An Explicit Derivation of the Möbius Function for Bruhat Order

Brant C. Jones

Abstract We give an explicit nonrecursive complete matching for the Hasse diagram of the strong Bruhat order of any interval in any Coxeter group. This yields a new derivation of the Möbius function, recovering a classical result due to Verma.

1 Introduction

The Bruhat partial order on the elements of a Coxeter group is a fundamental tool in algebraic combinatorics, representation theory and the geometry of Schubert varieties. In this work, we give a derivation of the Möbius function for this partial order based on an explicit nonrecursive matching of the Hasse diagram. The Möbius function is used to invert formulas defined by sums over Bruhat intervals, and gives the Euler characteristic in poset topology. Many proofs of the Möbius function have appeared in the literature; see [4, 7, 11, 12, 15, 16].

Our construction is closest to Verma's original argument, although it is phrased in terms of combinatorial objects called masks that are related to Kazhdan–Lusztig combinatorics. In [16], Verma constructs a complete matching of the Hasse diagram of the Bruhat interval $[x, w]$ in "half" the cases: when there exists a Coxeter generator s_i such that $xs_i > x$ and $ws_i < w$. In the other cases, he applies an inductive argument to prove the Möbius function formula, but this argument does not extend to give a complete matching of the Bruhat interval. The complete matching that we give below can be seen to agree with Verma's in the case that there exists s_i satisfying $xs_i > x$

B. C. Jones (✉)
Department of Mathematics, One Shields Avenue, University of California,
Davis, CA 95616, USA
e-mail: brant@math.ucdavis.edu
URL: http://math.ucdavis.edu/~brant/

and $ws_i < w$. This case is also an example of a special matching that has been used to compute R-polynomials in Kazhdan–Lusztig theory; see [2, Proposition 5.6.1]. In addition, [14] have used a complete matching of the intervals in finite Coxeter groups in order to apply discrete Morse theory to totally nonnegative flag varieties. We show that our construction also agrees with this matching for the case of finite Coxeter groups.

Our matching unifies these constructions, and has the advantage of being given explicitly and nonrecursively. It also extends to intervals in infinite Coxeter groups.

2 Construction

Let W be a Coxeter group with generating set S and relations of the form $(s_i s_j)^{m(i,j)} = 1$. An *expression* is any product of generators from S and the *length $l(w)$* is the minimum length of any expression for the element w. Such a minimum length expression is called *reduced*. Given $w \in W$, we represent reduced expressions for w in sans serif font, say $\mathsf{w} = \mathsf{w}_1 \mathsf{w}_2 \cdots \mathsf{w}_p$ where each $\mathsf{w}_i \in S$. For any $x, w \in W$, we say that $x \leq w$ *in Bruhat order* if a reduced expression for x appears as a subword (that is not necessarily consecutive) of some reduced expression for w. There are several other characterizations of this partial order on the elements of W; see [2, 10] for details. If s_i appears as the last (first, respectively) factor in some reduced expression for w, then we say that s_i is a *right (left, respectively) descent* for w; otherwise, s_i is an *right (left, respectively) ascent* for w. If s_i is a descent for an element w with reduced expression $\mathsf{w} = \mathsf{w}_1 \mathsf{w}_2 \cdots \mathsf{w}_p$ then the *Exchange Condition* implies that there exists an index i for which $ws_i = \mathsf{w}_1 \cdots \mathsf{w}_{i-1} \widehat{\mathsf{w}}_i \mathsf{w}_{i+1} \cdots \mathsf{w}_p$, where the hat indicates omission.

The following lemma gives a useful property of Bruhat order.

Lemma 2.1 (Lifting Lemma) [2, Proposition 2.2.7] *Suppose $x < w$, s_i is a right descent for w, and s_i is a right ascent for x. Then, $xs_i \leq w$ and $ws_i \geq x$.*

In this work, we will represent Bruhat relations using a combinatorial model inspired by Deodhar [8] and Billey–Warrington [5] for the purpose of studying Kazhdan–Lusztig polynomials. Fix a reduced expression $\mathsf{w} = \mathsf{w}_1 \mathsf{w}_2 \cdots \mathsf{w}_p$. Define a *mask σ* associated to the reduced expression w to be any binary vector $(\sigma_1, \ldots, \sigma_p)$ of length $p = l(w)$. Every mask corresponds to a subexpression of w defined by $\mathsf{w}^\sigma = \mathsf{w}_1^{\sigma_1} \cdots \mathsf{w}_p^{\sigma_p}$ where

$$\mathsf{w}_j^{\sigma_j} = \begin{cases} \mathsf{w}_j & \text{if } \sigma_j = 1 \\ 1 & \text{if } \sigma_j = 0. \end{cases}$$

Each w^σ is a product of generators so it determines an element of W. For $1 \leq j \leq p$, we also consider initial sequences of a mask denoted $\sigma[j] = (\sigma_1, \ldots, \sigma_j)$, and the corresponding initial subexpression $\mathsf{w}^{\sigma[j]} = \mathsf{w}_1^{\sigma_1} \cdots \mathsf{w}_j^{\sigma_j}$. In particular, we have $\mathsf{w}^{\sigma[p]} = \mathsf{w}^\sigma$. We also use this notation to denote initial sequences of expressions, so $\mathsf{w}[j] = \mathsf{w}_1 \cdots \mathsf{w}_j$.

