POSET TOPOLOGY, MOVES, AND BRUHAT INTERVAL POLYTOPE LATTICES

CHRISTIAN GAETZ AND PATRICIA HERSH

Abstract. We study the poset topology of lattices arising from orientations of 1-skeleta of *directionally simple* polytopes, with Bruhat interval polytopes $Q_{e,w}$ as our main example. We show that the order complex $\Delta((u,v)_w)$ of an interval therein is homotopy equivalent to a sphere if $Q_{u,v}$ is a face of $Q_{e,w}$ and is otherwise contractible. This significantly generalizes the known case of the permutahedron. We also show that saturated chains from u to v in such lattices are connected, and in fact highly connected, under moves corresponding to flipping across a 2-face. When w is a Grassmannian permutation, this implies a strengthening of the restriction of Postnikov's move-equivalence theorem to the class of BCFW bridge decomposable plabic graphs.

1. Introduction

Bruhat interval polytopes were introduced in [11] and further studied in [20], motivated by their connections to the Toda lattice and to the moment map on the totally positive part of the flag variety. Recently, these polytopes have attracted significant attention as the moment polytopes of generic torus orbit closures in Schubert varieties [12, 13], as distinguished flag positroid polytopes [7], and for their connection to Mirković–Vilonen polytopes [18].

It is natural to study the *1-skeleton poset* P_w of a Bruhat interval polytope $Q_w = Q_{e,w}$, obtained by orienting the 1-skeleton of Q_w with respect to a certain cost vector $\boldsymbol{\rho}$. This is a partial order on the lower Bruhat interval $[e, w]$ which is intermediate in strength between the weak order and the strong (Bruhat) order. The order complexes of intervals in both the weak and strong orders are known [3, 4, 6] to be homotopy equivalent to balls or spheres; indeed these are fundamental examples in poset topology. These results have also found many applications to the study of the weak and strong orders, and in particular they allow for the easy determination of their Möbius

2010 *Mathematics Subject Classification.* 05E45, 06A07.

Key words and phrases. Bruhat interval polytope, 1-skeleton poset, BCFW bridge decomposition, poset topology, directionally simple.

P.H. was supported by NSF grant DMS-1953931 and NSF-RTG grant DMS-2039316.

functions, which have long (see e.g. [21]) been of interest. In our first main result, we show that these results extend to intervals in the posets P_w.

We write $[u, v]_w$ for the closed interval between u and v in the poset P_w and $A_w(u, v)$ for the set of atoms of this interval; we write $(u, v)_w$ for the open interval.

Theorem 1.1. *The order complex $\Delta((u, v)_w)$ of any open interval in P_w is homotopy equivalent a sphere $\mathbb{S}^{|A_w(u,v)|-2}$ if $Q_{u,v}$ is a face of Q_w, and is otherwise contractible. Thus the Möbius function takes values $\mu_{P_w}(u, v) = (-1)^{|A_w(u,v)|}$ or $\mu_{P_w}(u, v) = 0$ accordingly.*

Theorem 1.1 answers a question[1] of Richard Stanley.

Our approach to Theorem 1.1 extends the techniques developed by the second named author in [10], but requires new arguments. It was proven there that a *simple* polytope whose 1-skeleton poset is the Hasse diagram of a lattice satisfies the conclusions of Theorem 1.1. The first named author proved in [8] that the posets P_w are indeed lattices. The polytopes Q_w are not in general simple, but they do have the weaker property of being *directionally simple* [8]. We give a new argument to show in Theorem 3.4 that these conditions suffice to imply that the lattice-theoretic join agrees with a convex-geometric *pseudojoin* operation. We use this together with a face non-revisiting property of Q_w to prove Theorem 1.1.

Our second main result concerns maximal chains in $[u, v]_w$ (that is, facets of $\Delta((u, v)_w)$) or more generally in any lattice which is the 1-skeleton poset of a directionally simple polytope. If a maximal chain γ_1 enters some 2-dimensional face F at its unique source (with respect to our fixed 1-skeleton orientation), continues along part of ∂F, and then exists at the unique sink, we can perform a *move* to produce a new maximal chain γ_2 which goes the other way around ∂F. Two maximal chains are *move-equivalent* if they can be connected by a sequence of such moves. We study the graph $\mathcal{M}(u, v)$ whose vertices are these maximal chains and whose edges correspond to such moves.

Theorem 1.2. *Let $\mathcal{O}$ be a facial orientation (see Definition 2.4) and suppose that a polytope Q is $\mathcal{O}$-directionally simple such that $\mathcal{O}$ is the Hasse diagram of a lattice L. Suppose that $[u, v]_L$ has a atoms. Then $\mathcal{M}(u, v)$ is connected, and if $a \geq 2$ is moreover $(a-1)$-connected.*

Theorem 1.2 is surprisingly powerful. For example, when $Q = Q_{w_0}$ is the *permutahedron* the lattice is the (right) weak order on the symmetric group. Maximal chains in weak order correspond to *reduced words* of w_0, and a move across a 2-face F corresponds to a *commutation move* or *braid move* according to whether F is a square or hexagon. Thus in this case Theorem 1.2 recovers Matsumoto's Theorem [14, 19] that all reduced words are connected under commutation or braid moves.

[1] Personal communication at FPSAC 2023.

