Séminaire Lotharingien de Combinatoire **78B** (2017)
Article #82, 12 pp.

Proceedings of the 29th Conference on Formal Power Series and Algebraic Combinatorics (London)

Whitney Duals of Geometric Lattices

Rafael S. González D'León[1] and Joshua Hallam[2]

[1]*Department of Mathematics, University of Kentucky, Lexington, KY 33124, USA*
[2]*Department of Mathematics and Statistics, Wake Forest University, Winston-Salem, NC 27109, USA*

Abstract. Given a graded partially ordered set P, let $w_k(P)$ and $W_k(P)$ denote its Whitney numbers of the first and second kind respectively. We call a graded partially ordered set Q a *Whitney Dual* of P if $|w_K(P)| = W_K(Q)$ and $W_k(P) = |w_k(Q)|$ for all k. In this extended abstract, we show that every geometric lattice has a Whitney dual. This is done constructively, using edge labelings and quotient posets.

Keywords: Posets, Möbius function, Whitney numbers, Edge labelings, Quotient Posets

1 Introduction

Throughout this abstract we will assume that all our partially ordered sets (or *posets*) are finite, graded, and contain a minimum element (denoted by $\hat{0}$). Moreover, ρ will be used to denote the rank function. Recall that for a poset P, the Möbius function is defined recursively by

$$\mu(x,y) = \begin{cases} 1 & \text{if } x = y, \\ -\sum_{x \leq z < y} \mu(x,z) & \text{if } x \neq y. \end{cases}$$

The *Whitney numbers of the first kind* $w_k(P)$ are defined by

$$w_k(P) = \sum_{\rho(x)=k} \mu(\hat{0}, x),$$

and the *Whitney numbers of the second kind* $W_k(P)$ are defined by

$$W_k(P) = \sum_{\rho(x)=k} 1.$$

We assume that the reader is familiar with poset and poset topology terminology, see [8, 9] for background and notation.

The authors of [4] noticed in their study of the poset of weighted partitions, Π_n^w, that the Whitney numbers of Π_n^w were closely related to the Whitney numbers for the poset $\mathcal{F}_n$ of rooted spanning forests on $[n]$ already studied by D. Reiner in [6] and B. Sagan in

[7]. The surprising fact was that the Whitney numbers of the first and second kind of Π_n^w were (up to a sign) the same as the Whitney numbers of the second and first kind of $\mathcal{F}_n$. That is $|w_K(\Pi_n^w)| = W_K(\mathcal{F}_n)$ and $W_k(\Pi_n^w) = |w_k(\mathcal{F}_n)|$. This phenomenon occurs for many other pairs of posets and motivates the following definition.

Definition 1. Let P and Q be graded posets. We say that P and Q are *Whitney Duals* if for all $k \geq 0$ we have that

$$|w_k(P)| = W_k(Q) \text{ and } |w_k(Q)| = W_k(P).$$

According to this definition Π_n^w and $\mathcal{F}_n$ are Whitney Duals. We now discuss another example.

Example 1. Let Π_n denote the poset whose underlying set is formed by the set partitions of $[n]$ with order relation given by $\pi \leq \pi'$ whenever every block of π is contained in some block of π'. We say that the partitions are ordered by *refinement* and we call Π_n the *partition lattice*. In Figure 1a we illustrate the Hasse diagram of Π_3 together with the Möbius values $\mu(\hat{0}, \pi)$ next to each element π.

Now let T be a tree with vertices labeled by distinct integers. We call the smallest vertex of T the *root*. We say T is an *increasing tree* if the sequence of vertex labels read along any path starting at the root of T is increasing. An *increasing spanning forest* is a collection of increasing trees whose vertex labels form a partition of $[n]$. The word "spanning" here indicates that these forests are spanning forests of the complete graph. Let $\mathcal{ISF}_n$ be the set of increasing spanning forests on $[n]$. We define a partial order on $\mathcal{ISF}_n$ by saying that $F_1 \lessdot F_2$ if exactly two trees in F_1 are replaced by the tree in F_2 that is obtained after joining their roots with an edge. Note that the root of the resulting tree is the smaller label among the roots of the two joined trees. Figure 1b illustrates the Hasse diagram of $\mathcal{ISF}_3$ together with its Möbius values. For more information about increasing spanning forests see [5].

k	$w_k(\Pi_n)$	$W_k(\mathcal{ISF}_n)$	$w_k(\mathcal{ISF}_n)$	$W_k(\Pi_n)$
0	1	1	1	1
1	-3	3	-3	3
2	2	2	1	1

Table 1: Whitney numbers of the first and second kind for Π_3 for $\mathcal{ISF}_3$.

When the Whitney numbers of Π_3 and $\mathcal{ISF}_3$ are listed side by side (see Table 1), we see that they are indeed Whitney duals. Later, we will see that for $n \geq 1$, Π_n and $\mathcal{ISF}_n$ are Whitney duals (see Proposition 4).