We say that a position j (for $2 \leq j \leq p$) of the fixed reduced expression w is a *defect* with respect to the mask σ if

$$\mathsf{w}^{\sigma[j-1]} \mathsf{w}_j < \mathsf{w}^{\sigma[j-1]}.$$

Note that the defect status of position j does not depend on the value of σ_j. We say that a defect position is a *0-defect* if it has mask-value 0, and call it a *1-defect* if it has mask-value 1. If a mask has no defect positions at all, then we say it is a *constant mask on the reduced expression* w *for the element* w^σ. This terminology arises from the fact that these masks correspond precisely to the unique constant term in the Kazhdan–Lusztig polynomial $P_{x,w}(q)$ in the combinatorial model mentioned above. Other authors [13, 14] have used the term "positive distinguished subexpression" to define an equivalent notion.

The following result is due to Deodhar [8, Proposition 2.3(iii)], and has also appeared in work of [13] related to totally nonnegative flag varieties, as well as [1] in the context of sorting algorithms on Coxeter groups. As the lemma is central to our work, we include a proof here for completeness.

Lemma 2.2 (Deodhar) *Let* $\mathsf{w} = \mathsf{w}_1 \cdots \mathsf{w}_p$ *be a reduced expression for an element* $w \in W$ *and let* $x \leq w$. *Then there is a unique constant mask* σ *on* w *for* x.

Proof We describe a greedy algorithm to construct such a mask. Let $r_{p+1}(x) = x$ and $i = p$. We inductively assign

$$\sigma_i := \begin{cases} 0 & \text{if } \mathsf{w}_i \text{ is a right ascent for } r_{i+1}(x) \\ 1 & \text{if } \mathsf{w}_i \text{ is a right descent for } r_{i+1}(x) \end{cases}$$

$$\text{and} \quad r_i(x) := \begin{cases} r_{i+1}(x) & \text{if } \mathsf{w}_i \text{ is a right ascent for } r_{i+1}(x) \\ r_{i+1}(x) \cdot \mathsf{w}_i & \text{if } \mathsf{w}_i \text{ is a right descent for } r_{i+1}(x) \end{cases}$$

for each i from p down to 1.

Note that the constraint that σ have no defects forces the choice of mask-value at each step. Hence, there can be at most one mask σ on w for x. In particular, the algorithm produces a constant mask on w for x if and only if $r_1(x)$ is the identity.

We claim that $r_1(x)$ is always the identity. Note that we have a constant mask consisting of all 1 entries for $x = w$. Hence, if $x < w$ and we run the algorithm for both elements simultaneously, we initially have $r_{p+1}(x) = x \leq w = r_{p+1}(w)$. Observe that for each $i \leq p$, whenever we have $r_{i+1}(x) \leq r_{i+1}(w)$ then $r_i(x) \leq r_i(w)$. This follows by definition when $r_i(w) = r_{i+1}(w)$, and by an application of the Lifting Lemma 2.1 in the case that $r_{i+1}(w)$ covers $r_i(w)$ using the fact that $r_i(x)$ always has w_i as a right ascent by construction. Since $r_1(w) = 1$, this implies by induction that $r_1(x) = 1$ so the algorithm produces a constant mask for all $x < w$.   □

Example 2.3 If $W = A_4$, $\mathsf{w} = s_2 s_3 s_4 s_1 s_2 s_3$ and $x = s_1 s_2 s_1$ then σ is

s_2	s_3	s_4	s_1	s_2	s_3
1	0	0	1	1	0

as a result of

$$r_7(x) = s_1 s_2 s_1 = r_6(x), r_5(x) = s_2 s_1, r_4(x) = s_2 = r_3(x) = r_2(x), r_1(x) = 1.$$

Remark 2.4 Suppose σ and τ are constant masks on a fixed reduced expression w. Then it can be shown that the mask $\nu = \sigma \vee \tau$ defined by

$$\nu_i := \begin{cases} 1 & \text{if } \sigma_i = 1 \text{ or } \tau_i = 1, \\ 0 & \text{otherwise.} \end{cases}$$

is a constant mask. Although Bruhat order is not a lattice, the operation $\vee$ can be used to define an associated join-semilattice that respects Bruhat order, once we fix a reduced expression w.

In fact, [1] has shown that this partial order on $[1, \mathsf{w}]$ is a lattice that lies maximally between the weak and strong Bruhat orders on W.

Continue to fix the reduced expression $\mathsf{w} = \mathsf{w}_1 \cdots \mathsf{w}_p$ for $w \in W$ and suppose $y \le x \le w$. We describe a notion of relative mask that captures this pair of Bruhat relations. Let τ be the unique constant mask on w for x. Then, w^τ is a reduced expression for x and we may let ν be the unique constant mask on w^τ for y. We combine these into a *relative mask* $\sigma = (\sigma_1, \ldots, \sigma_p)$ by

$$\sigma_j = \begin{cases} X & \text{if } \tau_j = 0 \\ \nu_j & \text{if } \tau_j = 1. \end{cases}$$

In this situation, we call τ the *X-mask* associated to (w, σ), also denoted $\Xi(\sigma)$. We denote $(\mathsf{w}^\tau)^\nu$ by w^σ. We say that position j is a *defect* in the relative mask σ if $\mathsf{w}^{\sigma[j-1]}\mathsf{w}_j < \mathsf{w}^{\sigma[j-1]}$. Note that only positions in σ with mask-value X can be defects, by definition. We will indicate these defect positions by X^d in our illustrations of relative masks.