Theorem 1.2 extends a similar result of Athanasiadis, Edelman, and Reiner who work in the setting of simple polytopes (and with the assumption that $u = \hat{0}, v = \hat{1}$) [2]. Our generalization to directionally simple polytopes is important because Q_w is almost never simple for w a Grassmannian permutation, and move-equivalence on these Bruhat interval polytopes is of particular interest: if w is Grassmannian, then, by a result of Williams [22], Q_w is isomorphic to a *bridge polytope* and maximal chains in P_w correspond to *BCFW bridge decompositions*. It is shown in [2] that the higher connectivity of Theorem 1.2 fails for general polytopes; it is an interesting question for future study to determine exactly how much further this result can be pushed.

BCFW bridge decompositions are physically motivated ways of building up reduced *plabic graphs* with trip permutation w [1]. Moves across a 2-face now correspond to applications of Postnikov's moves [15] to the resulting plabic graphs (see [22, Thm. 5.3]). A fundamental theorem of Postnikov [15] says that any two reduced plabic graphs with the same trip permutation are move-equivalent. We obtain the restriction of this result to the class of BCFW decomposable plabic graphs as an easy corollary of Theorem 1.2. Our result has the advantage of guaranteeing that two BCFW decomposable plabic graphs can be connected by moves, with all intermediate graphs also BCFW decomposable[2]. Furthermore, we are not aware of prior results on the higher connectivity of the plabic moves.

Corollary 1.3. *Let $w \in S_n$ be Grassmannian. Then any two BCFW bridge decomposable plabic graphs with trip permutation w can be connected by a sequence of moves, without leaving the class of BCFW decomposable plabic graphs. Moreover, the move graph is $(a-1)$-connected, where a is the number of distinct simple reflections appearing in a reduced word for w.*

2. Preliminaries

2.1. **Bruhat interval polytopes.** We refer the reader to [5] for background on Coxeter groups and Bruhat order. We view the symmetric group S_n as a Coxeter group, with adjacent transpositions as the simple reflections. We write $\preceq$ for the Bruhat order on S_n, and $\leq_R$ for the right weak order.

Definition 2.1 (Kodama–Williams [11])**.** For a permutation $w \in S_n$, denote by $\mathbf{w}$ the vector $(w^{-1}(1), \ldots, w^{-1}(n)) \in \mathbb{R}^n$. For $w' \preceq w \in S_n$, define the corresponding *Bruhat interval polytope* as the convex hull $Q_{w',w} := \mathrm{conv}\{\mathbf{u} \mid w' \preceq u \preceq w\}$. If $w' = e$ is the identity permutation, we write Q_w for $Q_{e,w}$. See Figure 1 for an example.

Definition 2.2. Let P_w denote the partially ordered set $([e, w], \leq_w)$ with cover relations $u \lessdot_w v$ whenever $\ell(v) = \ell(u) + 1$ (where ℓ denotes Coxeter length) and there is

[2]We thank Melissa Sherman-Bennett for pointing out that there exist reduced plabic graphs which are not BCFW decomposable.

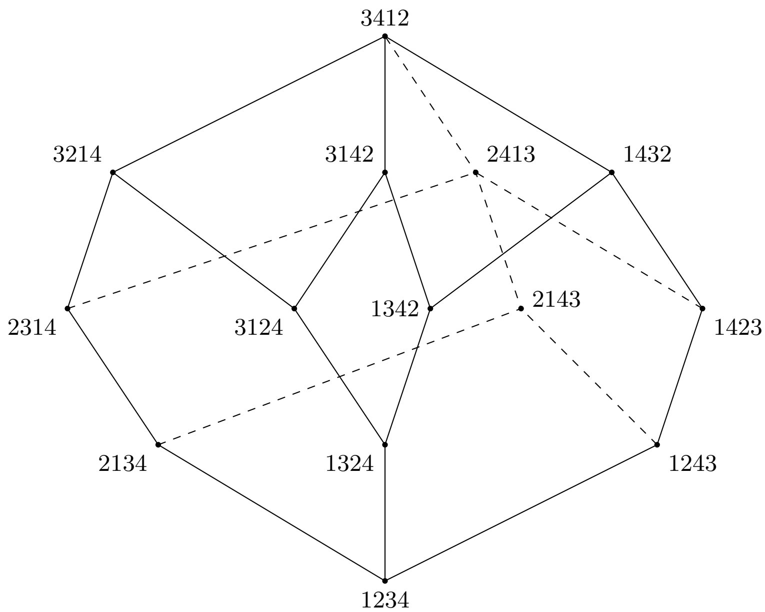

FIGURE 1. The Bruhat interval polytope Q_{3412}.

an edge of Q_w connecting $\mathbf{u}$ and $\mathbf{v}$. We often conflate elements u of P_w and vertices $\mathbf{u}$ of Q_w when no confusion can result.

A key property of Bruhat interval polytopes is the following.

Theorem 2.3 (Tsukerman-Williams)**.** *Every face of a Bruhat interval polytyope is a Bruhat interval polytope.*

Theorem 2.3 implies in particular that every cover relation in P_w is a cover relation in Bruhat order. This implies in turn that Bruhat interval polytopes are *generalized permutahedra* in the sense of Postnikov [16].

The poset P_w can alternatively be described as the 1-skeleton poset of the polytope Q_w with respect to a cost vector. Such posets are the subject of [10].

Definition 2.4. We say an acyclic orientation $\mathcal{O}$ of the 1-skeleton of a polytope Q is a *facial orientation* if the restriction of $\mathcal{O}$ to each face F of Q has a unique source and a unique sink; we write source(F) and sink(F) for these vertices. The *1-skeleton poset* $P(Q, \mathcal{O})$, a partial order on the vertices of Q, is defined as the transitive closure of $\mathcal{O}$. In this paper, we will be most interested in cases where $\mathcal{O}$ is the Hasse diagram of $P(Q, \mathcal{O})$ (that is, when $\mathcal{O}$ is transitively reduced); this is the *Hasse diagram property* of [10].