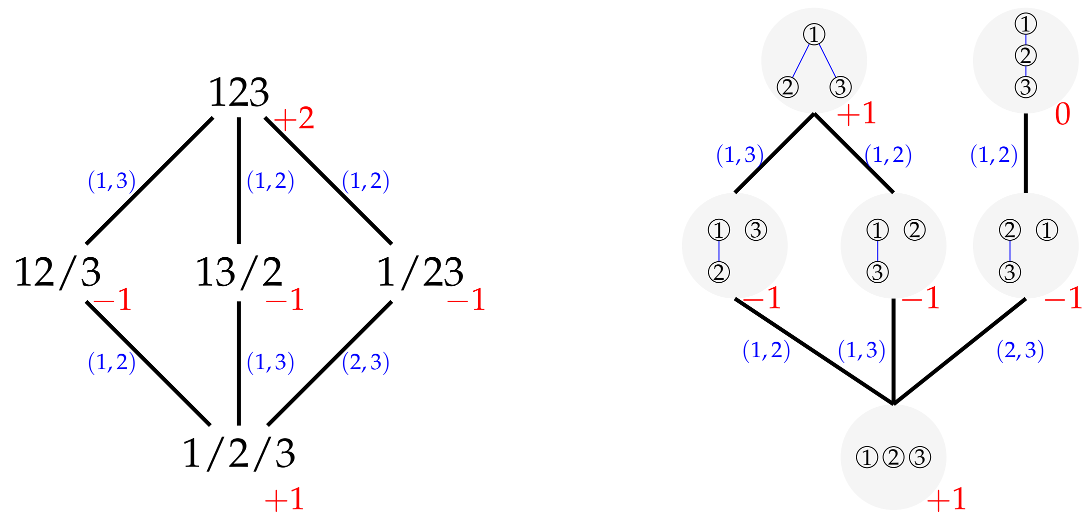

(a) Π_3 (b) $\mathcal{ISF}_3$

Figure 1

The rest of the paper is organized as follows. In Section 2, we consider some examples of Whitney duals. We pay particular attention to Π_n and $\mathcal{ISF}_n$. In Section 3, we explain how to use edge labelings and quotient posets to construct Whitney duals for posets which posses a certain type of edge labeling that we call an $\overline{\text{EW}}$-labeling. We also show in this section that every geometric lattice has an $\overline{\text{EW}}$-labeling and hence a Whitney dual. We finish with a section on future work.

2 Examples of Whitney Duals

2.1 Eulerian posets, self-duality, and posets without Whitney duals

A graded poset P is *Eulerian* if $\mu(x, y) = (-1)^{\rho(y)-\rho(x)}$ for all $x \leq y$ in P. Thus, in an Eulerian poset we have that $|w_k(P)| = W_k(P)$ for all k. Therefore, every Eulerian poset has a Whitney dual, namely itself. It is natural to ask if all posets which are their own Whitney dual are Eulerian. The poset in Figure 2a shows that this is not the case [1]. This leads to the question of whether there is a natural characterization of self Whitney-dual posets.

Not every ranked poset has a Whitney dual. For example, consider the three element chain C in Figure 2b. We have that $w_2(C) = 0$ and $W_2(C) = 1$. If Q was a Whitney dual of C, then $|w_2(Q)| = 1$ and $W_2(Q) = 0$, which is clearly impossible. This illustrates the fact that a poset P with $|w_k(P)| = 0$ for some k smaller than the rank of the poset cannot

[1] We thank Cyrus Hettle from University of Kentucky for pointing out this example to the authors.

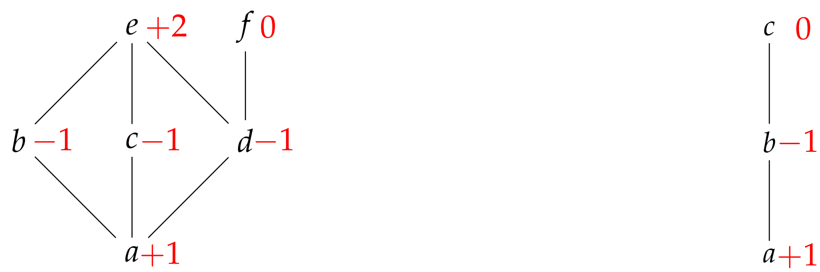

(a) Self Whitney-dual, but not Eulerian (b) No Whitney duals

Figure 2: Examples of self Whitney-duality and a poset without Whitney dual

have a Whitney dual.

2.2 Edge labelings and Whitney duality

We now discuss edge labelings and their relation with Whitney numbers. First, let us recall some basic facts about edge labelings. For complete treatments on the topic, see [1, 3, 8]. Let P be a poset, and let $\mathcal{E}(P)$ be the set of edges of the Hasse diagram of P. Moreover, let Λ be an arbitrary fixed poset that will be considered as the *poset of labels*. An *edge labeling* of P is a map $\lambda : \mathcal{E}(P) \to \Lambda$.