Example 2.5 The relative masks encoding the Bruhat interval $[s_2, s_2 s_1 s_3 s_2]$ in type A are given by

	s_2	s_1	s_3	s_2		s_2	s_1	s_3	s_2	subexpression for $x \in [s_2, s_2 s_1 s_3 s_2]$
$\sigma =$	0	0	0	1	$\tau =$	1	1	1	1	$s_2 s_1 s_3 s_2$
$\sigma =$	1	0	0	X^d	$\tau =$	1	1	1	0	$s_2 s_1 s_3$
$\sigma =$	0	0	X	1	$\tau =$	1	1	0	1	$s_2 s_1 s_2$
$\sigma =$	0	X	0	1	$\tau =$	1	0	1	1	$s_2 s_3 s_2$
$\sigma =$	X	0	0	1	$\tau =$	0	1	1	1	$s_1 s_3 s_2$
$\sigma =$	1	0	X	X^d	$\tau =$	1	1	0	0	$s_2 s_1$
$\sigma =$	1	X	0	X^d	$\tau =$	1	0	1	0	$s_2 s_3$
$\sigma =$	X	0	X	1	$\tau =$	0	1	0	1	$s_1 s_2$
$\sigma =$	X	X	0	1	$\tau =$	0	0	1	1	$s_3 s_2$
$\sigma =$	X	X	X	1	$\tau =$	0	0	0	1	s_2

Here, $\mathsf{w}^\sigma = s_2$ for all of these masks.

Our goal is to give a matching on the Hasse diagram of the Bruhat interval $[y, w]$ using the relative masks for y on a fixed reduced expression for w as an encoding. The following definition will allow us to define a procedure that is reversible.

Definition 2.6 Let $\mathsf{w} = \mathsf{w}_1 \cdots \mathsf{w}_p$ be a reduced expression for an element $w \in W$ and let σ be a relative mask on w with X-mask τ. We call position j a *shifted descent* of (w, σ) if $\mathsf{w}^{\tau[j-1]} \geq \mathsf{w}^{\sigma[j]}$ in Bruhat order.

Example 2.7 Consider

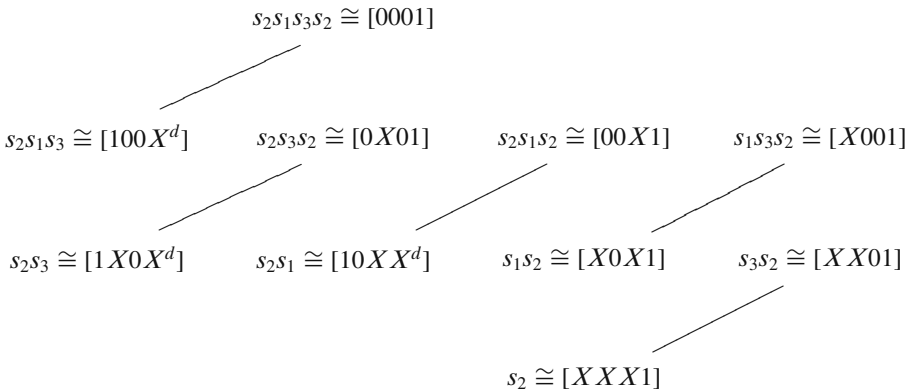|		s_2	s_1	s_3	s_2		s_2	s_1	s_3	s_2	subexpression for
											$x \in [s_2, s_2 s_1 s_3 s_2]$
$\sigma =$		0	X	0	1	$\tau =$	1	0	1	1	$s_2 s_3 s_2$
$\sigma =$		X	0	0	1	$\tau =$	0	1	1	1	$s_1 s_3 s_2$

The first mask has position 4 as a shifted descent because $s_2 s_3 \geq s_2$. The second mask does not have position 4 as a shifted descent because $s_1 s_3 \not\geq s_2$.

We are now in a position to define our matching. The rough idea that motivates the following definition is to remove the rightmost X from a relative mask σ in a way that preserves w^σ and is reversible.

