Minuscule posets from neighbourly graph sequences

N.J. Wildberger

School of Mathematics
UNSW
Sydney NSW 2052
Australia

February 7, 2001

1 Introduction

We construct minuscule posets, an interesting family of posets arising in Lie theory, algebraic geometry and combinatorics, from sequences of vertices of a graph with particular *neighbourly* properties.

We begin by associating to any sequence of vertices in a simple graph X, here always assumed connected, a partially ordered set called a *heap*. This terminology was introduced by Viennot ([11]) and used extensively by Stembridge in the context of fully commutative elements of Coxeter groups (see [8]), but our context is more general and graph-theoretic. The heap of a sequence of vertices is that partially ordered set whose total linear orders correspond to all possible sequences obtained from the original one by switching adjacent elements which are not neighbours in X. Furthermore sequences which are equivalent under such interchanges (of adjacent elements which are not neighbours in the graph) give rise to identical heaps.

A heap will be called *neighbourly* if the associated sequences have the property that between any two successive occurrences of a vertex x there occurs at least two occurrences of a neighbour of x.

Heaps arising from maximal neighbourly sequences which in fact have exactly two neighbours between any two occurrences of a vertex x are classified. In our main result, we prove that any graph X having such a maximal neighbourly heap which is in fact two-neighbourly must be one of the Dynkin - Coxeter diagrams A_n, D_n, or E_6, E_7, and that the corresponding heaps are exactly the minuscule posets defined and studied by Proctor in [4].

In the last section we briefly connect these interesting minuscule posets (actually they are all distributive lattices) to Lie theory, algebraic geometry, and combinatorics. This paper could be viewed as an elementary graph theoretic approach to their study. We were led to these posets in our attempt to construct Lie algebra representations directly from Dynkin diagrams, work which is described in [12].

2 Neighbourly heaps for a graph

Let X be a simple graph. By an X-*sequence* we mean a sequence $s = (x_1, \dots, x_n)$ of vertices of X. If we transform s to s' by switching x_i and x_{i+1} for some i then there are three possibilities:

1) x_i and x_{i+1} are neighbours in X – (an X-*switch*)

2) x_i and x_{i+1} are distinct and not neighbours – (a free *switch*)

3) $x_i = x_{i+1}$ – (a redundant *switch*).

Any $X-$sequence s' obtainable from s by free switches is defined to be *equivalent* to s; we write $s \simeq s'$ and let $[s]$ denote the equivalence class of s, which we call an X-*string*. We refer to the x_i in $s = (x_1, \dots, x_n)$ as the *occurrences* in s; as occurrences they are considered distinct even if as vertices of X there may be repetitions. We partially order the occurrences x_i in s by declaring $x_i < x_j$ if $i < j$ and x_i, x_j are neighbours or identical vertices in X. The resulting poset P_s is unchanged by free switches and so depends only on the X-string $[s]$. We refer to $P_s = P_{[s]}$ as the X-*heap* of $[s]$.

Proposition 2.1 *The X-string $[s]$ consists exactly of the total orderings of $P_{[s]}$ consistent with the partial order.*

Proof. Any sequence s' obtained from s by free switches has the same heap and so is an ordering of $P_{[s]}$ consistent with the partial ordering. Conversely suppose s' is an ordering of P consistent with the partial order. Let's show that we can free switch s' to obtain s. Suppose by induction that s and s' agree up to to the kth term so that

$$s = (x_1, x_2, \cdots, x_k, x_{k+1}, \cdots x_n)$$

$$s' = (x_1, x_2, \cdots, x_k, y_{k+1}, \cdots y_n)$$

and that $x_{k+1} = x$. Clearly there is a first occurrence of x in $y_{k+1}, \cdots y_n$, and if this first occurrence is preceded by a neighbour $y = y_j$ in X of x, then since any two neighbours are necessarily related, we must have $y_j < x_{k+1}$ in P. But this contradicts the fact that P is the heap of s, in which x_{k+1} occurs before y_j. $\square$

Example 1 Suppose $X = A_n$ labelled as shown.

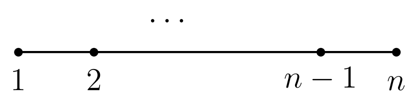

If we consider only X-sequences which are permutations of $\{1, \ldots, n\}$, the associated heaps are 'stock market graphs' where each successive node is either up or down from the previous. We get naturally a map from S_n to the set of sequences $\{(\eta, \ldots, \eta n - 1) \mid \eta = \pm 1\} = T$. It is natural to ask for the distribution of this map: how many permutations map to a given $t \in T$? When t is the zigzag sequence alternating plus and minus one, this is known as André's Problem, and the answer is given by Euler numbers, or Entringer numbers. The general case has been recently solved by G. Szekeres.

Example 2 Suppose $X = E_6$ labelled as shown

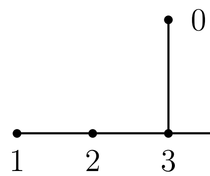

E_6

The X-sequence $s = (1, 2, 3, 0, 4, 5, 3, 2, 4, 3, 1, 0, 2, 3, 4, 5)$ has heap

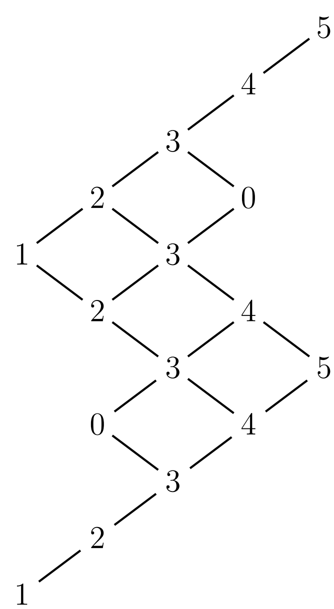

$F(E_6, 1)$

For future reference, we refer to this particular heap as $F(E_6, 1)$.

Definition An X-sequence $s = (x_1, \ldots, x_n)$ will be called *neighbourly* if between any two consecutive occurrences of a vertex x there are at least two occurrences of some neighbour or neighbours of x. This property is preserved by free switches, so we also speak of neighbourly X-strings and X-heaps.

A neighbourly X-sequence s will be called *maximal* if F cannot be extended by the addition of a vertex x in any position to a larger neighbourly X-sequence s', and similarly for X-strings and heaps. The neighbourly E_6-heap of Example 2 is maximal.

A neighbourly X-string or X-heap will be called *two-neighbourly* if there are exactly two occurrences of some neighbour or neighbours of x between any two consecutive occurrences of any vertex x. The heap $F(E_6, 1)$ of Example 2 is two-neighbourly.

Recall that a *lattice* is a poset such that for $a, b \in L$ the least upper bound $a \vee b$ and greatest lower bound $a \wedge b$ exist uniquely. When these operations satisfy the usual distributive laws, the lattice is called *distributive*. If P is any poset, an *ideal* of P is a subset I such that $x \in I,\ y \leq x$ implies $y \in I$. Let $J(P)$ denote the poset of all ideals of P ordered by inclusion. Then $J(P)$ is always a distributive lattice, and any distributive lattice is of the form $J(P)$ for some poset P.

Proposition 2.2 *If a graph X has a maximal neighbourly X-heap then X is a tree.*

Proof. If X is not a tree, consider the first occurrences of the elements of some fixed cycle in X. The last occurrence in this set is necessarily preceded by two neighbours, which contradicts maximality. $\square$

Proposition 2.3 *If F is a maximal neighbourly X-heap for some simple graph X, then F is a lattice.*

Proof. Let us suppose that P is a maximal neighbourly X-heap for some graph X and that $F = P_{[s]}$ for some X-sequence s. The previous Proposition shows that X must be a tree.

We first make a remark of general interest. Since the partial order on F is generated by the relations $x_i < x_j$ if $i < j$ and x_i, x_j are neighbours or identical in X, two occurrences x_i and x_j are related by $x_i < x_j$ if and only if there is a subsequence of s, $x_i = x_{i_1}, x_{i_2}, \cdots, x_{i_k} = x_j$ such that $i_1 < i_2 < \cdots i_k$ (this is what we mean by a subsequence) and such that any two successive elements in the subsequence are neighbours in X. That is, x_{i_j} and $x_{i_{j+1}}$ are neighbours, for all $j = 1, \cdots, k-1$.

It can be useful to imagine that the vertices of X are lights which are turned off and on in sequence according to s, so that the term x_i in s means that vertex x is lit up at time i. One is allowed to move from a vertex to a neighbouring vertex precisely when that neighbouring vertex is lit. Then to say that $x_i < y_j$ is just to say that you can get from vertex x at time i to vertex y at time j by a sequence of allowed moves.

Now suppose we have two occurences $x_i = x$ and $x_j = y$ in s with say $i < j$. To say that there is a unique z_k so that $x_i \leq z_k$ and $y_j \leq z_k$ is to say that there is unique vertex on which two players A and B can meet at the earliest possible time if they start at x and y at times i and j respectively.

Since X is a tree, if our two players want to meet as soon as possible they will have to approach each other along the unique path which separates them, say $x = x^0, x^1, \cdots, x^k = y$. This means that A will move to x^1 at the first opportunity, B will move to x^{k-1} at the first opportunity and so on. If they can meet in this way it is clear that there is a unique vertex and time when they will do so. Otherwise, they will reach a point when they are unable to decrease the distance between them. Without loss of generality let us assume this from the beginning. It means there is no occurence of x^1 past time i (and no occurence of x^{k-1} past time j).

But then by maximality there can be no occurrence of x^2 past time i either since then the previous occurence of x^1 (which must exist) will be followed by two occurences of its neighbours but not by another occurence of itself, which is impossible. So after time i there is no occurence of x^1, x^2 and so on. But we are told that $x^k = y$ does occur after time i so our assumption is impossible.

A similar argument shows that there is a unique occurence w_l with $w_l \leq x_i$ and $w_l \leq y_j$. $\square$

Recall the family of graphs $D_n, n \geq 4$ and E_7 and E_8 labelled as shown

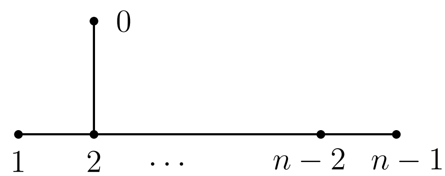

D_n

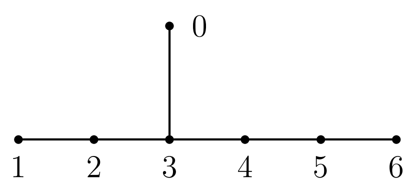

E_7

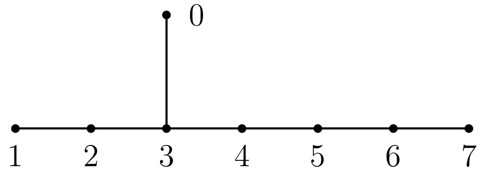

E_8

Theorem 2.1 *Let X be a simple graph for which there exists a maximal neighbourly X-heap F which is two-neighbourly. Then X is one of the graphs $A_n, n \geq 1,\ D_n,\ n \geq 4,\ E_6$ or E_7. There are exactly n such X-heaps for A_n, three for D_n, two for E_6 and one for E_7.*

The resulting X-heaps are precisely the set of *minuscule posets* defined and studied in Proctor [4]. Let us illustrate what these minuscule posets look like. Note that each is a *coloured* poset, where the colours are the vertices of the top tree in each poset.

a) The case A_n. We label the minuscule A_n-heaps $F(A_n, k)$ $k = 1, \ldots, n$. Hopefully the following example will make the general case clear.

For $n = 5$

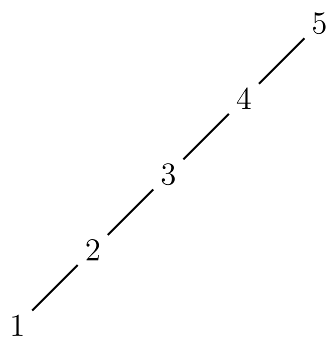

$F(A_5, 1)$

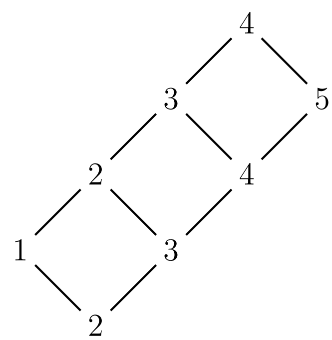

$F(A_5, 2)$

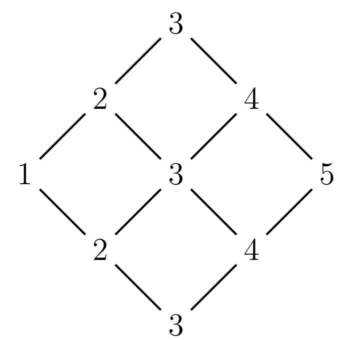

$F(A_5, 3)$

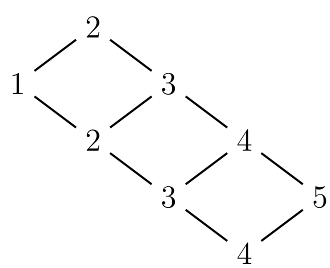

$F(A_5, 4)$

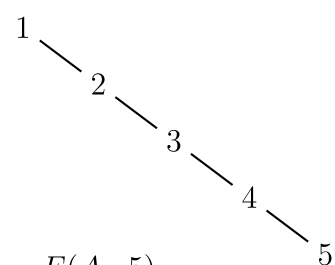

$F(A_5, 5)$

b) The case D_n. The minuscule D_n-heaps are labelled $F(D_n, 0),\ F(D_n, 1)$ and $F(D_n, n-1)$. The following example for $n = 5$ should make the general case clear.

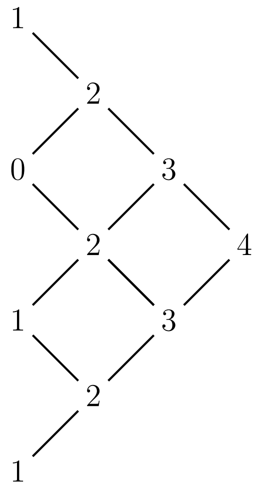

$F(D_5, 0)$

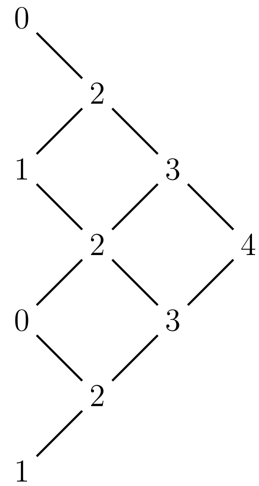

$F(D_5, 1)$

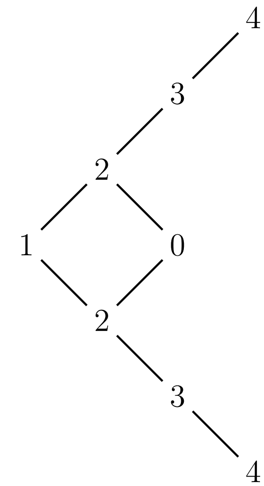

$F(D_5, 4)$

The heaps $F(D_n, 0)$ and $F(D_n, 1)$ have the same triangular shape with $n(n-1)/2$ elements, while $F(D_n, n-1)$ consists of a square symmetrically placed between two chains, and has $2(n-1)$ elements.

c) The case E_6. There are two minuscule E_6-heaps labelled $F(E_6, 1)$ and $F(E_6, 5)$. The heap $F(E_6, 1)$ appeared in Example 2. The heap $F(E_6, 5)$ has the same shape, and is the inverse of $F(E_6, 1)$.

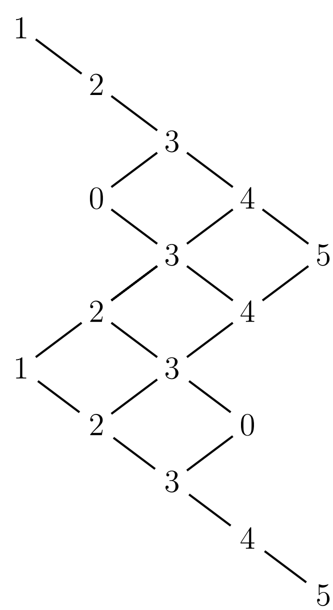

$F(E_6, 5)$

d) The case E_7. There is only one minuscule E_7 heap labelled $F(E_7, 6)$.

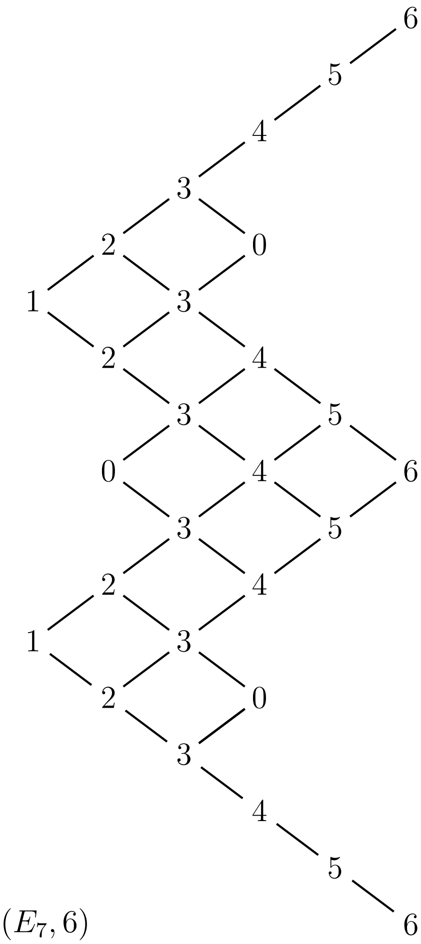

$F(E_7, 6)$

This lovely lattice, which we might call the *swallow*, is symmetric, spindle shaped, Sperner, Gaussian and enjoys other interesting combinatorial properties (see [7],[9],[12]).

Note that in each case the graph X is an ideal of the minuscule X-heap and that the minimal vertex appears in the label of that X-heap.

Proof of the Theorem. The proof will be broken down into several steps. We will show that the assumption on s implies that X must be a tree with no vertices of degree 4 or more and at most one vertex of degree 3. Then the possibilities for this latter case will be analysed by reducing it to study of triples of integers satisfying certain recursive properties. So let X and F be given as in the theorem and let s be some X-sequence with heap F.

Lemma 2.1 *X is a tree.*

Proof. Let s_0 be the sequence of initial occurrences of the vertices of X in s, and $P_{[s_0]}$ the associated heap, which clearly involves all vertices of X. Now if X has a cycle then the largest element of this cycle appearing in $P_{[s_0]}$ is necessarily preceded by two neighbours, so contradicting maximality. $\square$

Lemma 2.2 *X cannot have a vertex of degree 4 or more.*

Proof. Suppose X has a vertex e with neighbours a, b, c, d. Since each occurs in s, e must occur at least twice.

Between the first and second occurrences of e we can have at most 2 occurrences of neighbours of e – that means, say, that c and d do not occur. But then both c and d must occur before the first occurrence of e (if they didn't, we could add them, contradicting maximality) so we can add another e to the front of the sequence which is impossible. $\square$

Lemma 2.3 *X cannot have two vertices of degree 3.*

Proof. If X has at least two vertices of degree three then it has a subgraph Y of the following form

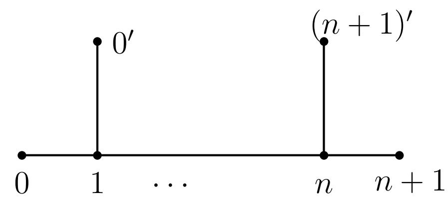

D_n

Consider the first occurrences in s of the vertices of the subgraph Y and the associated heap P_Y. If the occurrences of the vertices 1 and n are unrelated in P_Y then an easy argument shows that P_Y must have the following form for some k, $1 < k < n$.

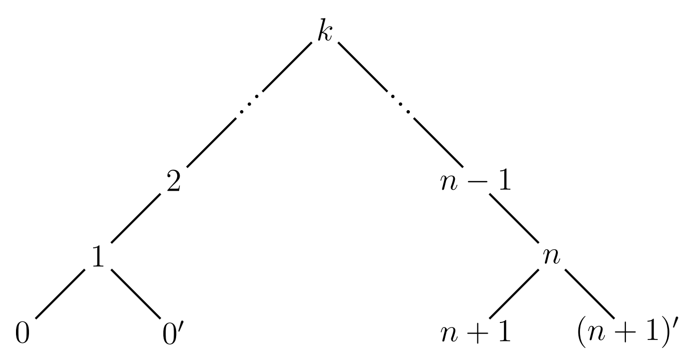

That means that the next occurrence of either 1 or n must precede the next occurrence of 2 or $n-1$, that then the next occurrence of 2 or $n-1$ must precede the next occurrence of 3 or $n-2$ etc. But that will imply that the next occurrence of k is preceded by more than two of its neighbours, a contradiction. $\square$

On the other hand if say $1 < n$ in P_Y then again an easy argument shows that the associated heap P_Y must have up to relabelling the following form.

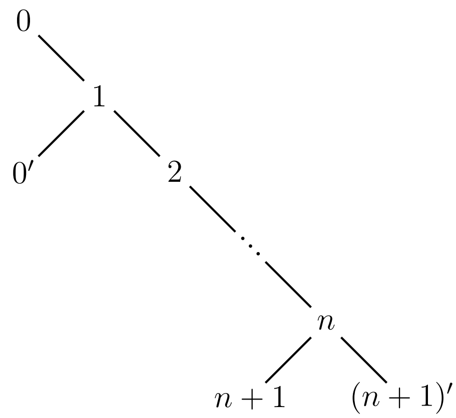

But then the next occurrence of n must precede the next occurrence of $n-1$, which must precede the next occurrence of $n-2$ and so on down to i, which is then preceded necessarily by three occurrences of neighbours of itself since its first occurrence, again a contradiction. $\square$

Now suppose that X has exactly one vertex, call it d, of degree 3, with chains of length $\alpha, \beta, \gamma > 0$ emanating from it, labelled $a_1, a_2, \ldots, a_\alpha$, $b_1, b_2, \ldots, b_\beta$ and $c_1, c_2, \ldots, c_\gamma$ as shown.

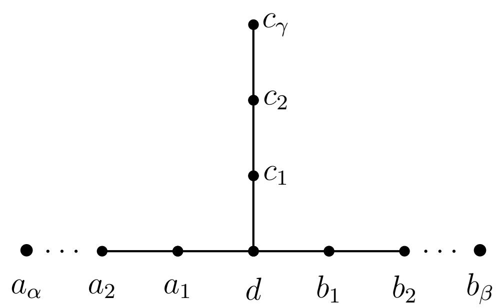

We imagine weighting the vertices linearly as follows

$$d > c_1 > c_2 > \cdots > c_\gamma > b_1 > b_2 > \ldots > b_\beta > a_1 > a_2 > \cdots > a_\alpha$$

and make the convention that wherever possible lighter elements move forward by free switches in a sequence s (and so upwards in the reverse Hasse diagram for $P_{[s]}$). In other words $a_i a_j$ is replaced by $a_j a_i$ if $i > j$ and $|i-j| \neq 1$, da_j is replaced by $a_j d$ if $j \neq 1$ (and similarly with b_i^s, c_i^s) and $b_i a_j$ is replaced by $a_j b_i$, etc. The weighting above then induces a partial order on elements of an X-string $[s]$ so that there is a unique minimal X-sequence t where no further free switches of the above type are possible.

Let us look in t at the successive occurrences of d and refer to the i^{th} interval of t as the segment following the i^{th} d and before the $(i+1)^{\text{st}}$ d (if it occurs), for $i = 1, \cdots, r$.

Lemma 2.4 *For any i, $1 \leq i \leq r$, there are non-negative integers α_i, β_i, γ_i such that the i^{th} interval has the form*

$$a_1 a_2 \, \dots \, a_{\alpha_i} \; b_1 b_2 \, \dots \, b_{\beta_i} \; c_1 c_2 \, \dots \, b_{\gamma_i}.$$

Proof. Since all the a_j can be freely switched with all the b_j and all the c_j and the b_j with the c_j, the fact that the a_j are lighter than the b_j which are lighter than the c_j means that the i^{th} interval will consist of a sequence of a_j followed by a sequence of b_j followed by a sequence of c_j with some of these sequences possibly empty.

The first a_j must be a_1, otherwise it would switch with d out of the i^{th} interval. The second a_j must be a_2 since it cannot be a_1 and any other a_j would freely switch to the left out of the interval. Continuing, we must start with a maximal sequence of a_j of the form $a_1 a_2 \, \dots \, a_{\alpha_i}$ for some $\alpha_i \leq \alpha$. But then the neighbourly condition ensures that no more α_j are possible. Since the b_j and c_j sequence are subject to the same analysis, the Lemma is proved. $\square$

Let us represent the sequence

$$a_1 a_2 \, \dots \, a_{\alpha_i}$$

by the shorthand symbol a^{α_i}.

Proposition 2.4 *If there are r intervals then t has the form*

$$t = \cdots \quad d_{(1)} a^{\alpha_1} b^{\beta_1} c^{\gamma_1} d_{(2)} a^{\alpha_2} b^{\beta_2} c^{\gamma_2} d_{(3)} \, \cdots \quad d_{(r)} a^{\alpha_r} b^{\beta_r} c^{\gamma_r},$$

where $d_{(k)}$ is the k^{th} occurrence of d and where the $\alpha_i, \beta_i, \gamma_i$ satisfy

1. *for $i = 1, \dots, r-1$ exactly one of $\alpha_i, \beta_i, \gamma_i$ is zero*
2. *for $i = r$ exactly two of $\alpha_i, \beta_i, \gamma_i$ is zero*
3. *if $\alpha_i > 0$ then $\alpha_{i+1} = \alpha_i - 1$ for $i = 1, \cdots, r-1$ (and similarly for β_i and γ_i)*
4. *if $\alpha_i = 0$ for $i = 1, \dots, r-1$ then $\alpha_{i+1} > 0$ (and similarly for β_i and γ_i).*

Proof. If there are r intervals then let us show that t cannot end in $d_{(r+1)}$. If two of $\alpha_r, \beta_r, \gamma_r$ were non-zero, say α_r and β_r, and there was an $(r+1)^{\text{st}}$ occurrence of d, then by maximality another c_1 could be added after this, contradicting the assumption of r intervals. This also proves 2). Statement 1) is a consequence of the two neighbourliness of t.

Lets prove 3). Suppose $\alpha_i > 0$ for some $i = 1, \dots, r-1$. Then $\alpha_{i+1} \geq \alpha_i$ is impossible since the element a_{α_i} in the i^{th} interval is then separated from the a_{α_i} in the $(i+1)^{\text{st}}$ interval by a single neighbour, namely $a_{\alpha_i - 1}$ if $\alpha_i > 1$ or d if $\alpha_i = 1$. Now if $\alpha_{i+1} < \alpha_i - 1$ then there must be a following occurrence (after the $(i+1)^{\text{st}}$ interval) of $a_{\alpha_{i+1}+1}$, since two neighbours of it have occurred. But when it does occur next it does so with $a_{\alpha_{i+1}}$ preceding it — meaning at least 3 neighbours between occurrences.

To prove 4), note that if $\alpha_i = 0$ and $\alpha_{i+1} = 0$ then 3 d's will have occurred between the previous a_1 and the following a_1. $\Box$

Without loss of generality we may assume that $\alpha_1 > 0$, $\beta_1 > 0$ and $\gamma_1 = 0$. This means there is necessarily by maximality an occurrence of c_1 before the first d.

Lemma 2.5 *The portion of t before the first occurrence of d is*

$$t = c_\gamma \, c_{\gamma-1}, \cdots c_1 d_{(1)} \, \cdots .$$

Proof. We first note that no a_j or b_j may precede $d_{(1)}$. Since c_1 does occur before $d_{(1)}$, neither a_1 of b_1 can otherwise we could add another occurrence of d to the beginning of the sequence. But then neither a_2 or b_2 can occur, because otherwise we could add an a_1 or b_1 before it, contradicting the previous statement. Continuing we obtain the claim.

To see that c_1 is necessarily immediately to the left of $d_{(1)}$, observe that any c_j, $j > 2$, is freely switched to the left of the c_1 occurrence immediately preceding $d_{(1)}$. If c_2 occurs between this c_1 and $d_{(1)}$ then since $\gamma_1 = 0$ (assumption) there are three neighbours of c_1 between it and the next occurrence of it after $d_{(2)}$, which is impossible. Similarly the next previous c_j must be c_2, then c_3 and so on. If as we proceed left from $d_{(1)}$ in t we find two occurrences of c_j then there must also be two occurrences of c_{j-1}, of c_{j-2}, and so on until two occurrences of c_1 mean another d can be added to the beginning, which is impossible. Thus t has the prescribed form. $\Box$

If we agree to write $c_\gamma \, c_{\gamma-1}, \cdots c_1$ as $c^{-\gamma}$ then we see that t has the form

$$t = c^{-\gamma} d_{(1)} a^{\alpha_1} b^{\beta_1} d_{(2)} a^{\alpha_2} b^{\beta_2} c^{\gamma_2} \, \cdots \, d_{(r)} a^{\alpha_r} b^{\beta_r} c^{\gamma_r},$$

where we now analyse the possibilities for the sequence of triples

$$(0, 0, -\gamma), \, (\alpha_1, \beta_1, 0), \, (\alpha_2, \beta_2, \gamma_2), \, \cdots \, (\alpha_r, \beta_r, \gamma_r).$$

We know $\alpha_1, \beta_1, \gamma_2 > 0$. Since at least one of $\alpha_2, \beta_2, \gamma_2$ is zero, without loss of generality we may assume that $\beta_2 = 0$ so that $\beta_1 = 1$ from 3) or 4) of Proposition. The above sequence of triples is then of the form

$$(0, 0, -\gamma)(\alpha_1, 1, 0)(\alpha_1 - 1, 0, \gamma_2) \, \cdots .$$

Lemma 2.6 $\beta = 1$.

Proof. If $\beta > 0$ consider first occurrence of b_2. It is then preceded by two b_1's, so we may add b_2 to the beginning of t contradicting the previous Lemma. $\square$

Suppose now that $r = 2$. Then since 2 of α_2, β_2, γ_2 are zero and γ_2 we know is not, we must have $\alpha_2 = 0$ so that $\alpha_1 = 1$. By maximality $\gamma_0 = \gamma_2 = \gamma$ and so the sequence of triples for t is

$$(0, 0, -\gamma),\ (1, 1, 0), (0, 0, \gamma).$$

This corresponds to $X = D_n$

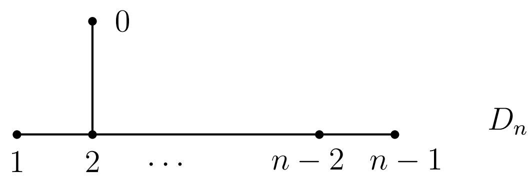

and the sequence

$$t = (n-1, n-2, \cdots 3, 2, 1, 0, 2, 3, \cdots, n-1).$$

In the case $n = 5$ the associated heap has the form

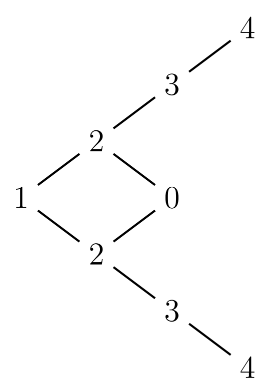

Suppose now that $r > 2$. Then exactly one of $(\alpha_2, \beta_2, \gamma_2) = (\alpha_1 - 1, 0, \gamma_2)$ is zero, so that $(\alpha_3, \beta_3, \gamma_3) = (\alpha_1 - 2, 1, \gamma_2 - 1)$. If $r = 3$ then both $\alpha_1 - 2$ and $\gamma_2 - 1$ must be 0, giving $\alpha_1 = 2$, $\gamma_2 = 1$ and the only possible maximal form for the sequence of triples being

$$(0, 0, -1)(2, 1, 0)(1, 0, 1)(0, 1, 0).$$

This corresponds to $X = D_5$ with sequence

$$t = (1\,2\,3\,4\,0\,2\,3\,1\,2\,0)$$

and heap

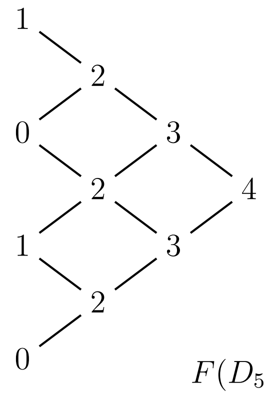

$F(D_5, 0)$

If $r > 3$ then exactly one of $(\alpha_3, \beta_3, \gamma_3) = (\alpha_1 - 2, 1, \gamma_2 - 1)$ is zero.

We consider the 2 cases $\alpha_1 = 2$ and $\gamma_2 = 1$ separately.

Case $\alpha_1 = 2$: If $\alpha_1 = 2$, $\gamma_2 > 1$ then the triple sequence for t must have the form

$$(0, 0, -\gamma)(2, 1, 0)(1, 0, \gamma_2)(0, 1, \gamma_2 - 1)(\alpha_4, 0, \gamma_2 - 2).$$

Now α_4 must be 2, since $\alpha_4 > 0$ by Proposition and if $\alpha_4 = 1$ then the next occurrence of a_2 (which must occur) will have (at least) three neighbours between it and the first, while if $\alpha_4 > 2$ then there ought to be an a_3 before $d_{(1)}$ which there is not. Thus the triple sequence for t is

$$(0, 0, -\gamma)(2, 1, 0)(1, 0, \gamma_2)(0, 1, \gamma_2 - 1)(2, 0, \gamma_2 - 2).$$

If $r = 4$ then $\gamma_2 = 2$ and we have

$$(0, 0, -2)(2, 1, 0)(1, 0, 2)(0, 1, 1)(2, 0, 0)$$

corresponding to $X = E_6$

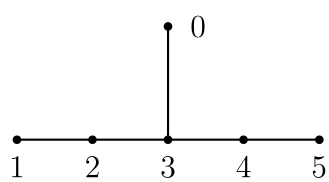

E_6

with

$$t = (1\,2\,3\,4\,5\,0\,3\,2\,1\,4\,3\,2\,0\,3\,4\,5)$$

.

The corresponding heap is one of the two minuscule posets for E_6.

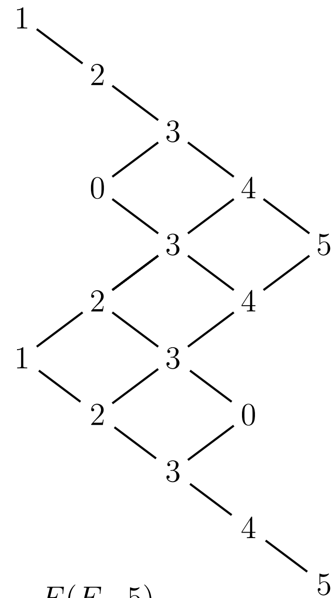

$F(E_6, 5)$

If $r > 4$ then $(\alpha_5, \beta_5, \gamma_5) = (1, 1, \gamma_2 - 3) = (1, 1, 0)$ which gives $\gamma_2 = 3 = \gamma$ and $(\alpha_6, \beta_6, \gamma_6) = (0, 0, 3)$ for maximality, yielding a final sequence

$$(0, 0, -3)(2, 1, 0)(1, 0, 3)(0, 1, 2)(2, 0, 1)(1, 1, 0)(0, 0, 3)$$

corresponding to $X = E_7$

with

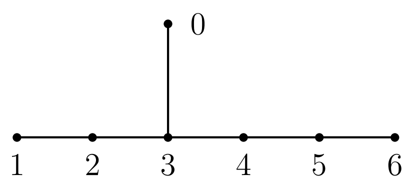

E_7

$$t = (6543210324563045321432034 56)$$

The corresponding heap is the unique minuscule poset for E_7, which we call the *swallow*. It is perhaps somewhat remarkable that this distributive lattice is precisely the lattice of order ideals in the either of the minuscule posets for E_6. This is part of a more general 'cascading' phenomenon which goes back to an observation of Steinberg noted and explained by Proctor in [P]. The minuscule posets for E_6 are themselves lattices of order ideals in the spin posets for D_5.

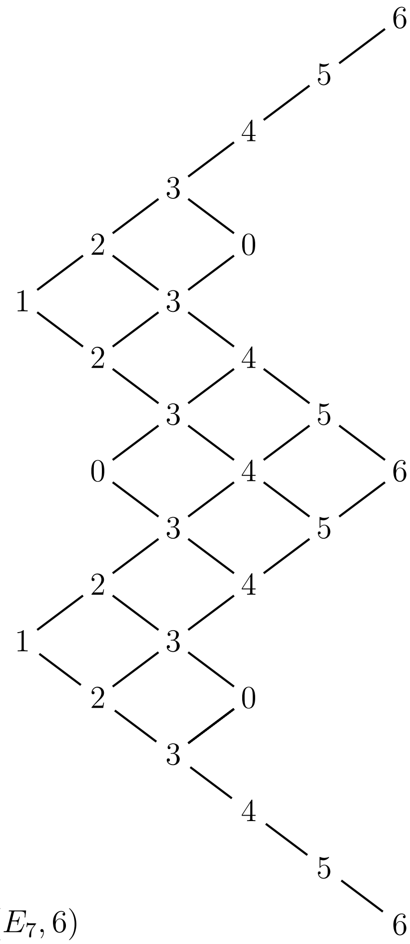

$F(E_7, 6)$

This completes the analysis of the case $\alpha_1 = 2$.

Case $\gamma_2 = 1$: We now examine the case $r > 3$ with $\gamma_2 = 1$ and triple sequence

$$(0, 0, -\gamma)(\alpha_1, 1, 0)(\alpha_1 - 1, 0, 1)(\alpha_1 - 2, 1, 0).$$

Then $\gamma = 1$ for if $\gamma > 1$ the first occurrence of c_2 must occur before $d_{(1)}$ by maximality (since we know c_1 occurs before $d_{(1)}$), while then the next occurrence follows at least three c_1's, which is impossible. Thus $\beta = \gamma = 1$ and the triple sequence must have the form

$$(0, 0, -1)(\alpha, 1, 0)(\alpha - 1, 0, 1)(\alpha - 2, 1, 0), \cdots, (0, 1, 0) \text{ or } (0, 0, 1)$$

depending on the parity of α. Thus $X = D_n$ and

$$t = 1\,2\,3\,4 \cdots n - 1\,0\,2\,3\,4 \cdots n - 2\,1\,2\,3\,0\,2\,1 \cdots 2\,3\,1\,2\,0.$$

These are the same kind of triangular heaps as the example of $F(D_5, 4)$ pictured earlier.

Finally suppose X has no vertices of degree 3 or more, and t begins with a vertex d which has two chains emanating from it as shown

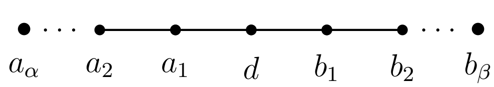

This is really a special case of the above where now $\gamma = 0$, and the same arguments show that t is of the form

$$t = d^{(1)}a^{\alpha_1}b^{\beta_1}d^{(2)}a^{\alpha_2}b^{\beta_2} \ \ldots \ d^{(r)}a^{\alpha_r}b^{\beta_r}.$$

Note that we have used the assumption that t begins with d. Now by neighbourliness, each $\alpha_i, \beta_i > 0$ for $i = 1, \cdots, r-1$, and since for $\alpha_i > 0$, $\alpha_{i+1} = \alpha_i - 1$ we see that the sequences $(\alpha_1, \alpha_2, \cdots, \alpha_r)$, $(\beta_1, \beta_2, \cdots, \beta_r)$ are decreasing incrementally and one must end it zero. It follows that $\alpha_1 = \alpha_1$ $\beta_1 = \beta$ and t is uniquely determined, namely

$$t = da^{\alpha}b^{\beta}da^{\alpha-1}b^{\beta-1}d \cdots da^{\max(0,|\alpha-\beta|)}b^{\max(0,|\alpha-\beta|)}.$$

Here for example is the case $\alpha = 3$, $\beta = 1$, corresponding to $X \simeq A_5$.

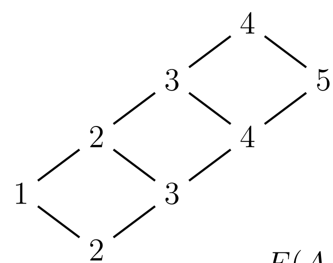

$F(A_5, 2)$

This completes the proof. $\square$

3 Connections

The heaps we have constructed are examples of coloured posets, since each vertex may be considered to be coloured by the corresponding vertex of the Coxeter graph. If we ignore the labels, these posets are just the irreducible 'minuscule' posets defined by Proctor in [4] and shown in figure 2 of Proctor [5]. As indicated in [4], these posets encapsulate the structure of some of the most important Bruhat orders on Weyl groups; in fact if an irreducible Bruhat poset is a lattice then either the Weyl group W is of type G_2 or the poset is isomorphic to the poset induced on the W-orbit of a minuscule weight with respect to the usual ordering of weights.

These posets play interesting roles in algebraic geometry and Lie theory, including describing the cohomology ring for minuscule flag manifolds including the Grassmanians. See for example Hiller [2] and Seshadri [7] for connections with the Schubert calculus of G/P where P is the stabilizer in a simple Lie group G of a maximal weight space in a minuscule representation.

Minuscule representations have the property that all weights are conjugate under the Weyl group. In this case, the geometry and order structure of this orbit of weights naturally determines much about the representation. All of the simply laced simple Lie algebras have minsucule representations with the sole exception of E_8 (which is why the latter does not appear in our main result). For connections with minuscule representations, see Stembridge [9], Parker and Rohrle [3], and recent work of Donnelly [1].

Some other combinatorial characterizations of minuscule posets appear on pp 344-345 of [4], including the fact that they constitute all known 'Gaussian' posets and that they are exactly the posets of join-irreducibles of the lattice of weights of minuscule representations of simple Lie algebras. It is also noted there that minscule posets are strongly Sperner, as well as being rank unimodal and rank symmetric.

More recently Proctor has shown that the minuscule posets are exactly the self-dual "d-complete" posets in [6]. Stembridge has found a new characterization of "coloured d-complete" posets which consists of (H1) and (H2) on p8 of [10]. In this language, the posets of this paper are those maximal amongst those satisfying (H1) and (H2*) which in addition satisfy (H2). Here (H2*) refers to having at least two elements whose labels are adjacent to i contained in every open subinterval between two elements labelled i.

References

[1] Donnelly, R. G., Solitary bases of Representations for semisimple Lie algebras, preprint (2000).

[2] Hiller H., Geometry of Coxeter groups , Research Notes in Mathematics, 54 Pitman(Advanced Publishing Program), Boston, Mass., 1982.

[3] Parker C. and Rohrle G., Minuscule Representations, preprint (1993).

[4] Proctor R. A., Bruhat Lattices, Plane Partition Generating Functions, and Minuscule Representations, Europ. J. Combinatorics 5, 331–350 (1984).

[5] Proctor R. A., A Dynkin diagram classification theorem arising from a combinatorial problem, Advances Math. 62, 103–117 (1986).

[6] Proctor R. A., Dynkin diagram classification of λ-minuscule Bruhat lattices and of d-complete posets, J. Algebraic Combin. 9, no.1, 61–94 (1999).

[7] Seshadri C. S., Geometry of G/P.I. Theory of standard monomials for minuscule representations, in C.P. Ramanujam –a tribute, 207–239, Tata Inst. Fund. Res. Studies in Math., 8, Springer, Berlin, 1978.

[8] Stembridge J. R., On the fully commutative elements of Coxeter groups, J. Alg. Combins. 5, no.4, 353–385, (1996).

[9] Stembridge J. R., On minuscule representations, plane partitions and involutions in complex Lie groups, Duke Math. J., no.2, 469–490 (1994).

[10] Stembridge J. R., Minuscule Elements of Weyl groups, preprint (1999).

[11] Viennot, G. X. Heaps of pieces I: Basic definitions and combinatorial lemmas, in 'Combinatoire Énumérative' (G. Labelle and P. Leroux, eds.), pp. 321–350, Lect. Notes in Math. Vol 1234, Springer–Verlag, Berlin, 1985.

[12] Wildberger N. J., A combinatorial construction for simply-laced Lie algebras, preprint (2000).