Generic cost vectors give an easy source of facial orientations. For $Q \subset \mathbb{R}^d$, a vector $\mathbf{c} \in \mathbb{R}^d$ is *generic* if no edge of Q is orthogonal to $\mathbf{c}$. Given such a vector, we obtain a facial orientation $\mathcal{O}_{\mathbf{c}}$ by orienting each edge according to increasing inner product with $\mathbf{c}$; in this case we may write $P(Q, \mathbf{c})$ for $P(Q, \mathcal{O}_{\mathbf{c}})$.

Example 2.5. Consider the permutahedron Q_{w_0} for S_n. It is observed in [10] that for cost vector $\boldsymbol{\rho} \coloneqq (n, n-1, \ldots, 1)$ the poset $P(Q_{w_0}, \boldsymbol{\rho})$ is the right weak order on S_n and that $\mathcal{O}_{\boldsymbol{\rho}}$ has the Hasse diagram property. It is further observed in [8] that for any Bruhat interval polytope Q_w with $w \in S_n$ and this same cost vector $\boldsymbol{\rho}$ that $\mathcal{O}_{\boldsymbol{\rho}}$ is the Hasse diagram of the poset $P_w \coloneqq P(Q_w, \boldsymbol{\rho})$.

Recall that a d-polytope Q is *simple* if each vertex is incident to exactly d edges, or, equivalently, if, for each vertex $v \in Q$, each subset of the edges incident to v spans a distinct face. Most of the results in [10] were proven for simple polytopes. In this paper, we work in the more general setting of *directionally simple* polytopes. We work with a more general notion of directional simplicity than that introduced in [8], which only allowed for orientations induced by a generic cost vector.

Definition 2.6. We say Q is *$\mathcal{O}$-directionally simple* with respect to a facial orientation $\mathcal{O}$ if for each vertex v of Q, each subset E of the edges outgoing from v (with respect to $\mathcal{O}$) spans a distinct face of Q. We write F_E for this face.

Directional simplicity is a weaker condition than simplicity since the existence of F_E is required only for subsets of edges outgoing from v, not for all subsets of edges incident to v.

Theorem 2.7 (Theorem 5.3, [8])**.** *For any $w \in S_n$, the Bruhat interval polytope Q_w is directionally simple with respect to the cost vector $\boldsymbol{\rho}$.*

It will also be very important for us that the posets P_w are lattices.

Theorem 2.8 (Theorem 4.5, [8])**.** *For any $w \in S_n$, the poset P_w is a lattice.*

Remark 2.9. The more general polytopes $Q_{w',w}$ do not always have the lattice or directional simplicity properties (see [8, §1.3.3]), so our results do not necessarily apply to these.

2.2. **Poset topology.** We briefly review some background on poset topology.

Definition 2.10. The *Möbius function* of a finite bounded partially ordered set $(P, \leq)$ is the function $\mu_P : P \times P \to \mathbb{Z}$ satisfying

$$\mu_P(u, v) = -\sum_{u \leq z < v} \mu_P(u, z),$$

for $u \neq v$, and $\mu_P(u, u) = 1$, for all $u \in P$.

Definition 2.11. The *order complex*, denoted $\Delta(P)$, of a poset P is the abstract simplicial complex whose i-faces are the chains $v_0 < \cdots < v_i$ of P.

We will use the following well known fact due to Philip Hall.

Proposition 2.12. *Let P be a finite bounded poset and let $u \leq_P v$. Then*

$$\mu_P(u,v) = \tilde{\chi}(\Delta((u,v)_P)),$$

where $\tilde{\chi}$ denotes reduced Euler characteristic.

Recall from [17] the Quillen fiber lemma:

Lemma 2.13 (Quillen fiber lemma)**.** *Suppose that $f : P \to P'$ is an ordering-preserving map such that for each $p' \in P'$ the subposet $P^{\leq p'} := \{p \in P | f(p) \leq p'\}$ has contractible order complex. Then $\Delta(P)$ and $\Delta(P')$ are homotopy equivalent.*

Since a poset and its dual have the same order complex, an immediate consequence is what we will call the "dual Quillen fiber lemma":

Corollary 2.14 (dual Quillen fiber lemma)**.** *Suppose that $f : P \to Q$ is an order-preserving map such that for each $p' \in P'$ the subposet $P^{\geq p'} := \{p \in P | f(p) \geq p'\}$ has contractible order complex. Then $\Delta(P)$ and $\Delta(P')$ are homotopy equivalent.*

2.3. **Topology of simple 1-skeleton lattices.** The following result of [10] helped motivate the present work.

Theorem 2.15 (Theorem 5.13, [10])**.** *Let Q be a simple polytope with facial orientation $\mathcal{O}$ such that $L := P(Q, \mathcal{O})$ is a lattice with Hasse diagram $\mathcal{O}$. Then any open interval $(u, v)_L$ has order complex homotopy equivalent to a sphere $\mathbb{S}^{a-2}$, where a is the number of atoms in $[u, v]_L$, if v equals the join of the atoms of $[u, v]$, and is otherwise contractible.*

It is natural to try to generalize this result to directionally simple polytopes. We will also need the next notion which was introduced in [10] for simple polytopes but makes equally good sense for directionally simple polytopes.

Definition 2.16. Suppose that a polytope Q is $\mathcal{O}$-directionally simple where $\mathcal{O}$ is the Hasse diagram of $P(Q, \mathcal{O})$. Then the *pseudo-join* $\mathrm{psj}(S)$ of a collection S of upper covers of an element u in $P(Q, \mathcal{O})$ is the unique sink of the unique $|S|$-face F_S containing u and the elements of S. The face F_S exists and is unique by directional simplicity and has a unique sink since $\mathcal{O}$ is a facial orientation.

Theorem 2.17 (Theorem 4.7, [10])**.** *Let Q be a simple polytope with facial orientation $\mathcal{O}$ such that $\mathcal{O}$ is the Hasse diagram of a lattice L. Then for any collection S of atoms of an interval $[u, v]_L$, we have $\bigvee S = \mathrm{psj}(S)$.*

3. Poset topology of Bruhat interval polytopes

In this section we prove Theorem 1.1, our first main theorem.

Recall that we write $\preceq$ for Bruhat order and $\leq_w$ for the order on P_w. We write $(u,v)_B$ and $[u,v]_B$ for open and closed intervals in Bruhat order and $(u,v)_w$ and $[u,v]_w$ for open and closed intervals in P_w. We also fix the notation $\mathcal{O}_w$ for the orientation of the 1-skeleton of Q_w induced by the cost vector $\boldsymbol{\rho}$; this is the Hasse diagram of P_w. Finally, we write $\wedge_w$ and $\vee_w$ for the meet and join operators in the lattice P_w.

Combining Theorems 2.8 and 2.15, it is easy to conclude that Theorem 1.1 holds whenever the Bruhat interval polytope Q_w is simple. Our goal is to extend this result to all Q_w.

We first establish a face non-revisiting property for the Q_w.

Proposition 3.1. *Let $v_0 \to v_1 \to \cdots \to v_k$ be a directed path in $\mathcal{O}_w$ and suppose that v_0 and v_k both lie in some face F of Q_w. Then v_i lies in F for all $i = 0, \ldots, k$.*

Proof. By Theorem 2.3, the vertices of the face F are the elements $[u,v]_B$ of a Bruhat interval. Note that $[u,v]_w \subset [u,v]_B$ since every cover relation in P_w is a cover relation in Bruhat order. Thus we have $v_i \in [v_0, v_k]_w \subset [u,v]_w \subset [u,v]_B$ for all i, so all v_i lie in F. □

Proposition 3.2. *Let $u \leq_w v$ and let $S, S' \subset A_w(u,v)$ be distinct subsets. Then $\bigvee S \neq \bigvee S'$.*

Proof. It suffices to prove the claim when $v = w$, so suppose we are in this case. Suppose for the sake of contradiction that $\bigvee S = \bigvee S' = x$ for S and S' distinct subsets of $A_w(u,v)$. For any $T \subset A_w(u,v)$, let F_T denote the unique smallest face of Q_w containing u and all elements of T. Notice that, by Proposition 3.1, $x \in F_S \cap F_{S'}$, since $\mathrm{sink}(F_S)$ and $\mathrm{sink}(F_{S'})$ are upper bounds in P_w for S and S', respectively. By Theorem 2.7, T is exactly the set of upper covers of u contained in F_T. Thus we have that $F_S \cap F_{S'} = F_{S \cap S'}$. Since $S \neq S'$, we may assume without loss of generality that $S \setminus S' \neq \emptyset$. Consider any $s \in S \setminus S'$; by the prior observation we know that $s \notin F_{S \cap S'}$. Since $u \lessdot_w s$, there is a directed edge $u \to s$ in $\mathcal{O}_w$. There is also a directed path $s \to \cdots \to x$ since $s \leq_w x = \bigvee S$. Concatenating these paths gives a directed path from u to x that includes s. But this contradicts Proposition 3.1 since $u, x \in F_{S \cap S'}$ but $s \notin F_{S \cap S'}$. □

Proposition 3.3. *For any $u \leq_w v$ there is a surjective order-preserving map $f : [u,v]_w \to \mathcal{P}(A_w(u,v))$, where $\mathcal{P}$ denotes the power set (partially ordered by inclusion).*

Proof. For each $v' \in [u,v]_w$, define $f(v') = \{a \in A_w(u,v) \mid a \leq_w v'\}$. The map f is clearly order-preserving and is surjective since $f(\bigvee S) = S$ by Proposition 3.2. □

We now prove a general theorem about directionally simple polytopes whose 1-skeleton poset is the Hasse diagram of a lattice, namely that joins equal pseudojoins. This generalizes [10, Thm. 4.7] which proved the case of simple polytopes.

Theorem 3.4. *Let Q be $\mathcal{O}$-directionally simple such that $\mathcal{O}$ is the Hasse diagram of a lattice $L = P(Q, \mathcal{O})$. Let $u \in L$ and suppose $u \lessdot a_1, \ldots, a_r$. Then $a_1 \vee \cdots \vee a_r = \mathrm{psj}(a_1, \ldots, a_r)$.*

Proof. First observe that $a_1 \vee \cdots \vee a_r \leq \mathrm{psj}(a_1, \ldots, a_r)$ by virtue of $\mathrm{psj}(a_1, \ldots, a_r)$ being an upper bound for $a_1, \ldots, a_r$. Therefore it suffices to show that $\mathrm{psj}(a_1, \ldots, a_r) \leq a_1 \vee \cdots \vee a_r$.

We proceed by induction on r, using the base case $r = 2$. This case was proven as [10, Thm. 4.6] (see [10, Thm. 5.13] for discussion of the case for general facial orientations $\mathcal{O}$) in the setting of simple polytopes, but we note that the proof applies equally well to directionally simple polytopes without modification. So assume now that $r \geq 3$; in this case we need a different proof than in [10].