Definition 2.8 Let σ be a relative mask on w with X-mask τ. Find the rightmost position j in (w, σ) where one of the following conditions holds, and apply the given transformation to obtain a new relative mask denoted $\varphi(\sigma)$:

(1) If $\sigma_j = X$ and σ_j is not a defect then change σ_j to 0.
(2) If $\sigma_j = 0$ then change σ_j to X. Note that by definition, σ_j cannot be a defect in this case.
(3) If $\sigma_j = X$ and σ_j is a defect then $\mathsf{w}^{\tau[j-1]} \geq \mathsf{w}^{\sigma[j-1]} > \mathsf{w}^{\sigma[j-1]}\mathsf{w}_j$. Hence, we may assign the unique constant mask for $\mathsf{w}^{\sigma[j-1]}\mathsf{w}_j$ on $\mathsf{w}^{\tau[j-1]}$ to the entries of $\mathsf{w}^{\tau[j-1]}$ and set σ_j to 1.
(4) If $\sigma_j = 1$ and σ_j is a shifted descent then $\mathsf{w}^{\tau[j-1]} \geq \mathsf{w}^{\sigma[j]}$ so we may assign the unique constant mask for $\mathsf{w}^{\sigma[j]}$ on $\mathsf{w}^{\tau[j-1]}$ to the entries of $\mathsf{w}^{\tau[j-1]}$ and set σ_j to X. Note that by definition, σ_j becomes a defect in this case.

Example 2.9 The matching given by φ on $[s_2, s_2 s_1 s_3 s_2]$ is:

$$s_2 s_1 s_3 s_2 \cong [0001]$$

$s_2 s_1 s_3 \cong [100 X^d]$ $s_2 s_3 s_2 \cong [0X01]$ $s_2 s_1 s_2 \cong [00X1]$ $s_1 s_3 s_2 \cong [X001]$

$s_2 s_3 \cong [1X0X^d]$ $s_2 s_1 \cong [10XX^d]$ $s_1 s_2 \cong [X0X1]$ $s_3 s_2 \cong [XX01]$

$$s_2 \cong [XXX1]$$

We now give our main results.

Lemma 2.10 *The function φ given in Definition 2.8 always produces a valid relative mask.*

Proof Let σ be a relative mask on w. Applying φ to σ interchanges exactly one X entry in position j with an entry that is either 0 or 1, and also possibly rearranges the 0 and 1 entries lying to the left of j. By definition, we never create a 0-defect nor 1-defect at j, and applying any of the rules at position j preserves the element $\mathsf{w}^{\sigma[j]}$, so the defect status of positions $k > j$ does not change. Hence, to show that $\varphi(\sigma)$ is valid it suffices to show that $\varphi(\sigma)$ has a constant X-mask. Let τ denote the X-mask of σ.

Suppose φ acts at position j by changing σ_j from X to 1 or 0. Then the X-mask τ' of $\varphi(\sigma)$ is the result of changing the rightmost 0 to a 1 in the X-mask τ of σ. Since τ is a constant mask, we have that τ' is constant by [2, Proposition 5.3.9].

Now consider the case where φ acts at position j by changing σ_j from 1 or 0 to X, and suppose for the sake of contradiction that the X-mask τ' of $\varphi(\sigma)$ is not constant. By Definition 2.8, we have $\sigma_i = 1$ for all $i > j$, for otherwise we would have applied φ to position i. Since τ' is not constant, there exists a leftmost position $k > j$ that becomes a 1-defect in τ'. Hence, we have the schematic shown below.

$$
\begin{array}{ccccccccccc}
& 1 & & \cdots & & j & & \cdots & & (k-1) & k \\
\sigma = & * & X & \cdots & * & * & 1 & \cdots & 1 & 1 & 1 \\
\tau' = & 1 & X & \cdots & 1 & X & 1 & \cdots & 1 & 1 & 1^d
\end{array}
$$

Here, the positions marked by $*$ are the non-X positions of σ, so these positions have mask-value 1 in τ'.

If $\sigma_j = 0$, then $\mathsf{w}^{\sigma[k-1]} \leq \mathsf{w}^{\tau'[k-1]}$ and w_k is a descent for $\mathsf{w}^{\tau'[k-1]}$ while w_k is an ascent for $\mathsf{w}^{\sigma[k-1]}$. Hence, by the Lifting Lemma 2.1 we have $\mathsf{w}^{\sigma[k]} \leq \mathsf{w}^{\tau'[k-1]} \leq \mathsf{w}^{\tau[k-1]}$ so k is a shifted descent in σ, contradicting the rightmost choice of move in Definition 2.8.

Next, suppose $\sigma_j = 1$, as shown in the schematic below.

$$
\begin{array}{ccccccccccc}
& 1 & & \cdots & & j & & \cdots & & (k-1) & k \\
\sigma = & * & X & \cdots & * & 1 & 1 & \cdots & 1 & 1 & 1 \\
\varphi(\sigma) = & * & X & \cdots & * & X^d & 1 & \cdots & 1 & 1 & 1 \\
\tau' = & 1 & X & \cdots & 1 & X & 1 & \cdots & 1 & 1 & 1^d
\end{array}
$$

Then, since φ operates on σ at position j, we have that j is a shifted descent. Therefore, position j in the mask $\varphi(\sigma)$ is X^d and $\mathsf{w}^{\varphi(\sigma)[j]} = \mathsf{w}^{\sigma[j]}$.

