldsymbol{\gamma}$ to be the partial order generated by taking the transitive closure of the relations

$$\gamma_p \prec \gamma_q \quad \text{whenever} \quad \langle \gamma_p, \gamma_q \rangle \neq 0.$$

The heap ordering of the co-root sequence $\boldsymbol{\gamma}^\vee$ is defined similarly.

If s_i and s_j are a pair of generators that occur in positions $k-1$ and k of the reduced expression, then $\gamma_k = x^{-1}\alpha_j$ and $\gamma_{k-1} = x^{-1}s_j\alpha_i$, where $x = s_{i_{k+1}} \cdots s_{i_l}$. Hence

$$\langle \gamma_{k-1}, \gamma_k^\vee \rangle = \langle s_j\alpha_i, \alpha_j^\vee \rangle = -\langle \alpha_i, \alpha_j^\vee \rangle. \tag{5.1}$$

In particular, s_i and s_j commute if and only if γ_{k-1} and γ_k are orthogonal. Furthermore, if s_i and s_j do commute, then we have $s_j\alpha_i = \alpha_i$ and $\gamma_{k-1} = x^{-1}\alpha_i$, so transposing two commuting generators corresponds to transposing two adjacent orthogonal roots in $\boldsymbol{\gamma}$ and hence creates a new root sequence with the same heap ordering.

If w is fully commutative, then any reduced expression for w can be obtained by a sequence of such operations, so all root sequences yield the same partial ordering of $\Phi(w)$.

The following result justifies the terminology.

PROPOSITION 5.2. *If P is the heap of a reduced word and $\boldsymbol{\gamma}$ is the corresponding root sequence, then $p \prec q$ in P if and only if $\gamma_p \prec \gamma_q$ in the heap ordering of $\boldsymbol{\gamma}$.*

Proof. If $\gamma_p \prec \gamma_q$ is a covering relation, then $\langle \gamma_p, \gamma_q \rangle \neq 0$, so γ_p precedes γ_q in every root sequence belonging to the commuting equivalence class of $\boldsymbol{\gamma}$, so p precedes q in every linear extension of P, so $p \prec q$. Conversely, if $p \prec q$ is a covering relation of P, then the corresponding terms of the reduced expression cannot commute, and there is a reduced expression for w in which γ_p and γ_q appear consecutively in the corresponding root sequence. Hence (5.1) implies $\langle \gamma_p, \gamma_q \rangle \neq 0$ and $\gamma_p \prec \gamma_q$. $\square$

The following result generalizes Theorem 2.4 of [**FS**] from the simply-laced case to any (symmetrizable Kac-Moody) Weyl group.

THEOREM 5.3. *Given $w \in W$, the following are equivalent:*

(a) *There is no reduced expression for w containing a subword of the form $s_is_js_i$, where s_j is short relative to s_i. (In particular, w is fully commutative.)*
(b) *There is no triple of co-roots $\alpha^\vee, \beta^\vee, \alpha^\vee + \beta^\vee \in \Phi^\vee(w)$.*
(c) *The heap ordering of some (equivalently, every) co-root sequence for w is consistent with the dual of the standard ordering of $\Phi^\vee$ (i.e., $\alpha^\vee \prec \beta^\vee$ implies $\alpha^\vee > \beta^\vee$).*

Proof. We argue that the negations of these conditions are equivalent.

$\neg$(a) $\Rightarrow$ $\neg$(b). If (a) fails, then w has a reduced expression that contains $s_is_js_i$, where $\langle \alpha_j, \alpha_i^\vee \rangle = -1$. In the corresponding co-root sequence there must be terms of the form $x^{-1}\alpha_i^\vee$, $x^{-1}s_i\alpha_j^\vee = x^{-1}(\alpha_j^\vee + k\alpha_i^\vee)$, and

$$x^{-1}s_is_j\alpha_i^\vee = x^{-1}s_i(\alpha_j^\vee + \alpha_i^\vee) = x^{-1}(\alpha_j^\vee + (k-1)\alpha_i^\vee)$$

for some $x \in W$, where $k = -\langle \alpha_i, \alpha_j^\vee \rangle \geqslant 1$. It follows that $\alpha^\vee, \beta^\vee, \alpha^\vee + \beta^\vee \in \Phi^\vee(w)$, where $\alpha^\vee = x^{-1}\alpha_i^\vee$ and $\beta^\vee = x^{-1}(\alpha_j^\vee + (k-1)\alpha_i^\vee)$.