Let U denote the set of vertices of the face $F_{a_1,\ldots,a_r}$. We will show that all elements of U, including $\mathrm{psj}(a_1, \ldots, a_r)$, lie below $a_1 \vee \cdots \vee a_r$. Let $u = u_0, u_1, \ldots, u_k$ be an ordering of U such that if $u_i < u_j$ then $i < j$. Each element of U lies in some facet of $F_{a_1,\ldots,a_r}$. We will prove by induction on the index i of the vertex $u_i = \mathrm{source}(G)$ that all vertices of each facet G lie weakly below $a_1 \vee \ldots \vee a_r$. If $i = 0$, so $\mathrm{source}(G) = u_0 = u$, then $G = F_{a_1,\ldots,\widehat{a_j},\ldots,a_r}$ for some j. By the inductive hypothesis on r, we have

$$\mathrm{sink}(G) = \mathrm{psj}(a_1, \ldots, \widehat{a_j}, \ldots, a_r) = a_1 \vee \cdots \vee \widehat{a_j} \vee \cdots \vee a_r \leq a_1 \vee \cdots \vee a_r.$$

Thus all vertices of G lie weakly below $a_1 \vee \cdots \vee a_r$.

Now suppose $\mathrm{source}(G) = u_i$ with $i > 0$. Let $b_1, \ldots, b_{r-1}$ be the upper covers of u_i, so $G = F_{b_1,\ldots,b_{r-1}}$. Since $\dim(F_{a_1,\ldots,a_r}) = r \geq 3$, the edge $\overline{u_i b_1}$ also lies in some other facet $G' \neq G$ of $F_{a_1,\ldots,a_r}$. Let $u_j = \mathrm{source}(G')$. Note that $u_j \neq u_i$, since otherwise u_i would have at least r upper covers, an impossibility since $u_i \neq u$ and since Q is directionally simple. Thus we have $u_j < u_i$ and, by our choice of ordering, $j < i$. We conclude by induction on i that all vertices of G', and in particular b_1, lie weakly below $a_1 \vee \cdots \vee a_r$. The same is true of $b_2, \ldots, b_{r-1}$, and hence $\mathrm{psj}(b_1, \ldots, b_{r-1}) = b_1 \vee \cdots \vee b_{r-1} \leq a_1 \vee \cdots \vee a_r$ by the inductive assumption on r. We conclude that all vertices of G lie weakly below $a_1 \vee \cdots \vee a_r$, and this completes the proof. □

Remark 3.5. One might be tempted to try to prove *the monotone Hirsch conjecture* for directionally simple polytopes whose 1-skeleton is a lattice, and thereby for all Bruhat interval polytopes, in light of Preuß's result in the appendix of [10] for the case of simple polytopes. However Preuß's proof does not seem to readily generalize to directionally simple polytopes.

Applying Theorems 2.7, 2.8 and 3.4, we have the following corollary.

Corollary 3.6. *For any* $u \lessdot_w a_1, \ldots, a_r$ *we have* $\mathrm{psj}(a_1, \ldots, a_r) = a_1 \vee_w \cdots \vee_w a_r$.

With a little further work, we also derive the following corollary.

Corollary 3.7. *The Bruhat interval polytope* $Q_{u,v}$ *is a face of* Q_w *if and only if* $v = \bigvee A_w(u, v)$.

Proof. If $v = \bigvee A_w(u, v)$, then by Corollary 3.6 we have $v = \mathrm{psj}(A_w(u, v))$. Thus the vertex set of $F_{A_w(u,v)}$ is exactly $[u, v]_w$ by Proposition 3.1. By Theorem 2.3, we must have $F_{A_w(u,v)} = Q_{u,v}$, so $Q_{u,v}$ is a face.

Otherwise we have $v >_w \bigvee A_w(u, v) = \mathrm{psj}(A_w(u, v))$, so v does not lie in $F_{A_w(u,v)}$. Then $Q_{u,v}$ cannot be a face of Q_w, since $F_{A_w(u,v)} \subsetneq Q_{u,v}$ would both be $|A_w(u, v)|$-dimensional faces. □

We are now ready to prove Theorem 1.1

Proof of Theorem 1.1. The order-preserving f from Proposition 3.3 restricts to a surjection from $(u, v)_w$ either to $\mathcal{S}_1 := \mathcal{P}(A_w(u, v))\backslash\{\emptyset, A_w(u, v)\}$ or to $\mathcal{S}_2 := \mathcal{P}(A_w(u, v))\backslash \{\emptyset\}$ according to whether $A_w(u, v) \in \mathrm{im}(f|_{(u,v)_w})$, that is, according to whether $\bigvee A_w(u, v) <_w v$. Notice that, since P_w is a lattice by Theorem 2.8, the dual fiber $(u, v)_w^{\geq S}$ has a unique minimal element, namely $\bigvee S$, for each $\emptyset \neq S \subseteq A_w(u, v)$. This is a cone point in $\Delta((u, v)_w^{\geq S})$, so this complex is contractible. Thus we may apply apply the dual Quillen Fiber Lemma (Corollary 2.14) to deduce that $\Delta((u, v)_w)$ is homotopy equivalent to $\Delta(\mathcal{S}_1)$ or to $\Delta(\mathcal{S}_2)$. The complex $\Delta(\mathcal{S}_2)$ is contractible due to having the cone point $A_w(u, v)$. The poset $\mathcal{S}_1$ is the face poset of the boundary of a simplex, so its order complex is the barycentric subdivision of this sphere. Thus $\Delta(\mathcal{S}_1)$ is a sphere $\mathbb{S}^{|A_w(u,v)|-2}$. We apply Corollary 3.7 to deduce that we are in the case of a sphere exactly when $Q_{u,v}$ is a face of Q_w. Finally, the claim about μ_{P_w} follows easily by applying Proposition 2.12. □

Remark 3.8. Our proof of Theorem 1.1 could be adapted to apply to any directionally simple polytope Q with facial orientation $\mathcal{O}$ such that $\mathcal{O}$ is the Hasse diagram of a lattice and such that Q satisfies a face non-revisiting property as in Proposition 3.1. It would be interesting to find other interesting examples of such polytopes.