Hence, we have $\mathsf{w}^{\varphi(\sigma)[k-1]} \leq \mathsf{w}^{\tau'[k-1]}$ because we have exhibited one as a submask of the other without 1-defects. Moreover, w_k is a descent for $\mathsf{w}^{\tau'[k-1]}$ and an ascent for $\mathsf{w}^{\varphi(\sigma)[k-1]}$, so by the Lifting Lemma 2.1, we have $\mathsf{w}^{\varphi(\sigma)[k-1]}\mathsf{w}_k \leq \mathsf{w}^{\tau'[k-1]}$. Hence,

$$
\mathsf{w}^{\sigma[k]} = \mathsf{w}^{\sigma[k-1]}\mathsf{w}_k = \mathsf{w}^{\varphi(\sigma)[k-1]}\mathsf{w}_k \leq \mathsf{w}^{\tau'[k-1]} \leq \mathsf{w}^{\tau[k-1]}
$$

so k was a shifted descent in σ to begin with, contradicting the rightmost choice of move in Definition 2.8.

Since all cases where τ' is not constant lead to a contradiction, we have completed the proof that $\varphi(\sigma)$ is a valid relative mask.   □

Lemma 2.11 *The function φ given in Definition* 2.8 *is an involution on the set of relative masks on* w.

Proof To see that φ is an involution, we observe that rule (2) inverts rule (1) in Definition 2.8, and it is straightforward to verify that rule (4) inverts rule (3). Moreover, applying any of the rules at position j preserves the element $\mathsf{w}^{\sigma[j]}$, so the mask-value and defect status of positions $k > j$ does not change.

Since the rules in Definition 2.8 depend only on mask-value, defect status and shifted descent status, the only way in which applying a move at position j can create a new move at position $k > j$ is if k becomes a shifted descent in $\varphi(\sigma)$. Moreover, this can only occur as a result of applying rules (1) or (3) to σ.

For the sake of contradiction suppose this occurs and among all counterexamples, consider one such that $l(w)$ is minimal. Then, $\varphi(\sigma)$ operates at position j and $\varphi(\varphi(\sigma))$ operates at position $k > j$. Let τ be the X-mask of σ and τ' be the X-mask of $\varphi^2(\sigma)$. Thus, we have the following schematic.

$$
\begin{array}{ccccccccc}
 & 1 & \cdots & j & \cdots & i & \cdots & k-1 & k \\
\sigma = & * & \cdots & X^{(d)} & \cdots & 1 & \cdots & 1 & 1 \\
\varphi(\sigma) = & * & \cdots & * & \cdots & 1 & \cdots & 1 & 1 \\
\varphi^2(\sigma) = & * & \cdots & 1 & \cdots & 1 & \cdots & 1 & X^d
\end{array}
$$

We begin by justifying the main points of this schematic. Note that we do not assume the defect status of σ_j is known. If σ_j is a defect then $\varphi(\sigma)_j = 1$, and if it is not then $\varphi(\sigma)_j = 0$. In any case, observe that $\varphi^2(\sigma)_j$ must be 1, for otherwise the mask $\varphi^2(\sigma)$ shows that position k is already a shifted descent in σ. Also, observe that all the entries between j and k have mask-value 1 in σ and $\varphi(\sigma)$, for otherwise we contradict that j is the rightmost move in σ. By assuming that w is minimal length, we have that there are no other X entries in any of the masks because if there exists an X entry in one of the masks, it exists in all three of the masks, by virtue of the fact that we only adjust the X-masks at positions j and k as shown. Hence, any X-positions could be removed from all three masks simultaneously.

Next, we consider all possible cases of mask-values for σ and $\varphi^2(\sigma)$ on w_1.

Case $(\sigma_1 = 0$ and $\varphi^2(\sigma)_1 = 0)$ or $(\sigma_1 = 1$ and $\varphi^2(\sigma)_1 = 1)$. Here, we have

$$
\begin{array}{ccccccccc}
 & 1 & \cdots & j & j+1 & \cdots & k-1 & k \\
\sigma = & 0 & \cdots & X^{(d)} & 1 & \cdots & 1 & 1 \\
\varphi^2(\sigma) = & 0 & \cdots & 1 & 1 & \cdots & 1 & X^d
\end{array}
\qquad \text{or}
$$

$$
\begin{array}{ccccccccc}
 & 1 & \cdots & j & j+1 & \cdots & k-1 & k \\
\sigma = & 1 & \cdots & X^{(d)} & 1 & \cdots & 1 & 1 \\
\varphi^2(\sigma) = & 1 & \cdots & 1 & 1 & \cdots & 1 & X^d
\end{array}
$$

Let ν denote the restriction of σ to $\mathsf{w}_2 \cdots \mathsf{w}_k$. Then the restriction of $\varphi^2(\sigma)$ to $\mathsf{w}_2 \cdots \mathsf{w}_k$ shows that k becomes a shifted descent in $\varphi(\nu)$. Therefore, we obtain a counterexample on $\mathsf{w}_2 \cdots \mathsf{w}_k$, contradicting our minimal length choice of w.

Case $\sigma_1 = 0$ and $\varphi^2(\sigma)_1 = 1$. In this case, w_1 is a left descent for w^σ.

$$
\begin{array}{ccccccccc}
 & 1 & \cdots & j & j+1 & \cdots & k-1 & k \\
\sigma = & 0 & \cdots & X^{(d)} & 1 & \cdots & 1 & 1 \\
\varphi^2(\sigma) = & 1 & \cdots & 1 & 1 & \cdots & 1 & X^d
\end{array}
$$

By the Exchange Condition, there exists some position i such that $\mathsf{w}_1 \mathsf{w}^\sigma = \mathsf{w}^\nu$ where ν is obtained from σ by changing a single mask-value 1 entry at w_i to have mask-value 0. Observe that if $i > j$ then $\mathsf{w}_1 \mathsf{w}^{\sigma[i-1]} = \mathsf{w}^{\sigma[i]}$ so changing σ_1 to 1 would witness that i was a shifted descent in σ, a contradiction.