Let P be a poset with edge labeling λ. For any saturated chain

$$\mathbf{c} : (x = x_0 \lessdot x_1 \lessdot \cdots \lessdot x_{\ell-1} \lessdot x_\ell = y)$$

there is a corresponding *word of labels*

$$\lambda(\mathbf{c}) = \lambda(x_0, x_1)\lambda(x_1, x_2) \cdots \lambda(x_{\ell-1}, x_\ell).$$

We say that $\mathbf{c}$ is *increasing* if its word of labels $\lambda(\mathbf{c})$ is *strictly* increasing, that is, $\mathbf{c}$ is increasing if

$$\lambda(x_0, x_1) < \lambda(x_1, x_2) < \cdots < \lambda(x_{\ell-1}, x_\ell).$$

We say that $\mathbf{c}$ is *ascent-free* if its word of labels $\lambda(\mathbf{c})$ has no ascents, i.e. $\lambda(x_i, x_{i+1}) \not< \lambda(x_{i+1}, x_{i+2})$, for all $i = 0, \ldots, \ell - 2$. Clearly there can be chains that are neither increasing nor ascent-free.

Definition 2. An edge labeling is an *ER-labeling* if in each closed interval $[x, y]$ of P, there is a unique increasing maximal chain (in [8] this type of labeling is called an R-labeling). By analogy, we say that an edge labeling is an ER*-labeling if in each closed interval $[x, y]$ of P, there is a unique ascent-free maximal chain.

The following theorem due to R. Stanley provides a relation between ER and ER*-labelings and the Möbius function in a poset P.

Theorem 1 (c.f. Theorem 3.14.2 in [8]). *Let P be a poset with an ER-labeling (ER*-labeling). Then*

$$\mu(x,y) = (-1)^{\rho(y)-\rho(x)} |\{\mathbf{c} \mid \mathbf{c} \text{ is an ascent-free (increasing) maximal chain in } [x,y]\}|.$$

Using Definition 2 and Theorem 1, we can describe the Whitney numbers of a poset with an ER-labeling (ER*-labeling) by the enumeration of saturated chains as follows.

Proposition 1. *Let P be a poset with an ER-labeling (ER*-labeling). Then $|w_k(P)|$ is the number of ascent-free (increasing) saturated chains starting at $\hat{0}$ of length k. Moreover, $|W_k(P)|$ is the number of increasing (ascent-free) saturated chains starting at $\hat{0}$ of length k.*

With this proposition in mind, it seems natural to use edge labelings to understand Whitney duals. We illustrate this concept with Π_n and $\mathcal{ISF}_n$ next.

2.3 Π_n and $\mathcal{ISF}_n$

In order to prove that Π_n and $\mathcal{ISF}_n$ are Whitney duals, we are going to describe an ER-labeling on Π_n and an ER*-labeling on $\mathcal{ISF}_n$. Recall that if $\pi \lessdot \sigma$ is a covering relation in Π_n , then σ is obtained from π by merging two blocks of π.

Definition 3. Let $\lambda : \mathcal{E}(\Pi_n) \to [n] \times [n]$ be the edge labeling defined by $\lambda(\pi \lessdot \sigma) = (i,j)$ where $i < j$ and i and j are the minimum elements of the two blocks of π that were merged to obtain σ and $[n] \times [n]$ has the lexicographic order induced by the natural order of $[n]$. This labeling is a special case of Björner's minimum labeling for geometric lattices described in [2]. In Figure 1a the labeling λ of Π_3 is depicted.

Proposition 2 (c.f. [2]). *The labeling $\lambda : \mathcal{E}(\Pi_n) \to [n] \times [n]$ of Definition 3 is an ER-labeling.*

Now consider an ER*-labeling of $\mathcal{ISF}_n$. We will use the convention that an edge of the complete graph is an ordered pair (i,j) with $i < j$.

Definition 4. We define an edge labeling $\lambda^* : \mathcal{E}(\mathcal{ISF}_n) \to [n] \times [n]$ by setting $\lambda^*(F_1 \lessdot F_2)$ to be the unique edge in F_2 which is not in F_1. Figure 1b depicts the labeling λ^* of $\mathcal{ISF}_3$.

The labels for each cover relation of $\mathcal{ISF}_n$ indicate the two roots of the trees that get connected. Similarly, in each cover relation of Π_n the labels indicate what are the smallest elements of the two blocks that were merged. Note that the two roots of an increasing forest have the smallest labels in their connected components. If we think of the connected components as blocks of a partition, we see λ and λ^* are closely related. We also have that any maximal chain in a fixed interval, in both of these posets, is completely determined by its word of labels.