$\neg$(b) $\Rightarrow$ $\neg$(c). Suppose $\alpha^\vee, \beta^\vee, \alpha^\vee + \beta^\vee \in \Phi^\vee(w)$. These terms must appear in every co-root sequence for w in the order $\alpha^\vee, \alpha^\vee + \beta^\vee, \beta^\vee$ or the reverse. Otherwise, by suitable truncation, there would exist $x \in W$ such that $x(\alpha^\vee + \beta^\vee)$ is negative while $x\alpha^\vee$ and $x\beta^\vee$ are both positive, or vice-versa. If both orderings occurred among the set of co-root sequences for w within a given commuting equivalence class, then the three co-roots would have to be pairwise incomparable in the heap ordering, and thus pairwise orthogonal. However, this contradicts the fact that they are linearly dependent. Therefore, within a given commuting equivalence class, $\alpha^\vee$ always precedes $\alpha^\vee + \beta^\vee$ or $\beta^\vee$ always precedes $\alpha^\vee + \beta^\vee$, whence $\alpha^\vee \prec \alpha^\vee + \beta^\vee$ or $\beta^\vee \prec \alpha^\vee + \beta^\vee$, contradicting (c).

$\neg$(c) $\Rightarrow$ $\neg$(a). If the heap ordering of some co-root sequence fails to be consistent with the dual of the standard ordering, then there must be a covering pair $\alpha^\vee \prec \beta^\vee$ in the heap such that $\alpha^\vee \not> \beta^\vee$. If there is more than one such pair available, choose one so that $\beta^\vee$ is maximal among all such pairs, relative to the heap ordering.

By choosing a suitable reduced expression $w = s_{i_1} \cdots s_{i_l}$ in the same equivalence class, we may obtain a co-root sequence in which $\alpha^\vee$ and $\beta^\vee$ appear in positions $p-1$ and p, for some p. Setting $i = i_{p-1}$ and $j = i_p$, it follows that $\beta^\vee = x^{-1}\alpha_j^\vee$ and $\alpha^\vee = x^{-1}s_j\alpha_i^\vee = x^{-1}(\alpha_i^\vee + k\alpha_j^\vee)$, where $x = s_{i_{p+1}} \cdots s_{i_l}$ and $k = -\langle \alpha_j, \alpha_i^\vee \rangle \geqslant 1$. Since $\alpha^\vee \not> \beta^\vee$ and

$$\alpha^\vee - \beta^\vee = x^{-1}\alpha_i^\vee + (k-1)x^{-1}\alpha_j^\vee,$$

it cannot be the case that $x^{-1}\alpha_i^\vee$ and $x^{-1}\alpha_j^\vee$ are both positive. However $\beta^\vee = x^{-1}\alpha_j^\vee$, so $x^{-1}\alpha_i^\vee$ must be negative. Thus x has a reduced expression that begins with s_i (e.g., see [**H**, §5.4]), so we may assume $i_{p+1} = i$ and the term immediately following $\beta^\vee$ in the co-root sequence is $\gamma^\vee = x^{-1}s_i\alpha_i^\vee = -x^{-1}\alpha_i^\vee$.

If $k = 1$, we are done. Otherwise, $k > 1$ and

$$\beta^\vee - \gamma^\vee = x^{-1}\alpha_i^\vee + x^{-1}\alpha_j^\vee = \frac{k-2}{k-1}x^{-1}\alpha_i^\vee + \frac{1}{k-1}(\alpha^\vee - \beta^\vee)$$

is in the positive linear span of $x^{-1}\alpha_i^\vee$ and $\alpha^\vee - \beta^\vee$. Since the former is a negative co-root, and the latter is not a sum of positive co-roots, it follows that $\beta^\vee \not> \gamma^\vee$, contradicting our choice of $\beta^\vee$. $\square$

Remark 5.4. (a) Either of Propositions 2.1 or 2.3 show that minuscule elements satisfy the equivalent conditions of the above result. But there are non-minuscule elements that also satisfy the conditions (e.g., $w = s_3s_1s_2s_4s_3$ in D_4).

(b) Fan proves the equivalence of parts (a) and (b) for simply-laced finite Weyl groups and briefly discusses (in a dual form) the multiply-laced case in [**F**].

Given that Theorem 5.3 provides circumstances where the heap ordering and the dual of the standard ordering of $\Phi^\vee(w)$ are related, it is natural to investigate circumstances where the two orderings coincide.