The mask ν may be not be a constant mask on w^τ, but there exists a unique constant mask γ for the element w^ν on $\mathsf{w}_2^{\tau_2} \cdots \mathsf{w}_k^{\tau_k}$ and γ still has mask-value 1 on $\mathsf{w}_{j+1} \cdots \mathsf{w}_k$ by the algorithm from Lemma 2.2. By abuse of notation, let γ denote the corresponding relative mask on $\mathsf{w}_2 \cdots \mathsf{w}_k$.

If there exists a shifted descent in position m of γ where $j < m \le k$, then m must have been a shifted descent in σ, a contradiction. To see this, observe that $\mathsf{w}^{\sigma[m]} = \mathsf{w}_1 \mathsf{w}^{\gamma[m]}$ because $\mathsf{w}^\sigma = \mathsf{w}_1 \mathsf{w}^\gamma$ and these reduced expressions agree in positions $m, \ldots, k$. Therefore, if $\mathsf{w}^{\gamma[m]} \le \mathsf{w}^{\tau[m-1]}$ then $\mathsf{w}_1 \mathsf{w}^{\sigma[m]} \le \mathsf{w}^{\tau[m-1]}$ and by the Lifting Lemma, we have $\mathsf{w}^{\sigma[m]} \le \mathsf{w}^{\tau[m-1]}$.

Then, the restriction of $\varphi^2(\sigma)$ to $\mathsf{w}_2 \cdots \mathsf{w}_k$ shows that k becomes a shifted descent in $\varphi(\gamma)$. Therefore, we obtain a counterexample on $\mathsf{w}_2 \cdots \mathsf{w}_k$, contradicting our minimal length choice of w.

Case $\sigma_1 = 1$ and $\varphi^2(\sigma)_1 = 0$. In this case, w_1 is a left descent for $\mathsf{w}^{\varphi^2(\sigma)}$.

$$
\begin{array}{ccccccccc}
 & 1 & \cdots & j & j+1 & \cdots & k-1 & k \\
\sigma = & 1 & \cdots & X^{(d)} & 1 & \cdots & 1 & 1 \\
\varphi^2(\sigma) = & 0 & \cdots & 1 & 1 & \cdots & 1 & X^d
\end{array}
$$

By the Exchange Condition, there exists some position i such that $\mathsf{w}_1 \mathsf{w}^{\varphi^2(\sigma)} = \mathsf{w}^\nu$ where ν is obtained from $\varphi^2(\sigma)$ by changing a single mask-value 1 entry to have mask-value 0. The mask ν may be not be a constant mask on $\mathsf{w}_2 \cdots \mathsf{w}_{k-1}$, but there exists a constant mask γ on $\mathsf{w}_2 \cdots \mathsf{w}_{k-1}$ for the element w^ν by Lemma 2.2. Let ρ denote σ restricted to $\mathsf{w}_2 \cdots \mathsf{w}_k$. Then, γ shows that k becomes a shifted descent in $\varphi(\rho)$. Therefore, we obtain a counterexample on $\mathsf{w}_2 \cdots \mathsf{w}_k$, contradicting our minimal length choice of w.

This in all cases, we have shown that applying a move at position j cannot create a new move at some position $k > j$. Hence, φ is an involution. $\square$

Theorem 2.12 *The function φ given in Definition 2.8 is a complete matching of the Hasse diagram of Bruhat order on $[y, w]$ whenever $y < w$.*

Proof Encode $[y, w]$ as a set of relative masks for y on a fixed reduced expression
w for w, so each $x \in [y, w]$ is given by $\mathsf{w}^{\Xi(\sigma)}$ where $\Xi(\sigma)$ is the X-mask of a
relative mask σ. By Lemmas 2.10 and 2.11, the function φ given in Definition 2.8
is an involution that interchanges exactly one X entry in each relative mask for an
entry that is either 0 or 1. Since the X-masks of both elements are constant masks,
$l(w) - l(x)$ is given by the number of X entries in the relative mask, so this operation
represents a cover relation in Bruhat order. Hence, we have that φ is a matching on
the Hasse diagram of $[y, w]$. Unmatched relative masks must contain no 0 entries nor
X entries at all, so consist of all 1 entries, and this occurs only if $w = y$. Hence, the
matching is complete when $y < w$, and the result follows.   □

Corollary 2.13 *The Möbius function of the Bruhat interval* $[y, w]$ *is* $\mu(x, w) =$
$(-1)^{l(w)-l(x)}$.

Proof Following Verma [16], it suffices to show that there exists a complete matching
of the Hasse diagram of the Bruhat interval $[y, w]$ whenever $y < w$. This matching
can then be interpreted as a sign-reversing involution on

$$\sum_{y \leq x \leq w} (-1)^{l(w)-l(x)}$$

proving that the sum is equal to the Kronecker function $\delta_{y,w}$. This follows from
Theorem 2.12.   □

In work related to totally nonnegative flag varieties, [14] have given another
complete matching on Bruhat intervals $[y, w]$ of a finite Coxeter group W. This
matching M is defined recursively, starting from an EL-labeling of the interval and
a chosen reduced expression w for w. To describe this matching, we begin with a
reduced expression w_0 for the longest element w_0 of W having $\mathsf{w}^{-1} = \mathsf{w}_p \mathsf{w}_{p-1} \cdots \mathsf{w}_1$
as a left factor. Then we obtain a total ordering on the reflections of W using the
inversion sequence constructed from the reduced expression w_0 by

$$\mathsf{w}_p > \mathsf{w}_p \mathsf{w}_{p-1} \mathsf{w}_p > \cdots > \mathsf{w}_p \mathsf{w}_{p-1} \cdots \mathsf{w}_{p-i+1} \mathsf{w}_{p-i} \mathsf{w}_{p-i+1} \cdots \mathsf{w}_{p-1} \mathsf{w}_p > \cdots . \quad (2.1)$$

We label all of the Bruhat cover relations $x' \lessdot x$ in $[y, w]$ by the unique right
reflection t such that $x' = xt$. Then, Dyer [9] has shown that this is an EL-labeling.
Rietsch and Williams construct a matching M from this EL-labeling using a result of
Chari [6].

Remark 2.14 For finite Coxeter groups, we show that the matching M is the same as
the matching given in Theorem 2.12, working by downward induction on the ranks
of the partial order $[y, w]$. We begin at the top rank r containing w.

Let x be an unmatched element on the current rank r. Consider the relative mask
σ associated to x. Since x is a maximal unmatched element, when we apply φ to
σ, we operate by placing an X in the rightmost position i such that the element
w^τ associated to the resulting X-mask $\tau = \Xi(\varphi(\sigma))$ still contains y in Bruhat order.
Moreover, observe that none of the entries to the right of i are shifted descents, nor

do they have mask values X or 0, for otherwise σ would already have been matched. Hence, w^τ is the element

$$x \cdot \left(\mathsf{w}_p \mathsf{w}_{p-1} \cdots \mathsf{w}_{i+1} \mathsf{w}_i \mathsf{w}_{i+1} \cdots \mathsf{w}_{p-1} \mathsf{w}_p \right),$$

and any element $x \cdot \left(\mathsf{w}_p \mathsf{w}_{p-1} \cdots \mathsf{w}_{j+1} \mathsf{w}_j \mathsf{w}_{j+1} \cdots \mathsf{w}_{p-1} \mathsf{w}_p \right)$ for $j > i$ does not contain y in Bruhat order, for otherwise j would be a shifted descent in σ.

In comparison, [14, Corollary 7.8] states that the matched edge $x' \lessdot x$ in M has the largest EL-label in the sense of Eq. 2.1 among all of the edges descending from x in $[y, w]$. But this is equivalent to the rightmost condition that we used to choose i. Hence, we see that w^τ is equal to the element x' that is matched to x in M. This proves that the matchings agree on all elements down to rank r, and we can proceed to apply the argument to the unmatched elements on rank $r - 1$. Continuing in this fashion, we find that the matchings agree on $[y, w]$.

The description given in Theorem 2.12 has the advantage of being nonrecursive and also permits some observations that are perhaps less clear in the other language. For example, we see that the matched edges of M are always labeled by one of the reflections that represent inversions in w, so the matching does not depend on how w is completed to a reduced expression for w_0.

3 Further Questions

Bruhat order extends to parabolic quotients of Coxeter groups as described in [2, Section 2.5]. Deodhar has given a parabolic version of the Möbius function formula in Deodhar [7, Theorem 1.2], and it would be interesting to extend the mask matching given above to recover his result.

Also, the order complex associated to a Bruhat interval $[x, w]$ is a topological space known to be homeomorphic to the $(l(w) - l(x) - 2)$-sphere. It would be interesting to recover the poset topology of the Bruhat intervals from the combinatorial matching we have given above.

As a preliminary step in this direction, we have observed that our matching is acyclic, in the sense used in discrete Morse theory. When W is a finite Coxeter group, this could also be inferred from [14] by Remark 2.14.

Definition 3.1 Consider the Hasse diagram of Bruhat order as a directed graph with an edge $w \to x$ if w covers x. Given a matching, reverse the direction of each edge in the Hasse diagram corresponding to a matched edge. We say the matching is *acyclic* if there are no directed cycles in the resulting directed graph.

Theorem 3.2 *The function φ given in Definition 2.8 is an acyclic matching of the Hasse diagram of Bruhat order on $[y, w]$ whenever $y < w$.*

Proof Let w be a reduced expression for w and consider the relative masks on w for y. Every directed cycle has at least two pairs of up-down edges. Observe that each edge pointing up corresponds to a matched edge, so is obtained by removing the rightmost X entry. Each edge pointing down corresponds to a non-matched edge, and this is a Bruhat cover on the elements encoded by the X-masks.

Recall that $\Xi(\sigma)$ denotes the X-mask of a relative mask σ on w. Suppose we have a pair of up-down edges in a directed cycle

$$x := \mathsf{w}^{\Xi(\sigma)} \to z := \mathsf{w}^{\Xi(\gamma)} \to x' := \mathsf{w}^{\Xi(\sigma')}.$$

Here, z covers x and x' with $x \neq x'$, and $\gamma = \varphi(\sigma)$. We claim that the rightmost X-entry in σ' occurs strictly left of the rightmost X-entry in σ, which implies that there are no directed cycles.

Let i denote the position in w where φ acts on σ. Since i is the rightmost move and φ does not alter the mask-values to the right of i, we must have $\gamma_j = 1$ for all $j > i$, and none of the positions γ_j are shifted descents for $j > i$.

Consider the rightmost X-entry in σ' and suppose for the sake of contradiction that it occurs in position $j \geq i$. Then, the relative masks γ and σ' agree on all positions strictly right of j according to the algorithm given in Lemma 2.2 since we always encode the same element y. At position j, we have $\Xi(\gamma)[j] = 1$, and $\Xi(\sigma')[j] = 0$. Hence,

$$\mathsf{w}^{\gamma[j]} = \mathsf{w}^{\sigma'[j]} \leq \mathsf{w}^{\Xi(\sigma')[j-1]} = \mathsf{w}^{\Xi(\gamma)[j-1]}. \tag{3.1}$$

To see the last equality, we use a Lifting Lemma argument. We have $\mathsf{w}^{\Xi(\sigma')[j-1]} = \mathsf{w}^{\Xi(\sigma')[j]} \leq \mathsf{w}^{\Xi(\gamma)[j]}$, and w_j is a right descent for $\mathsf{w}^{\Xi(\gamma)[j]}$, but w_j is a right ascent for $\mathsf{w}^{\Xi(\sigma')[j-1]}$. So, $\mathsf{w}^{\Xi(\sigma')[j-1]} \leq \mathsf{w}^{\Xi(\gamma)[j]}\mathsf{w}_j = \mathsf{w}^{\Xi(\gamma)[j-1]}$. However, $\mathsf{w}^{\Xi(\sigma')[j-1]}$ and $\mathsf{w}^{\Xi(\gamma)[j-1]}$ have the same length so they must be equal.

Equation 3.1 proves that j is a shifted descent in γ. It also shows that if $j = i$, then $x = x'$. In any case, we reach a contradiction. Hence, the rightmost X-entry in σ' occurs strictly left of the rightmost X-entry in σ, so the matching is acyclic. □

This is consistent with the main result of [3].

References

1. Armstrong, D.: The sorting order on a Coxeter group. J. Comb. Theory Ser. A. (2009). doi:10.1016/j.jcta.2009.03.009
2. Björner, A., Brenti, F.: Combinatorics of Coxeter Groups. Graduate Texts in Mathematics, vol. 231. Springer, New York (2005)
3. Björner, A.: Posets, regular CW complexes and Bruhat order. Eur. J. Comb. **5**(1), 7–16 (1984)
4. Björner, A., Wachs, M.: Bruhat order of Coxeter groups and shellability. Adv. Math. **43**(1), 87–100 (1982)
5. Billey, S., Warrington, G.S.: Kazhdan–Lusztig polynomials for 321-hexagon-avoiding permutations. J. Algebr. Comb. **13**(2), 111–136 (2001)
6. Chari, M.K.: On discrete Morse functions and combinatorial decompositions. Discrete Math. **217**(1–3), 101–113 (2000), Formal Power Series and Algebraic Combinatorics (Vienna, 1997)
7. Deodhar, V.V.: Some characterizations of Bruhat ordering on a Coxeter group and determination of the relative Möbius function. Invent. Math **39**(2), 187–198 (1977)

8. Deodhar, V.V.: A combinatorial setting for questions in Kazhdan–Lusztig theory. Geom. Dedic. **36**(1), 95–119 (1990)
9. Dyer, M.J. Hecke algebras and shellings of Bruhat intervals. Compos. Math. **89**(1), 91–115 (1993)
10. Humphreys, J.E.: Reflection Groups and Coxeter Groups. Cambridge Studies in Advanced Mathematics, vol. 29. Cambridge University Press, Cambridge (1990)
11. Kazhdan, D., Lusztig, G.: Representations of Coxeter groups and Hecke algebras. Invent. Math **53**(2), 165–184 (1979)
12. Marietti, M.: Algebraic and combinatorial properties of zircons. J. Algebr. Comb. **26**(3), 363–382 (2007)
13. Marsh, R.J., Rietsch, K.: Parametrizations of flag varieties. Representat. Theory **8**, 212–242 (2004) (electronic)
14. Rietsch, K., Williams, L.: Discrete Morse theory for totally non-negative flag varieties. arXiv:0810.4314 [math.CO] (2008)
15. Stembridge, J.R.: A short derivation of the Möbius function for the Bruhat order. J. Algebr. Comb. **25**(2), 141–148 (2007)
16. Verma, D.N.: Möbius inversion for the Bruhat ordering on a Weyl group. Ann. Sci. Éc. Norm. Super. (4) **4**, 393–398 (1971)