4. Higher connectivity of move graphs

In this section we study move graphs of chains in 1-skeleton posets and prove Theorem 1.2.

Definition 4.1. Let Q be $\mathcal{O}$-directionally simple. Let $\mathcal{M}(Q, \mathcal{O})$ denote the graph whose vertices are the directed paths from source to sink in $\mathcal{O}$ and whose edges connect pairs of such paths differing by a single move across a 2-face of Q.

When $\mathcal{O}$ is the Hasse diagram of a poset P, for each $u \leq_P v$ we define $\mathcal{M}(u, v, \mathcal{O})$ as the graph whose vertices are the saturated chains from u to v and whose edges connect pairs of saturated chains differing by a single move across a 2-face of Q. In this case the vertices of $\mathcal{M}(Q, \mathcal{O})$ are the maximal chains of P.

The arguments below are inspired by the proof of [2, Thm. 2.1] in the setting of simple polytopes, but some of our adaptations to the directionally simple setting are a bit delicate. This more general setting is important for application to BCFW bridge decompositions, however, as Bruhat interval polytopes Q_w of Grassmannian permutations w are almost never simple ([9, Thm. 1.6] implies a pattern avoidance characterization for such w). We have retained notation from [2] when possible.

Lemma 4.2. *Let Q be an $\mathcal{O}$-directionally simple polytope such that $\mathcal{O}$ is the Hasse diagram of a lattice L. Then, for any $u \leq_L v$, and any chains $\gamma_1, \gamma_2 \in \mathcal{M}(u, v, \mathcal{O})$, there exists a path from γ_1 to γ_2 such that all intermediate chains γ contain all vertices in $\gamma_1 \cap \gamma_2$.*

Proof. We construct such a path, denoted $\gamma_1 * \gamma_2$, as follows.

Suppose that $\gamma_1 \neq \gamma_2$ and let x be the lowest element of $\gamma_1 \cap \gamma_2$ such that x is covered by distinct elements in γ_1 and γ_2. Let x' be the next lowest element of $\gamma_1 \cap \gamma_2$. Let e_1 (resp. e_2) be the edge upward from x in γ_1 (resp. γ_2). By directional simplicity, there exists a 2-face F of Q containing e_1 and e_2. By [10, Thms. 4.6 & 5.13], $\text{sink}(F) = a_1 \vee a_2$, where a_1, a_2 are the vertices connected to x by e_1, e_2. Since $x' \geq a_1, a_2$, the lattice property implies that there exists a saturated chain p from $\text{sink}(F)$ to $\text{sink}(Q)$ passing through x'. If $\text{sink}(F) \in \gamma_1$, then choose p to coincide with γ_1 on $[\text{sink}(F), \text{sink}(Q)]$. In any event, choose p to coincide with γ_1 on $[x', \text{sink}(Q)]$.

Let γ_1^F be the chain that coincides with γ_1 on $[\text{source}(Q), a_1]$, then follows the unique path in the boundary of F proceeding upward from a_1 to $\text{sink}(F)$, then coincides with p on $[\text{sink}(F), \text{sink}(Q)]$. Likewise, let γ_2^F be the maximal chain that coincides with γ_2 on $[\text{source}(Q), a_2]$, then follows the unique path in the boundary of F proceeding upward from a_2 to $\text{sink}(F)$, then coincides with p on $[\text{sink}(F), \text{sink}(Q)]$.

We let $\gamma_1 * \gamma_2$ be the path in $\mathcal{M}(u, v, \mathcal{O})$ determined by the following sequence of moves: first apply the inductively-defined series of moves in $\gamma_1 * \gamma_1^F$; then apply the move across F from γ_1^F to γ_2^F; finally, apply the inductively-defined series of moves in $\gamma_2^F * \gamma_2$. Note that this construction depends on the choices of the path p in each iteration. It is clear by the choices of x, x' that all intermediate chains γ in $\gamma_1 * \gamma_2$ contain all vertices of $\gamma_1 \cap \gamma_2$. □

Lemma 4.2 implies in particular that $\mathcal{M}(u, v, \mathcal{O})$ is connected. Our goal in the remainder of the section is to study the higher vertex connectivity of $\mathcal{M}(u, v, \mathcal{O})$. We will make use of special properties of the paths $\gamma_1 * \gamma_2$ constructed in the proof of Lemma 4.2.

Lemma 4.3. *Let Q be an $\mathcal{O}$-directionally simple polytope such that $\mathcal{O}$ is the Hasse diagram of a lattice L. Let γ_1, γ_2 be maximal chains in L. Suppose there exist $x <_L x'$*

with $x, x' \in \gamma_1 \cap \gamma_2$, and with x covered by distinct elements in γ_1 and γ_2. Then the path $\gamma_1 * \gamma_2$ holds fixed the segment $\gamma_1|_{[x',\mathrm{sink}(Q)]}$ until after it has transformed $\gamma_1|_{[\mathrm{source}(Q),x']}$ into $\gamma_2|_{[\mathrm{source}(Q),x']}$.

Proof. This follows from our convention, in the construction of $\gamma_1 * \gamma_2$, of choosing p to coincide with γ_1 on the interval $[\mathrm{sink}(F), \mathrm{sink}(Q)]$ in the event that $\mathrm{sink}(F)$ is an element of γ_1. □

Lemma 4.4. *Let Q be an $\mathcal{O}$-directionally simple polytope such that $\mathcal{O}$ is the Hasse diagram of a lattice L. Suppose γ and γ_1 are saturated chains in an interval $[u, v]_L$ that include distinct cover relations upward from u. Let γ_2 be any node on the path $\gamma * \gamma_1$. Denote by e_1 the lowest edge in γ_1, and denote by e_2 the lowest edge in γ_2 that is not in γ.*

Then there exists an alternating sequence $(\epsilon_0, F_1, \epsilon_1, F_2, \ldots, F_r, \epsilon_r)$ such that $\epsilon_0 = e_1$, $\epsilon_r = e_2$, each ϵ_i is an edge of Q whose lower vertex x_i is in γ, and each F_i is a 2-face of Q containing $\epsilon_{i-1}, \epsilon_i$, and all edges and vertices of γ between x_{i-1} and x_i.

Proof. By directional simplicity we may choose F_1 to be the unique 2-face containing both e_1 and the lowest edge in γ. If e_2 is an edge in F_1, let $\epsilon_1 = e_2$, in which case $r = 1$ and we are done. Otherwise, define ϵ_1 to be the lowest edge not contained in γ that is in the path, in the boundary of F_1, from $\mathrm{source}(F_1)$ to $\mathrm{sink}(F_1)$ that does not contain e_1. Such an edge ϵ_1 exists, since otherwise $\mathrm{sink}(F_1)$ would be in γ with a path upward from it to the edge e_2, but by Lemma 4.3 this would imply $\gamma_2 \notin \gamma * \gamma_1$. For $\epsilon_1 \neq e_2$ let F_2 be the unique 2-face containing both ϵ_1 and the edge in γ that proceeds upward from a vertex in ϵ_1. Repeat in this manner until reaching a 2-face that contains e_2; this process terminates by virtue of the construction in Lemma 4.2. □

Lemma 4.5. *The sequence $(\epsilon_0, F_1, \epsilon_1, F_2, \ldots, F_r, \epsilon_r)$ from Lemma 4.4 is uniquely determined by γ and e_2. In particular, γ and e_2 determine e_1.*

Proof. By the requirements for F_r in Lemma 4.4, the face F_r includes both the edge ϵ_r and an edge in γ leading upward to ϵ_r. We showed such an F_r exists in the proof of Lemma 4.4. There can only be one 2-face in a polytope containing a specified pair of edges, so this 2-face is uniquely determined. By virtue of the edge ϵ_{r-1} having its lower vertex in γ and the entire path from x_{r-1} to x_r also being in γ, the edge ϵ_{r-1} must be the unique lowest edge not contained in γ that is in the boundary of F_r. We likewise deduce that F_{r-1} is uniquely determined as the only 2-face that includes both ϵ_{r-1} and the edge in γ leading up to ϵ_{r-1}. Repeating in this manner shows that F_1 is likewise uniquely determined, and F_1 uniquely determines the edge ϵ_0 in the boundary of F_1 that proceeds upward from the source of F_1 and is not in γ. □

Lemma 4.6. *Let Q be an $\mathcal{O}$-directionally simple polytope such that $\mathcal{O}$ is the Hasse diagram of a lattice L. If γ, γ_1, and γ_1' are saturated chains in an interval $[u, v]_L$ using distinct edges upward from u, then $(\gamma * \gamma_1) \cap (\gamma * \gamma_1') = \{\gamma\}$.*

Proof. Consider any node γ_2 on the path $\gamma * \gamma_1$ and any node γ_2' on the path $\gamma * \gamma_1'$. Let e_2 be the lowest edge in γ_2 not shared with γ, and let e_2' be the lowest edge in γ_2' not shared with γ. To prove $\gamma_2 \neq \gamma_2'$, it suffices to prove $e_2 \neq e_2'$. If $e_2 = e_2'$, then Lemma 4.5 shows that γ and e_2 would give rise to the same alternating sequence of edges and 2-faces that γ and e_2' would produce. In particular, they would give rise to the same edges upward from u, contradicting γ_1 and γ_1' having distinct edges upward from u. □

We are now ready to prove Theorem 1.2, our second main theorem.

Proof of Theorem 1.2. The case with $a = 1$ is handled by Lemma 4.2. The proof of the case with $a \geq 2$ is via the same argument as in [2], with the dimension d of a simple polytope replaced by the number a of atoms of $[u, v]_L$ throughout, using Lemmas 4.3 to 4.6 to justify various claims under our modified hypotheses. We recall this argument now, as it applies in our setting.

Let Γ be any subset of size at most $a - 2$ of the set of vertices of $\mathcal{M}(u, v, \mathcal{O})$. Let E be the set of atoms of $[u, v]_L$ that are contained in chains from Γ. Let M be the induced subgraph of $\mathcal{M}(u, v, \mathcal{O}) \setminus \Gamma$ whose vertices are those chains γ that do not use any cover relation from E. Notice that $|M| \geq 2$, since at least 2 atoms of $[u, v]_L$ lie outside E and there is at least one saturated chain containing each such atom.

Given any $\gamma_1, \gamma_2 \in M$, observe that no vertex on the path $\gamma_1 * \gamma_2$ uses any atom from E, since all $\gamma \in \gamma_1 * \gamma_2$ use an atom used by either γ_1 or γ_2. Thus M is connected.

It remains to prove that for any chain $\gamma \notin \Gamma$ whose lowest cover relation e uses an atom from E there is a path from γ to an element of M that stays within $\mathcal{M}(u, v, \mathcal{O}) \setminus \Gamma$. Let k be the size of the subset of Γ consisting of chains using the edge e. The number of cover relations upward from u in $[u, v]_L$ that are not in E equals $a - |E| \geq a - |\Gamma| + (k - 1) \geq k + 1$. This implies the existence of $k + 1$ saturated chains in M all involving distinct cover relations upward from u. By Lemma 4.6, the paths in $\mathcal{M}(u, v, \mathcal{O})$ from γ to these $k + 1$ saturated chains are vertex disjoint except for sharing the initial node γ. Thus at least one of these paths avoids all k elements of Γ having initial edge e. By the construction from Lemma 4.2, this path stays within $\mathcal{M}(u, v, \mathcal{O}) \setminus \Gamma$. □

Acknowledgements

We are grateful to Richard Stanley for suggesting the investigation of the topology of Bruhat interval polytope posets. We also thank Melissa Sherman-Bennett and Lauren Williams for helpful conversations about BCFW bridge decompositions.

References

[1] Nima Arkani-Hamed, Jacob Bourjaily, Freddy Cachazo, Alexander Goncharov, Alexander Postnikov, and Jaroslav Trnka. *Grassmannian geometry of scattering amplitudes*. Cambridge University Press, Cambridge, 2016.

[2] Christos A. Athanasiadis, Paul H. Edelman, and Victor Reiner. Monotone paths on polytopes. *Math. Z.*, 235(2):315–334, 2000.

[3] A. Björner. Posets, regular CW complexes and Bruhat order. *European J. Combin.*, 5(1):7–16, 1984.

[4] Anders Björner. Orderings of Coxeter groups. In *Combinatorics and algebra (Boulder, Colo., 1983)*, volume 34 of *Contemp. Math.*, pages 175–195. Amer. Math. Soc., Providence, RI, 1984.

[5] Anders Björner and Francesco Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.

[6] Anders Björner and Michelle Wachs. Bruhat order of Coxeter groups and shellability. *Adv. in Math.*, 43(1):87–100, 1982.

[7] Jonathan Boretsky, Christopher Eur, and Lauren Williams. Polyhedral and tropical geometry of flag positroids. *Algebra Number Theory*, 18(7):1333–1374, 2024.

[8] Christian Gaetz. One-skeleton posets of Bruhat interval polytopes. *Adv. Math.*, 430:Paper No. 109216, 22, 2023.

[9] Christian Gaetz and Yibo Gao. On the minimal power of q in a Kazhdan–Lusztig polynomial. *Adv. Math.*, 457:Paper No. 109941, 2024.

[10] Patricia Hersh. Posets arising as 1-skeleta of simple polytopes, the nonrevisiting path conjecture, and poset topology. *Discete Comput. Geom.*, Published online in 2023.

[11] Yuji Kodama and Lauren Williams. The full Kostant-Toda hierarchy on the positive flag variety. *Comm. Math. Phys.*, 335(1):247–283, 2015.

[12] Eunjeong Lee and Mikiya Masuda. Generic torus orbit closures in Schubert varieties. *J. Combin. Theory Ser. A*, 170:105143, 44, 2020.

[13] Eunjeong Lee, Mikiya Masuda, Seonjeong Park, and Jongbaek Song. Poincaré polynomials of generic torus orbit closures in Schubert varieties. In *Topology, geometry, and dynamics—V. A. Rokhlin-Memorial*, volume 772 of *Contemp. Math.*, pages 189–208. Amer. Math. Soc., [Providence], RI, [2021] ©2021.

[14] Hideya Matsumoto. Générateurs et relations des groupes de Weyl généralisés. *C. R. Acad. Sci. Paris*, 258:3419–3422, 1964.

[15] Alexander Postnikov. Total positivity, Grassmannians, and networks, 2006. `arxiv:math/0609764`.

[16] Alexander Postnikov. Permutohedra, Associahedra, and Beyond. *International Mathematics Research Notices*, 2009(6):1026–1106, 01 2009.

[17] Daniel Quillen. Homotopy properties of the poset of nontrivial p-subgroups of a group. *Adv. Math.*, 28(2):101–128, 1978.

[18] Mario Sanchez. Mirković-Vilonen polytopes from combinatorics, 2023. arxiv:2311.16979.

[19] Jacques Tits. Le problème des mots dans les groupes de Coxeter. In *Symposia Mathematica (INDAM, Rome, 1967/68), Vol. 1*, pages 175–185. Academic Press, London-New York, 1969.

[20] E. Tsukerman and L. Williams. Bruhat interval polytopes. *Adv. Math.*, 285:766–810, 2015.

[21] Daya-Nand Verma. Möbius inversion for the Bruhat ordering on a Weyl group. *Ann. Sci. École Norm. Sup. (4)*, 4:393–398, 1971.

[22] Lauren K. Williams. A positive Grassmannian analogue of the permutohedron. *Proc. Amer. Math. Soc.*, 144(6):2419–2436, 2016.

(Gaetz) Department of Mathematics, University of California, Berkeley, CA 94720
E-mail address: gaetz@berkeley.edu

(Hersh) Department of Mathematics, University of Oregon, Eugene, OR 97403
E-mail address: plhersh@uoregon.edu