THEOREM 5.5. *If w is dominant minuscule, then the heap of w is dual-isomorphic to the standard ordering of $\Phi^\vee(w)$. In fact, $(\Phi^\vee(w), \prec) = (\Phi^\vee(w), >)$.*

Proof. Assume that w is minuscule relative to the dominant weight λ. By Theorem 5.3, we know that $\alpha^\vee \prec \beta^\vee$ in $\Phi^\vee(w)$ implies $\alpha^\vee > \beta^\vee$, so it suffices to prove the converse. Given that $\alpha^\vee > \beta^\vee$, we have $\alpha^\vee = \beta^\vee + \sum_{i\in J} c_i\alpha_i^\vee$, where $J \subseteq [n]$ and the c_i's are positive integers. By Proposition 5.1, it follows that

$$\sum_{i\in J} c_i\langle\lambda, \alpha_i^\vee\rangle = \langle\lambda\,, \alpha^\vee - \beta^\vee\rangle = 0.$$

However λ is dominant, so $\langle\lambda, \alpha_i^\vee\rangle = 0$ for all $i \in J$; in particular, $w\alpha_i^\vee$ must be positive for all $i \in J$ (again by Proposition 5.1).

If $\alpha^\vee \not\prec \beta^\vee$, then there would exist a co-root sequence in which $\beta^\vee$ precedes $\alpha^\vee$. Hence by truncation, there would exist some $x \in W$ with an inversion set $\Phi^\vee(x) \subset \Phi^\vee(w)$ that contains $\alpha^\vee$ but not $\beta^\vee$. However if $x\alpha^\vee$ is negative and $x\beta^\vee$ is positive, then $x\alpha_i^\vee$ must be negative for some $i \in J$, contradicting the fact that $\Phi^\vee(x) \subset \Phi^\vee(w)$. □

REMARK 5.6. (a) The above argument shows that if $\alpha^\vee > \beta^\vee$, $\alpha^\vee \in \Phi^\vee(w)$ and $\langle\lambda, \beta^\vee\rangle = 1$, then $\beta^\vee \in \Phi^\vee(w)$. In other words, if w is (dominant) λ-minuscule, then $\Phi^\vee(w)$ is an order ideal of $\Phi^\vee_\lambda := \{\alpha^\vee \in \Phi^\vee : \langle\lambda, \alpha^\vee\rangle = 1\}$, relative to $<$. In particular, since $\Phi^\vee_\lambda$ (and therefore $\Phi^\vee(w)$) is an order-convex subset of $(\Phi^\vee, >)$, it follows that every dominant minuscule heap is isomorphic to a convex subposet of $(\Phi^\vee, >)$.

(b) If w is minuscule but not dominant, then the heap ordering and dual-standard ordering of $\Phi^\vee(w)$ need not coincide. For example, in the affine root system of type $A_3^{(1)}$ (with the nodes numbered so that $1, 2, 3, 4$ form a circuit), the element $w = s_3s_1s_2s_4s_1$ is easily seen to be minuscule, and the corresponding root sequence is

$$\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_5) = (2\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4,\ \alpha_1 + \alpha_2 + \alpha_4,\ \alpha_1 + \alpha_2,\ \alpha_1 + \alpha_4,\ \alpha_1).$$

However, the relation $\gamma_1 > \gamma_2$ has no counterpart in the heap ordering. See Figure 7, where the posets are displayed using label i for the vertex γ_i.

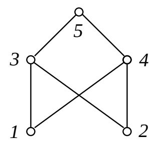

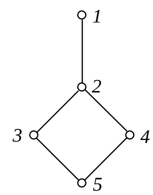

FIGURE 7: $(\Phi^\vee(w), \prec)$ and $(\Phi^\vee(w), <)$ for a non-dominant minuscule w.

It is plausible that if the Dynkin diagram is acyclic, then every element satisfying the conditions of Theorem 5.3 has an inversion set whose heap ordering and standard ordering are dual. For example, although we omit the proof, we claim that this is true at least when W is finite. It is also plausible that the standard ordering is the "wrong" ordering for this purpose.

QUESTION 5.7. *Is there a partial ordering $\lhd$ of $\Phi^\vee$ so that for every $w \in W$ satisfying the conditions of Theorem 5.3, $(\Phi^\vee(w), \lhd)$ is the heap ordering of $\Phi^\vee(w)$?*

In case the Cartan matrix satisfies (cf. Remark 2.7(a))

$$a_{ij} < 0 \quad \Rightarrow \quad a_{ij} = -1 \text{ or } a_{ji} = -1, \tag{5.2}$$

computer searches suggest that a candidate for $\lhd$ is the partial order obtained by taking the transitive closure of the relations

$$\alpha^\vee \lhd \beta^\vee \quad \text{whenever} \quad \alpha^\vee - \beta^\vee \in (\Phi^\vee)^+.$$

In the finite case, this coincides with the dual of the standard ordering. It also eliminates the extraneous relation in the example discussed in Remark 5.6(b).

The necessity of (5.2) can be seen as follows. The element $w = s_i s_j$ satisfies the conditions of Theorem 5.3, and has the co-root sequence $(s_j\alpha_i^\vee, \alpha_j^\vee)$, so $s_j\alpha_i^\vee \prec \alpha_j^\vee$ in the heap (assuming $a_{ij} < 0$). However if $a_{ij}, a_{ji} \leqslant -2$, it is not hard to show that there is no pair of positive co-roots whose sum is $s_j\alpha_i^\vee$.

The following result shows that this ordering is at least consistent with the heap.

PROPOSITION 5.8. *Assume that the Cartan matrix satisfies (5.2) and that w satisfies the conditions of Theorem 5.3.*

(a) *If $\alpha^\vee \prec \beta^\vee$ is a covering relation in the heap of $\Phi^\vee(w)$, then $\alpha^\vee - \beta^\vee \in (\Phi^\vee)^+$.*
(b) *If $\alpha^\vee, \beta^\vee \in \Phi^\vee(w)$ and $\alpha^\vee - \beta^\vee \in (\Phi^\vee)^+$, then $\alpha^\vee \prec \beta^\vee$ in the heap.*

Proof. (a) If $\alpha^\vee \prec \beta^\vee$ is a covering relation, then there is a co-root sequence in which $\alpha^\vee$ and $\beta^\vee$ occur consecutively, and hence (5.1) and (5.2) imply $\langle \alpha, \beta^\vee \rangle = 1$ or $\langle \beta, \alpha^\vee \rangle = 1$.

Thus the reflection of $\alpha^\vee$ through the hyperplane perpendicular to β is $\alpha^\vee - \beta^\vee$, or vice-versa. Either way, $\alpha^\vee - \beta^\vee$ is a co-root, necessarily positive by Theorem 5.3(c).

(b) If $\alpha^\vee$ and $\beta^\vee$ are unrelated in the heap ordering, then they must be orthogonal and occur consecutively in some co-root sequence for w. It follows that there is an orthogonal pair of simple roots α_i, α_j such that $\alpha^\vee = x^{-1}\alpha_i^\vee$ and $\beta^\vee = x^{-1}\alpha_j^\vee$ for some $x \in W$ (cf. (5.1)). However since $\alpha_i^\vee - \alpha_j^\vee$ cannot be a co-root (it is neither positive nor negative), this contradicts the fact that $\alpha^\vee - \beta^\vee$ is a co-root. Hence $\alpha^\vee$ and $\beta^\vee$ must be related in the heap, and since $\alpha^\vee - \beta^\vee$ is positive, the relation must be $\alpha^\vee \prec \beta^\vee$, by Theorem 5.3(c). $\square$

References

[**B**] N. Bourbaki, "Groupes et Algèbres de Lie, Chp. IV–VI," Masson, Paris, 1981.

[**F**] C. K. Fan, "A Hecke Algebra Quotient and Properties of Commutative Elements of a Weyl Group," Ph. D. thesis, MIT, 1995.

[**FS**] C. K. Fan and J. R. Stembridge, Nilpotent orbits and commutative elements, *J. Algebra* **196** (1997), 490–498.

[**H**] J. E. Humphreys, "Reflection Groups and Coxeter Groups," Cambridge Univ. Press, Cambridge, 1990.

[**K**] V. G. Kac, "Infinite Dimensional Lie Algebras," Cambridge Univ. Press, Cambridge, 1990.

[**P1**] R. A. Proctor, Minuscule elements of Weyl groups, the numbers game, and d-complete posets, *J. Algebra* **213** (1999), 272–303.

[**P2**] R. A. Proctor, Dynkin diagram classification of λ-minuscule Bruhat lattices and of d-complete posets, *J. Algebraic Combin.* **9** (1999), 61–94.

[**P3**] R. A. Proctor, Generalized Young diagrams with well-defined jeu de taquin emptying procedures, preprint.

[**St1**] J. R. Stembridge, On the fully commutative elements of Coxeter groups, *J. Algebraic Combin.* **5** (1996), 353–385.

[**St2**] J. R. Stembridge, Quasi-minuscule quotients and reduced words for reflections, preprint.