Proposition 3. *The labeling $\lambda^* : \mathcal{E}(\mathcal{ISF}_n) \to [n] \times [n]$ of Definition 4 is an ER*-labeling.*

Proof. Let $F \in \mathcal{ISF}_n$ and suppose the edges of F are $(i_m, j_m) > (i_{m-1}, j_{m-1}) > \cdots > (i_1, j_1)$, ordered lexicographically. Let F_k be the forest with edges $(i_m, j_m), (i_{m-1}, j_{m-1}), \cdots, (i_{m-k+1}, j_{m-k+1})$ (note that F_0 is the graph with no edges). Then for all k, F_k is a subforest of F. Since subforests of increasing spanning forests are increasing spanning forests (see [5]), each F_k is an increasing spanning forest. Note that there is unique edge in F_k not in F_{k-1}. We claim that this edge connects two roots. Suppose we are creating F_k from F_{k-1}. Then we add the edge (i_{m-k+1}, j_{m-k+1}). Let T be the tree containing i_{m-k+1} and let r be the root of T. If $r \neq i_{m-k+1}$, then at some point as we built up T, we added an edge containing r. Since $i_{m-k+1} > r$ this edge preceded (i_{m-k+1}, j_{m-k+1}) in lexicographic order. This is impossible by how we constructed the F_k's. Thus i_{m-k+1} is the root. Also, if j_{m-k+1} were not a root, then adding the edge (i_{m-k+1}, j_{m-k+1}) would create a subforest of F which is not increasing. Thus, $F_0 \lessdot F_1 \lessdot \cdots \lessdot F_m$ is a maximal chain in the interval $[\hat{0}, F]$ which has an ascent-free word of labels. Since every chain in $[\hat{0}, F]$ is labeled by $(i_m, j_m), (i_{m-1}, i_{j-1}), \cdots, (i_1, j_1)$ and since the word of labels completely determines the chain, we see that $[\hat{0}, F]$ has a unique ascent-free maximal chain. Now observe that every interval $[F, G]$ is isomorphic to an interval of the form $[\hat{0}, F']$, where F' is obtained from F by contracting the vertices from each connected component of F and labeling the contracted set of vertices with the label of the root of the component. Therefore, the above argument shows that every interval has a unique ascent-free maximal chain and hence λ^* is an ER*-labeling. □

Proposition 4. *The posets Π_n and $\mathcal{ISF}_n$ are Whitney duals.*

Proof. First we explain why using λ and λ^* there are the same number of increasing saturated chains starting at $\hat{0}$ in Π_n and $\mathcal{ISF}_n$ and the same number of ascent-free saturated chains starting at $\hat{0}$ in Π_n and $\mathcal{ISF}_n$. Let $\mathbf{c}$ be a saturated chain starting at $\hat{0}$ in Π_n. As we move up the chain, the labels indicate the minimum elements of the two blocks being merged. The sequence of labels as we move up also describes a way to create an increasing spanning forest. Here the blocks of the partitions correspond to the trees in the forest. When we merge two blocks we connect the roots of the two corresponding trees. It follows that for every saturated chain in Π_n starting at $\hat{0}$, there is a saturated chain in $\mathcal{ISF}_n$ starting at $\hat{0}$ with the same word of labels. This chain is unique since saturated chains in $\mathcal{ISF}_n$ starting at $\hat{0}$ are completely determined by their word of labels. This correspondence between saturated chains in Π_n starting at $\hat{0}$ and saturated chains starting at $\hat{0}$ in $\mathcal{ISF}_n$ is therefore a bijection. Moreover, this bijection preserves the word of labels. Therefore there are the same number of increasing (ascent-free) saturated chains in Π_n starting at $\hat{0}$ and increasing (ascent-free) saturated chains in $\mathcal{ISF}_n$ starting at $\hat{0}$. Proposition 1, implies that Π_n and $\mathcal{ISF}_n$ are Whitney duals. □

As we will see in the next section, the situation for constructing Whitney duals for general posets follows the same principle that we just described for Π_n and $\mathcal{ISF}_n$.

Suppose we want to show that P and Q are Whitney duals. We define edge labelings on P and Q in such a way that the labeling on P is an ER-labeling and the one on Q is an ER*-labeling. If these posets, under the two labelings, have the same number of increasing (ascent-free) saturated chains starting at $\hat{0}$, then Proposition 1 implies that P and Q are Whitney duals.

3 Constructing Whitney Duals

3.1 $\overline{\text{EW}}$-labelings

In this section, we will show how to construct Whitney duals for certain posets. We will need edge labelings with the following property.

Definition 5. Let λ be an ER-labeling. We say λ has the *rank two switching property* provided that for every interval $[x, y]$ with $\rho(y) - \rho(x) = 2$, if ab is the word of labels of the unique increasing chain in the interval, then there exists a unique chain in $[x, y]$ whose word of labels is ba.

In Figure 1a, one can see that the labeling of Π_3 given in Definition 3 has the rank two switching property. Indeed, the increasing chain in the unique rank two interval of Π_3 is labeled by $(1, 2), (1, 3)$ and there is a unique chain labeled by $(1, 3), (1, 2)$. In fact, Π_n has the rank two switching property for all $n \geq 1$. One can verify this using the fact that every rank two interval is either isomorphic to Π_3 or a boolean algebra of rank two. More generally, for every geometric lattice there is a labeling which has the rank two switching property (see Proposition 10). By repeatedly applying the rank two switching property we get the following lemma.

Lemma 1. *Let P be a poset and let λ be an ER-labeling with the rank two switching property. For every interval $[x, y]$ and every maximal chain $\mathbf{c}$ in $[x, y]$ there exists an ascent-free maximal chain $\mathbf{c}'$ in $[x, y]$ with the same multiset of labels as $\mathbf{c}$.*

We now introduce the notion of an $\overline{\text{EW}}$-labeling, where the letter "W" comes from the fact that it will provide sufficient conditions to construct a Whitney dual.

Definition 6. Let $\lambda : \mathcal{E}(P) \to \Lambda$ be an ER-labeling of P with Λ a total order. We say λ is an *$\overline{EW}$-labeling* if the following holds.

1. λ has the rank two switching property.
2. For each interval $[x, y]$, if $\mathbf{c}$ and $\mathbf{c}'$ are distinct ascent-free maximal chains in $[x, y]$, then the multisets of labels for $\mathbf{c}$ and $\mathbf{c}'$ are different.

Remark 1. The reason for using an overline on $\overline{\text{EW}}$ is that in a forthcoming paper on Whitney duals we define an EW-labeling by relaxing condition 2 in Definition 6 and the requirement of Λ to be a totally ordered set. In the case of geometric lattices these two stronger conditions greatly simplify the discussion and proofs.

As we commented after Definition 5, the labeling of Π_n given in Definition 3 has the rank two switching property. Moreover, the edge labels come from a lexicographical order which is a total order. Finally, the sequence along any maximal chain uniquely identifies that chain and so any two different ascent-free maximal chains in any interval, cannot have the same set of labels. It follows that the labeling in Definition 3 is an $\overline{\text{EW}}$-labeling.

We now turn our attention to quotient posets which is the next piece of the puzzle we need to construct Whitney duals.

3.2 Quotient posets

Definition 7. Let P be a graded poset and let $\sim$ be an equivalence relation on P such that if $x \sim y$, then $\rho(x) = \rho(y)$. We define the *quotient poset* $P/\sim$ to be the set of equivalence classes ordered by $X \leq Y$ if and only if there exists $x \in X$, $y \in Y$ and $z_1, z_2, \ldots, z_k \in P$ such that

$$x \sim z_1 \leq z_2 \sim z_3 \leq \cdots \leq z_{k-2} \sim z_{k-1} \leq z_k \sim y.$$

The next proposition follows from Definition 7.

Proposition 5. *Let P be a graded poset and let $\sim$ be an equivalence relation on P such that if $x \sim y$, then $\rho(x) = \rho(y)$. Then we have the following.*

1. *$P/\sim$ is a graded poset.*
2. *For $X \in P/\sim$, we have $\rho(X) = \rho(x)$ for all $x \in X$.*
3. *For $X, Y \in P/\sim$, $X \lessdot Y$ if and only if $x \lessdot y$ for some $x \in X$ and $y \in Y$.*

Given a poset P, let $C(P)$ denote the poset whose elements are saturated chains of P starting at $\hat{0}$ ordered by inclusion. We call $C(P)$ the *chain poset of P*. Figure 3 depicts Π_3 and $C(\Pi_3)$. Suppose that λ is an edge labeling of P. If $\mathbf{c} \in C(P)$, we write $S(\mathbf{c})$ to denote the underlying multiset of labels along $\mathbf{c}$. Additionally, we write $e(\mathbf{c})$ for the element of P where $\mathbf{c}$ terminates. Let $\sim_\lambda$ be the equivalence relation on $C(P)$ defined by $\mathbf{c}_1 \sim_\lambda \mathbf{c}_2$ whenever, $e(\mathbf{c}_1) = e(\mathbf{c}_2)$ and $S(\mathbf{c}_1) = S(\mathbf{c}_2)$. We use $Q_\lambda(P)$ to denote $C(P)/\sim$. In our example of Π_3, every element of $C(\Pi_3)$ is in its own equivalence class except for the chains $1/2/3 \lessdot 12/3 \lessdot 123$ and $1/2/3 \lessdot 13/2 \lessdot 123$ since they both terminate at 123 and have the same underlying (multi)set of labels. Taking the quotient to obtain $Q_\lambda(\Pi_3)$, we see that we get a Whitney dual of Π_3 isomorphic to $\mathcal{ISF}_3$. Figure 3c depicts $Q_\lambda(\Pi_3)$

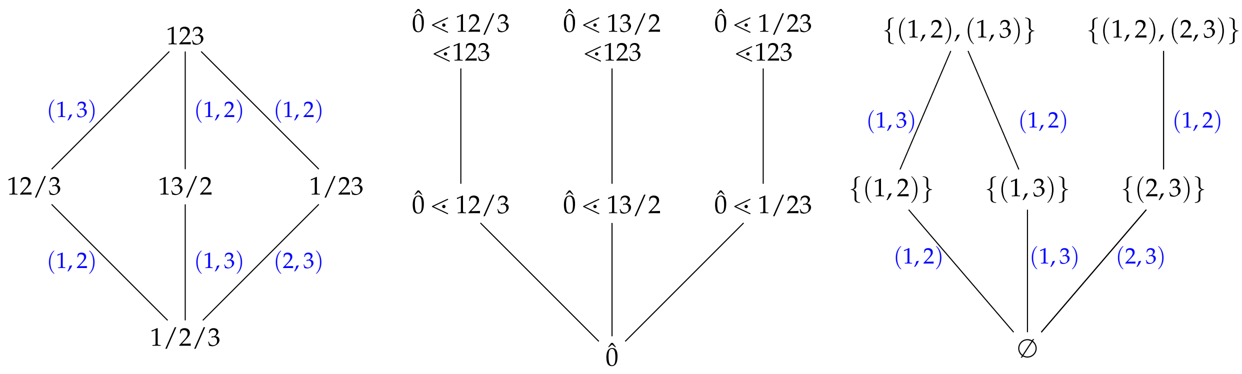

(a) Π_3 with labeling λ (b) $C(\Pi_3)$ (c) $Q_\lambda(\Pi_3)$ with labeling λ^*.

Figure 3

where we have identified the equivalence classes by the underlying set of labels on the chains.

Note that by the definition of $\sim_\lambda$ and Lemma 1, each equivalence class $X \in Q_\lambda(P)$ corresponds to a unique ascent-free maximal chain in $[\hat{0}, e(X)]$. In fact this correspondence is a bijection between ascent-free saturated chains starting at $\hat{0}$ in P of length k and equivalence classes in $Q_\lambda(P)$ of rank k. The following proposition immediately follows.

Proposition 6. *Let λ be an $\overline{EW}$-labeling. Then $|w_K(P)| = W_k(Q_\lambda(P))$.*

Suppose that λ is an edge labeling of P. We will define an edge labeling λ^* on $Q_\lambda(P)$ that depends on λ. Note that by definition, $\mathbf{c} \sim_\lambda \mathbf{c}'$ implies $S(\mathbf{c}) = S(\mathbf{c}')$. In light of this, we will use $S(X)$ to denote the multiset of labels in any chain in X. Moreover, if $X \lessdot Y$ in $Q_\lambda(P)$ then there exists a unique element in $S(Y) \setminus S(X)$. Define an edge labeling λ^* on $Q_\lambda(P)$ by

$$\lambda^*(X \lessdot Y) = S(Y) \setminus S(X). \tag{3.1}$$

This edge labeling for $Q_\lambda(\Pi_3)$ appears in Figure 3c.

Since all the saturated chains inside an equivalence class X of $Q_\lambda(P)$ terminate at the same element of P, we can define $e(X)$ to be $e(\mathbf{c})$ for any $\mathbf{c} \in X$. Additionally, we will use $\mathcal{C}(x,y)$ for the set of maximal chains in $[x,y]$ and $\mathcal{C}_S(x,y)$ for the set of maximal chains in $[x,y]$ whose multiset of labels is S. For example, for Π_3 with labeling λ depicted in Figure 1a, $\mathcal{C}_{\{(1,2),(1,3)\}}(1/2/3, 123) = \{1/2/3 \lessdot 12/3 \lessdot 123, 1/2/3 \lessdot 13/2 \lessdot 123\}$.

Proposition 7. *Let λ be an $\overline{EW}$-labeling of P, $[X,Y]$ an interval in $Q_\lambda(P)$, and $S = S(Y) \setminus S(X)$. There is a bijection $\varphi : \mathcal{C}(X,Y) \to \mathcal{C}_S(e(X), e(Y))$ that preserves the word of labels on the chains.*

Remark 2. It is important to first fix the interval $[X,Y]$. In general it is not true that chains in $[x,y]$ map to unique chains in $Q_\lambda(P)$.

Proof. Let $\varphi : \mathcal{C}(X,Y) \to \mathcal{C}_S(e(X),e(Y))$ be defined by sending $X = X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k = Y$ to $e(X) = e(X_1) \lessdot e(X_2) \lessdot \cdots \lessdot e(X_k) = e(Y)$. We claim φ is well-defined. To see this, note that if $X_i \lessdot X_{i+1}$ then by Proposition 5 there are chains $\mathbf{c} \in X_i$ and $\mathbf{c}' \in X_{i+1}$ with $\mathbf{c} \lessdot \mathbf{c}'$ in $C(P)$. By definition of $C(P)$ we have that $e(\mathbf{c}) \lessdot e(\mathbf{c}')$ and $\mathbf{c}' = \mathbf{c} \cup \{e(\mathbf{c}')\}$. Therefore $e(X_1) \lessdot e(X_2) \lessdot \cdots \lessdot e(X_k)$ is saturated chain in $[e(X),e(Y)]$. Recall that $\lambda^*(X_i \lessdot X_{i+1})$ is the unique element of $S(X_{i+1}) \setminus S(X_i)$. But then this must also be $S(\mathbf{c}') \setminus S(\mathbf{c})$ which is exactly the value of $\lambda(e(X_i) \lessdot e(X_{i+1}))$. It follows that $\varphi(X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k)$ is a maximal chain in $[e(X),e(Y)]$ with label multiset $S(Y) \setminus S(X)$, and hence φ is well-defined. This argument also shows that φ preserves the word of labels.

Next, we show φ is injective. Suppose that $X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k$ and $X_1' \lessdot X_2' \lessdot \cdots \lessdot X_k'$ are maximal chains in $[X,Y]$ which have the same image under φ. Then $e(X_i) = e(X_i')$ for all i. If $X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k \neq X_1' \lessdot X_2' \lessdot \cdots \lessdot X_k'$, then there is a minimal i such that $X_i \neq X_i'$. Since $X_1 = X = X_1'$, it must be that $i > 1$. But then $S(X_i) \setminus S(X_{i-1}) = \lambda(e(X_{i-1}) \lessdot e(X_i)) = S(X_i') \setminus S(X_{i-1}')$. Since $X_{i-1} = X_{i-1}'$, this implies that $S(X_i) = S(X_i')$. However, if $e(X_i) = e(X_i')$ and $S(X_i) = S(X_i')$, then $X_i = X_i'$, a contradiction. It follows that φ is injective.

Finally, we show φ is surjective. Let $e(X) = x_1 \lessdot x_2 \lessdot \cdots \lessdot x_k = e(Y)$ be a maximal chain in $[e(X),e(Y)]$ with label multiset $S = S(Y) \setminus S(X)$. Additionally, let $\mathbf{c}$ be any element of X. Define $\mathbf{c}_i$ to be the chain $\mathbf{c} \cup \{x_1, x_2, \ldots, x_i\}$. Then $\mathbf{c} = \mathbf{c}_1 \lessdot \mathbf{c}_2 \lessdot \mathbf{c}_2 \lessdot \cdots \lessdot \mathbf{c}_k$ is a saturated chain in $C(P)$ with $e(\mathbf{c}_i) = x_i$. Note that by construction, $S(\mathbf{c}_k) = S(Y)$. Let X_i be the equivalence class in $Q_\lambda(P)$ containing $\mathbf{c}_i$, then $X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k$ is a saturated chain in $Q_\lambda(P)$ with $e(X_i) = x_i$. Clearly $X_1 = X$. Moreover, we have that $e(X_k) = x_k = e(Y)$ and $S(X_k) = S(\mathbf{c}_k) = S(Y)$. Therefore, $X_k = Y$. So we have that $X = X_1 \lessdot X_2 \lessdot \cdots \lessdot X_k = Y$ is a chain in $[X,Y]$ and it is a preimage of $x_1 \lessdot x_2 \lessdot \cdots \lessdot x_k$ implying φ is surjective. □

By Lemma 1 and Definition 6, if P has an $\overline{\mathrm{EW}}$-labeling, then for every interval $[x,y]$ in P and every maximal chain $\mathbf{c}$ in $[x,y]$ with label multiset S, there is a unique ascent-free maximal chain in $[x,y]$ with label multiset S. Therefore, Proposition 7 implies that every interval $[X,Y]$ in $Q_\lambda(P)$ has a unique ascent-free chain. Thus we have the following.

Proposition 8. *Let P be a poset and let λ be an $\overline{\mathrm{EW}}$-labeling. Then λ^* defined as in* (3.1) *is an ER*-labeling of $Q_\lambda(P)$.*

Assume now that λ is an $\overline{\mathrm{EW}}$-labeling and that $X^1, X^2, \ldots, X^n$ are the different elements of $Q_\lambda(P)$ that satisfy $e(X^i) = x$ for all i. If there is an increasing maximal chain in $[\hat{0}, X^i]$, then by Proposition 7 there is an increasing maximal chain in $[\hat{0}, x]$ with label multiset $S(X^i)$. Since there is exactly one increasing maximal chain in $[\hat{0}, x]$ and since (by the definition of $\sim_\lambda$) $S(X^i) \neq S(X^j)$ for all $i \neq j$, there is exactly one X^i such that $[\hat{0}, X^i]$ has an increasing maximal chain. We then have that the number of increasing saturated chains starting at $\hat{0}$ and of length k in $Q_\lambda(P)$ is equal to the number of elements $x \in P$

with $\rho(x) = k$ and hence equal to the number of increasing saturated chains starting at $\hat{0}$ and of length k in P. Proposition 1 then implies the following proposition.

Proposition 9. *Let P be a poset and λ be an $\overline{\text{EW}}$-labeling of P. Then $W_k(P) = |w_K(Q_\lambda(P))|$.*

Combining Propositions 6 and 9 together we obtain our main theorem.

Theorem 2. *Let P be a poset with an $\overline{\text{EW}}$-labeling λ. Then $Q_\lambda(P)$ is a Whitney dual of P.*

3.3 Whitney duals of geometric lattices

Next we show that every geometric lattice has an $\overline{\text{EW}}$-labeling. The labeling we discuss was originally introduced by Björner [2].

Definition 8. Let L be a geometric lattice with set of atoms $A(L)$. Fix a total order on $A(L)$. Now define $\lambda : \mathcal{E}(L) \to A(L)$ by setting $\lambda(x \lessdot y) = a$ where a is the smallest atom such that $x \vee a = y$. We call λ a *minimum labeling of* L.

Proposition 10. *For any geometric lattice L a minimum labeling of L is an $\overline{\text{EW}}$-labeling.*

Proof. It was shown in [2] that a minimum labeling is an ER-labeling. Also, for any interval $[x, y]$ the labels along any maximal chain uniquely determine the chain since one can read off the elements of the chain by taking joins of x with the labels along the chain. Thus it suffices to show that a minimum labeling has the rank two switching property.

Let λ be a minimum labeling of L, let $[x, y]$ be an interval of rank two and suppose that ij is the word of labels of the increasing chain, $x \lessdot x \vee i \lessdot x \vee i \vee j = y$. Since L is geometric and j is an atom not underneath x, $x \lessdot x \vee j \lessdot y$. Observe that $\lambda(x \lessdot x \vee j) = j$, since if this was not the case, this would imply $\lambda(x \vee i \lessdot y) < j$ which is a contradiction. Moreover, i is not below $x \vee j$ and i is below y. Since there is a unique increasing chain in $[x, y]$, i is the smallest atom that appears as a label in $[x, y]$. It follows that $\lambda(x \vee j \lessdot y) = i$. We conclude that the chain $x \lessdot x \vee j \lessdot y$ has the word of labels ji. Since joins are unique, there is only one chain in $[x, y]$ with word of labels ji and thus any minimum labeling satisfies the rank two switching property. □

We have the following theorem as a corollary.

Theorem 3. *Every geometric lattice has a Whitney dual.*

4 Future Work

The conditions necessary for an $\overline{\text{EW}}$-labeling are somewhat restrictive. While the rank two switching property is fundamental to our construction of Whitney duals, the other conditions can be relaxed. In upcoming work, the authors will explore posets which have a more flexible version of an $\overline{\mathrm{EW}}$-labeling. This allows one to construct Whitney duals for a larger family of posets. The weighted partition poset, which motivated this work, is part of this family.

Acknowledgements

We thank the anonymous referees for their very useful comments. In particular, we are thankful for the suggested simpler proof of Proposition 10.

References

[1] A. Björner. "Shellable and Cohen-Macaulay partially ordered sets". *Trans. Amer. Math. Soc.* **260** (1980), pp. 159–183. DOI.

[2] A. Björner. "On the homology of geometric lattices". *Algebra Universalis* **14** (1982), pp. 107–128. DOI.

[3] A. Björner and M. L. Wachs. "On lexicographically shellable posets". *Trans. Amer. Math. Soc.* **277** (1983), pp. 323–341. DOI.

[4] R. S. González D'León and M. L. Wachs. "On the (co)homology of the poset of weighted partitions". *Trans. Amer. Math. Soc.* **368** (2016), pp. 6779–6818. DOI.

[5] J. Hallam, J. L. Martin, and B. E. Sagan. "Increasing spanning forests in graphs and simplicial complexes". 2016. arXiv:1610.05093.

[6] D. L. Reiner. "The combinatorics of polynomial sequences". *Stud. Appl. Math.* **58** (1978), pp. 95–117. DOI.

[7] B. E. Sagan. "A note on Abel polynomials and rooted labeled forests". *Discrete Math.* **44** (1983), pp. 293–298. DOI.

[8] R. P. Stanley. *Enumerative Combinatorics, Vol. 1*. 2nd ed. Cambridge Studies in Advanced Mathematics, Vol. 49. Cambridge University Press, 2012.

[9] M. L. Wachs. "Poset topology: tools and applications". *Geometric Combinatorics*. IAS/Park City Math. Ser., Vol. 13. Amer. Math. Soc., 2007, pp. 497–615.