J Algebr Comb (2010) 31: 169–216
DOI 10.1007/s10801-009-0181-4

Smooth and palindromic Schubert varieties in affine Grassmannians

Sara C. Billey · Stephen A. Mitchell

Received: 28 May 2008 / Accepted: 5 May 2009 / Published online: 23 May 2009

Abstract We completely determine the smooth and palindromic Schubert varieties in affine Grassmannians, in all Lie types. We show that an affine Schubert variety is smooth if and only if it is a closed parabolic orbit. In particular, there are only finitely many smooth affine Schubert varieties in a given Lie type. An affine Schubert variety is palindromic if and only if it is a closed parabolic orbit, a chain, one of an infinite family of "spiral" varieties in type A, or a certain 9-dimensional singular variety in type B_3. In particular, except in type A there are only finitely many palindromic affine Schubert varieties in a fixed Lie type. Moreover, in types D and E an affine Schubert variety is smooth if and only if it is palindromic; in all other types there are singular palindromics.

The proofs are for the most part combinatorial. The main tool is a variant of Mozes' numbers game, which we use to analyze the Bruhat order on the coroot lattice. In the proof of the smoothness theorem we also use Chevalley's cup product formula.

Keywords Affine Grassmannians · Schubert varieties

1 Introduction

Let G be a simply-connected simple compact Lie group, with complexification $G_{\mathbb{C}}$. The affine Grassmannian $\mathcal{L}_G$ is a projective ind-variety, homotopy-equivalent to the

The first author was supported by the Royalty Research Fund and NSF grant DMS-0800978.
The second author was supported by the National Science Foundation grant DMS-0504795.

S.C. Billey · S.A. Mitchell (✉)
Department of Mathematics, University of Washington, Seattle, WA 98195, USA
e-mail: mitchell@math.washington.edu

S.C. Billey
e-mail: billey@math.washington.edu

loop space ΩG and closely analogous to a maximal flag variety of $G_{\mathbb{C}}$. It has a Schubert cell decomposition

$$\mathcal{L}_G = \coprod_{\lambda \in \mathcal{Q}^\vee} e_\lambda,$$

where $\mathcal{Q}^\vee$ is the coroot lattice and e_λ is the corresponding Schubert cell. The closure X_λ of e_λ is a finite dimensional projective variety that we call an *affine Schubert variety*. In this paper we completely determine the smooth and palindromic affine Schubert varieties.

Theorem 1.1 *Let X_λ be an affine Schubert variety. Then the following are equivalent*:

a) *X_λ is smooth*;
b) *X_λ satisfies Poincaré duality integrally*;
c) *X_λ is a closed parabolic orbit*.

Of course (a) $\Rightarrow$(b) and (c) $\Rightarrow$ (a) are immediate; the significant point is (b) $\Rightarrow$ (c).

Corollary 1.2 *There are only finitely many smooth Schubert varieties in a fixed $\mathcal{L}_G$.*

In fact, it is easy to see that the non-trivial closed parabolic orbits are in bijective correspondence with connected subgraphs of the affine Dynkin diagram containing the special node s_0 (Proposition 4.1).

A node of the Dynkin diagram is *minuscule* if there is an automorphism of the affine diagram carrying it to the special node s_0. A *minuscule flag variety* is a flag variety whose parabolic isotropy group is the maximal parabolic obtained by deleting a minuscule node. (*Warning:* Our "minuscule" flag varieties would be called "co-minuscule" in some sources.) Similarly, if a variety X is isomorphic to G'/P for some reductive algebraic group G' and maximal parabolic subgroup P, we will say X is a *maximal flag variety*.

It turns out that every closed parabolic orbit in $\mathcal{L}_G$ is a minuscule flag variety of some simple algebraic group, and that every minuscule flag variety occurs as a closed parabolic orbit in some affine Grassmannian (see Proposition 4.5). Hence we obtain as a by-product:

Corollary 1.3 *Let X be a Schubert variety in a minuscule flag variety. Then the following are equivalent*:

a) *X is smooth*;
b) *X satisfies Poincaré duality integrally*;
c) *X is a closed parabolic orbit*.

This corollary generalizes the fact that in the type A Grassmannian, $G_k\mathbb{C}^n$, the smooth Schubert varieties are the ones that are themselves Grassmannians [1, Cor. 9.3.3].

A Schubert variety X_λ is *palindromic* if it has palindromic Poincaré polynomial $|X_\lambda|(t) = 1 + a_1 t + \ldots + a_{d-1}t^{d-1} + t^d$, where d is the complex dimension of X_λ.

In other words, X_λ satisfies Poincaré duality additively: $a_k = a_{d-k}$ for all $0 < k < d$. We say that X_λ is a *chain* if $|X_\lambda|(t) = 1 + t + t^2 + \ldots + t^d$. In type A_n, there are two infinite families of palindromic Schubert varieties (one family if $n = 1$) $X_{n,k}$, $X'_{n,k}$ of dimension kn, introduced by the second author in [16]. We call these *spiral varieties*, for reasons to be explained in Section 11. The two families are conjugate under the automorphism of the affine Dynkin diagram fixing the special node s_0. The spiral varieties have a number of interesting properties (the first three are proved in [16]):

(1) $H_*X_{n,k}$ realizes the "degree filtration" on $H_*\Omega SU(n+1)$;

(2) $X_{n,k}$ is the variety of k-dimensional submodules in a free module of rank $n+1$ over the truncated polynomial ring $\mathbb{C}[z]/z^k$;

(3) $H^*X_{n,k}$ and $H^*G_k\mathbb{C}^{n+k}$ are isomorphic as graded abelian groups, but not as rings unless $k = 1$;

(4) (Cohen-Lupercio-Segal [8]) $X_{n,k}$ is homotopy-equivalent to $Hol_k(\mathbb{P}^1, G_{n+1}\mathbb{C}^\infty)$, the space of holomorphic maps of degree k.

Theorem 1.4 *X_λ is palindromic if and only if one of the following conditions holds*:

a) X_λ is a closed parabolic orbit (in particular, X_λ is smooth).

b) X_λ is a chain.

c) G has type A_n and X_λ is spiral.

d) G has type B_3 and $\lambda = (3, 0, -1)$.

There is some overlap in conditions a)-d). For example, in the simply-laced case a chain is a closed parabolic orbit if and only if it is a projective space, and these occur frequently. In type A_n the two spiral classes of minimal dimension n are projective spaces, but the others are neither smooth nor chains. The peculiar exception in type B_3 is a singular 9-dimensional variety with Poincaré polynomial $1 + t + t^2 + 2t^3 + 2t^4 + 2t^5 + 2t^6 + t^7 + t^8 + t^9$.

Corollary 1.5 *If G is not of type A, there are only finitely many palindromic Schubert varieties in $\mathcal{L}_G$.*

It is easy to see that in all types there are only finitely many chains (Corollary 7.2), so the corollary follows immediately from the theorem.

By a special case of a theorem of Carrell and Peterson ([6]; see also [13], XII, §2), an affine Schubert variety is palindromic if and only if it is rationally smooth. This yields the corollary:

Corollary 1.6 *Schubert variety. Then the following are equivalent*:

a) X_λ is palindromic;

b) X_λ is rationally smooth;

c) X_λ satisfies rational Poincaré duality;

d) X_λ satisfies one of the conditions (a)-(d) of Theorem 1.4.

Since we have enumerated all the palindromic Schubert varieties, the corollary can be proved *ad hoc* by checking that the singular ones satisfy rational Poincaré duality.

The equivalence of rational Poincaré duality and rational smoothness is well-known [15].

In the simply-laced case (this excludes affine type A_1, which should not be regarded as simply-laced) every chain is a projective space (Corollary 7.6), and so in particular is smooth. Hence:

Corollary 1.7 *In types D and E, an affine Schubert variety is smooth if and only if it is palindromic. In all other types there are singular palindromics.*

This contrasts with an unpublished theorem of Dale Peterson, which asserts that for ordinary Schubert varieties the corollary holds in all simply-laced types ADE. Combining Peterson's result with Corollary 1.7, we get evidence for the following conjecture.

Conjecture 1.8 *In any flag manifold $\mathcal{G}/\mathcal{Q}$ of affine or classical type, smoothness is equivalent to rational smoothness for all non-cyclic simply laced types.*

Theorem 1.4 gives a second proof of Theorem 1.1: Having listed all the palindromic Schubert varieties, we need only run through the list and show that only the closed parabolic orbits satisfy Poincaré duality integrally.

In a second forthcoming article [2], the authors consider an alternative approach to the proof of Theorem 1.4. In this work, we associate a natural family of bounded partitions to each element in the coroot lattice in such a way that the relations in Young's lattice on partitions imply relations in Bruhat order on coroot lattice elements. These relations are sufficient to differentiate all palindromic affine Schubert varieties from the non-palindromic ones.

Several tools developed in this paper may be of interest in their own right:

- A simple type-independent description for the covering relation in Bruhat order on coroot lattice elements corresponding to linear reflections is given in Proposition 9.1. A similar description for the covering relations corresponding to certain affine reflections $r_{1,\beta}$ (Proposition 9.3).
- A type-independent characterization is given of rigid elements (elements with a unique reduced expression) and chains in $\tilde{W}^S$. In order to be as explicit as possible, we give a type by type description of the corresponding coroot lattice elements in addition to the general characterization.
- The Poincaré polynomials for all smooth and palindromic affine Schubert varieties are explicitly determined.
- A necessary condition for smoothness in chains is given in terms of Chevalley's rule for multiplication in the cohomology ring.

Let us outline the proofs of Theorems 1.1 and Theorem 1.4. We use a variant of the Mozes *numbers game* [18] to observe the affine Weyl group action on the coroot lattice elements. The proof of Theorem 1.1 begins by considering some elementary obstructions to palindromy. We call this the "palindromy game" in analogy with the Mozes numbers game. The point is simply that $\mathcal{L}_G$ has only one cell of each of the first few dimensions, and hence a palindromic X_λ can't have too many cells near

the top. This already narrows down the possibilities for smooth or palindromic affine Schubert varieties considerably. In particular, in any affine Grassmannian there is a unique 2-cell, and hence a palindromic X_λ of complex dimension d can only have one $(2d-2)$-cell. If X_λ satisfies Poincaré duality, then multiplication by the generator of H^2X_λ induces an isomorphism $H^{2d-2}X_\lambda \xrightarrow{\cong} H^{2d}X_\lambda$. It is known that the classical formula of Chevalley for this cup product generalizes to the affine case (indeed to any Kac-Moody flag variety), putting further severe restrictions on which λ can occur. Along the way we also classify the chains together with their cup product structure. Each lemma, theorem and proposition is stated independent of type. A few of the lemmas employ a type by type description of the Dynkin diagram and coroot lattice elements.

The proof of Theorem 1.1 uses only the weak order on the coroot lattice; in other words, it only uses descents of the form $\lambda \downarrow s\lambda$ with s one of the Coxeter generators of the affine Weyl group. Theorem 1.4, on the other hand, requires a more elaborate version of the palindromy game incorporating the full Bruhat order; in other words, it requires descents of the form $\lambda \downarrow r\lambda$ in which r is an affine reflection associated to a non-simple root.

We note that there are useful tools for identifying smooth and rationally smooth affine Schubert varieties already in the literature. Kumar's [12] test for smoothness and rational smoothness of Schubert varieties can be applied to the affine Grassmannian [12] by testing each Schubert variety at the identity. Carrell-Kuttlers' [7] algorithm for identifying the singular locus outside of type G_2 could also be used by showing that each Schubert variety not corresponding to a closed parabolic orbit has a non-empty singular locus. However, at this time, we do not see how to use these tests to give the simple description of the smooth and rationally smooth Schubert varieties in our main theorems.

2 Notation

We follow the notation from [5] whenever possible.

G: simple, simply-connected compact Lie group of rank n
T: maximal torus, with Lie algebra $\mathfrak{t}$
W: Weyl group
$S = \{s_1, s_2, \ldots, s_n\}$: set of Coxeter generators for W
Φ, Φ^+: root system, positive roots
α_s, $s \in S$: simple positive roots. If $s = s_i$, we also denote α_s by α_i.
$m_s(\alpha)$, $\alpha \in \Phi$: $\alpha = \sum_{s\in S} m_s(\alpha)\alpha_s$
α_0: highest root; set $m_s = m_s(\alpha_0)$
D: Dynkin diagram with S as set of nodes
$\mathcal{Q}^\vee$, $\mathcal{P}^\vee$: Coroot lattice, coweight lattice
$G_\mathbb{C}$, $T_\mathbb{C}$: complexification of G, T.
B, B^-: Borel subgroup containing $T_\mathbb{C}$, opposite Borel subgroup.

$\tilde{G}_\mathbb{C}$: $G_\mathbb{C}(\mathbb{C}[z, z^{-1}])$, or regular maps $\mathbb{C}^\times \longrightarrow G_\mathbb{C}$
$\tilde{P}$: $G_\mathbb{C}(\mathbb{C}[z]) \subset \tilde{G}_\mathbb{C}$, or regular maps $\mathbb{C} \longrightarrow G_\mathbb{C}$

$\mathcal{L}_G$: $\tilde{G}_{\mathbb{C}}/\tilde{P}$, the affine Grassmannian
$\tilde{B}$: $\{f \in \tilde{P} : f(0) \in B^-\}$
$\tilde{W}$: affine Weyl group
$\tilde{S}$: $S \cup \{s_0\}$, the Coxeter generators for $\tilde{W}$
$\tilde{D}$: affine Dynkin diagram with $\tilde{S}$ as set of nodes
$\tilde{\Phi}$: affine root system $\mathbb{Z} \times \Phi$
$\tilde{W}^S$: set of minimal length representatives for $\tilde{W}/W$
ℓ, ℓ^S: length function on $\tilde{W}$, length function relative to S

Bruhat coverings. If $\sigma \in \tilde{W}/W$, and r is an affine reflection, we write $\sigma \downarrow r\sigma$ if $\ell^S(r\sigma) = \ell^S(\sigma) - 1$. Thus σ covers $r\sigma$ in the Bruhat order. If $r \in \tilde{S}$ we call this a *left descent*. The partial order generated by the left descents is the *left weak order* or just *weak order* for short. Note the right weak order isn't useful on $\tilde{W}/W$ since every non-trivial element has only s_0 as a right descent. If $\ell^S(r\sigma) = \ell^S(\sigma) + 1$, we write $\sigma \uparrow r\sigma$.

Minuscule nodes: We call a node s of D *minuscule* if it satisfies the equivalent conditions: (i) There is an automorphism of $\tilde{D}$ carrying s_0 to s; (ii) $m_s = 1$, where m_s is the coefficient of α_s in the highest root α_0. We call a flag variety of $G_{\mathbb{C}}$ *minuscule* if its parabolic isotropy group is the maximal parabolic obtained by deleting a minuscule node. Note, the minuscule fundamental coweights form a set of distinct representatives for $\mathcal{P}^\vee/\mathcal{Q}^\vee$. We caution the reader that these coweights would be called "co-minuscule" in some sources such as [1]. However, our *minuscule nodes* correspond with their fundamental cominuscule coweights.

Long and short nodes: In the simply-laced case we regard all roots as long. A node s of the Dynkin diagram is regarded as long/short according as the corresponding simple root α_s is long/short.

Poincaré series: If A is a suitable graded object—a ranked poset, a graded abelian group, etc.—we write $|A|(t)$ for the Poincaré series of A. The spaces X considered in this paper invariably have their homology groups concentrated in even dimensions, and as a slight variant of this notation we write $|X|(t) = \sum_i a_i t^i$, where $a_i = rank\, H_{2i} X$. Similarly, the CW-complexes we consider have only even-dimensional cells, and it will be convenient to call the $2k$-skeleton of X the *complex k-skeleton*. Here the *k-skeleton* of X is the union of all cells of dimension up to and including k.

3 The coroot lattice

In this section we set down some basic facts and notation concerning the coroot lattice.

3.1 The coroot lattice and the affine Weyl group

The affine Weyl group $\tilde{W}$ is the group of affine transformations of $\mathfrak{t}$ generated by all reflections across hyperplanes $\alpha = k$, where $\alpha \in \Phi$, $k \in \mathbb{Z}$. It fits into a split extension

$$\mathcal{Q}^\vee \longrightarrow \tilde{W} \longrightarrow W.$$

Hence there are bijections

$$\mathcal{Q}^\vee \xrightarrow{\cong} \tilde{W}/W \xrightarrow{\cong} \tilde{W}^S,$$

mapping $\lambda \in \mathcal{Q}^\vee$ to $w_\lambda \in \tilde{W}^S$ if λ and w_λ are in same coset of $\tilde{W}/W$. Thus we have two distinct canonical sets of coset representatives for $\tilde{W}/W$. Define

$$\ell^S(\lambda) = \min\{\ell(\lambda w) : w \in W\} = \ell(w_\lambda).$$

Both length functions can be computed using the formulas of [11]:

$$\ell(\lambda) = \sum_{\alpha\in\Phi^+} |\alpha(\lambda)|, \qquad \ell^S(\lambda) = \ell(\lambda) - q(\lambda), \tag{1}$$

where

$$q(\lambda) = |\{\alpha \in \Phi^+ : \alpha(\lambda) > 0\}|.$$

The equivalent length generating functions $|\tilde{W}^S|(t) = \sum_{w\in\tilde{W}^S} t^{\ell(w)}$ and $|\mathcal{Q}^\vee|(t) = \sum_{\lambda\in\mathcal{Q}^\vee} t^{\ell^S(\lambda)}$ can be obtained from the following beautiful formula due to Bott.

Theorem 3.1 [4] *Let $e_1, e_2, \ldots, e_n$ be the exponents of W. Then*

$$|\mathcal{L}_G|(t) = \prod_{i=1}^{n} \frac{1}{(1-t^{e_i})}.$$

Now recall that $\lambda \in \mathcal{Q}^\vee$ is *dominant* (resp. *anti-dominant*) if $\alpha(\lambda) \geq 0$ (resp. $\alpha(\lambda) \leq 0$) for all $\alpha \in \Phi^+$. It follows that $\lambda = w_\lambda \in \tilde{W}^S$ if and only if λ is anti-dominant.

The coroot lattice also inherits a Bruhat order and a left weak order from $\tilde{W}/W$. For $s \in S$ and $\lambda \in \mathcal{Q}^\vee$ we have

$$\begin{aligned} \lambda \downarrow s\lambda &\iff \alpha_s(\lambda) < 0 \\ \lambda \uparrow s\lambda &\iff \alpha_s(\lambda) > 0 \\ \lambda = s\lambda &\iff \alpha_s(\lambda) = 0 \end{aligned} \tag{2}$$

If $s = s_0$, the same three conditions hold with $\alpha_s(\lambda)$ replaced by $1 - \alpha_0(\lambda)$.

We view the set of double cosets $W\backslash\tilde{W}/W$ asymmetrically, regarding it as the orbit set of the left W action on $\tilde{W}/W$. Note that $\lambda \in \mathcal{Q}^\vee$ is dominant (resp. anti-dominant) if and only if it is the unique minimal (resp. maximal) element of its left W-orbit in $\tilde{W}/W$.

We will almost always denote elements $\lambda \in \mathcal{Q}^\vee$ using the expansion of λ in terms of the fundamental coweights $\omega_s^\vee$, $s \in S$, or equivalently as $\mathbb{Z}$-valued functions on D. Again, we freely interchange the notation $\omega_s^\vee$ and $\omega_i^\vee$ if $s = s_i$. In a given type Φ the elements of S are ordered as on page 214 following [5]; note the slightly odd ordering there in type E, in which the "off-line" node is labelled as s_2. If Φ has rank

n and $\lambda = \sum a_i \omega_i^\vee$ then we write $\lambda = (a_1, \ldots, a_n)$, where $a_i = \alpha_i(\lambda)$. We use the symbol $\underline{0}$ to denote a sequence of zeros whose length is irrelevant or determined by the context. Occasionally, however, we use what we call "standard notation", meaning the customary explicit representation of the root systems in some $\mathbb{R}^n$ as in [5]; we write e_i for the standard basis where Bourbaki writes ϵ_i.

Remark: We think of $\mathcal{Q}^\vee$ in several different ways: (1) as a lattice in $\mathfrak{t}$; (2) as a group of translations acting on $\mathfrak{t}$; (3) as a set of coset representatives for $\tilde{W}/W$; (4) as the subgroup $Hom(S^1, T) \subset \tilde{G}_{\mathbb{C}}$. This last identification uses the fact that G is simply-connected, so that the coroot lattice and the integral lattice $Ker(exp : \mathfrak{t} \longrightarrow T) = Hom(S^1, T)$ coincide. It should be clear from the context which of these interpretations is intended.

3.2 Comparison with the coweight lattice

The coroot lattice is a subgroup of finite index in the coweight lattice

$$\mathcal{P}^\vee = \{v \in \mathfrak{t} : \alpha(v) \in \mathbb{Z} \, \forall \alpha \in \Phi\}.$$

The fundamental coweights $\omega_s^\vee$, $s \in S$, are defined by $\alpha_t(\omega_s^\vee) = \delta_{st}$. Below we summarize criteria to determine if an element in $\mathcal{P}^\vee$ is actually an element of $\mathcal{Q}^\vee$.

We write $\pi_1\Phi$ for $\mathcal{P}^\vee/\mathcal{Q}^\vee$, since the latter depends only on Φ and is the fundamental group of the adjoint form of G. The minuscule fundamental coweights $\omega_s^\vee$ form a complete set of representatives for the non-trivial cosets in $\pi_1\Phi$. It is possible, however, for non-minuscule fundamental coweights to represent non-trivial elements of $\pi_1\Phi$. For the convenience of the reader, and because we need to know all such non-trivial fundamental coweights and the relations between them, we will describe $\pi_1\Phi$ type by type. The reference is [5], where some of the data is left implicit in the description of the fundamental weights. Since the data in [5] is in terms of weights rather than coweights, some translation is necessary; in particular, the weights in type B are the coweights in type C and vice-versa. For simplicity we identify the fundamental coweights with the nodes of D, and write $s \sim t$ if $s = t$ in $\pi_1\Phi$. If $\delta \in \mathcal{P}^\vee$, we write $\delta = (a_1, \ldots, a_n)$ if $\delta = \sum a_i \omega_i^\vee$.

A_n: We have $\pi_1\Phi = \mathbb{Z}/(n+1)$, with the elements of $\{\omega_i^\vee : 1 \le i \le n\}$ representing the distinct non-trivial classes. Equivalently, the nodes of D represent the non-trivial classes. Hence, $(a_1, \ldots, a_n) \in \mathcal{Q}^\vee \Leftrightarrow \sum i a_i \equiv 0 \bmod (n+1)$.

B_n: Here $\pi_1\Phi = \mathbb{Z}/2$; the non-trivial nodes are the odd nodes and these are all identified. Hence, $(a_1, \ldots, a_n) \in \mathcal{Q}^\vee \Leftrightarrow \sum a_{odd} \equiv 0 \bmod 2$.

C_n: Here $\pi_1\Phi = \mathbb{Z}/2$; s_n is the only non-trivial node. Hence, $(a_1, \ldots, a_n) \in \mathcal{Q}^\vee \Leftrightarrow a_n \equiv 0 \bmod 2$.

D_n: Here we have

$$\pi_1\Phi \cong \begin{cases} \mathbb{Z}/2 \times \mathbb{Z}/2 & \text{if } n \text{ even} \\ \mathbb{Z}/4 & \text{if } n \text{ odd} \end{cases}$$

The non-trivial classes are represented by s_1, s_{n-1}, s_n. If $k < n-1$, then s_k is non-trivial if and only if k is odd, in which case $s_k \sim s_1$. Now let Σ_{odd} denote the sum of

the a_i's with i odd, $i < n-1$. If n is odd, $(a_1, \ldots, a_n) \in \mathcal{Q}^\vee \Leftrightarrow a_{n-1} - a_n + 2\Sigma_{odd} \equiv 0 \bmod 4$. If n is even, $(a_1, \ldots, a_n) \in \mathcal{Q}^\vee \Leftrightarrow \Sigma_{odd} + a_{n-1} \equiv 0 \equiv \Sigma_{odd} + a_n \bmod 2$.

E_6: $\pi_1\Phi = \mathbb{Z}/3$, with s_1 and s_6 representing the non-trivial nodes. The other non-trivial nodes are $s_3 \sim s_6$ and $s_5 \sim s_1$. Hence, $(a_1, \ldots, a_6) \in \mathcal{Q}^\vee \Leftrightarrow a_1 - a_3 + a_5 - a_6 \equiv 0 \bmod 3$.

E_7: $\pi_1\Phi = \mathbb{Z}/2$, with s_2, s_5, s_7 the non-trivial nodes. Hence, $(a_1, \ldots, a_7) \in \mathcal{Q}^\vee \Leftrightarrow a_2 + a_5 + a_7 \equiv 0 \bmod 2$.

E_8, F_4, G_2: $\pi_1\Phi$ is trivial, so $\mathcal{Q}^\vee = \mathcal{P}^\vee$.

4 Parabolic orbits

4.1 Closed parabolic orbits

Let $I \subset \tilde{S}$ be a proper subset that contains s_0. The corresponding parabolic subgroup $P_I \subset \tilde{G}_{\mathbb{C}}$ is generated by $\tilde{B}$ and the simple reflections in I lifted to $\tilde{G}_{\mathbb{C}}$. Note that P_I has a unique closed orbit in $\mathcal{L}_G$, namely

$$Y_I = P_I \tilde{P}/\tilde{P} = P_I/P_{I-\{s_0\}}.$$

Note that Y_I depends only on the component of s_0 in the subgraph of $\tilde{D}$ defined by I. Hence:

Proposition 4.1 *The non-trivial closed parabolic orbits are in bijective correspondence with connected subgraphs of $\tilde{D}$ containing s_0.*

From now on we will assume that I is connected, and let p_G denote the number of subgraphs as in the proposition. (We usually ignore the trivial orbit, which corresponds to $P_\emptyset = \tilde{B}$).

Let L_I denote the Levi factor of P_I. Since any proper sub-Coxeter system of $(\tilde{W}, \tilde{S})$ is finite, L_I is a finite dimensional algebraic group whose commutator subgroup $G_{I,\mathbb{C}}$ is the simple algebraic group associated to the Dynkin diagram defined by I. It is clear that $Y_I = G_{I,\mathbb{C}}/Q$, where Q is the maximal parabolic subgroup of $G_{I,\mathbb{C}}$ associated to $I - \{s_0\}$. Hence Y_I is isomorphic to an ordinary maximal flag variety, and in particular is irreducible and smooth. Since Y_I is also $\tilde{B}$-invariant, it is therefore a smooth Schubert variety in $\mathcal{L}_G$. Hence $Y_I = X_\sigma$, where σ is the maximal element of $(W_I)^{I-\{s_0\}} \subset \tilde{W}^S$.

Note that using Proposition 4.1, one can easily compute p_G and list all the closed parabolic orbits explicitly. For example, in type A_n there are $\binom{n+1}{2}$ non-trivial closed parabolic orbits, all of which are Grassmannians. In type E_8 there are ten non-trivial closed parabolic orbits, nine of which are projective spaces. The exception is the maximal closed parabolic orbit Y_I obtained by deleting the node s_1; from the affine Dynkin diagram we see that Y_I has type D_8/D_7, a nonsingular quadric hypersurface of dimension 14.

The dimension of any ordinary flag variety $G_{\mathbb{C}}/Q_J$ can be computed as follows: Let Φ_J denote the root subsystem corresponding to J. Then

$$dim\,(G_{\mathbb{C}}/Q_J) = |\Phi^+| - |\Phi_J^+|.$$

In particular, we have

$$dim\, Y_I = |\Theta_I^+| - |\Theta_{I-\{s_0\}}^+|.$$

where Θ_I is the root system associated to the subgraph I. We identify Θ_I with the root subsystem of Φ having $\{\alpha_0\} \cup \{-\alpha_s : s \in I - \{s_0\}\}$ as a base.

Now let $\mathcal{N}(I)$ denote the set of neighbor nodes of I; note that $\mathcal{N}(I)$ uniquely determines I given that $s_0 \in I$ and I is connected. Let

$$A_I = \{\alpha \in \Phi^+ : m_s(\alpha) = m_s(\alpha_0)\, \forall s \in \mathcal{N}(I)\}.$$

Lemma 4.2 *$A_I = \Theta_I^+ - \Theta_{I-\{s_0\}}^+$. Hence $dim\, Y_I = |A_I|$.*

Proof: It is clear that $\Theta_I^+ - \Theta_{I-\{s_0\}}^+ \subset A_I$. For the reverse inclusion, note that any positive root $\alpha \in \Phi^+$ can be obtained from α_0 by successively subtracting simple roots α_s for various $s \in S$. If $\alpha \in A_I$, then no such α_s can have $s \in \mathcal{N}(I)$. Now write $\tilde{S}$ as a disjoint union $\tilde{S} = I \coprod \mathcal{N}(I) \coprod K$. Then if $\beta \in \Theta_I$ and $\gamma \in \Phi_K$, $\beta + \gamma$ is not a root. Hence no such α_s can have $s \in K$, and it follows that $\alpha \in \Theta_I^+ - \Theta_{I-\{s_0\}}^+$.

Now let $\lambda_I \in \mathcal{Q}^\vee$ denote the coroot lattice representative for the top cell in Y_I. Then

$$\ell^S(\lambda_I) = dim\, Y_I = |A_I|.$$

Furthermore, setting $S^+(\lambda) = \{s \in S : \alpha_s(\lambda) > 0\}$, we have:

Lemma 4.3 *$S^+(\lambda_I) = \mathcal{N}(I)$.*

Proof: If $s \in I$ then $s\lambda_I \leq \lambda_I$ in $\tilde{W}/W$, so $\alpha_s(\lambda_I) \leq 0$. If $s \in K$ (where K is as in the proof of Lemma 4.2), then s commutes with the elements of I; hence $\alpha_s(\lambda_I) = 0$. Finally, if $s \in \mathcal{N}(I)$ then $\lambda_I \uparrow s\lambda_I$; hence $\alpha_s(\lambda_I) > 0$.

There is a simple way to recognize a dominant closed parabolic orbit:

Proposition 4.4 *Suppose λ is non-trivial and dominant. Then X_λ is a closed parabolic orbit if and only if $\alpha_0(\lambda) = 2$.*

Proof: Suppose $\alpha_0(\lambda) = 2$. Equivalently, λ has the form (i) $\omega_i^\vee$ with $m_i = 2$; or (ii) $2\omega_i^\vee$ with $m_i = 1$ (i.e. i minuscule); or (iii) $\omega_i^\vee + \omega_j^\vee$ with i, j minuscule. It follows that there is a unique connected subset $I \subset \tilde{S}$ containing s_0 such that $S^+(\lambda) = \mathcal{N}(I)$, where in case (iii) this uses the fact that every minuscule node of S is a leaf node. Moreover, X_λ is P_I-invariant, and inspection of cases (i)-(iii) shows that $\ell^S(\lambda) = |A_I|$. Hence $dim\, X_\lambda = dim\, Y_I$ by Lemma 4.2, forcing $X_\lambda = Y_I$.

Conversely, suppose λ_I is dominant. Then $\lambda_I \downarrow s_0\lambda_I$, and hence $\alpha_0(\lambda_I) \geq 2$. Moreover,

$$|A_I| = \ell^S(\lambda_I) = \sum_{\alpha(\lambda_I)>0} (\alpha(\lambda_I) - 1) \geq \sum_{\alpha \in A_I} (\alpha(\lambda_I) - 1) = (\alpha_0(\lambda_I) - 1)|A_I|.$$

Hence $\alpha_0(\lambda_I) = 2$.

We conclude this section with a proof of Corollary 1.3.

Proposition 4.5 *Every closed parabolic orbit is a minuscule flag variety of some simple algebraic group. Moreover every minuscule flag variety occurs as a closed parabolic orbit in some affine Grassmannian.*

Proof: The first assertion is immediate, since α_0 can occur at most once in a positive root of the system Θ. Conversely, suppose s is a minuscule node of D, and let $Z = G_{\mathbb{C}}/Q_{S-s}$ denote the corresponding flag variety. Then there is an automorphism ϕ of $\tilde{D}$ taking s to s_0. Hence $Z \cong Y_I$, where $I = \tilde{S} - \{s\}$.

Corollary 1.3 is now clear from Theorem 1.1.

4.2 General parabolic orbits

In this section we fix a proper subset $I \subset \tilde{S}$ and consider arbitrary P_I orbits in $\mathcal{L}_G$. The propositions here are well-known (see [17] for a detailed exposition), so we omit the proofs.

Call an element $\lambda \in \mathcal{Q}^\vee$ *I-minimal* if it is the minimal element of its left $\tilde{W}_I$-orbit. Then λ is I-minimal if and only if $\alpha_s(\lambda) \geq 0$ for every $s \in I$ (or $1 - \alpha_0(\lambda) \geq 0$ when $s = s_0$). Every P_I -orbit contains a unique I-minimal λ, and from now on we assume λ is I-minimal unless otherwise specified. Let $\mathcal{O}_\lambda = P_I\lambda \subset \mathcal{L}_G$, and let $M_\lambda = L_I\lambda \subset \mathcal{O}_\lambda$ denote the corresponding Levi orbit.

Proposition 4.6 *$\mathcal{O}_\lambda$ is isomorphic as an algebraic variety to the total space of a vector bundle ξ_λ over M_λ, with fiber dimension $\ell^S(\lambda)$. (Here λ is the I-minimal representative of the orbit.)*

The bundle ξ_λ can be described explicitly in terms of a certain representation of L_I arising from the adjoint representation of L_I on the the Lie algebra of the unipotent radical of P_I. We will not need this description here.

The closure relations on the P_I orbits are given by the Bruhat order on the set of I-minimal λ. We then have a filtration of $\mathcal{L}_G$ for which the quotients are the Thom spaces $T(\xi_\lambda)$. Now let $p_\lambda(t)$ denote the generating function for the cells that lie in $\mathcal{O}_\lambda$. Note that $p_\lambda(t)$ is counting cells, not homology groups (although $1 + p_\lambda(t) = |T(\xi_\lambda)|(t)$). In fact

$$p_\lambda(t) = t^{\ell^S(\lambda)}|M_\lambda|(t).$$

Now let $K_\lambda \subset I$ denote those s such that $\alpha_s(\lambda) = 0$ (or $1 - \alpha_0(\lambda) = 0$ if $s = s_0$; in the language of the next section, these are the "zero nodes" of λ). Then M_λ is a flag variety of type Φ_I/Φ_{K_λ}.

Proposition 4.7 *We have*

$$|\mathcal{L}_G|(t) = \sum_{\lambda} t^{\ell^S(\lambda)} |M_\lambda|(t),$$

where the sum is over all I-minimal λ. Moreover,

$$|M_\lambda|(t) = \frac{\prod_{i=1}^{|I|}(1-t^{e_i})}{(1-t)^{|I|-|K_\lambda|}\prod_{j=1}^{|K_\lambda|}(1-t^{f_j})},$$

where $e_1, ..., e_{|I|}$ (resp. $f_1, ..., f_{|K_\lambda|}$) are the degrees of $\tilde{W}_I$ (resp. $\tilde{W}_{K_\lambda}$).

We apply this proposition to the exceptional Schubert variety in type B_3.

Corollary 4.8 *Let G have type B_3. Then $X_{(3,0,-1)}$ is a singular palindromic of dimension* 9, *with Poincaré polynomial* 1112222111. *It satisfies Poincaré duality over $\mathbb{Q}$ but not over $\mathbb{Z}$, and has singular locus $X_{(2,0,0)}$.*

Proof: Note that $\ell^S(3,0,-1) = 9$. It is not difficult to compute the Bruhat order on $\tilde{W}^S$ through dimension 9 (by hand or by computer); then one can read off the Poincaré polynomial. Another approach is as follows: Let $I = \{s_0, s_2, s_3\}$ and note that $X_{(3,0,-1)}$ is P_I invariant. We will now use the proposition to analyze the P_I-orbit decomposition through dimension 9.

Let $\lambda_0 = (0,0,0)$, $\lambda_1 = (-1,0,1)$, $\lambda_3 = (-2,1,0)$. Note that these elements are I-minimal of S-length 0, 3, and 5, respectively. The Levi orbits M_{λ_i} have types B_3/B_2, B_3/A_2, B_3/A_1 respectively. Since B_n has exponents $2, 4, \ldots 2n$, while A_n has exponents $2, 3, \ldots n+1$, we conclude from Proposition 4.7 that

1. $|M_{\lambda_0}|(t) = 111111$ (dimension 5); note that $M_{\lambda_0} = \mathcal{O}_{\lambda_0}$ is the closed parabolic orbit $X_{(2,0,0)}$.
2. $|M_{\lambda_1}|(t) = 1112111$ (dimension 6)
3. $|M_{\lambda_2}|(t) = 123444321$ (dimension 8).

The generating function for the cells of $\mathcal{O}_{\lambda_0} \cup \mathcal{O}_{\lambda_1} \cup \mathcal{O}_{\lambda_2}$ is then $|M_{\lambda_0}|(t) + t^3|M_{\lambda_1}|(t) + t^5|M_{\lambda_2}|(t)$. Comparing with $|\mathcal{L}_G(t)| = \frac{1}{(1-t)(1-t^3)(1-t^5)}$, we conclude that the complex 8-skeleton of $\mathcal{L}_G$ is contained in $\mathcal{O}_{\lambda_0} \cup \mathcal{O}_{\lambda_1} \cup \mathcal{O}_{\lambda_2}$. Since the top cell of $\mathcal{O}_{\lambda_3}$ is in dimension $5 + 8 = 13 > 9$, it follows that $X_{(3,0,-1)} = \mathcal{O}_{\lambda_0} \cup \mathcal{O}_{\lambda_1}$. Hence $|X_{(3,0,-1)}|$ is as claimed, and in particular $X_{(3,0,-1)}$ is palindromic.

By Theorem 1.1 (or by direct application of the Chevalley formula; see §6), $X_{(3,0,-1)}$ does not satisfy Poincaré duality integrally, and in particular is singular. Since the open orbit $\mathcal{O}_{\lambda_1}$ is smooth, the singular locus can only be $\mathcal{O}_{\lambda_0} = X_{(2,0,0)}$. Finally, $X_{(3,0,-1)}$ satisfies Poincaré duality rationally by the Carrell-Peterson theorem.

5 The palindromy game I: weak order and the coroot lattice

In §5.1 we give an informal overview of the palindromy game, in its simpler form using only the weak order. More details are given in §5.2.

5.1 The node-firing game

In this section we describe a variation on Mozes *numbers game* [18] which we call the *node-firing game*. Our description of the game follows [3, 9]. The purpose of this game is to make the bijection between $\tilde{W}^S$ and $\mathcal{Q}^\vee$ explicit and to highlight the left-weak order.

We identify the coweight lattice $\mathcal{P}^\vee$ with the group of $\mathbb{Z}$-valued functions on the Dynkin diagram D, where $s \mapsto \alpha_s(\lambda)$. We extend this labeling to the affine diagram $\tilde{D}$ by putting the value $1 - \alpha_0(\lambda)$ on the node s_0. This latter value is of course determined by the others, but it is important to include it as part of the picture. The coroot lattice $\mathcal{Q}^\vee$ can then be identified with a subgroup of finite index in the group of all labelled diagrams using the criteria in Section 3.2. For example, suppose Φ has type D_n with $n \geq 6$, and $\lambda = -\omega_2^\vee + 2\omega_3^\vee + \omega_{n-2}^\vee$. Then the labelled diagram of λ is

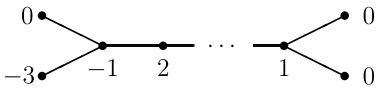

Here s_0 is the lower left node, and the nodes not shown are all labeled zero. Note, this labeled diagram is not in $\mathcal{Q}^\vee$ if n is odd. In fact, $\lambda = \omega_{n-2}^\vee \bmod \mathcal{Q}^\vee$, and hence $\lambda \in \mathcal{Q}^\vee$ if and only if n is even.

Now λ has a left descent $\lambda \downarrow s\lambda$ precisely when s is a negative node; i.e., $\alpha_s(\lambda) < 0$, or $1 - \alpha_0(\lambda) < 0$ when $s = s_0$. We refer to this descent as *firing* the node s. The effect of such a firing on the labeled diagram is as follows:

(1) The value at s is replaced by its negative;
(2) $k\alpha_s(\lambda)$ (or $1 - \alpha_0(\lambda)$, when $s = s_0$) is added to each adjacent node t, where

$$k = \begin{cases} 1 & \text{if } s \text{ is long or } s, t \text{ are joined by a single edge} \\ 2 & \text{if } s, t \text{ are joined by a double edge with } s \text{ at the short end, or } \Phi \text{ has type } A_1 \\ 3 & \text{if } s, t \text{ are joined by a triple edge with } s \text{ at the short end} \end{cases}$$

If s is at the short end of a multiple bond, we call the firing *back-firing*. As an example in type F_4, let λ be given by the following diagram:

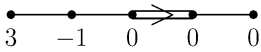

Firing twice yields

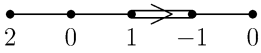

If we were firing along a type A subgraph, the configuration of adjacent 1,-1 surrounded by zeros would simply continue moving steadily to the right. Here, however, firing s_3 back-fires against the arrow to produce

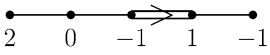

which has two negative nodes and hence covers a pair of elements in the weak order. Thus we have reached a fork in the Hasse diagram of the order ideal of λ.

We remark that the node-firing game yields a simple algorithm for computing the bijections $\mathcal{Q}^\vee \cong \tilde{W}^S$ defined by the diagram

$$\mathcal{Q}^\vee \xrightarrow{\cong} \tilde{W}/W \xrightarrow{\cong} \tilde{W}^S.$$

Suppose first that we are given $\lambda \in \mathcal{Q}^\vee$. If $\lambda = 0$, then every node has label 0 except for the node corresponding to s_0 which is labeled 1. If $\lambda \neq 0$, then λ has at least one negative node $t_1 \in \tilde{S}$. Then $\lambda \downarrow t_1\lambda$. Repeating the process yields

$$\lambda \downarrow t_1\lambda \downarrow t_2t_1\lambda \downarrow \ldots \downarrow t_mt_{m-1}\ldots t_1\lambda = 0,$$

where $m = \ell^S(\lambda)$. Taking $\sigma = t_1t_2\cdots t_m$, we have (i) $\sigma \in \tilde{W}^S$ (in particular, $t_m = s_0$), (ii) the product is reduced, and (iii) $\lambda W = \sigma W$. Hence $\lambda \mapsto \sigma$.

In the reverse direction, suppose we are given $\sigma \in \tilde{W}^S$. The corresponding $\lambda \in \mathcal{Q}^\vee$ is obtained by letting σ act on $0 \in \mathcal{Q}^\vee$, and is computed explicitly as follows: Choose a reduced decomposition $\sigma = t_1t_2\cdots t_m$, where necessarily $t_m = s_0$. Then fire up starting from 0:

$$0 \uparrow t_m \cdot 0 \uparrow t_{m-1}t_m \cdot 0 \uparrow \ldots \uparrow t_1t_2\ldots t_m \cdot 0 = \lambda.$$

(By "'firing up" we mean simply reverse the node-firing procedure described above, firing positive nodes to raise length rather than negative nodes to lower length.) Many examples of these node-firings can be found below.

5.2 Elementary obstructions to palindromy

In order to show that a given λ is *not* palindromic, we show that it has too many cells near the top dimension. Often these violations of palindromy can be detected by merely firing negative nodes. In general, however, we must consider more general Bruhat descents $\lambda \downarrow r\lambda$ defined by non-simple reflections $r \in \tilde{W}$. Whenever possible, we arrange things so that the necessary information can be read off directly from the labelled diagram. Indeed the reader may find it helpful—or at least amusing—to think of this process as a game, called the *palindromy game*, in which the object of one player is to find a palindromic λ satisfying given initial conditions, while the object of the other is to prevent it by finding an excess of Bruhat descents $\lambda \downarrow r\lambda$. In fact we will often refer to such descents as "moves".

We emphasize that the "palindromy game" is much easier to play than to write down. In many cases, the reader may prefer to draw the pictures and work out the moves for herself, rather than wade through the verbiage required to explain them in print. To get started, here is an informal discussion of the most basic principles of the game. Suppose that λ is nonzero and palindromic. Then there are the following *palindromy rules*. These rules give necessary conditions for X_λ to be palindromic.

Rule 1. λ has exactly one negative node $s \in \tilde{D}$ (the case $s = s_0$ corresponds to λ dominant).

Rule 2. Except in type A, there cannot be two zero nodes adjacent to s.

Rule 3. If $s \neq s_0$, then $\alpha_0(\lambda) \leq 1$.

Rule 4. More generally, λ cannot "fork too soon".

Note that Rule 3 is a special case of rule 1. If an arbitrary λ violates Rule 3 or if $s = s_0$ and $\alpha_0(\lambda) > 2$ then we say λ is *overweight*.

See below for the precise definition of "forks too soon". Informally, this just means that the Hasse diagram of λ (coming down from the top) reaches a fork sooner than the Hasse diagram of $\tilde{W}^S$ (coming up from the bottom), thereby violating palindromy. Often one can see this instantly from the diagram.

Example: In the F_4 example above, $\lambda = (-1, 0, 0, 0)$ clearly forks too soon because s_1 has a "head-start" on s_0.

Example: Surprisingly, E_8 is in many ways the simplest type. One reason for this is that it is the unique simply-laced type with no minuscule nodes; another reason will be given below. Firing up from $0 \in \mathcal{Q}^\vee$ in E_8 it takes six steps to reach the fork. Hence if a given λ is to have any chance at palindromy, the fork at s_4 must be suitably protected. For example, suppose $s = s_1$ is the unique negative node of λ. Then either s must be blocked away from s_4 by an intermediate positive node (picture the negative value moving to the right under repeated firings), or at least one of the exit nodes s_2, s_5 of the fork must be positive (and in fact must be at least as large as $|\alpha_1(\lambda)|$, but we ignore this refinement for the moment).

Now suppose in addition that $\alpha_1(\lambda) = -1$. Then by exploiting Rule 1 and Rule 3 and the coefficients m_s in the expansion $\alpha_0 = \sum_{s \in S} m_s(\alpha)_s$ (see §12), we see at once that λ can only have the form

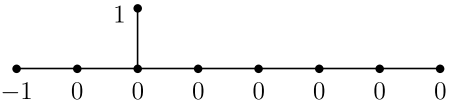

Here X_λ turns out to be the closed parabolic orbit Y_I with $I = \tilde{S} - \{s_2\}$, which has has type $\mathbb{P}^8$. This is easily seen by firing the negative node all the way down to the bottom. Note that the 1 serves to protect the fork, and is killed by the -1 as it passes by. This example illustrates why all three of the E types are actually easier than the classical types: There are few repetitions among the coefficients m_s.

Consider the length generating function in Theorem 3.1 expanded out

$$|\tilde{W}^S|(t) = \sum_{\sigma \in \tilde{W}^S} t^{\ell(\sigma)} = 1 + t + \ldots + t^{k-1} + a_k t^k + \ldots$$

where $k = k_G$ is minimal such that $a_k > 1$. If no such k exists, we set $k_G = \infty$ (this happens only in type A_1).

Pictorially, the first fork (going up) in the Hasse diagram for $\tilde{W}^S$ occurs at height $k_G - 1$ (see the diagrams at the end of the paper for examples). To compute k_G we start at s_0 on the affine Dynkin diagram, and follow the only possible path until a node of degree 3 or higher is reached. The number of nodes in such a path is $k_G - 1$.

Here we allow doubling back along a multiple edge; for example, in type F_4 we reach a fork at s_3, where we have the option of continuing to s_4 or following the unused edge back to s_2:

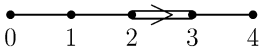

The coefficient a_{k_G} is just the number of options at the fork. Thus k_G and a_{k_G} are easily determined by inspecting the affine Dynkin diagrams.

$$k_G = \begin{cases} 2 & \text{type } A_n, n > 1 \\ 3 & \text{type } B, C, D \\ 4 & \text{type } E_6 \\ 5 & \text{type } E_7, F_4, G_2 \\ 7 & \text{type } E_8 \\ \infty & \text{type } A_1 \end{cases} \tag{3}$$

$$a_{k_G} = \begin{cases} 3 & \text{in type } D_4 \\ 2 & \text{otherwise} \end{cases} \tag{4}$$

Remark: This result can also be proved by computing the rational cohomology of the loop group: It is well known that

$$H^*(BG;\mathbb{Q}) \xrightarrow{\cong} (H^*(BT;\mathbb{Q}))^W,$$

and that the ring of invariants $(H^*(BT:\mathbb{Q}))^W$ is a polynomial algebra on generators of complex dimension $d_1 \le d_2 \le \ldots \le d_n$. The degrees d_i can be computed explicitly in each Lie type; see [10], p. 59. Since ΩG is the double-loop space of BG, we have $k_G = d_2 - 2$, yielding the table above. From this point of view, the exceptional value $a_{k_G} = 3$ in type D_4 can be traced to the fact that $H^8 BSpin(8)$ has rank 3, with generators the Pontrjagin classes p_1^2, p_2 plus the Euler class.

Now for all X_λ, $|X_\lambda|(t) \le |\tilde{W}^S|(t)$ coefficient by coefficient. In addition, if X_λ is palindromic of dimension d, then

$$D|X_\lambda|(t) \le |\tilde{W}^S|(t)$$

where if $f(t) = a_0 + a_1 t + \ldots + a_n t^n$ is a polynomial of degree n, the *dual polynomial* is $Df(t) = t^n f(t^{-1}) = a_n + a_{n-1}t + \ldots + a_0 t^n$. In a range of dimensions (depending on λ), the inequality will actually be an equality. In any case, (3) forces restrictions on $|X_\lambda|(t)$ near the top dimension d. Consider for example the configurations in a Hasse diagram

which we call a *pair*, a *fork*, and a *trident* respectively, with λ sitting at the top. If X_λ is palindromic then λ cannot cover a pair or a trident, and if Φ is not of type A then

it cannot cover a fork. In type B we will encounter an (upside down) *scepter*

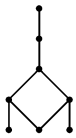

A palindromic λ cannot cover a scepter.

In general, we say that λ *forks too soon* if the Hasse diagram of its order ideal (coming down from the top) reaches a fork sooner than the Hasse diagram of $\tilde{W}^S$ (coming up from the bottom). More precisely: Say $|X_\lambda|(t) = 1 + a_1 t + \ldots + a_k t^k + t^{k+1} + \ldots + t^m$, where $m = \ell^S(\lambda)$ and k is maximal such that $a_k > 1$. Then λ forks too soon if $0 \leq m - k < k_G$. Hence, this proves Rule 4 in the palindromy game.

Example: We show that if G has type E_8 and λ is anti-dominant and non-trivial, then X_λ is not palindromic. Assume there is a unique negative node s (otherwise λ covers a pair), and all other nodes are zero. Assume further that s is a leaf node of D (otherwise λ covers a fork). Finally, if s is one of the three leaf nodes s_1, s_2, s_8, then by repeated firing we reach the fork in the Dynkin diagram in 3, 2 or 5 steps respectively. But it takes 6 steps to reach the fork from s_0. Thus λ forks too soon and hence is not palindromic.

6 Poincaré duality and the affine Chevalley formula

We state an affine version of the Chevalley formula, Proposition 6.1, and record its implications for Poincaré duality. Proposition 6.1 is a special case of a vastly more general cup product formula in equivariant cohomology, valid for arbitrary Kac-Moody flag varieties ([13], Corollary 11.3.17 and Remark 11.3.18).

Let $[X_\lambda] \in H_{2d}\mathcal{L}_G$ denote the homology class carried by X_λ, where $d = \ell^S(\lambda)$. These classes form the *Schubert basis* of $H_*\mathcal{L}_G$. Of course we can equally well regard $[X_\lambda]$ as a homology class in any Schubert variety containing X_λ. Let $y_\lambda \in H^{2d}\mathcal{L}_G$ be Kronecker dual to $[X_\lambda]$ with respect to the Schubert basis. We will use the abbreviation y for the special class $y = y_{\alpha_0^\vee} = y_{s_0}$, the generator of $H^2\mathcal{L}_G$.

We have seen that if X_λ is palindromic of dimension d, then it has just one $2d-2$ cell, and hence $H^{2d-2}X_\lambda \cong \mathbb{Z}$. If in addition X_λ satisfies Poincaré duality, then cup product with y defines an isomorphism $H^{2d-2}X_\lambda \xrightarrow{\cong} H^{2d}X_\lambda$. Hence the following *affine Chevalley formula* puts severe restrictions on the possible such λ.

Proposition 6.1 *If $\lambda \uparrow s\lambda$ for $s \in \tilde{S}$, then*

$$< yy_\lambda, [X_{s\lambda}] >= \begin{cases} c\alpha_s(\lambda) & \text{if } s \neq s_0 \\ 1 - \alpha_0(\lambda) & \text{if } s = s_0 \end{cases}$$

where

$$c = \begin{cases} 1 & \text{if } \alpha_s \text{ is long} \\ 2 & \text{if } \alpha_s \text{ is short in type } B,C,F \\ 3 & \text{if } \alpha_s \text{ is short in type } G_2 \end{cases}$$

Note that the assumption $\lambda \uparrow s\lambda$ is equivalent to the positivity of $\alpha_s(\lambda)$ or $1 - \alpha_0(\lambda)$ in the node-firing game. It will be convenient to reformulate the Chevalley formula in terms of the cap product.

Proposition 6.2 *If $\lambda \downarrow s\lambda$ for $s \in \tilde{S}$, then*

$$< y_{s\lambda}, y \cap [X_\lambda] >= \begin{cases} -c\alpha_s(\lambda) & \text{if } s \neq s_0 \\ \alpha_0(\lambda) - 1 & \text{if } s = s_0 \end{cases}$$

where $c = 1, 2, 3$ is the constant defined in Proposition 6.1.

Proof: This follows by simply reversing the roles of λ and $s\lambda$ in Proposition 6.1.

We then have at once:

Proposition 6.3 *If X_λ satisfies Poincaré duality and $\lambda \downarrow s\lambda$ for $s \in \tilde{S}$, then $c = 1$ and $\alpha_s(\lambda) = -1$ (or $1 - \alpha_0(\lambda) = -1$). In particular α_s must be long.*

Remark: This proposition already suffices to show that there are only finitely many Schubert varieties satisfying Poincaré duality in a fixed $\mathcal{L}_G$, since it bounds the values $\alpha_s(\lambda)$ for $s \in S$.

7 Chains

In this section we study the chains in $\mathcal{L}_G$. In particular, we will show that the theorems of the introduction hold for chains. We begin with some simple observations.

7.1 General observations

Proposition 7.1 *Every infinite subset of $\tilde{W}^S$ is cofinal in the Bruhat order.*

Proof: If $I \subset \tilde{S}$ is a proper subset, then $\tilde{W}_I$ is a finite Coxeter group. Let k_I denote the maximal length of an element of $\tilde{W}_I$, and let $k = max\, k_I$, where I ranges over all such proper subsets. Then if $w \in \tilde{W}$ and $\ell(w) > k$, every reduced expression for w contains every $s \in \tilde{S}$ at least once.

Now let $V \subset \tilde{W}^S$ be an infinite subset, and let $\sigma \in \tilde{W}^S$. Since $\tilde{W}^S$ has only finitely many elements of any fixed length, we can choose $v \in V$ with $\ell(v) \geq (k+1)\ell(\sigma)$. It then follows from the preceding paragraph that $\sigma \leq v$, proving that V is cofinal.

Corollary 7.2 *If W is not of type A_1, there are only finitely many chains X_λ in $\mathcal{L}$.*

Proof: If there are infinitely many chains, then it follows from Proposition 7.1 that $\tilde{W}^S$ itself is a chain. But this is the case only in type A_1.

Following Stembridge's terminology [19], call an element of a Coxeter group *rigid* if it has a unique reduced expression. Note that if w is rigid, then so is any element obtained by taking a factor of the reduced expression for σ. The next result follows by an easy induction on length.

Proposition 7.3 *Every chain is rigid.*

We will classify the chains by first classifying all the rigid elements. In fact the rigid elements are easily determined from the affine Dynkin diagram. Suppose $s, t \in \tilde{S}$ with $s \neq t$, and let m_{st} denote the order of st. Thus $m_{st} = \{2, 3, 4, 6\}$. Then a rigid element cannot contain subwords of the form st, sts, $stst$, $ststst$ respectively in these four cases. Therefore, rigid elements are fully commutative [19]. We interpret these restrictions on the Dynkin diagram as follows:

Let $\sigma = t_k t_{k-1} \cdots t_1$ be the reduced expression for a rigid element, where $\sigma \in \tilde{W}^S$ and hence $t_1 = s_0$. Then every pair of adjacent nodes in this expression must be adjacent in $\tilde{D}$. Hence σ determines and is determined by a path in $\tilde{D}$ starting at s_0. Furthermore the path in question can reverse direction only along a multiple edge, and if the multiple edge is a double edge then it can reverse direction only once. If it is a triple edge then the path can reverse direction at most three times. Let us call such a path an *admissible path*. Then every rigid element is associated to an admissible path in this way and vice versa. This allows us to read off the rigid elements directly from $\tilde{D}$.

Examples: 1. Type A. There are two infinite families of rigid elements, obtained in the evident way by starting at s_0 and following an admissible path of arbitrary length clockwise or counterclockwise around the diagram. These are the only rigid elements. Note that the two families are conjugate under the involution of $\tilde{D}$ fixing s_0.

2. Type C. Consider an admissible path (starting at s_0).

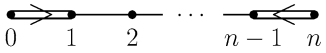

At s_1 there are two options: We can reverse direction to obtain a maximal rigid element $s_0 s_1 s_0$, or we can continue to the right. In the latter case we obtain an infinite family of rigid elements by running back and forth along $\tilde{D}$ in the evident way. These are the only rigid elements.

3. If W is not of type A or C, then there are only finitely many rigid elements. This is also clear, because then $\tilde{D}$ has no cycles and at most one multiple edge; hence every maximal admissible path eventually terminates at a leaf node.

Let σ be a chain of length m and let $y_0 = 1, y_1, \ldots, y_m$ denote the Schubert basis for $H^* X_\sigma$. Define integers a_k by $y_1 y_{k-1} = a_k y_k$ for $1 \leq k \leq m$. Note that these integers are positive by the Chevalley formula, and $a_1 = 1$. In particular, $H^*(X_\sigma; \mathbb{Q})$ is a truncated polynomial algebra $\mathbb{Q}[y_1]/y_1^{m+1}$, and hence X_σ satisfies rational Poincaré duality. Call $(a_1, \ldots, a_m)$ the *cup sequence* of σ.

Proposition 7.4 *Let σ be a chain of length m. Then X_σ satisfies Poincaré duality over $\mathbb{Z}$ if and only if the cup sequence of σ is palindromic, in the sense that $a_k = a_{m-k+1}$ for all k.*

Proof: Define c_k by $y_1^k = c_k y_k$, and note that $c_k \neq 0$. Then Poincaré duality holds if and only if $c_k c_{m-k} = c_m$ for all k. But $c_k = a_1 \dots a_k$, and the result follows by induction on k.

7.2 Simply-laced types

We show that the chains in the simply laced types all lead to smooth affine Schubert varieties.

Proposition 7.5 *Let $\sigma = t_k t_{k-1} \cdots t_1$ be a chain in $\tilde{W}^S$ whose associated admissible path has no multiple edges. Then X_σ is a closed parabolic orbit isomorphic to $\mathbb{P}^k$, and hence X_σ is smooth.*

Proof: Suppose that W is not of type A. Then the t_i's are distinct, since $\tilde{D}$ has no cycles and the path cannot reverse direction. Hence the path of σ is just a type A_k subgraph I of $\tilde{D}$, and $X_\sigma = Y_I \cong \mathbb{P}^k$.

If W has type A_n, then $n > 1$ and σ belongs to one of the two infinite families of rigid elements described above. The two families are conjugate, so we may as well suppose the path of σ runs counterclockwise, so $\sigma = s_{k-1} \cdots s_1 s_0$ for some k, the subscripts being interpreted mod $n+1$. Then $k - 1 < n$: For if $k - 1 = n$ then $\sigma \downarrow s_{k-1} s_{k-3} \cdots s_1 s_0$, so that σ covers a pair; hence for $k - 1 \geq n$, σ is not a chain. Therefore $k - 1 < n$, in which case the argument used above shows that X_σ is a closed parabolic orbit isomorphic to $\mathbb{P}^k$.

Corollary 7.6 *In the simply-laced case ADE (excluding A_1), every chain is a closed parabolic orbit isomorphic to $\mathbb{P}^k$.*

To describe the chains explicitly, it suffices to list the maximal chains. In type A_n for $n > 1$ there are two maximal chains namely $s_2 s_3 \cdots s_n s_0$ and $s_{n-1} \cdots s_1 s_0$ with coroot lattice representatives $(1, -1, \underline{0})$ and $(\underline{0}, -1, 1)$ respectively. The corresponding affine Schubert varieties are both isomorphic to $\mathbb{P}^n$. In types DE every rigid element is a chain. In type D_n there are three maximal chains, two $\mathbb{P}^n$'s and one $\mathbb{P}^3$. Finally there are two $\mathbb{P}^5$'s in E_6, a $\mathbb{P}^7$ and a $\mathbb{P}^5$ in E_7, and a $\mathbb{P}^8$ and a $\mathbb{P}^7$ in E_8. We leave it to the interested reader to write down the minimal length and coroot lattice representatives for these chains.

7.3 Non-simply-laced types

In this subsection, we classify the chains in each of the non-simply-laced types and identify which ones index smooth affine Schubert varieties.

A_1: Every element is a chain; indeed $\tilde{W}^S$ itself forms an infinite chain with cup sequence $a_k = k$. To see this, note that the coroot lattice representative for the element

of length k is $\lambda = (2j)$ if $k = 2j - 1$ and $\lambda = (-2j)$ if $k = 2j$. In the first case we have $\alpha(\lambda) = 2j$ and in the second $1 - \alpha(\lambda) = 2j + 1$, where in this case $\alpha = \alpha_1 = \alpha_0$ is the unique positive root. Hence our claim follows from the Chevalley formula. We also conclude that the only non-trivial smooth Schubert variety in type A_1 is the unique closed parabolic orbit, $X_{s_0} \approx \mathbb{P}^1$.

This discussion also recovers the well-known fact in type A_1 that $H^*\mathcal{L}_G$ is a divided power algebra, a fact normally deduced from the equivalence $\mathcal{L}_G \cong \Omega SU(2)$ and the Serre spectral sequence.

B_n: There are two maximal chains for $n \geq 3$: (1) $\sigma = s_1 s_2 s_0$ with coroot lattice representative $(-1, 0, 1, \underline{0})$; and (2) $\tau = s_0 s_2 s_3 \cdots s_n s_{n-1} \cdots s_3 s_2 s_0$, with coroot lattice representative $(2, \underline{0})$.

Note that both X_σ and X_τ are closed parabolic orbits Y_I, where $I = \{s_0, s_1, s_2\}$ and $I = \tilde{S} - \{s_1\}$ respectively. $X_\sigma \approx \mathbb{P}^3$ is of the simply laced type covered above. X_τ is a flag variety of type B_n/B_{n-1}; that is, a nonsingular quadric hypersurface of dimension $2n - 1$. A standard calculation shows that the cup sequence is $(1, 1, \ldots, 1, 2, 1, \ldots, 1, 1)$; here this follows at once from the Chevalley formula. Thus the complex k-skeleton of X_τ satisfies Poincaré duality if and only if $k \leq n - 1$, in which case it is just $\mathbb{P}^k$. We observe that firing down from τ yields the chains below τ, namely starting with

$\tau =$ 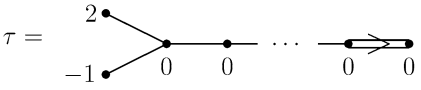

we get

$$(2, \underline{0}) \downarrow (2, -1, \underline{0}) \downarrow (1, 1, -1, \underline{0}) \downarrow (1, 0, 1, -1, \underline{0}) \downarrow \ldots.$$

There is one additional rigid element $\zeta = s_1 s_2 \cdots s_n \cdots s_2 s_0$. Note that after omitting the s_2 on the left we still have an element of $\tilde{W}^S$; this shows that ζ covers a pair and so is not a chain.

Remark: In type D_n, the closed parabolic orbit $Y = Y_{\tilde{S}-\{s_1\}}$ has type D_n/D_{n-1}, a nonsingular quadric hypersurface of dimension $2n - 2$. Its coroot lattice representative is $(2, \underline{0})$ as in type B_n. In this case, however, the middle homology has rank two, by a standard calculation or inspection of the Bruhat poset. Hence the Schubert subvarieties of Y of dimension k, $n - 1 < k < 2n - 2$, are not even palindromic. For future reference, we note that firing down from the top yields

$$(2, \underline{0}) \downarrow (2, -1, \underline{0}) \downarrow (1, 1, -1, \underline{0}) \downarrow \cdots, \tag{5}$$

exactly as in type B_n except that just above the middle dimension we reach the element $(1, \underline{0}, 1, -1, -1)$, which covers a pair.

C_n: Recall that there is an infinite family of rigid elements, obtained by running back and forth along $\tilde{D}$. The maximal chain in this family is $\sigma = s_1 s_2 \cdots s_n s_{n-1} \cdots s_1 s_0$ (note that $s_0\sigma$ covers a pair). There is one other maximal chain: $\tau = s_0 s_1 s_0$.

Note that X_τ is a closed parabolic orbit and is not $\mathbb{P}^3$ but rather the symplectic Grassmannian $Sp(2)/U(2)$. By the Chevalley formula its cup sequence is $(1, 2, 1)$ and hence its complex 2-skeleton is not smooth. The coroot lattice representative for τ is $(0, 1, \underline{0}) = \omega_2^\vee$.

The coroot lattice representative of σ is $\lambda = (-1, \underline{0}) = -\omega_1^\vee = -\alpha_0^\vee$, hence λ is anti-dominant and $\lambda = \sigma$ as elements of $\tilde{W}$. The cup sequence is $(1, 2, 2, \ldots, 2)$. Indeed, the complete list of coroot lattice representatives below σ is obtained as follows, starting from the top:

$$(-1, \underline{0}) \downarrow (1, -1, \underline{0}) \downarrow \ldots \downarrow (\underline{0}, 1, -1, 0) \overset{s_{n-1}}{\downarrow} (\underline{0}, 1, -2) \overset{s_n}{\downarrow} (\underline{0}, -1, 2)$$
$$\overset{s_{n-1}}{\downarrow} (\underline{0}, -1, 1, 0) \downarrow \ldots \downarrow (1, \underline{0}).$$

All but one of the factors of 2 in the cup sequence occurs because of a short node; the application of s_n also yields a factor of 2 because $\alpha_n(\underline{0}, -1, 2) = 2$. It follows that none of the complex skeleta are smooth, except for the $\mathbb{P}^1$ at the bottom.

An alternative way to identify the cup sequence of σ is to note that X_σ is the closure of the lowest non-trivial P-orbit, and it can be shown that it is therefore the Thom space of the line bundle over $\mathbb{P}^{2n-1}$ associated to the highest root [14]. This line bundle is just $(\gamma^*)^2$, where $\gamma^* \downarrow \mathbb{P}^{2n-1}$ is the hyperplane section bundle. Hence if u is the Thom class, by a general formula we have $u^2 = c_1(\gamma^{*2})u = 2yu$ where $y = c_1(\gamma^*)$, yielding the cup sequence above.

F_4: There are two maximal chains. The first is $s_0s_1s_2s_3s_2s_1s_0$, a closed parabolic orbit of type B_4/B_3 with coroot lattice representative $(0, 0, 0, 1)$. The coroot lattice representatives of its skeleta are given by

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

From our analysis of the type B case, we know that the cup sequence is $(1, 1, 1, 2, 1, 1, 1)$ and that only the complex k skeleta with $k \leq 3$ or $k = 7$ are Poincaré duality spaces.

The second maximal chain is $s_4s_3s_2s_1s_0$, which has coroot lattice representative $(0, 1, 0, -1)$ and cup sequence $(1, 1, 1, 2, 2)$. Hence it is not a Poincaré duality space.

Note that every rigid element is a chain.

G_2: There are two maximal chains. The first is $s_2s_1s_2s_1s_2s_0$, which has coroot lattice representative the anti-dominant class $(0, -1) = -\omega_2^\vee$. Its cup sequence is $(1, 1, 3, 2, 3, 1)$ and hence it is not a Poincaré duality space. To see this, we write down the coroot lattice representatives of its skeleta:

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

where we recall that firing the short node α_1 in type G_2 adds three times the short value to its neighbor. From the Chevalley formula we see that the short node firings produce cup product coefficients of 3, while the long node firing $(1, -2) \downarrow (-1, 2)$

yields a coefficient of 2 because of the -2 in the second position. We conclude that only the complex 1 and 2 skeleta are smooth; these are the two closed parabolic orbits $\mathbb{P}^1$, $\mathbb{P}^2$. It can be shown that $X_{-\omega_2^\vee}$ is the Thom space of a line bundle over the maximal flag variety of G_2 omitting the long node [14].

The second maximal chain is $s_0s_2s_1s_2s_0$, with coroot lattice representative $(1, 0)$. Its cup sequence is $(1, 1, 3, 2, 2)$, and hence it is not a Poincaré duality space.

There is one more rigid element $\zeta = s_0s_2s_1s_2s_1s_2s_0$. Note that omitting the s_2 on the left yields another element of $\tilde{W}^S$; hence ζ covers a pair and is not a chain.

7.4 Conclusions

The results of this section imply that Theorem 1.1 is true for all chains. More precisely:

Proposition 7.7 *Let X_λ be a chain. Then X_λ is smooth if and only if it is a closed parabolic orbit, in which case X_λ is either a projective space, a quadric hypersurface of type B_k/B_{k-1} for some k, or a symplectic Grassmannian of type C_2/A_1.*

8 Proof of the Smoothness Theorem 1.1

It is well known that a closed parabolic orbit is smooth and every smooth affine Schubert variety satisfies Poincaré duality over $\mathbb{Z}$. Therefore, to prove Theorem 1.1 it is only necessary to show that if X_λ satisfies Poincaré duality integrally then it is a closed parabolic orbit. Call λ *admissible* if

1. λ has a unique negative node $s \in \tilde{S}$;
2. s is a long node;
3. $\alpha_s(\lambda) = -1$ (or $1 - \alpha_0(\lambda) = -1$, if $s = s_0$).

Note that if X_λ satisfies Poincaré duality, then λ is an admissible palindromic by Proposition 6.3. In Section 7 we have classified all smooth chains. Hence to finish the main proof it suffices to prove the following key lemma:

Lemma 8.1 *Suppose λ is an admissible palindromic. Then either (i) X_λ is a closed parabolic orbit; or (ii) X_λ is a singular chain.*

We first dispose of the dominant and anti-dominant cases.

Lemma 8.2 *Suppose λ is admissible. Then*

(a) *if λ is dominant, then X_λ is a closed parabolic orbit;*
(b) *if λ is anti-dominant and palindromic, then G has type G_2 and $\lambda = -\omega_2^\vee$. Hence X_λ is a singular chain.*

Proof: (a) If λ is dominant and admissible, then $\alpha_0(\lambda) = 2$, and the assertion follows immediately from Proposition 4.4. (b) Suppose λ is antidominant, admissible, and palindromic. Then $\lambda = -\omega_s^\vee$, where s is a long node and $\omega_s^\vee \in Q^\vee$. In particular, s is

not minuscule. This rules out type A, since then all nodes are minuscule. Since in all other types λ cannot cover a fork, we conclude that s is a leaf node. This eliminates types BCD at once, since every leaf node is either short or minuscule. In type E every leaf node is either minuscule or forks too soon. In type F_4 the long leaf node s_1 forks too soon. This leaves type G_2 with $\lambda = -\omega_2^\vee$. This element is a chain, and is singular.

Recall from Section 4 that p_G is the number of non-trivial closed parabolic orbits, and that we have shown p_G is just the number of connected subdiagrams of the affine diagram containing s_0. Recall also that λ is *overweight* if it does not satisfy

$$\alpha_0(\lambda) \leq \begin{cases} 2 & \text{for all } \lambda \\ 1 & \text{if } \lambda \text{ is not dominant} \end{cases}$$

An admissible palindromic cannot be overweight.

Lemma 8.3 *If G is simply-laced, there are at most p_G admissible palindromics.*

Proof: Assume λ is admissible and palindromic. Let $s \in \tilde{D}$ denote the unique negative node of λ. By Lemma 8.2(b), we may assume that λ has at least one positive node. Note that "positive node" always refers to a node of D, while "negative node" refers to a node of $\tilde{D}$. The proof now proceeds type by type.

A_n: $p_G = \binom{n+1}{2}$. We may assume $n > 1$, since the case $n = 1$ was already settled in our study of chains. Note that λ can have at most two positive nodes (otherwise λ is overweight). Moreover if s_i, s_j are the positive nodes, with $i \leq j$, then $\alpha_i(\lambda) = 1 = \alpha_j(\lambda)$, where in the case $i = j$ this is to be interpreted as $\alpha_i(\lambda) = 2$. Now if s_k is the negative node, then $i + j - k = 0 \, mod \, (n + 1)$ (otherwise $\lambda \notin \mathcal{Q}^\vee$). Since there is a unique such k, and $k \neq i, j$, this shows that there are at most $\binom{n+1}{2}$ admissible palindromics.

D_n: $p_G = 2n$. There are at most three positive nodes (otherwise λ is overweight). If there are exactly three, then they are all minuscule and take the value 1 on λ (otherwise λ is overweight). Furthermore $s = s_{n-2}$ (otherwise λ covers a fork). Hence $\lambda = \omega_1^\vee + \omega_{n-1}^\vee + \omega_n^\vee - \omega_{n-2}^\vee$.

If there are two positive nodes t, u then at least one of them, say t, is minuscule (otherwise λ is overweight). If u is also minuscule, then using the characterization of the coroot lattice elements, one can check that λ has one of the following forms:

$$\lambda = \begin{cases} \omega_1^\vee \pm (\omega_{n-1}^\vee - \omega_n^\vee) & \\ \omega_{n-1}^\vee + \omega_n^\vee - \omega_1^\vee & (n \text{ even}) \\ \omega_{n-1}^\vee + \omega_n^\vee & (n \text{ odd}) \end{cases}$$

If u is not minuscule, then $s \neq s_0$, s is not minuscule, and $\alpha_t(\lambda) = 1 = \alpha_u(\lambda)$ (otherwise λ is overweight). Furthermore, $t = s_1$ (otherwise $\lambda \notin \mathcal{Q}^\vee$) and s, u are adjacent (otherwise λ covers a fork). Hence $\lambda = \omega_1^\vee \pm (w_k^\vee - w_{k-1}^\vee)$, where $2 < k < n - 1$. But $\omega_1^\vee - w_k^\vee + w_{k-1}^\vee$ is a non-palindromic skeleton of the quadric in (5), so we must have $\lambda = \omega_1^\vee + w_k^\vee - w_{k-1}^\vee$.

Hence there are at most $n - 1$ admissible palindromics with two positive nodes.

Suppose λ has exactly one positive node i. Then we claim

$$\lambda = \begin{cases} \omega_i^\vee & \text{if } i \text{ even}, i \neq n-1, n \\ \omega_i^\vee - \omega_1^\vee & \text{if } i \text{ odd}, i \neq 1, n-1, n \\ 2\omega_1^\vee & \text{if } i = 1 \\ 2\omega_i^\vee & \text{if } n \text{ even}, i = n-1, n \\ 2\omega_i^\vee - \omega_1^\vee & \text{if } n \text{ odd}, i = n-1, n \end{cases}$$

To prove one case of the claim, suppose $i = 1$. Then the negative node s can only be s_0 or s_2 (otherwise either λ covers a fork or $\lambda \notin \mathcal{Q}^\vee$). This forces $\alpha_1(\lambda) = 2$ (otherwise either $\lambda \notin \mathcal{Q}^\vee$ or λ is overweight). But if $s = s_2$ then $\lambda = (2, -1, \underline{0})$, a non-palindromic skeleton of the quadric. Hence $s = s_0$ and $\lambda = 2\omega_1^\vee$.

To prove another case, suppose n odd and $i = n$. Then the negative node s must be s_0, s_1 or s_{n-1} (otherwise λ covers a fork). If $s = s_0$ then $\alpha_n(\lambda) = 0 \, mod \, 4$ (otherwise $\lambda \notin \mathcal{Q}^\vee$), and hence λ is overweight. If $s = s_{n-1}$ then $\lambda = (\underline{0}, -1, a)$ for some $a > 0$. Then $a + 1 = 0 \, mod \, 4$ (otherwise $\lambda \notin \mathcal{Q}^\vee$), and hence λ is overweight. Hence $j = 1$b and $\alpha_n(\lambda)$ is even (otherwise $\lambda \notin \mathcal{Q}^\vee$), and then $\alpha_n(\lambda) = 2$ (otherwise λ is overweight). Hence $\lambda = 2\omega_n^\vee - \omega_1^\vee$.

The remaining cases are left to the reader. Thus there are at most $1+n-1+n = 2n$ admissible palindromics, as desired.

E_n: $p_G = 10$. For each $n = 6, 7, 8$, we will show that there are at most n admissible palindromics with one positive node, at most $10-n$ with two positive nodes, and none with more than two positive nodes.

E_6: There are at most three positive nodes (otherwise λ is overweight). If there are exactly three, then the negative node must be s_4 and only one of the adjacent nodes is occupied (otherwise λ is overweight). But then λ covers a fork, a contradiction. So there are at most two positive nodes.

Suppose there are exactly two. Then there are four possibilities: $(1, 0, 0, 0, 0, 1)$, $(0, 0, 1, -1, 1, 0)$, $(0, 0, 1, 0, -1, 1)$ and $(1, 0, -1, 0, 0, 1)$. For example, suppose that s_4 is the negative node. Then the positive nodes must be adjacent to it (otherwise λ covers a fork), and hence must each take the value 1 (otherwise λ is overweight). But this forces s_3, s_5 as the positive nodes (otherwise $\lambda \notin \mathcal{Q}^\vee$). Hence $\lambda = (0, 0, 1, -1, 1, 0, 0)$.

If there is just one positive node t, then for each choice of t there is only one possibility for λ. For example, suppose the positive node is s_6. Then the negative node must be s_1: For if s_0 is negative, then $\alpha_6(\lambda) = 0 \, mod \, 3$ (otherwise $\lambda \notin \mathcal{Q}^\vee$); but then λ is overweight. The other four choices of negative node all fork too soon. It follows that $\alpha_6(\lambda)$ is even (otherwise $\lambda \notin \mathcal{Q}^\vee$), and hence $\alpha_6(\lambda) = 2$ (otherwise λ is overweight). Hence $\lambda = (-1, 0, 0, 0, 0, 2)$.

E_7: There are at most three positive nodes (otherwise λ is overweight). If there are exactly three, then the negative node must be s_4 and the adjacent nodes s_3, s_5 must be zero (otherwise λ is overweight). But then λ covers a fork, a contradiction. So there are at most two positive nodes.

If there are exactly two, we find that there are three possibilities: $(0, 1, 0, -1, 1, 0, 0)$, $(0, 1, 0, 0, -1, 0)$,$(0, 1, 0, 0, 0, -1, 1)$.

If there is just one positive node t, then for each choice of t there is only one possibility for λ. For example, suppose s_4 is positive. Then $\alpha_4(\lambda) = 1$ and the negative node must be s_3 or s_5 (otherwise λ is overweight). But s_5 can't occur, since then $\lambda \notin \mathcal{Q}^\vee$. Hence $\lambda = \omega_4^\vee - \omega_3^\vee$.

E_8: There are at most three positive nodes (otherwise λ is overweight). If there are exactly three, then λ is still overweight unless $\lambda = \omega_1^\vee + \omega_2^\vee + \omega_8^\vee - \omega_4^\vee$. But then λ covers a fork, a contradiction. So there are at most two positive nodes.

If there are exactly two, then λ must be either $(1, 1, -1, \underline{0})$ or $(0, 1, 1, -1, \underline{0})$. If there is only one positive node t, then for each such t there is at most one possibility for λ. In all cases one finds that the alternatives fork too soon or are overweight; details are left to the reader.

We can now prove Lemma 8.1. In the simply laced case, we know that every closed parabolic orbit satisfies Poincaré duality and hence is admissible and palindromic. Lemma 8.3 implies the converse holds also so Lemma 8.1 holds.

We now turn to the non-simply laced types.

B_n: $p_G = 2n - 2$. We will show that (1) there are at most $2n - 2$ admissible palindromics that satisfies Poincaré duality, and (2) all other admissible palindromics are singular skeletons of the quadric $X_{(2,\underline{0})}$ (and in particular, are chains).

There are at most two positive nodes (otherwise λ is overweight). If there are exactly two, then one of them is s_1 (otherwise λ is overweight). If the other is s_2 then $\lambda = (1, 1, -1)$ (otherwise λ covers a fork). This element is a singular skeleton of the quadric. So suppose that the two positive nodes are s_1, s_j, where $j > 2$. Then we claim $\lambda = \omega_1^\vee + \omega_j^\vee - \omega_{j-1}^\vee$. To prove the claim, let $s = s_i$. If $i \neq 2$ then $j = i \pm 1$ (otherwise λ covers a fork). Then there are two palindromic solutions: $\lambda = (1, \underline{0}, -1, 1, \underline{0})$ and $\lambda = (1, \underline{0}, 1, -1, \underline{0})$. In the second case λ is a singular skeleton of the quadric, hence a singular chain. If $i = 2$ then $j = 3$ and $\lambda = (1, -1, 1, \underline{0})$ (otherwise λ covers a fork, since s_0 is necessarily a zero node). Hence there are at most $n - 2$ admissible palindromics that satisfy Poincaré duality.

Suppose there is one positive node s_j. Then we claim that

$$\lambda = \begin{cases} \omega_j^\vee & \text{if } j \text{ even} \\ \omega_j^\vee - \omega_1^\vee & \text{if } j \text{ odd, } j > 1 \\ 2\omega_1^\vee & \text{if } j = 1 \end{cases}$$

If j is even then i is also even (otherwise $\lambda \notin \mathcal{Q}^\vee$). Since $i \neq n$, this forces $i = 0$ (otherwise λ covers a fork). Hence $\lambda = \omega_j^\vee$.

If j is odd and $j > 1$, then $\alpha_j(\lambda) = 1$ (otherwise λ is overweight). Hence i is odd (otherwise λ is overweight). Since $i \neq n$, this forces $i = 1$ (otherwise λ covers a fork).

Now suppose $j = 1$. If $\alpha_1(\lambda)$ is odd, then i is odd (otherwise $\lambda \notin \mathcal{Q}^\vee$), in which case λ covers a fork. So i is even. Hence $\alpha_j(\lambda)$ is even (otherwise $\lambda \notin \mathcal{Q}^\vee$), forcing $\alpha_j(\lambda) = 2$ and $i = 0, 1$. If $i = 1$ then $\lambda = (2, -1, \underline{0})$, a singular skeleton of the quadric. Hence $i = 0$ and $\lambda = 2\omega_1^\vee$. This completes the proof of our claim.

C_n: $p_G = n$. We must have $s = s_i$ for $i = 0, n$. But if $i = n$ then $\lambda \notin \mathcal{Q}^\vee$. Hence $i = 0$ and λ is dominant. This forces $\lambda = \omega_i^\vee$ if $i < n$, and $\lambda = 2\omega_n^\vee$ if $i = n$, as shown earlier.

F_4: $p_G = 4$. If λ has more than one positive node, then $\lambda = (1, -1, 0, 1)$, a singular chain.

Now suppose there is one positive node s_j. For each of $j = 2, 3$, it is easy to check that there is only one corresponding admissible palindromic, namely $(-1, 1, 0, 0)$ and $(0, -1, 1, 0)$. For $j = 1, 4$ there are two admissible palindromics: $(1, 0, 0, 0)$ and $(1, -1, 0, 0)$ for $j = 1$, and $(0, 0, 0, 1)$ and $(-1, 0, 0, 1)$ for $j = 4$. The dominant classes are singular chains.

G_2: $p_G = 2$. Suppose λ is an admissible palindromic (and is not anti-dominant). Then λ has exactly one positive node s_j (otherwise λ is overweight), . $j = 1, 2$. If $j = 2$ there are three admissible palindromics: $(-1, 1)$, $(-1, 2)$ and $(0, 1)$. The first two are singular chains. If $j = 1$ there are two: $(1, 0)$ and $(1, -1)$. The first is a singular chain.

This completes the proof of Lemma 8.1 and Theorem 1.1.

Remark: Let I be a connected subgraph of $\tilde{D}$ containing s_0. Then $S^+(\lambda) = \mathcal{N}(I)$ by Lemma 4.3. From this point of view, the proof of Theorem 1.1 amounts to showing that (i) If X_λ satisfies Poincaré duality, then $S^+(\lambda) = \mathcal{N}(I)$ for some (unique) I, and (ii) for each I there is a unique λ such that $S^+(\lambda) = \mathcal{N}(I)$ and X_λ satisfies Poincaré duality.

9 The palindromy game II: Bruhat order and the coroot lattice

In the characterization of the smooth Schubert varieties, it turned out (somewhat surprisingly) that we only needed the weak order. For the palindromy theorem, however, we will need more general Bruhat descents of the form $\lambda \downarrow r\lambda$, where r is an affine reflection associated to a non-simple root. As it happens, we will only need two kinds of such reflections: The linear reflection $s_\beta \in W$ associated to a positive root β, and the affine reflection $r_\beta = r_{1,\beta}$ associated to the affine root $(1, \beta)$, where again β is a positive root. In the spirit of the palindromy game, we will often refer to such descents as "moves". Whenever possible we describe these moves $\lambda \downarrow s_\beta\lambda$, $\lambda \downarrow r_\beta\lambda$ in terms of the Dynkin diagram $\tilde{D}$.

Note that if r_1 and r_2 are any two distinct reflections (linear or affine), then by elementary geometry we have that $r_1\lambda = r_2\lambda \Rightarrow r_1\lambda = \lambda = r_2\lambda$. Hence if $\lambda \downarrow r_1\lambda$ and $\lambda \downarrow r_2\lambda$, it follows that λ covers a pair.

9.1 β-positive and β-negative pairs

Let β be a positive root and let $\lambda \in Q^\vee$. We want to determine when the reflections s_β and r_β lower the S-length of λ by 1. Note that $r_\beta\lambda = s_\beta\lambda + \beta^\vee$. For all $\alpha \in \Phi^+$ we have the formulas

$$\alpha(s_\beta\lambda) = \alpha(\lambda) - \alpha(\beta^\vee)\beta(\lambda)$$

and

$$\alpha(r_\beta\lambda) = \alpha(\lambda) + \alpha(\beta^\vee)(1 - \beta(\lambda)).$$

In order to evaluate the change in S-length after reflection, it is convenient to partition the positive roots into the following sets:

1. β itself;
2. The β-null roots; i.e., $\{\alpha : \alpha(\beta^\vee) = 0\}$. Thus $s_\beta\alpha = \alpha$.
3. β-positive pairs α, α': These are characterized by $s_\beta\alpha = \alpha'$.
4. β-negative pairs α, α': These are characterized by $s_\beta\alpha = -\alpha'$.

The apparent symmetry of α and α' is misleading. In the β-positive case, we will always take α to have $\alpha(\beta^\vee) < 0$. Then $\alpha'(\beta^\vee) > 0$, and $\alpha + k\beta = \alpha'$, where $k = -\alpha(\beta^\vee) = 1, 2$, or 3.

In the β-negative case $\alpha'(\beta^\vee) = \alpha(\beta^\vee)$, and $\alpha + \alpha' = k\beta$. When $k = 1$ there is no way to distinguish α, α', but for $k = 2, 3$ we can and will always choose α so that $\beta - \alpha$ is a positive root. