(Accepted for publication in the Journal of Algebra as of 7/13/98.)

Minuscule Elements of Weyl Groups, the Numbers Game, and d-Complete Posets

Robert A. Proctor*

Department of Mathematics
University of North Carolina
Chapel Hill, North Carolina 27599

Abstract. Certain posets associated to a restricted version of the numbers game of Mozes are shown to be distributive lattices. The posets of join irreducibles of these distributive lattices are characterized by a collection of local structural properties, which form the definition of d-complete poset. In representation theoretic language, the top "minuscule portions" of weight diagrams for integrable representations of simply laced Kac-Moody algebras are shown to be distributive lattices. These lattices form a certain family of intervals of weak Bruhat orders. These Bruhat lattices are useful in studying reduced decompositions of λ-minuscule elements of Weyl groups and their associated Schubert varieties. The d-complete posets have recently been proved to possess both the hook length and the jeu de taquin properties.

*Supported in part by NSA Grant No. MDA 904-95-H-1018.

Corrected Page Proofs Version, 1/14/99

1. Introduction

Except for some motivating comments, the only background needed to read this paper is familiarity with basic poset concepts. There are combinatorial motivations for this material which are independent of Lie theory. In fact, the "numbers game" viewpoint of this paper may be explained to lay people. Readers who are unfamiliar with Weyl groups should skip the next six paragraphs and resume reading with the description of the numbers game.

In the context of the highest weight theory of finite dimensional representations for simple Lie algebras, the simplest representations are the minuscule representations. A finite dimensional representation is *minuscule* if all of its weights lie in the Weyl group orbit of the highest weight. Although there are no minuscule representations for Kac-Moody algebras, the weight diagrams of many integrable Kac-Moody representations begin with small "minuscule portions."

This paper combinatorially characterizes the minuscule portions of the weight diagrams for dominant integral highest weight representations of simply laced Kac-Moody algebras. The "d-complete posets" which arise in this characterization have applications to representations, to reduced decompositions of Weyl group elements, and to two seemingly unrelated combinatorial issues.

Let λ be a dominant integral weight for a simply laced Kac-Moody algebra, and let W be its Weyl group. Dale Peterson defines $w \in W$ to be *λ-minuscule* if there exists some decomposition $s_{i_k} \cdots s_{i_1}$ of w such that $s_{i_j}(s_{i_{j-1}} \cdots s_{i_1}\lambda) = (s_{i_{j-1}} \cdots s_{i_1}\lambda) - \alpha_{i_j}$ for $1 \leq j \leq k$, where α_i is the simple root associated to s_i. The portion of the weight diagram for the representation with highest weight λ from λ to $w\lambda$ will be seen to have properties similar to those enjoyed by the entire weight diagram for a minuscule representation of a simple Lie algebra. Let L_w denote the poset of weights between λ and $w\lambda$ with respect to the reverse of the usual ordering by positive simple roots. It turns out that L_w is the same poset as the interval $[e,w]$ in the Bruhat and weak Bruhat orders on the set of cosets W/W_λ, where W_λ is the stabilizer of λ.

Our first theorem, Theorem A, states that L_w is a distributive lattice. In this paper we will call such lattices "wave lattices." (In title of the companion paper [Pr1] to this paper, these lattices were called "λ-minuscule Bruhat lattices.") Any finite distributive lattice L may be more succinctly described by its subposet P of join irreducible elements. If a poset P can arise as the poset of join irreducibles for a wave lattice L, then we will call it a "wave poset." The main result of this paper is a complete characterization of wave lattices in terms of the corresponding wave posets. The characterization of wave posets is given in terms of certain local structural properties. A poset P is said to be "d-complete" if it possesses this collection of local structural properties (Section 3). So the main result of this paper, Theorem B, may be stated as: A poset is a wave poset if and only if it is a d-complete poset. A concept closely related to the notion of d-complete poset is that of "colored d-complete poset." The weight generation process also produces a coloring of the elements of the wave poset P. Proposition 8.6 implies that the notions of d-complete and of colored d-complete are equivalent.

The notion of d-complete poset first arose as follows. Let $\mathfrak{g}$ be a simply laced Kac-Moody algebra. Let λ be a dominant integral weight, let w be λ-minuscule, and let P be the colored wave poset associated to w as above. In [Pr4], we produced a combinatorial analog of Lakshmibai's basis [Lak] for the Demazure modules of the positive Borel subalgebra of $\mathfrak{g}$ with lowest weights $m(w\lambda)$ for $m \geq 1$. This construction was stated independently of conventional Lie theory in the context of arbitrary colored posets. It was shown that a family of modules associated to a colored poset P possessed bases like those of Lakshmibai and Seshadri if and only if P was colored d-complete. The present paper completes the program of [Pr4] in the following sense: Theorem B guarantees that the methods of [Pr4] can be applied to all of the Demazure modules described above. A corollary to the methods of [Pr4] is a combinatorial proof of the Demazure character formula for this small family of Demazure modules.

Any λ-minuscule Weyl group element w has the "fully commutative" property introduced by Stembridge [Ste]. The wave/Bruhat lattices L_w indexed by such λ-minuscule w describe the Schubert subcell structure of certain Schubert varieties X_w. These Schubert varieties are the simplest Schubert varieties in a certain sense, i.e. they are the analogs in a certain sense of the Schubert varieties of a Grassmannian. Further Weyl group, representation theory, and geometric comments appear in Section 10.

The numbers game of Mozes [Moz] can be explained to any lay person. Let G be a finite simple graph: no loops or multiple edges are allowed. Here is the restricted version of this one player game which will be considered in this paper. The game begins at an initial state which consists of an assignment of integer labels to the nodes of G. One move consists of the following: Choose a node d which currently has a $+1$ label. Change that label to -1 and add $+1$ to the labels of each of the nodes adjacent to d. We say that we have "fired" the node d. From this new state, one is again allowed to perform any such move. We organize the various tallies of "node firings" which can arise from a sequence of moves away from the fixed initial state into a partially ordered set. The general version of our first result states that any principal order ideal of this poset is a distributive lattice. For our main result, we require that all of the initial labels be non-negative. The distributive lattices which arise in that case are just the wave distributive lattices mentioned three paragraphs above. That paragraph describes the main result of this paper, Theorem B, and should be read by all readers at this point.

The posets which are produced by Theorem B, the d-complete posets, form a class of posets which are of interest for combinatorial reasons. Shapes (Ferrers diagrams) and shifted shapes are diagrams upon which tableaux and plane partitions are formed. Such diagrams may be viewed as posets. Tableaux and plane partitions on shapes and on shifted shapes have attracted much attention over the last 25 years or so. Is there a more general class of posets in which some of the nice properties enjoyed by shapes and shifted shapes continue to hold? The class of d-complete posets is a good generalization with respect to two such properties.

In [Pr1] we give an explicit Dynkin diagram classification of all possible d-complete posets. There it is shown that any connected d-complete poset has a globally tree-like structure such that each of the local "slant irreducible" components falls into one of 15 possible classes. Each slant irreducible component is indexed by a Dynkin diagram which is embedded in the bottom of its order diagram. For 14 of 15 classes, the Dynkin diagrams are of "general type E." This classification may be regarded as a classification of simply laced λ-minuscule Weyl group elements or of the "simplest" Schubert varieties. (As is explained in Section 3, the definition of d-complete poset used in the present paper is the order dual of the definition used in [Pr1] and in the other papers of this series.)

The first two of the 15 classes of irreducible d-complete posets are the classes of shapes and of shifted shapes. Rooted tree posets are shown in [Pr1] to form a family of d-complete posets which are trivial in a certain sense. These posets are defined in Section 3 below, as are the members of a fourth family of d-complete posets, the "double-tailed diamonds." These four families of d-complete posets were the only infinite families of posets known to have the "hook length" property [Sag]: A poset P is said to be "hook length" if its associated P-partition generating function factors in a certain nice fashion analogous to identities discovered by Euler and Stanley. The first three of these families were the only infinite families of posets known to have the "jeu de taquin" property: A poset P is said to be jeu de taquin if the result of playing Schützenberger's sliding game is independent of the order in which the "empty labels" are slid out. In 1994 we conjectured that any d-complete poset possesses both the hook length and jeu de taquin properties. After this paper was refereed, Dale Peterson and we proved the hook length conjecture by combining facts from algebraic geometry and representation theory with the wave viewpoint of this paper. A corollary to this result is a generalization of the hook product formula for the number of standard Young tableaux on an ordinary shape to a product formula for the number of order extensions of any d-complete poset. Given the relationships described in Section 10, this corollary can be viewed as a conversion of Dale Peterson's hook formula for the number of reduced decompositions of a λ-minuscule element into a combinatorial form analogous to the original Frame-Robinson-Thrall form. More recently, we have proved the jeu de taquin conjecture as well, by modifying and then extending Kimmo Eriksson's proof of the jeu de taquin property for shapes to the other 14 classes of irreducible d-complete posets.

There are three or four differences between the numbers game process of Mozes and our process. Relative to [Moz], we are multiplying all labels by -1 in order to agree with the traditional Lie theoretic approach employed in [Pr2]. Mozes allowed real numbers instead of integers. More importantly, relative to our sign convention, Mozes required that only one label for the initial state be positive. For our main result, we require that all initial labels be non-negative. Most importantly, Mozes's process allows one move for each node with a positive label x: This label is replaced by $-x$, and x is added to the labels of each of the neighboring nodes. We require $x = +1$ to make a move based at that node. The general numbers game has also been studied in [AKP] [Bjö] [Eri].

If one forms a square matrix from the set of vectors corresponding to our moves at the various nodes, one obtains the "discretized Laplacian" operator for the graph G. Visualize a portion of G which is simply a chain. Suppose that one node v has $+1$ assigned to it and that the other nodes have been assigned 0. View this situation as a guitar string which has been pulled up at v but not yet released. The states locally attainable from this initial situation can be visualized as snapshots of pulses propagating away from v as waves after the string is released. This square matrix is the generalized Cartan matrix for the simply laced Kac-Moody algebra whose Dynkin diagram is G. Executing a move based at the ith node corresponds to subtracting the positive simple root α_i. Further details relating aspects of the numbers game to aspects of representation theory are given in Section 10.

In [Pr2] we used the (as yet unnamed) numbers game to determine which of the Bruhat orders on parabolic quotients W^J of finite Weyl groups were distributive lattices. The answer was that the weight diagrams of the minuscule representations constituted almost all such "Bruhat lattices." We defined a poset P to be *minuscule* if it was the poset of join irreducibles for the weight diagram of a minuscule representation. Ideals of these posets are the simple Lie algebra antecedents for the "wave posets" studied in this paper in the context of simply laced Kac-Moody algebras. The minuscule posets were pictured in Figure 2 of [Pr3]. In [Pr1] we show that a d-complete poset is minuscule if and only if it is order self-dual. In [Pr2], with R. Stanley, we showed that every minuscule poset is "Gaussian" [Sta, p. 288].

The two main results of this paper are stated at the end of Section 2, after our restricted version of the numbers game and the poset of tallies are defined. Section 3 presents the definition of d-complete poset. In Section 4, we obtain the distributive lattice result in the general context where some of the initial labels may be negative. Section 5 constructs the posets of join irreducibles for those distributive lattices. Beginning with Section 6, we assume that all of the initial labels are non-negative. Sections 6 and 7 prove that wave posets are d-complete. Sections 8 and 9 prove that any d-complete poset can arise as a wave poset.

2. A Restricted Numbers Game

Readers who are familiar with Lie representations or Weyl groups should read Section 10 in tandem with this section. Let G be a simple graph with node set N. Let $\Lambda := \mathbb{Z}^N$. Elements of Λ are *wave states*. For each $d \in N$, let $\alpha_d \in \Lambda$ be defined by $\alpha_d := (a_b)_{b \in N}$

with $a_d = +2$, $a_c = -1$ when c is adjacent to d, and $a_b = 0$ otherwise. For each $d \in N$, the operator $\mathcal{S}_d$ is partially defined on Λ by $\mathcal{S}_d\mu = \mu - \alpha_d$ if $\mu = (m_b)_{b\in N}$ with $m_d = +1$. (If $m_d \neq 1$, then $\mathcal{S}_d\mu$ is undefined.) If $m_d = +1$, applying $\mathcal{S}_d$ to μ *fires* node d. Let $\lambda = (l_b)_{b\in N}$ be a given fixed element of Λ; it is the *initial state*. Our restricted numbers game begins with this initial state: The various play sequences of the game consist of applying the operators $\mathcal{S}_d$ to λ and its successors, provided that the actions are defined at each step. For example, Figure 1 presents a graph G (which is the Dynkin diagram E_6). Figure 2 shows some states generated by beginning to play the restricted numbers game starting with a particular λ. In the labels appearing in this figure, the entries are arranged to correspond with the geometry of G as it is depicted in Figure 1. The symbols "1," "0," and "–" indicate component values of states of +1, 0, and –1 respectively. The initial state λ is at the bottom of Figure 2; it has a +1 at the leftmost node of G and 0's elsewhere. The result of firing the +1 is shown above λ. Thirteen later states produced with various choices of firings are shown higher up in Figure 2.

For some choices of G and λ, it is possible for this game to "loop" indefinitely. For example, if G is the affine Dynkin diagram $E_6^{(1)}$ [Kac, p. 54] with its central node numbered 0, let $\lambda = (0,1,0,-1,0;0,0)$. Then $\mathcal{S}_{1'}\mathcal{S}_0\mathcal{S}_{1''}\mathcal{S}_{2''}\mathcal{S}_{2'}\mathcal{S}_{1'}\mathcal{S}_0\mathcal{S}_{1''}\mathcal{S}_1\mathcal{S}_0\mathcal{S}_2\mathcal{S}_1\lambda = \lambda - \xi$, where $\xi = 1\alpha_2 + 2\alpha_1 + 3\alpha_0 + 2\alpha_{1'} + 1\alpha_{2'} + 2\alpha_{1''} + 1\alpha_{2''}$. But note that $\xi = 0$ in Λ. Hence $\mathcal{S}_{1'}\mathcal{S}_0\mathcal{S}_{1''}\mathcal{S}_{2''}\mathcal{S}_{2'}\mathcal{S}_{1'}\mathcal{S}_0\mathcal{S}_{1''}\mathcal{S}_1\mathcal{S}_0\mathcal{S}_2\mathcal{S}_1\lambda = \lambda$. (This example was inspired by Figure 7 of [Ste] and considerations from [Pr1].)

We would like to introduce a partial ordering on the set of wave states produced by the numbers game beginning with the initial state λ such that the earliest states generated occur at the lowest elements of the poset. The example just presented indicates that this is impossible in general. It turns out that it *is* possible when all of the initial labels are non-negative. However, we can show this only after a considerable amount of work. The additional objects introduced next, "tallies," can immediately be organized into a poset. Although we will continue to refer to wave states, the tallies will become the primary objects for stating and deriving our results.

Continue to work in the context of a fixed simple graph G with a given fixed initial state $\lambda \in \Lambda$. Now we track the number of times that each node is fired. Let $\Xi := \mathbb{N}^N$. Elements of

Ξ are called *tallies*. Denote the *initial tally* $(0,0,\ldots,0) \in \Xi$ by θ. The *wave state function* Ψ: $\Xi \to \Lambda$ is defined by $\Psi(\pi) = \lambda - \sum_{b\in N} p_b\alpha_b$, if $\pi = (p_b)_{b\in N} \in \Xi$. Note that $\Psi(\theta) = \lambda$, the initial state. An operator s_d will act on Ξ in parallel with the action of $\mathcal{S}_d$ on Λ to record the firings of the node d. Let $\varepsilon_d \in \Xi$ be defined by $\varepsilon_d := (e_b)_{b\in N}$ with $e_d = +1$ and $e_b = 0$ elsewhere. For each $d \in N$, the operator s_d is partially defined on Ξ by $s_d\pi = \pi + \varepsilon_d$ if $\Psi(\pi) = (m_b)_{b\in N}$ with $m_d = +1$. (If $m_d \neq 1$, then $s_d\tau$ is undefined.) Let $\pi \in \Xi$ be such that $\Psi(\pi) = \mu = (m_b)_{b\in N}$ with $m_d = +1$. Then $\mathcal{S}_d\mu$ is defined; let $\nu := \mathcal{S}_d\mu$. Note that $\Psi(s_d\pi) = \Psi(\pi+\varepsilon_d) = \nu$: By adding +1 to the d*th* component of π, the operator s_d is tallying the firing of the d*th* node in passing from μ to ν.

We recursively build up a subset Ω of Ξ beginning with $\Omega := \{\theta\}$. If $\pi \in \Omega$ and $s_d\pi = \rho$ for some $d \in N$, then adjoin ρ to Ω. Figure 3 shows 14 tallies generated by working up from θ in Ω for the example shown in Figure 2. It is best to think of Figure 3 as the primary object; now Figure 2 merely depicts the wave states $\Psi(\pi)$ corresponding to the tallies π. From now on we assume that for the given (G,λ), the (possibly infinite) set Ω of all possible tallies has been generated. Suppose that $\pi, \rho \in \Omega$, with $\pi = (p_b)_{b\in N}$ and $\rho = (r_b)_{b\in N}$. We define a partial order on Ω by: $\pi \leq \rho$ if $p_b \leq r_b$ for every $b \in N$. Let $\zeta \in \Omega$. Define L_ζ to be the subposet of Ω consisting of all $\pi \in \Omega$ such that $\pi \leq \zeta$. The topmost tally shown in Figure 3 is one possible choice for a $\zeta \in \Omega$ in this example; the rest of Figure 3 gives the order diagram for the poset L_ζ.

A lattice L is *distributive* if for every $x, y, z \in L$, the identities $x\wedge(y\vee z) = (x\wedge y)\vee(x\wedge z)$ and $x\vee(y\wedge z) = (x\vee y)\wedge(x\vee z)$ hold. Let P be a poset. A subset $I \subseteq P$ is an *ideal* if it is closed below: $y \in I$, $x \leq y \Rightarrow x \in I$. Given P, the poset J(P) is *the poset of ideals of* P ordered by containment. It is always a distributive lattice. Let L be a distributive lattice. An element $x \in L$ is *join irreducible* if it covers exactly one other element. The poset shown in Figure 3 is a distributive lattice. The symbol "#" has been placed above each of its join irreducible elements. Given a distributive lattice L, let j(L) denote its *poset of join irreducibles*. The poset j(L) for the lattice L appearing in Figure 3 is displayed separately in Figure 4. The fundamental theorem of finite distributive lattices [Sta, p. 106] states that if P is a finite poset and L is a finite distributive lattice, then $j(J(P)) \cong P$ and $J(j(L)) \cong L$. So a finite distributive lattice L is completely determined by its poset of join irreducibles P.

Assuming the definition of "d-complete poset" which is given in the next section, we can now state our main results.

Theorem A. *Let* G *be a simple graph with node set* N *and let* $\lambda \in \mathbb{Z}^N$. *Let* Ω *be the associated poset of tallies. Then for any* $\zeta \in \Omega$, *the poset* L_ζ *is a distributive lattice.*

Theorem B. *Let* G *be a simple graph with node set* N *and let* $\lambda \in \mathbb{N}^N$. *Let* Ω *be the associated poset of tallies. A poset* P *may arise as* $j(L_\zeta)$ *for some such* G *and* λ *and some* $\zeta \in \Omega$ *if and only if it is d-complete.*

Although we will not use the following reduction, we mention it for the sake of exposition. There is no point in allowing initial labels which are +2 or larger, since those nodes of G would just be ignored by the restricted numbers game. So for Theorem B, every initial label can be assumed to be 0 or +1. Firing one of the initial +1's will enable adjacent nodes to be fired, provided that their initial labels were 0. Any new +1 which arises during a firing sequence can have its heritage traced back to a particular original +1. If a wave propagating outwardly from one original +1 collides with a wave propagating outwardly from another original +1, a +2 label will be created and that node will henceforth be inactive. Fix some $\zeta \in \Omega$. If we want to study only L_ζ, it can be seen that we may as well replace G at the outset with a disjoint union of rooted trees in which each root has the initial label +1 and all other nodes have initial label 0. (Each node which appears at least once in ζ can be unambiguously assigned to one of the original +1 nodes according to the heritage of the first wave to reach it. The "unambiguous" claim can be verified by considering the set of unfired nodes in ζ; this will be independent of the chain chosen to reach ζ from θ.) Since everything in this paper is well-behaved under disjoint union, one might as well restrict attention to the (G,λ) cases in which G is a rooted tree and λ has only a +1 at the root of G.

The paper [Pr1] takes up where Theorem B leaves off by classifying all d-complete posets. As a consequence of that classification, it can be seen that there is an upper bound on the length of any firing sequence for a given finite G and initial $\lambda \in \mathbb{N}^N$. This is in contrast to the example above which had $\lambda \in \mathbb{Z}^N$. We do not know of any overlap between the finiteness result just stated and the various terminating numbers games results of [AKP] [Bjö] [Eri] [Moz].

3. Definitions of d-Complete and Colored d-Complete Posets

In this paper, all *posets* (i.e. partially ordered sets) are assumed to be finite. Let P be a poset. The order-dual poset P^* is defined on the same set as P, but with order defined by $x \leq y$ in P^* if $x \geq y$ in P. The definitions presented in this section are the order duals of the analogous definitions presented in the other papers of this series. In those other papers, a connected d-complete poset has a unique maximal element. This agrees with the majority of historical precedents: weight diagrams of representations, and ordinary and shifted tableaux for the hook length and jeu de taquin properties. In this paper only, connected d-complete posets will have unique minimal elements. We could not bear to define the order on the tallies in Ω by $\rho \leq \pi$ if $r_b \geq p_b$ for every $b \in N$. Staying with the maximal convention would have required considering the "meet irreducibles" of L_ζ rather than the traditional "join irreducibles." Finally, the convention used in this paper means that the lattices L_ζ will be intervals in weak Bruhat orders, rather than their order duals, since the convention for Bruhat orders is to start with the identity as the unique minimal element.

Let P be a poset. If x is covered by y in P, then we write $x \to y$. A *chain* is a subset of P of the form $x_0 \to x_1 \to \dots \to x_n$. If $w,z \in P$, then the *interval* $[w,z] := \{x \in P: w \leq x \leq z\}$. Also, $(w,z] := [w,z]-\{w\}$. And $[w,z)$ and (w,z) are analogously defined. Given $x \in P$, the *principal ideal* (x) is defined to be $\{y \in P: y \leq x\}$. Given $x_1,\dots,x_m \in P$, the *ideal* $(x_1,\dots x_m)$ generated by these elements is defined to be the union of the ideals (x_i) for $1 \leq i \leq m$. From the context at hand, it will be clear whether (x,y) refers to an doubly open interval or to an ideal generated by two elements.

A subset $\{w,x,y,z\}$ of P is a *diamond* if z covers x and y, and each of x and y cover w. The poset $d_3(1)$ is the four element poset consisting of one diamond $\{w,x,y,z\}$.

The poset $d_4(1)$ is the six element poset formed by adjoining one element above the maximal element of $d_3(1)$ and one element below the minimal element of $d_3(1)$. Generalizing, for $k \geq 3$, the *double-tailed diamond* poset $d_k(1)$ has $2k-2$ elements, of which two are incomparable elements in a middle rank and $k-2$ apiece form chains above and below the two incomparable elements. (In [Pr3], the poset $d_k(1)$ was defined to be the poset of join irreducibles of the Bruhat lattice $D_k(1)$ for the ω_1 representation of the simple Lie algebra of type D_k.)

An interval $[w,z]$ in P is a d_3-*interval* if it is a diamond $\{w,x,y,z\}$ for some x and y, or in other words, if $[w,z] \cong d_3(1)$. More generally, for $k \geq 3$, we say that an interval $[w,z]$ is a d_k-*interval* if is isomorphic to $d_k(1)$. A d_3^--*interval* $[x,y;z]$ consists of three elements x, y, and z such that z covers both x and y. For $k \geq 4$, we say that an interval $[x,z]$ is a d_k^--*interval* if is isomorphic to $d_k(1)-\{b\}$, where b is the minimal element of $d_k(1)$.

Suppose that $[x,y;z]$ is a d_3^--interval in a poset P. If there is no $w \in P$ such that $\{w,x,y,z\}$ is a d_3-interval, then $[x,y;z]$ is an *incomplete* d_3^--interval. If there exists $z' \neq z$ such that $[x,y;z']$ is also a d_3^--interval, then we say that $[x,y;z]$ and $[x,y;z']$ *overlap*. A poset P is *d_3-complete* if it contains no incomplete d_3^--intervals, if the minimal element of each d_3-interval is not covered by any elements outside of that interval, and if it contains no overlapping d_3^--intervals. We have just required:

(D1) Anytime two elements x and y are covered by a third element z, there must exist a fourth element w which is covered by each of x and y,

(D2) If $\{w,x,y,z\}$ is a diamond in P, then w is covered by only x and y in P, and

(D3) No two elements z and z' can cover each of two other elements x and y.

Let $k \geq 4$. Suppose $[x,z]$ is a d_k^--interval in which y is the unique element covered by z. If there is no $w \in P$ covered by x such that $[w,z]$ is a d_k-interval, then $[x,z]$ is an *incomplete* d_k^--interval. If there exists $z' \neq z$ covering y such that $[x,z']$ is also a d_k^--interval, then we say that $[x,z]$ and $[x,z']$ *overlap*. For any $k \geq 4$, a poset P is *d_k-complete* if:

(D4) There are no incomplete d_k^--intervals,

(D5) If $[w,z]$ is a d_k-interval, then w is covered by only one element in P, and

(D6) There are no overlapping d_k^--intervals.

A poset P is d-*complete* if it is d_k-complete for every $k \geq 3$.

A *colored poset* (P,κ) consists of a poset together with a coloring map κ to a set of colors N. It is *properly colored* if no two incomparable elements are colored the same and no element is colored the same as an element it covers. It is *simply colored* if in addition: whenever an interval is a chain, the colors of the elements in that interval are distinct. For any $k \geq 3$, a colored poset (P,κ) is *colored* d_k*-complete* if it is simply colored and d_k-complete in such a way that $\kappa(w) = \kappa(z)$ in every invocation of D1 or D4. It is *colored d-complete* if it is colored d_k-complete for every $k \geq 3$. The colored definitions could be simplified by dropping conditions D3 and D6: Proposition 8.1 states that D1 (or D4) together with simple coloring implies D3 (or D6). Trivially, any d-complete colored poset (P,κ) becomes a d-complete poset when the coloring function κ is ignored. Proposition 8.6 states that any d-complete poset can be colored so that it becomes colored d-complete, and that there is essentially only one coloring which will work.

A *rooted tree* is a connected poset with a unique minimal element in which every other element covers at most one element. Rooted trees are obviously d-complete.

4. Proof of Theorem A

Let G be a fixed simple graph with node set N and let λ be a fixed element of $\Lambda = \mathbb{Z}^N$. Let Ω be the corresponding poset of tallies defined in Section 2.

For elements $\mu \in \Lambda$ and $\pi \in \Omega$, we will write $\mu = (m_b)$ and $\pi = (p_b)$ instead of $\mu = (m_b)_{b\in N}$ and $\pi = (p_b)_{b\in N}$. Let $d \in N$. If $\pi \in \Omega$ is $\pi = (p_b)$ with $\Psi(\pi) = \mu = (m_b)$, we will define $C_d(\pi)$ to be the d*th* coefficient of $\Psi(\pi)$: namely, $C_d(\pi) := m_d$. Within proofs, we will write $a_q \ldots a_1$ instead of $s_{a_q} \ldots s_{a_1}$.

Lemma 4.1. *Let* $\pi,\rho \in \Omega$. *If* $\pi \leq \rho$, *then there exist* $c_1,\ldots,c_r \in N$ *such that* $\rho = s_{c_r} \ldots s_{c_1} \pi$.

Proof. Let $\pi = a_p \ldots a_1\theta$ and $\rho = b_q \ldots b_1\theta$. Let t be maximal such that $a_1 = b_1, \ldots, a_t = b_t$. If $t = p$, we are done. Suppose that $t < p$. Note that $\mathcal{C}_{a_{t+1}}(a_t \ldots a_1\theta) = +1$. Since $\pi \leq \rho$, there must be some $i > t+1$ such that $b_i = a_{t+1}$. Letting h be the minimal such i, we have $b_j \neq a_{t+1}$ for $h > j \geq t+1$. Viewing the b's as a's, it is obvious that $b_h(b_t \ldots b_1\theta)$ is defined. We claim that $\sigma = b_{h-1}b_{h-2} \ldots b_{t+1}b_hb_t \ldots b_1\theta$ is defined. Since $b_h \notin \{b_{t+1}, \ldots, b_{h-1}\}$, the decrease of $\mathcal{C}_{b_h}(b_t \ldots b_1\theta)$ caused by firing b_h cannot adversely effect the original program of firing $b_{t+1}, \ldots, b_{h-1}$ after reaching $b_t \ldots b_1\lambda$. It is conceivable that one of $b_{t+1}, \ldots, b_{h-1}$, say b_j, is adjacent to b_h and that firing b_h earlier could increase the value of $\mathcal{C}_{b_j}(b_{j-1} \ldots b_1\theta)$ to $+2$ for some $h-1 \geq j \geq t+1$, and thereby interfere with that original program. But $\mathcal{C}_{b_h}(b_t \ldots b_1\theta) = \mathcal{C}_{a_{t+1}}(a_t \ldots a_1\theta) = +1$. If b_j is adjacent to b_h, then firing j would increase $\mathcal{C}_{b_h}(b_j \ldots b_1\theta)$ to $+2$ before b_h in its original position was reached. Then $b_q \ldots b_1\theta$ would not be defined. Hence b_h is not adjacent to any of $b_{t+1}, \ldots, b_{h-1}$, and so σ is defined. Clearly $\sigma = b_hb_{h-1}b_{h-2} \ldots b_{t+1}b_t \ldots b_1\theta$, and so $\rho = b_q \ldots b_{h+1}b_{h-1} \ldots b_{t+1}b_hb_t \ldots b_1\theta$ with $b_1 = a_1, \ldots, b_t = a_t, b_h = a_{t+1}$. By induction on t we can eventually reach ρ with a sequence of operators which passes through π. ■

Lemma 4.2. *Let* $\pi, \rho \in \Omega$. *Then* ρ *covers* π *in* Ω *if and only if there exists some* $d \in N$ *such that* $\rho = s_d\pi$.

Proof. One direction is immediate from the definition of $\leq$ on Ω and the other is immediate from Lemma 4.1. ■

Now we are ready to prove Theorem A. We continue to work in the context of our fixed G, N, λ, etc. Extend the partial order on Ω to all of Ξ. If $\pi \in \Xi$ is $\pi = (p_b)$, then we define the *height* of π by $ht(\pi) := \sum_{b \in N} p_b$.

Proof of Theorem A. Let $\zeta = (z_b)$ be a fixed tally in Ω. Then L_ζ is the interval $[\theta,\zeta]$ in Ω. Let $\pi = (p_b)$ and $\rho = (r_b)$ denote two elements of L_ζ. For each $b \in N$, let $g_b := \min(p_b,r_b)$ and $h_b := \max(p_b,r_b)$. Define two elements of Ξ as follows: the *weight meet* φ of π and ρ is defined by $\varphi := (g_b)$ and the *weight join* ψ of π and ρ is defined by $\psi := (h_b)$. In order to show that L_ζ is a lattice, we need only show that φ and ψ are always elements of Ω. In fact, distributivity will also be known at that point: The operations of componentwise max and componentwise min always satisfy the distributive axioms.

Note that $\psi \leq \zeta$. We inductively prove the following statements for $k \geq 0$:

(i) If $ht(\varphi) = k$, then $\varphi \in \Omega$.

(ii) If $ht(\varphi) = k$, then $\psi \in \Omega$.

Suppose $k = 0$. Then $\varphi = \theta \in \Omega$. The proof of (ii) for $k = 0$ is obtained by taking $k = 0$ in the second paragraph below.

Now suppose $k > 0$. Here we show $\varphi \in \Omega$ by assuming that (i) and (ii) have been shown for smaller values $k' < k$. Let τ be maximal in L_ζ such that $\tau \leq \varphi$. If $\tau = \varphi$, we are done. Suppose $\tau < \varphi$. Since $\tau \leq \pi$ and $\tau \leq \rho$, Lemma 4.1 implies that there exist $a_1, \ldots, a_p$ such that $a_p \ldots a_1 \tau = \pi$ and $b_1, \ldots, b_q$ such that $b_q \ldots b_1 \tau = \rho$. Let i be minimal such that $\tau + \varepsilon_{a_i} \leq \varphi$. Let j be minimal such that $b_j = a_i$. (Such a j must exist, since $\varepsilon_{a_i} \leq \varphi - \tau \leq \rho - \tau$.) Maximality of τ implies both that $i \neq 1$ and that $\mathcal{C}_{a_i}(\tau) \neq +1$. Note that $\mathcal{C}_{a_i}(b_{j-1} \ldots b_1 \tau) = +1$, which implies that the actions of $b_1, \ldots, b_{j-1}$ add at least 1 to $\mathcal{C}_{a_i}(\tau)$ when passing from τ to $b_{j-1} \ldots b_1 \tau$. By the minimality of i and the definition of φ, the intersection $\{a_1, \ldots, a_{i-1}\} \cap \{b_1, \ldots, b_{j-1}\}$ is empty. Hence the meet of $a_{i-1} \ldots a_1 \tau$ and $b_{j-1} \ldots b_1 \tau$ is τ. And $ht(\tau) < k$. Apply the induction assumption to conclude that the join γ of $a_{i-1} \ldots a_1 \tau$ and $b_{j-1} \ldots b_1 \tau$ is in Ω. The proof of the existence of the join indicates that $\gamma = b_{j-1} \ldots b_1 a_{i-1} \ldots a_1 \tau$. Note that $\mathcal{C}_{a_i}(a_{i-1} \ldots a_1 \tau) = +1$. Since a_j does not occur in $b_1, \ldots, b_{j-1}$, by the statement above concerning the actions of $b_1, \ldots, b_{j-1}$ we deduce that $\mathcal{C}_{a_i}(\gamma) \geq +2$. But $\gamma \leq \psi \leq \zeta$, and so Lemma 4.1 implies that there exist $d_1, \ldots, d_r$ such that $\zeta = d_r \ldots d_1 \gamma$. Now $\tau \leq \varphi \leq \zeta$ implies that at least one ε_{a_i} is added when passing from τ to ζ. Note that none of the operators in the sequence $b_{j-1} \ldots b_1 a_{i-1} \ldots a_1$ is equal to a_i. Hence at least one of $d_1, \ldots, d_r$ must be equal to a_i. But $\mathcal{C}_{a_i}(\gamma) \geq +2$ implies that none of $d_1, \ldots, d_r$ can equal a_i. This contradiction implies that $\tau = \varphi$.

Now we show the kth instance of (ii) using the kth instance of (i); i.e. that $\varphi \in \Omega$ implies $\psi \in \Omega$. Since $\varphi \leq \pi$ and $\varphi \leq \rho$, by Lemma 4.1 we know that there exist $a_1, \ldots, a_p$ such that $\pi = a_p \ldots a_1 \varphi$ and $b_1, \ldots, b_q$ such that $\rho = b_q \ldots b_1 \varphi$. So $\mathcal{C}_{b_1}(\varphi) = +1$. We claim that $b_1(a_p\ldots a_1\varphi)$ is defined. Since φ was the meet of π and ρ, we know that $b_1 \notin \{a_1,\ldots,a_p\}$. So $\mathcal{C}_{b_1}(\pi)$ is not decreased by firing any of $a_1, \ldots, a_p$. It could be increased to $+2$ by firing one of these, if one was adjacent to b_1. Since $\pi \leq \zeta$, Lemma 4.1 implies that ζ can be reached from π. Note that $\pi + \varepsilon_{b_1} \leq \zeta$. Hence b_1 must be fired at some point in passing from π to ζ. Therefore $\mathcal{C}_{b_1}(\varphi)$ cannot be increased to $+2$ by firing $a_1, \ldots, a_p$ in succession. So $b_1(a_p\ldots a_1\varphi)$ is defined. Note that the meet of $b_1a_p\ldots a_1\varphi$ and ρ is $\varphi + \varepsilon_{b_1}$, which is just $b_1\varphi \in \Omega$. Set $\varphi^{(1)} := b_1\varphi$ and $\pi^{(1)} := b_1a_p\ldots a_1\varphi$. Note that $\varphi^{(1)} \leq \pi^{(1)} \leq \zeta$ and $\varphi^{(1)} \leq \rho$. Therefore the reasoning above can be repeated until we have attained $\pi^{(q)} = \psi$, implying that $\psi \in \Omega$. ■

5. The Posets P_ζ of Join Irreducibles of L_ζ

Here we continue to work in the environment of Section 4 described by G, N, $\lambda \in \Lambda = \mathbb{Z}^N$, Ω, ζ, and L_ζ. We now identify the join irreducible elements of the distributive lattice L_ζ and derive some basic facts concerning the subposet P_ζ of L_ζ consisting of those elements.

Let $d \in N$. Suppose that $\pi = (p_b) \in \Xi$ is such that $\mathcal{C}_d(\pi) = -1$ and $p_d > 0$. Then define $r_d\pi := \pi - \varepsilon_d$. If $\mathcal{C}_d(\pi) \neq -1$ or if $p_d = 0$, then leave $r_d\pi$ undefined. Note that if $\Psi(\pi) = \mu$, then $\Psi(r_d\pi) = \mu + \alpha_d$. Hence $\mathcal{C}_d(r_d\pi) = +1$. Also note that if $s_d\tau$ is defined, then $r_ds_d\tau = \tau$, and if $r_d\pi$ is defined, then $s_dr_d\pi = \pi$.

Lemma 5.1. *If* $\pi \in \Omega$ *and* $r_d\pi$ *is defined, then* $r_d\pi \in \Omega$. *So* π *covers* τ *in* Ω *if and only if there exists some* $d \in N$ *such that* $r_d\pi = \tau$.

Proof. Let $\pi = b_k\ldots b_1\theta = (p_b)$. Since $r_d\pi$ is defined, we have $\mathcal{C}_d(\pi) = -1$ and $p_d > 0$. Let j be the largest i such that $b_i = d$. Consider $b_k\ldots b_{j+1}b_{j-1}\ldots b_1\theta$. We have $\mathcal{C}_d(b_j\ldots b_1\theta) = -1$. Since $\mathcal{C}_d(\pi) = -1$ and no $b_i = d$ for $i \geq j+1$, none of $b_{j+1}, \ldots, b_k$ can be adjacent to d. Hence firing d at stage j cannot affect the firability of any of $b_{j+1}, \ldots, b_k$. Hence $b_k\ldots b_{j+1}b_{j-1}\ldots b_1\theta$ is defined. It is clearly equal to $r_d\pi$, and so $r_d\pi \in \Omega$. The second statement now follows from Lemma 4.2. ■

Given $d \in N$ and $m \geq 0$, define $\eta(d,m)$ to be the meet of all $\pi = (p_b)$ in L_ζ such that

$p_d = m$. The empty meet in L_ζ is θ. Clearly $\eta(d,0) = \theta$ for any $d \in N$. We say that $\eta(d,m)$ is non-trivial if $\eta(d,m) \neq \theta$.

Lemma 5.2. *An element of* L_ζ *is join irreducible if and only if it is a non-trivial* $\eta(d,m)$ *for some* $d \in N$ *and some* $m \geq 1$. *If* $\eta(d,m)$ *is non-trivial, then* $r_d\eta(d,m)$ *is the unique element in* L_ζ *which it covers.*

Proof: Let $d \in N$ and $m \geq 1$ be such that $\eta(d,m) =: \pi = (p_b)$ is non-trivial. Since $\pi \neq \theta$, it must cover something in L_ζ. If $\pi = s_c\tau = (t_b)$ for some $c \neq d$, then $t_d = p_d = m$, contradicting the definition of $\eta(d,m)$. So we must have $\pi = s_d\tau$ for some $\tau \in \Omega$. Hence $\tau = r_d\pi$ is the unique element in L_ζ covered by $\eta(d,m)$.

Suppose $\pi = (p_b) \in L_\zeta$ covers only one element. Let d be the unique node such that $\pi = s_d\rho$ for some $\rho \in \Omega$. Let $m = p_d$ and set $\tau := \eta(d,m) = (t_b)$. So $t_d = m$. Clearly $\tau \leq \pi$. Suppose that $\tau < \pi$. Then $\pi = s_{a_q}\ldots s_{a_1}\tau$ for some $a_1, \ldots, a_q$. None of the a_i's can be d, since $p_d = t_d = m$. And $r_{a_q}\pi = s_{i_{q-1}}\ldots s_{i_1}\tau$ is defined, a contradiction since $a_q \neq d$. Hence $\pi = \tau = \eta(d,m)$. ■

Let P be an arbitrary finite poset and consider the distributive lattice $L := J(P)$. Each covering relation in L describes augmenting an ideal $I \subseteq P$ by an element $x \in P$ such that $I \cup \{x\} \subseteq P$ is also an ideal. Coloring the elements of P induces a coloring of the edges of L. (As T. Brylawski notes, a poset P is properly colored if and only if no two incident covering edges in the the order diagram of $J(P)$ have the same induced color.) Let L be a distributive lattice and set $P := j(L)$. The standard isomorphism map φ from $J(P)$ to L takes each principal ideal (y) to the join irreducible element $y \in L$.

Returning to the context established at the beginning of this section, we now define P_ζ to be the *poset of join irreducibles of* L_ζ. So P_ζ is the poset formed by putting the induced order from L_ζ on the set of elements π of L_ζ of the form $\eta(d,m) \neq \theta$ for some $d \in N$ and some $m \geq 1$. In Sections 5 - 8, the symbol $\eta(d,m)$ will almost always denote an element of P_ζ. (But it must be a non-trivial $\eta(d,m)$ when viewed as an element of L_ζ in order to be an element of P_ζ.)

We now want to transfer properties of elements of L_ζ (and their covering relations) to the corresponding ideals of P (and maximal containments thereof). We will now begin to often refer to elements of N as *colors*. If ρ covers π in L_ζ, then $\rho = s_d\pi$ for some color d. Assign the color $d \in N$ to this covering relation in L_ζ. Also define a *coloring function* κ: $P_\zeta \to N$ as follows: If $x \in P_\zeta$ is $\eta(d,m)$ for some $m \geq 1$, then define $\kappa(x) := d \in N$.

Proposition 5.3. *The coloring of the edges of* L_ζ *induced by the isomorphism* $L_\zeta \cong J(P_\zeta)$ *from the coloring of* P_ζ *agrees with the coloring of the edges of* L_ζ *defined by the operators* s_b.

Proof: Let π be a join irreducible element of L_ζ and let $x = \eta(d,m)$ be the corresponding element of P_ζ. Let I be the ideal (x) and let H be the ideal $(x)-\{x\}$. Let $\tau = (t_b) := \varphi[H]$ be the element of L_ζ corresponding to H. Note that $\varphi[I] = \varphi[(x)] = \pi =: (p_b)$. Since $\varphi[I]$ covers $\varphi[H]$, Lemma 5.2 implies that $\tau = r_d\pi$. Hence $\pi = s_d\tau$. Either way, the edge in L_ζ from $\varphi[H]$ to $\varphi[I]$ is colored by d.

More generally, the ideals of P_ζ which can be augmented by x to induce a coloring of an edge in $J(P_\zeta) \cong L_\zeta$ by d are precisely the ideals J containing H but not containing x. For such an ideal J, let $\rho = (r_b) := \varphi[J]$. In L_ζ we have $\rho \geq \tau$ and $\rho \not\geq \pi$. Hence $r_b \geq t_b$ for all $b \in N$. And $p_b = t_b$ when $b \neq d$ with $p_d = m$ and $t_d = m-1$. We can conclude that $r_d = m-1$. Now augmenting J by x corresponds under φ to taking the join of ρ with π, which corresponds to increasing r_d by 1. Letting $\sigma := \rho \vee \pi$ in L_ζ, we have by Lemma 4.1 that $\sigma = s_d\rho$. So the operator coloring of this edge in L_ζ is d, which agrees with the coloring induced by adjoining x. Since every edge of L_ζ is produced by some such augmentation, we are done. ■

Given an ideal $I \subseteq P_\zeta$, define the *color count* $\Phi[I]$ to be (u_b), where u_b is the number of elements $x \in I$ of color b. Any I can be built up from $\varnothing$ one colored element at a time, and any such building up will correspond to a chain in $J(P_\zeta)$ from $\varnothing$ consisting of colored edges. Let $\pi = (p_b) := \varphi[I] \in L_\zeta$. By Proposition 5.3, for each $d \in N$, the number p_d of operators s_d needed to reach π in Ω is equal to u_d. In other words, $(u_b) \equiv \Phi[I] = \varphi[I] \equiv (p_b)$. So given P_ζ, the lattice $L_\zeta \subseteq \Omega$ can be reconstructed by forming the color counts of all ideals in P_ζ and then ordering by componentwise comparison. Summarizing this section, we have:

Proposition 5.4. *Let* $\zeta = (z_b)$. *Let* $I \subseteq P_\zeta$ *be an ideal and let* $\pi = \Phi[I] = (p_b) = \varphi[I] \in L_\zeta$. *Then* π *is covered by an element* $s_d\pi$ *in* L_ζ *if and only if there exists a minimal element* x *in* $P_\zeta - I$ *of color* d *if and only if* $\mathcal{C}_d(\varphi[I]) = +1$ *and* $p_d < z_d$. *Then* $\varphi[I \cup \{x\}] = s_d\pi$. *And* π *covers an element* $r_d\pi$ *in* L_ζ *if and only if there exists a maximal element* x *in* I *of color* d *if and only if* $\mathcal{C}_d(\varphi[I]) = -1$ *and* $p_d > 0$. *Then* $\varphi[I - \{x\}] = r_d\pi$.

If P is a poset, then an *order extension of* P is an order preserving bijection ξ: $P \to \{1 < 2 < \ldots < |P|\}$. The preceding result implies that to each order extension ξ of P_ζ there corresponds one chain (or firing sequence) in L_ζ from θ to ζ. If, following Section 10, one prefers to think of ζ as a simply laced λ-minuscule element w of W, then the order extensions of the colored wave poset P_ζ correspond to the reduced decompositions of w. Hence:

Corollary 5.5. *Let* W *be a simply laced general Weyl group and let* λ *be a dominant integral weight. Let* w *be* λ*-minuscule. Then the number of reduced decompositions of* w *is equal to the number of extensions of the associated wave poset* P_ζ *to a total order.*

6. First Properties of Wave Posets

Let G be a simple graph with node set N. From now on we will assume that λ is a fixed element of $\mathbb{N}^N \subseteq \Lambda = \mathbb{Z}^N$; i.e. that all of the initial integer labels are non-negative. Let Ω, ζ, L_ζ, and P_ζ be as before. Any distributive lattice which arises as a L_ζ for some G and some non-negative $\lambda \in \mathbb{N}^N$ will be called a *wave lattice*. Any poset which arises as a $P_\zeta = j(L_\zeta)$ for some non-negative $\lambda \in \mathbb{N}^N$ will be called a *wave poset*. In this section we establish several facts concerning the local coloring structure of wave posets.

Lemma 6.1. *Let* $d \in N$ *and* $m > 1$ *be such that* $\eta(d,m)$ *is non-trivial. Then* $\eta(d,m-1) < \eta(d,m)$ *in* P_ζ. *Therefore all elements of a given color in* P_ζ *are comparable.*

Proof. Let $\pi = \eta(d,m-1)$ and $\rho = \eta(d,m)$ in L_ζ. Let $\tau := r_d\rho =: (t_b)$ as in Lemma 5.2. Then $t_d = m-1$, and so τ enters into the meet defining π. Hence $\pi \le \tau < \rho$ in L_ζ, and so $\pi < \rho$ in P_ζ. ■

When $\lambda \in \mathbb{N}^N$, we have $\mathcal{C}_b(\pi) \geq -1$ for every $\pi \in \Omega$ and every $b \in N$. If we know that $\ldots s_d \ldots \pi$ is defined, then $\mathcal{C}_d(\pi) < +2$, implying that $\mathcal{C}_d(\pi) \in \{-1,0,+1\}$. Any -1 can have arisen only from a firing of a $+1$. So if $\pi = (p_b)$, then $\mathcal{C}_d(\pi) = -1$ implies $p_d > 0$. Revisiting Proposition 5.4 with $\lambda \in \mathbb{N}^N$, we see that the requirement $p_d > 0$ can be dropped since it is now implied by $\mathcal{C}_d(\pi) = -1$. Hence an important observation is that the join irreducible elements π of L_ζ are those for which $\mathcal{C}_b(\pi) = -1$ for exactly one $b \in N$.

Lemma 6.2. *Let* $\pi \in \Omega$. *Then* $\Psi(\pi)$ *cannot have* -1*'s at adjacent nodes of* G.

Proof. Suppose there are -1's in $\Psi(\pi)$ at adjacent nodes c and d. Suppose that $\pi = a_n \ldots a_1 \theta$, with $a_q = c$ being (without loss of generality) the leftmost occurrence of c or d. Then $-1 = \mathcal{C}_d(a_n \ldots a_1 \theta) = \mathcal{C}_d(a_q \ldots a_1 \theta) = \mathcal{C}_d(a_{q-1} \ldots a_1 \theta) + 1$. But this is impossible, since $\mathcal{C}_d(a_{q-1} \ldots a_1 \theta) \geq -1$ when $\lambda \in \mathbb{N}^N$. ■

Lemma 6.3. *Let* $x,y \in P_\zeta$. *If* $\kappa(x)$ *is next to* $\kappa(y)$ *in* G, *then either* $x < y$ *or* $y < x$ *in* P_ζ.

Proof. If x and y are incomparable, they would be maximal elements of the ideal (x,y). Then by Proposition 5.4 there would be -1's at both $\kappa(x)$ and $\kappa(y)$ in $\Psi(\varphi[(x,y)])$, contradicting Lemma 6.2. ■

Proposition 6.4. *Let* x *be covered by* y *in* P_ζ, *with* $c := \kappa(x)$ *and* $d := \kappa(y)$. *Let* $\pi := \varphi[(x)]$ *and* $\rho := \varphi[(y)]$ *be the corresponding elements of* L_ζ. *There exists a unique chain in* L_ζ *from* π *up to* ρ. *If this chain arises from* $\rho = s_{b_1} \ldots s_{b_n} \pi$, *then* $b_n, \ldots, b_1$ *form a path in* G *consisting of distinct successively adjacent nodes. Here* $b_1 = d$ *is adjacent to* c, *and* $c \neq d$. *None of the other* b_i*'s are equal to* c *or adjacent to* c.

Proof. Since $\pi < \rho$ in L_ζ, we have $\rho = a_q \ldots a_1 \pi$ for some $a_1, \ldots, a_q \in N$. This corresponds in P_ζ to successively adjoining elements to (x) to produce (y). Since x is covered by y, all but the last of these elements (which is y itself) must be incomparable to x. By Lemma 6.1, we see that $a_i \neq c$ for $i < q$. By Lemma 6.3, we see that a_i is not adjacent to c for $i < q$. Since ρ is join irreducible, we have $a_q = d$. Let b_1 be the rightmost a_i which is adjacent to c: There must be at least one such a_i, or else there would be two or more -1's in $\Psi(a_q \ldots a_1\pi)$. If there was a $+1$ at b_1 in $\Psi(\pi)$, then stop. Otherwise, let b_2 be the righmost a_i to the right of b_1 in $a_q \ldots a_1\pi$ which is adjacent to b_1. There must be at least one such a_i somewhere in $a_1, \ldots, a_q$ before b_1, or else no $+1$ at b_1 would be created. If there was a $+1$ at b_2 in $\Psi(\pi)$, then stop. Repeat until some b_n is produced which has a $+1$ in $\Psi(\pi)$. There must be such a b_n, since $a_q \ldots a_1\pi$ is defined. The only b_j which could be equal to c would be b_1, if $b_1 = a_q$. But b_1 was chosen to be adjacent to c. Hence $b_j \neq c$ for $1 \leq j \leq n$. Then there must be 0's at each of $b_1, \ldots, b_{n-1}$ in $\Psi(\pi)$, since the only -1 label in $\Psi(\pi)$ was at c. Certainly $b_n\pi$ is defined. For $h \geq 1$, the definition of b_{h+1} implies that it is the only b_j in $b_{h+1} \ldots b_n\pi$ which is adjacent to b_h. Hence $\mathcal{C}_{b_h}(b_{h+2} \ldots b_n\pi) = 0$ and so $\mathcal{C}_{b_h}(b_{h+1} \ldots b_n\pi) = +1$, implying that $b_h \ldots b_n\pi$ is defined if $b_{h+1} \ldots b_n\pi$ is defined. So $b_1 \ldots b_n\pi$ is defined. Clearly $\Psi(b_1 \ldots b_n\pi)$ has only one -1 at b_1. So $b_1 \ldots b_n\pi$ is join irreducible in L_ζ. Since $\pi < b_1 \ldots b_n\pi \leq \rho$ and ρ is a minimal join irreducible above π, we must have $b_1 \ldots b_n\pi = \rho$. Hence $b_1 = a_q, \ldots, b_n = a_1$. The b_j's are distinct by definition, and b_j is adjacent only to b_{j+1} and b_{j-1} for $2 \leq j \leq m-1$. The node b_n is the unique b_j with a $+1$ in $\Psi(\pi)$. We know $b_1 = a_q = d$. It is clear that $b_n, \ldots, b_1$ is the unique sequence such that $\rho = s_{b_1} \ldots s_{b_n}\pi$, since $\pi - \rho$ determines the number of times each node of N will be used to reach ρ from π. We chose $b_1 = d$ to be adjacent to c, so $c \neq d$. ■

Corollary 6.5. *Let* P_ζ *be a wave poset. If* y *covers* x *in* P_ζ, *then the unique* -1 *in* $\Psi(\varphi[(y)])$ *is adjacent to the unique* -1 *in* $\Psi(\varphi[(x)])$. *Corresponding to any chain* $x = x_o \to x_1 \to \ldots \to x_n = y$ *in* P_ζ, *there is a path in* N *traced out by the unique* -1*'s associated to* $(x_o), (x_1), \ldots$.

An immediate consequence of Lemma 6.1 and Proposition 6.4 is:

Corollary 6.6. *Wave posets* P_ζ *are properly colored.*

Lemma 6.7. *Let* P_ζ *be a wave poset and let* $z \in P_\zeta$. *Then* z *cannot cover three or more distinct elements in* P_ζ.

Proof. Suppose that z covers w, x, and y and let the corresponding colors be d, c_1, c_2, and c_3. Since the w, x, and y are incomparable and P_ζ is properly colored, we see that d, c_1, c_2, and c_3 are distinct. Each of c_1, c_2, and c_3 is adjacent to d. Set $\pi := \varphi[(w,x,y)-\{w,x,y\}]$ and $\rho := \varphi[(w,x,y)]$. Note that $\upsilon = s_{c_3}s_{c_2}s_{c_1}\tau$. Hence the value at d in the wave picture for (w,x,y) is at least $+2$. But the node d must be fired in order to reach (z) from (w,x,y), which is now impossible. So z cannot cover three distinct elements. ■

A *cycle* in a simple graph G with node set N is a subset $C \subseteq N$ such that every element of C is adjacent to exactly two other elements of C. We say that G is *acylic* if it contains no cycles. Given $\zeta = (z_b) \in \Omega$, define the *support* of ζ to be $N_\zeta := \{d \in N\colon\ z_d > 0\}$. Let G_ζ denote the subgraph of G induced by N_ζ.

Lemma 6.8. *Let* $\zeta \in \Omega$. *Then* G_ζ *is acyclic.*

Proof. Let $\zeta = a_n \ldots a_1\theta$. Suppose that $C \subseteq N_\zeta$ is a cycle. Let d be the last node in C to be fired for the first time in $a_n \ldots a_1\theta$, and suppose that j is minimal such that $a_j = d$. Initially $\mathcal{C}_d(\theta) \geq 0$, and two neighbors of d will be fired before d is fired for the first time. Hence $\mathcal{C}_d(a_{j-1} \ldots a_1\theta) \geq +2$, implying that $a_j(a_{j-1} \ldots a_1\theta)$ is undefined. So there is no cycle $C \subseteq N_\zeta$. ■

Proposition 6.9. *Wave posets* P_ζ *are simply colored.*

Proof. Let $x,y \in P_\zeta$ be such that the interval $[x,y]$ is a chain. Denote this chain by $x = x_0 \to x_1 \to \ldots \to x_n = y$, and set $a_i := \kappa(x_i)$ for $1 \leq i \leq n$. Let $J := (y)$ and $I := J - \{x_1,\ldots,x_n\}$ be ideals in P. Set $\pi := \varphi[I]$ and $\rho := \varphi[J]$. Then $\rho = a_n \ldots a_1\pi$ in L_ζ. By Corollary 6.5, the a_i's trace out a path in G_ζ. Let $q < r$ be such that $a_q = a_r$ and such that there are no repeated colors in $a_{q+1}, \ldots, a_{r-1}$ or other occurrences of a_q. In order to recharge the -1 produced at a_q, there must be two distinct colors in $a_{q+1}, \ldots, a_{r-1}$ which are adjacent to a_q. These must be a_{q+1} and a_{r-1}. So we would have a cycle in G_ζ, which is impossible. So the colors occuring in this chain must be distinct. ■

Lemma 6.10. *Let* P_ζ *be a wave poset. Suppose that* $x < y$ *in* P_ζ *with* $\kappa(x) =: c$ *adjacent to* $\kappa(y) =: d$ *in* G. *Let* $(p_b) := \varphi[(x)]$ *and* $(r_b) := \varphi[(y)]$. *If* $p_c = r_c$, *then* y *covers* x *in* P_ζ.

Proof. The unique -1 is at c in $\varphi[(x)]$ and at d in $\Psi(\varphi[(y)])$. Suppose that y does not cover x; let $x = x_0 \to x_1 \to \dots \to x_n = y$ be a chain in P_ζ from x to y. By Corollary 6.5, the unique -1 will trace a corresponding path in G_ζ from c to d. If this path does not pass through c again after leaving it, then we would observe a cycle in G_ζ. But G_ζ is acylic. And the node c cannot be fired after leaving $\varphi[(x)]$ since $p_c = r_c$. So y must cover x. ■

7. Wave Posets are Colored d_k-Complete

We continue to work in the environment established at the beginning of Section 6. We begin by confirming that any wave poset P_ζ satisfies the axiom D1 for colored d_3-completeness.

Proposition 7.1. *Suppose that in a wave poset* P_ζ *there exists some* z *which covers two elements* x *and* y. *Then there exists a fourth element* $w \in P_\zeta$ *which is covered by both* x *and* y, *and in fact* $\kappa(w) = \kappa(z)$.

Proof: Let $\kappa(x) = b$ and $\kappa(y) = c$. By Lemma 6.7, we have $(z)-\{z\} = (x,y)$. Suppose $z = \eta(d,m)$ for some $d \in N$ and some $m \geq 1$. Let $(r_a) := \varphi[(x)]$. Proposition 6.4 implies that $r_d = m - 1$ and that b and c are each adjacent to d. Let $\upsilon = (u_a) := \varphi[(x,y)-\{x,y\}]$ in L_ζ. In passing from υ to $\varphi[(x,y)]$, the nodes b and c are fired once apiece without d being fired at all. Hence $u_d = -1$, or else the value at d would become too large to fire to reach $\varphi[(z)]$ from $\varphi[(z)-\{z\}]$. So d was fired at least once before υ, and so $m \geq 2$. Set $w := \eta(d,m-1)$. Obviously $x > w$ and $y > w$. Note that $r_d = m-1$. Lemma 6.10 implies that x covers w, and similarly it can be seen that y covers w. ■

Proposition 7.2. *Let* P_ζ *be a wave poset. If* $w,z \in P_\zeta$ *are such that* $w = \eta(d,m-1)$ *and* $z = \eta(d,m)$ *for some* $m \geq 2$, *then* $[w,z] \cong d_k(1)$ *for some* $k \geq 3$.

Proof. Let $w = x_0 \to x_1 \to \dots \to x_n = z$ be a chain in P_ζ from w to z of maximal length. By Proposition 6.4, we must have $n \geq 2$. Suppose that $n = 2$. By simple coloring, there must be another element y besides x_1 such that $x < y < z$. By the maximality of n, we have

$x \to y \to z$. By Lemma 6.7, we have $[w,z] \cong d_3(1)$.

Now suppose that we have verified this proposition for $2 \le n' < n$. By Corollary 6.5, the unique -1 traces out a path in N_ζ starting at d and ending at d as we observe $\Psi(\varphi[(x_0)])$, $\Psi(\varphi[(x_1)])$, ... in L_ζ. Since G_ζ is acyclic, we must have $\kappa(x_1) = \kappa(x_{n-1}) =: c$. Assume $n > 2$. Choose a firing sequence in L_ζ from $\varphi[(w))]$ to $\varphi[(z)]$ which passes through each $\varphi[(x_i)]$. Clearly the node d can be fired only as the very last step when reaching $\varphi[(z)]$ in L_ζ. There cannot be any firings at c beyond the arrivals at x_1 and x_{n-1}, or else the wave coefficient at d would be raised too high. Hence $x_1 = \eta(c,q)$ and $x_{n-1} = \eta(c,q+1)$ for some $q \ge 1$, and we can apply our induction hypothesis to conclude that $[x_1,x_{n-1}] \cong d_k(1)$ for some $k \ge 3$. Now the open interval (x_0,x_n) can contain more than $[x_1,x_{n-1}]$ only if there is either some $u \ne x_1$ covering $w = x_0$ such that $u < z$ or some $v \ne x_{n-1}$ covered by $z = x_n$ such that $v > w$. In either case, there would be a node $b \ne c$ adjacent to d in N which would have to be fired at least once when passing from $\varphi[(w)]$ to $\varphi[(z)]$ in L_ζ. With the two firings at c, this would raise $\mathcal{C}_d(\varphi[I])$ to at least $+2$ by the time we reached $I = (z)-\{z\}$. So $[w,z] \cong d_{k+1}(1)$. ■

Proposition 7.3. *Wave posets* P_ζ *satisfy the D4 colored* d_k*-complete axiom for any* $k \ge 4$.

Proof. The proof will be by induction on k, with the $k = 3$ case (Proposition 7.1) serving as the base case. Assume that $k \ge 4$ and that the result has been shown for $k' < k$.

Suppose that $x,z \in P_\zeta$ are such that $[x,z]$ is a d_k^--interval. Suppose $z = \sigma(d,m)$ for some $d \in N$ and some $m \ge 1$. Let y be the element covered by z which is in this interval. Let $c := \kappa(y)$. The colors c and d are adjacent in G. By induction, axiom D4 is satisfied at $k-1$ (or the axiom D1 when $k-1 = 3$), and so $[x,y] \cong d_{k-1}(1)$ implies that $\kappa(x) = c$. Induction and Lemma 8.2 below imply that the color c appears nowhere else in $[x,y]$. Set $\pi = (p_b) := \varphi[(x)-\{x\}]$. In passing from π to $\varphi[(z)-\{z\}]$, the node c is fired twice: the elements of color c form a totally ordered set, and so any other occurrences of c in this ideal build-up would have to occur in $[x,y]$. Accepting for now that it exists (i.e. that $m \ge 2$), set $w := \eta(d,m-1)$. We have $w < z$. Since d is adjacent to c, either $w > y$ or $w < y$. The former is impossible since z covers y. So $w < y$. And $w > x$ or $w < x$. Suppose $w > x$. Then w

$\in$ [x,y]. If w was a middle or higher element in [x,y], then [w,z] would be a chain and the simple coloring condition would be violated. If w were below the middle in [x,y], then a "sister" element above the middle would violate the simple coloring condition with z. So w < x, if it exists. So there are no elements of color d when passing from π to $\varphi[(z)]$, except for the very last step. This means that $p_d = -1$ to start, or else the value at d would become too large to fire from the two firings at c. This means that d was fired at least once in reaching π, and so $m-1 \geq 1$ and w exists. In passing from $\varphi[(w)]$ to $\varphi[(x)]$, the node d cannot be fired, since $w < x < z$ implies that $\varphi[(x)]$ has had d fired m–1 times. So Lemma 6.10 implies that x covers w. Since $\kappa(w) = d = \kappa(z)$, we have shown D4 at k. ■

Proposition 8.1 below will indicate that we now can infer from proper coloring, colored axioms D1/D4, and finiteness that any wave poset P_ζ satisfies axioms D3/D6 for d_k-completeness.

Proposition 7.4. *Wave posets* P_ζ *satisfy the D2/D5 axioms for* d_k*-completeness for* $k \geq 3$.

Proof. Suppose that $[w,z] \cong d_k(1)$ for some $k \geq 3$, and that there is an element u not below z which covers w. Then $\kappa(u)$ is adjacent to $\kappa(w)$ in G. Since both u and z are maximal elements of the ideal (u,z), the adjacent –1's in $\Psi(\varphi[(u,z)])$ would violate Lemma 6.2. So no such u can exist. ■

Since we have confirmed that any wave poset P_ζ satisfies the colored axioms D1-D6, we now know that it is colored d-complete. Recall that uncolored d-completenss follows trivially from colored d-completeness. Hence we have confirmed the forward direction of Theorem B, since the posets $j(L_\zeta)$ there are precisely the wave posets P_ζ.

8. Properties of d-Complete Posets

We begin by presenting the two lemmas which were used in Section 7.

Proposition 8.1. *Let* P *be a properly colored poset which satisfies the colored D1/D4 axioms for every* $k \geq 3$. *Then* P *satisfies the colored D3/D6 axioms for every* $k \geq 3$.

Proof. Let $k = 3$ and suppose that two elements z and z' each cover both of two other elements x and y. Let $\kappa(z) = c_1$ and $\kappa(z') = d_1$. Then applying D1 to z, we see that there must exist an element w with $\kappa(w) = c_1$. Similarly there exists w' with $\kappa(w') = d_1$. Since P is properly colored, $c_1 \neq d_1$. Hence we must have $w \neq w'$. This reasoning would need to be applied forever in order to satisfy D1. But P is finite. Hence having four such elements is impossible and D3 holds. Let $k \geq 4$ and suppose that two elements z and z' both cover an element y for which there exists an x such that $[x,y] \cong d_{k-1}(1)$, and that each of z and z' do not cover any other elements above x. Let $\kappa(z) = c_1$ and $\kappa(z') = d_1$. Applying D4 implies the existence of w and w' each covered by x, with (by proper coloring) distinct colors c_1 and d_1. So $w \neq w'$ again. Applying the D1/D4 axioms imply the existence of further elements below w and w', which together with repeating the above argument leads to a contradiction of the finiteness of P. So D6 is true as well, for $k \geq 4$. ■

Lemma 8.2. *Let* $k \geq 3$. *Let* P *be a colored* $d_{k'}$*-complete poset for every* $3 \leq k' \leq k$. *If* $w,z \in P$ *are such that* $[w,z] \cong d_k(1)$, *then no elements in* $[w,z]$ *other than* w *and* z *have color* $\kappa(w) = \kappa(z)$.

Proof. Let u and v be the two "middle" elements of $[w,z]$. Then $[w,u]$ is a chain. Since P is simply colored, all of the elements in $(w,u]$ have colors distinct from $\kappa(w)$. By colored $d_{k'}$-completeness for $k' < k$, we see that the colors in the chain $(u,z]$ are respectively the same as those in $[w,u)$. Repeat this reasoning for v and conclude that $\kappa(w)$ never occurs in (w,z). ■

Lemma 8.3. *No element can cover three or more other elements in a d-complete poset. "Strict neck" elements in* d_k*-intervals (necessarily* $k \geq 4$*) cannot cover any elements outside of the interval.*

Proof. Let P be d-complete. Suppose that $v \in P$ covers x, y, and z. Then by axiom D1, there exist elements w_1, w_2, and w_3 which are covered by x and y, y and z, and x and z respectively. If any two of these elements coincide, then axiom D2 would be violated. So they are distinct. Now apply D1 three more times to produce elements v_1, v_2, and v_3 which are covered respectively by pairs of the w_i's. This would continue forever, but we are assuming that P is finite. If a "strict neck" element covered an element outside of the interval, axiom D1 could be used to propagate diamonds downward until an element covering three other elements was produced. ■

Let P be any poset. The *bottom forest* F of P is the subposet consisting of every $y \in P$ such that $[x,y]$ is a (possibly empty) chain for every minimal element x of P.

Lemma 8.4. *Let* P *be a* d*-complete poset and let* z *be an element of its bottom forest* F. *Then there is exactly one minimal element* x *of* P *which is below* z.

Proof. Let w be one minimal element of P below z. Then there is a chain $w \to \ldots \to y \to z$. If any of these elements covers another element outside of this chain, then repeated applications of D1 would produce a downward sequence of diamonds which would eventually contradict the fact that w is minimal. So z can have only one minimal element below itself. ■

Lemma 8.5. *Let* P *be a* d*-complete poset with bottom forest* F *and let* z *be a maximal element of* P. *Then exactly one of the following is true:*

(i) *The element* z *is a maximal element of* F.

(ii) *There exists some* $k \geq 3$ *and some* $w < z$ *such that* $[w,z] \cong d_k(1)$.

Proof. If $z \in F$, then it is clearly maximal in F. By Lemma 8.4, the element z cannot be as in (ii). Suppose that $z \notin F$. If $z = z_1$ covers only one element, denote that element z_2. Repeat until some element z_n covers two elements x and y. This must happen eventually, or else $z \in F$. Then by D1, the elements x and y must cover some element w_n. Apply D4 for $k = 4, 5, \ldots, n+2$ to detect elements $w_{n-1}, \ldots, w_1$ such that $[w_i,z_i] \cong d_{n+3-i}(1)$. When $i = 1$, we have $[w_1,z_1] \cong d_k(1)$ for $k = n+2$. ■

Proposition 8.6. *Let* P *be a* d*-complete poset with bottom forest* F. *Let* N *be a set of colors such that* $|N| = |F|$. *Let* κ: $F \to N$ *be a bijection. Then there is a unique extension of* κ *to all of* P *which turns* P *into a colored* d*-complete poset.*

Proof: Use induction on $|P|$, with the base cases being the "forest" posets P such that $P = F$. If $P \neq F$, let z be a maximal element of P not in F. Let $P' := P - \{z\}$. Since P' is an ideal of P, it is d-complete. It will have the same bottom forest F. By induction, the coloring κ has been uniquely extended to $P' := P - \{z\}$. By (ii) of Lemma 8.5, the element z is the maximal element of a d_k-interval $[w,z]$ for some $k \geq 3$ and some $w < z$. Colored axioms D1/D4 force the extension of κ by $\kappa(z) := \kappa(w) =: d$. The sequence of choices of elements for the induction is immaterial, since these extensions of κ are determined only by the colors of elements below z, and not by colors of elements incomparable to z. Note that z does not cover w. Suppose that the element w is not the unique maximal element of color d in P', and that the adjunction of z is the first time for color d that the existing reference element w is not the unique maximal element of color d. Then at some earlier stage an element $z' \neq z$ was adjoined such that $[w,z'] \cong d_h(1)$ for some $h \geq 3$. The existence of both $[w,z']$ and $[w,z]$ in P would lead to a violation of D2/D5. So w was the unique maximal element of color d in P', and this implies that P is properly colored. The interval between two consecutive occurrences of one color will never be a chain when this procedure is followed. So P is simply colored. ■

9. d-Complete Posets are Wave Posets

Throughout this section let P be a fixed d-complete poset. Let F be the bottom forest of P. Fix a bijection κ from the elements of F to the colors in a set N such that $|N| = |F|$. Use Proposition 8.6 to color the elements of P so that it becomes a colored d-complete poset. Define a simple graph G on the node set N by taking its set of undirected edges to consist of the covering relations amongst the corresponding elements of F. As in Section 2, define $\Lambda := \mathbb{Z}^N$. Let $\lambda = (\ell_b)$ be the element of Λ such that $\ell_f = 1$ if $f \in N$ corresponds to a minimal element of F, and such that $\ell_f = 0$ otherwise. Construct Ω with respect to G and λ as in Section 2. Taking $I = P$ and ignoring the colors in the statement of the following proposition yields the reverse direction of Theorem B:

Proposition 9.1. *Let* P *be a colored d-complete poset with bottom forest* F. *Let* G, λ, *and* Ω *be as just defined. Let* I *be an ideal of* P. *Then there exists an element* $\zeta \in \Omega$ *such that* $\mathrm{P}_\zeta = \mathrm{I}$ *as colored posets.*

Proof. The ideal I is colored d-complete. Its colors come from N, and its bottom forest F_I is an ideal of F. The proof is by induction on $|\mathrm{I}|$. When $|\mathrm{I}| = 1$, it consists of one minimal element of P (or of F). It can be seen that a good ζ for I can be produced by firing the $+1$ at the corresponding node of G. Let $|\mathrm{I}| > 1$. Let z be a maximal element of I. Let $d := \kappa(z)$. Set $\mathrm{I}' = \mathrm{I} - \{z\}$. By induction, let $\zeta' \in \Omega$ be such that $\mathrm{I}' = \mathrm{P}_{\zeta'} = j(\mathrm{L}_{\zeta'})$ as colored posets. Below we will argue in each case that $s_d\zeta' =: \zeta$ is defined in Ω. Since L_ζ will have one more rank than $\mathrm{L}_{\zeta'}$ as a distributive lattice, it will be the case that $j(\mathrm{L}_\zeta)$ will consist of $j(\mathrm{L}_{\zeta'})$ with one element adjoined as a new maximal element. Proposition 5.4 will be used often in this proof.

There are two possibilities for z listed in Lemma 8.5, where P is to be replaced by I and F by F_I. First suppose that z is a maximal element of F_I. If z is also minimal in P, then I is the disjoint union of I' and z. Consider defining $\zeta := s_d\zeta'$. From the procedure of Proposition 8.6, any color appearing in I' (i.e. any node fired to reach ζ') must appear in the bottom forest of I'. Since the colors of elements of F_I are distinct, the node d cannot have been fired yet in ζ', and so the original $+1$ at d in $\Psi(\theta) = \lambda$ is still present. So ζ is defined, and $\eta(d,1)$ is a non-trivial element of L_ζ. Clearly $\eta(d,1) = s_d\theta$. Let t denote $\eta(d,1)$ viewed as an element of $j(\mathrm{L}_\zeta)$; it is of color d. Clearly $j(\mathrm{L}_\zeta)$ is the disjoint union of $\mathrm{I}' = j(\mathrm{L}_{\zeta'})$ and $\{t\}$, and so $j(\mathrm{L}_\zeta) = \mathrm{I}$ as desired.

Next suppose that z is maximal in F_I, but not minimal in P. Hence $\ell_d = 0$ in λ. Since z is non-minimal in F_I, it covers exactly one element of F_I, say y with $\kappa(y) := c$. Now $y \in \mathrm{I}'$, and so by induction the node at c is fired in reaching $\mathrm{L}_{\zeta'}$. As before, any color appearing in I' must appear in $\mathrm{F}_{\mathrm{I}'}$. So any such color cannot label an element of F above z. The only other node adjacent to d in G is c. So c is the only node adjacent to d which was fired in reaching ζ'. We now show that it was fired only once in this process. The element y is the minimal element of I' of color c, and there are no others: Otherwise by Proposition 7.2 for $\mathrm{P}_{\zeta'} = \mathrm{I}'$, there is an interval $[y,\eta(c,2)] \cong d_k(1)$ for some k in I'. Since $z \notin \mathrm{I}'$, it cannot be in this interval of I', and then in I it would be an element violating D2/D5 for that $d_k(1)$ interval. So the node c is fired only once in I', and thus $\mathcal{C}_d(\zeta') = +1$. Hence we may define $\zeta := s_d\zeta'$. We want to locate the non-trivial $\eta(d,1)$ in L_ζ. Let $y_n \to y_{n-1} \to \ldots \to y_1 = y$ describe $(y) \subseteq F$, and let $c_n, c_{n-1}, \ldots, c_1 = c$ be the corresponding colors. Then clearly $\eta(c,1) = c_1\ldots c_n\theta$ and $\eta(d,1) = dc_1\ldots c_n\theta$ in L_ζ. Let t denote $\eta(d,1)$ viewed as an element of $j(L_\zeta)$. Observe that $t \notin j(L_{\zeta'})$, and it covers only $\eta(c,1)$ in $j(L_\zeta)$ since $\eta(d,1) = s_d\eta(c,1)$ in L_ζ. But $\varphi[(y)] = \eta(c,1) \in j(L_{\zeta'})$. So $j(L_\zeta)$ consists of $j(L_{\zeta'}) = I'$ together with an element t of color d adjoined which covers only y. But z covered only y in I. Hence $j(L_\zeta) = I$ as desired.

Now we consider possibility (ii) of Lemma 8.5. Let $k \geq 3$ and $w \in I$ be such that z is the maximal element of a d_k-interval $[w,z]$.

Suppose $k = 3$. Then there exist $x, y, w \in I'$ such that in I we have z covering x and y, and in I' each of these elements covers w. Let b and c be the colors of x and y. Note that $\kappa(w) = \kappa(z) = d$ is adjacent to each of b and c. Let $z' > w$ be minimal in I' such that $\kappa(z') = d$. Then by Proposition 7.2 we have $[w,z'] \cong d_h$ for some $h \geq 3$. But in I we also have $[w,z] \cong d_3(1)$. This would lead to some kind of violation of D2/D5, since $z \neq z'$. Hence w is the maximal element of color d in I'. There is some $m \geq 1$ such that $w = \eta(d,m)$.

Note that there is a unique -1 at d in $\Psi(\varphi[(w)])$. Pass from $\varphi[(w)]$ to $\varphi[(x,y)]$ in $L_{\zeta'}$ by using edges specified by building up to (x,y) from (w) in I'. If $v \in (x,y) - (w)$ is not above w, then it is incomparable to w. Then Lemma 6.3 implies that $\kappa(v)$ is not adjacent to $\kappa(w) = d$. So no colors for elements used in the build-up besides those for x and y can be adjacent to d. So the firings at b and c are the only firings which occur adjacent to d in passing from (w) to (x,y). Hence $\mathcal{C}_d(\varphi[(x,y)]) = +1$. Let $a_1, \ldots, a_q$ be a sequence of nodes in $G_{\zeta'}$ such that $\zeta' = a_q\ldots a_1\varphi[(x,y)]$. These are just the colors of the elements in $I' - (x,y)$. Since w is the maximal element in I' of color d, none of the a_i's are equal to d. Suppose that one of them is adjacent to d. Let v be an element in $I' - (x,y)$ of that color, and let $\varphi[(v)] = (u_b)$. Since reaching w was the last time that d was fired, we have $u_d = m$. Then by Lemma 6.10, we see that v covers w in I'. So then v covers w in I. Since $[w,z] \cong$

$d_3(1)$ and I is d-complete, by D2 we see that v must be either x or y. But v was chosen from $I' - (x,y)$. So none of the $a_1, \ldots, a_q$ are adjacent to d. Hence $\mathcal{C}_d(\zeta') = +1$. Define $\zeta := s_d\zeta$.

We now know that $\eta(d,m+1)$ is non-trivial in L_ζ. By the $\lambda \in \mathbb{N}^N$ version of Proposition 5.4, we see that the only -1's in $\Psi(\varphi[(x,y)])$ are at b and c. Define $\pi := s_d\varphi[(x,y)]$. Then there is only one -1 in $\Psi[\pi]$, and it is at d. Hence $\eta(d,m+1) = \pi$. Let t denote $\eta(d,m+1)$ when viewed as an element of $j(L_\zeta)$. Since $w = \eta(d,r)$ was the maximal occurrence of d in $j(L_{\zeta'})$, the element $t \notin j(L_{\zeta'})$. So $j(L_\zeta)$ consists of $j(L_{\zeta'})$ together with the one element t of color d. By Lemma 8.3, we see that z covered exactly x and y in I. It is clear from the work above and Lemma 6.10 that t covers x and y in $j(L_\zeta)$. By Lemma 6.7, we see that these are the only elements that t covers in $j(L_\zeta)$. So $j(L_\zeta) = I$ when $k = 3$.

Now suppose that $k \geq 4$. Then there exists $w \in I'$ such that $[w,z] \cong d_k(1)$ in I. Note that $\kappa(w) = \kappa(z) = d$. Let x be the unique (by D5) element of I' which covers w and let y be the unique (by Lemma 8.3) element of I' which is covered by z. Note that $[x,y] \cong d_{k-1}(1)$. Let $c := \kappa(x) = \kappa(y)$. Note that c is adjacent to d. As with $k = 3$, we can see that the element w is the maximal element of color d in I'. Let $w = \eta(d,m)$.

Note that there is a unique -1 at d in $\Psi(\varphi[(w)])$. Pass from $\varphi[(w)]$ to $\varphi[(y)]$ in $L_{\zeta'}$ by using edges specified by building up to (y) from (w) in I'. If $v \in (y) - (w)$ is not above w, then it is incomparable to w. Then Lemma 6.3 implies that $\kappa(v)$ is not adjacent to $\kappa(w) = d$. So in the build-up, only colors for elements in $[x,y]$ could be adjacent to d. If a color b appears in $[x,y]$, then there exists a $u \in [x,y]$ such that $\kappa(u) = b$ and $[w,u]$ is a chain. Corollary 6.5 implies that the sequence of colors corresponding to this chain form a path in $G_{\zeta'}$. Since this graph is acyclic, none of these colors is adjacent to d except for the first, viz. $c = \kappa(x)$. So two firings at c are the only firings which occur adjacent to d in passing from (w) to (y). Hence $\mathcal{C}_d(\varphi[(y)]) = +1$. Let $a_1, \ldots, a_q$ be a sequence of nodes in $G_{\zeta'}$ such that $\zeta' = a_q \ldots a_1\varphi[(y)]$. These are just the colors of the elements in $I' - (y)$. An argument similar to the one used here for $k = 3$ shows that $\mathcal{C}_d(\zeta') = +1$. Again define $\zeta := s_d\zeta$.

We now know that $\eta(d,m+1)$ is non-trivial in L_ζ. Note that the unique -1 for $\varphi[(y)]$ is at c. Define $\pi := s_d\varphi[(y)]$. Then there is only one -1 in $\Psi[\pi]$, and it is at d. Hence $\eta(d,m+1) = \pi$. Let t denote $\eta(d,m+1)$ when viewed as an element of $j(L_\zeta)$. As above, the poset $j(L_\zeta)$ consists of $j(L_{\zeta'})$ together with the one element t of color d. By Lemma 8.3, we see that z covered only y in I. It is clear from the work above and Lemma 6.10 that t covers y in $j(L_\zeta)$. Note that $(y) = \eta(c,r)$ for some $r \geq 1$. Since $\eta(d,m+1) = s_d\eta(c,r)$ in L_ζ, it must be the case that t covers only y in $j(L_\zeta)$. So $j(L_\zeta) = I$ when $k \geq 4$ and we have finished the proof of Proposition 9.1. ■

10. Group, Representation, and Geometric Comments

Let G be fixed simple graph with finite node set N. View G as a simply laced Coxeter diagram, and let W be the corresponding Coxeter group. Let Λ be the set of labellings $(m_i)_{i\in N}$ of the nodes of G with integers. For each $k \in N$, let $\mathcal{S}_k$ be the operation on elements of Λ which multiplies the kth label m_k by -1 and adds m_k to the labels of each of the adjacent nodes. This action is linear, and it is easy to check that the Coxeter relations $\mathcal{S}_k^{\,2} = e$, $\mathcal{S}_j\mathcal{S}_k = \mathcal{S}_k\mathcal{S}_j$ (for j and k non-adjacent), and $\mathcal{S}_k\mathcal{S}_j\mathcal{S}_k = \mathcal{S}_j\mathcal{S}_k\mathcal{S}_j$ (for j and k adjacent) are preserved. Hence Λ is a W-module. Our restricted numbers game generates initial portions of the orbits of initial vectors under this action.

Let $\mathfrak{g}$ denote the simply laced Kac-Moody algebra with Dynkin diagram G, and let $\mathfrak{g}'$ denote its derived subalgebra. Let $\mathfrak{h}$ and $\mathfrak{h}'$ denote the Cartan subalgebras of $\mathfrak{g}$ and $\mathfrak{g}'$. The space $\mathfrak{h}'$ has a basis $\{\alpha_i^\vee\}_{i\in N}$ consisting of the simple coroots. Let $\{\omega_i\}_{i\in N}$ denote the fundamental weight basis of $\mathfrak{h}'^*$; here $\omega_i(\alpha_j^\vee) = \delta_{ij}$ is the definition of ω_i. For $i \in N$, let α_i denote the restriction of the ith simple root of $\mathfrak{g}$ to $\mathfrak{h}'$. With this notational convention we have $\alpha_i \in \mathfrak{h}'^*$, and the α_i are no longer linearly independent if $\mathfrak{g}$ is not finite dimensional. When one expresses a restricted root α_k in terms of the ω_i, one obtains $\alpha_k = 2\omega_k - \sum\omega_j$, where the sum is over all $j \in N$ which are adjacent to k in G. Let Λ denote the $\mathbb{Z}$-lattice in $\mathfrak{h}'^*$ generated by the ω_i: this is the set of the restrictions of the elements of the usual lattice $P \subseteq \mathfrak{h}^*$ to $\mathfrak{h}'$. Let Λ^+ denote the set of non-negative integral sums of the ω_i. Let W be the Weyl group of $\mathfrak{g}$. The action [Kac, p. 35] of W on $\mathfrak{h}^*$ restricts to $\mathfrak{h}'^*$. It is well known that the Weyl group W of $\mathfrak{g}$ is the Coxeter group of the preceding paragraph, and it is easy to see that the the Kac-Moody action of W on Λ becomes the action of the preceding paragraph when all vectors are written as N-tuples with respect to the ω_i basis.

We now relate the restricted numbers game of Section 2 to this action of W on Λ. The wave state space Λ becomes the lattice Λ. A potential initial state $\lambda \in \mathbb{N}^N$ for Theorem B is just a dominant integral weight $\lambda \in \Lambda^+$. The space of tallies Ξ is the set of non-negative sums of positive simple roots restricted to $\mathfrak{h}'$. The initial tally θ is the empty sum of positive simple roots, and $\Psi(\theta) = \lambda$. Other tallies are to be subtracted from λ after being converted to the ω_i basis. The wave state corresponding to the tally $s_{i_k}\ldots s_{i_1}\theta$ is $\Psi(s_{i_k}\ldots s_{i_1}\theta) = \mathcal{S}_{i_k}\ldots\mathcal{S}_{i_1}\lambda$. A further tally can be generated with the action of s_d if $[s_{i_k}\ldots s_{i_1}\lambda](\alpha_d^\vee) = +1$, i.e. if the d*th* coefficient with respect to the ω_i basis is $+1$.

Put the reverse of the usual order on Λ: for $\mu,\nu \in \Lambda$, we say that $\mu \geq \nu$ if $\nu - \mu$ is a sum of positive roots. Lemma 3.11 of [Kac] can be used to prove the following lemma: Let $\lambda \in \Lambda^+$ and $w \in W$. Then for $k \in N$, $\ell(\mathcal{S}_k w) > \ell(w)$ if and only if $\mathcal{S}_k w\lambda > w\lambda$.

Fix $\lambda \in \Lambda^+$. We say that $w \in W$ is *λ-minuscule* if there exists some decomposition $w = \mathcal{S}_{i_k}\ldots\mathcal{S}_{i_1}$ such that $\mathcal{S}_{i_j}(\mathcal{S}_{i_{j-1}}\ldots\mathcal{S}_{i_1}\lambda) = (\mathcal{S}_{i_{j-1}}\ldots\mathcal{S}_{i_1}\lambda) - \alpha_{i_j}$ for $1 \leq j \leq k$. Now think of λ as being an initial state. Applying the procedures of the restricted numbers game to λ will obviously produce all possible such decompositions $\mathcal{S}_{i_k}\ldots\mathcal{S}_{i_1}$ for all λ-minuscule elements w in W. Fix one such game play sequence $i_1, \ldots, i_k$ and set $w := \mathcal{S}_{i_k}\ldots\mathcal{S}_{i_1}$. Then $\zeta := \lambda - w\lambda$ is the corresponding tally $s_{i_k}\ldots s_{i_1}\theta$. Theorem A stated that the poset L_ζ defined in Section 2 is a distributive lattice. Under the map Ψ, this lattice becomes the subposet $\Psi(L_\zeta)$ of Λ. Let W_λ denote the stabilizer of λ in W. Using the results of this paper, one can associate a well-defined element $v \in W/W_\lambda$ to each tally $\pi \in L_\zeta$. Let the symbol w also denote the element of W/W_λ corresponding to ζ. Using the lemma just noted above, it can be seen that the order inherited from L_ζ by the subset of W/W_λ generated in this fashion is the same as the order obtained from the notion of weak Bruhat order. Let $[e,w]$ be the initial interval in the weak Bruhat order on W/W_λ determined by w. The lemma above can also be used to conversely show that any $v \in [e,w]$ is represented in L_ζ, and that $[e,w]$ is order isomorphic to L_ζ. When the full Bruhat order on W/W_λ is realized on the orbit $W\lambda$, the only covering relations which are not already present in the weak Bruhat order correspond to subtraction of multiples of non-simple roots. The fact that the ranks of $\Psi(L_\zeta)$ are spaced apart by the subtraction of just one positive simple root at each step implies that the initial interval $[e,w]$ in the full Bruhat order on W/W_λ is also isomorphic to L_ζ.

Fix $\lambda \in \Lambda^+$. Let $\Pi(\lambda)$ denote the set of weights of the integrable $\mathfrak{g}'$ module with highest weight λ. Let $w = s_{i_k} \ldots s_{i_1} \in W$ be λ-minuscule. Set $\zeta := s_{i_k} \ldots s_{i_1} \theta$ and $\mu := w\lambda$ $(= \Psi(\zeta))$. The set of weights corresponding to the elements of L_ζ, i.e. all weights of the form $\lambda - \pi$ for $\pi \in L_\zeta$, is $\Psi(L_\zeta)$. Let $\Pi(\lambda)_\mu$ denote the subset of $\Pi(\lambda)$ consisting of weights which are weakly below μ in our order on Λ. It will be shown in a future version of [Pr4] that $\Pi(\lambda)_\mu = \Psi(L_\zeta)$, and that each weight in this set occurs with multiplicity one for the representation. The set $\Pi(\lambda)_\mu$ is what we meant in the introduction by a "minuscule portion" of the weight diagram for the representation.

For $f \in N$, let $\omega_f = (\ell_b)_{b \in N}$ denote the element of Λ with $\ell_f = +1$ and $\ell_b = 0$ otherwise. The reduction given near the end of Section 2 amounts to saying: If we want to study a particular λ-minuscule w or the corresponding portion $[e,w]$ of a weight lattice, then we might as well study several simpler situations of the form (G',λ), where G' is a rooted tree Dynkin diagram with root f and $\lambda = \omega_f$.

Let w be a λ-minuscule element of a simply laced Weyl group. Let ζ be the corresponding tally constructed above, and let P_ζ be the corresponding colored wave poset constructed in Section 5. Then as noted just before Corollary 5.5, the reduced decompositions of w correspond to the order extensions of P_ζ. The $x = 1$ specialization of our recent hook length generating function result for d-complete posets (joint work with Dale Peterson) yields a hook product formula for the number of order extensions of P_ζ and for the number of reduced decompositions of w. Propositions 8.6 and 9.1 imply that one can obtain a hook product count of order extensions for any (uncolored) d-complete poset.

The classification of d-complete posets [Pr1] can be used to list all simply laced λ-minuscule elements and all of their reduced decompositions.

Let $\mathcal{G}$ be the simply laced Kac-Moody group corresponding to $\mathfrak{g}$. Consider $\lambda = \omega_f$ for some $f \in N$. Let $\mathcal{P}$ be the corresponding maximal parabolic subgroup of $\mathcal{G}$. Let w be λ-minuscule. Then the Schubert subvariety X_w of the flag manifold $\mathcal{G}/\mathcal{P}$ has a particularly simple structure which is described to some extent (Schubert cellular decomposition) by the distributive lattice L_ζ, where ζ corresponds to w as above. These subvarieties are closely analogous to the Schubert subvarieties X_w of Grassmannians. In those cases, the wave posets $P_\zeta = j(L_\zeta)$ arising for a fixed $\mathcal{G}/\mathcal{P}$ are precisely the shapes contained within a fixed rectangular poset. It is well known that each such shape P_ζ contains geometrically useful information concerning the corresponding variety X_w. The general simply laced "λ-minuscule" Schubert subvarieties X_w above are automatically classified at the same time that all d-complete posets are classfied in [Pr1], and it is hoped that the irreducible d-complete posets P drawn there will be as helpful in studying these varieties as the shapes P_ζ have been for the Grassmannian Schubert varieties. A small example of this was the re-derivation in [Pr2] of the degrees of the minuscule flag manifolds by counting the number of order extensions of the associated d-complete posets (then called "minuscule" posets).

Let W be an arbitrary Coxeter group. Stembridge has studied the elements $w \in W$ for which the weak Bruhat interval $[e,w]$ is a distributive lattice [Ste]. He showed that this property is equivalent to w being "fully commutative," viz. any reduced decomposition for w can be converted into any other reduced decomposition for w using only relations of the form $s_is_j = s_js_i$. Let w be λ-minuscule. By relating $[e,w]$ in the weak Bruhat order to the poset L_ζ as above, we can use Theorem A to conclude that $[e,w]$ is a distributive lattice. Then Theorem 2.2 of [Ste] implies that w is fully commutative. The wave posets P_ζ of this paper arise as "heaps" in [Ste]. Our Theorem B can be thought of as characterizing the heaps of some particularly nice fully commutative elements from simply laced Weyl groups.

I would like to thank Shrawan Kumar for some helpful conversations.

References

[AKP] N. Alon, I. Krasikov, Y. Peres, Reflection sequences, American Math. Monthly 96 (1989), 820 - 822.

[Bjö] A. Björner, On a combinatorial game of S. Mozes, preprint, 1988.

[Eri] K. Eriksson, Strongly convergent games and Coxeter groups, Ph.D. Thesis, Kungl Tekniska Högskolan, 1993.

[Kac] V. Kac, *Infinite dimensional Lie algebras*, Cambridge University Press, 1990.

[Lak] V. Lakshmibai, Bases for Demazure modules for symmetrizable Kac-Moody algebras, in *Linear Algebraic Groups and Their Representations,* 59 - 78, Contemporary Math. 153, AMS, Providence, 1993.

[Moz] S. Mozes, Reflection processes on graphs and Weyl groups, J. Combinatorial Theory A 53 (1990), 128 - 142.

[Pr1] R. Proctor, Dynkin diagram classification of λ-minuscule Bruhat lattices and of d-complete posets, J. of Algebraic Combinatorics, to appear.

[Pr2] R. Proctor, Bruhat lattices, plane partition generating functions, and minuscule representations, European J. of Combinatorics 5 (1984), 331-350.

[Pr3] R. Proctor, A Dynkin diagram classification theorem arising from a combinatorial problem, Advances in Math. 62 (1986), 103-117.

[Pr4] R. Proctor, Poset partitions and minuscule representations: External construction of Lie representations, Part I, preliminary manuscript.

[Sag] B. Sagan, Enumeration of partitions with hooklengths, European J. Combinatorics 3 (1982), 85-94.

[Sta] R. Stanley, *Enumerative Combinatorics, Vol. I*, Wadsworth & Brooks/Cole, Monterey, 1986.

[Ste] J. Stembridge, On the fully commutative elements of Coxeter groups, J. of Algebraic Combinatorics 5 (1996), 353 - 385.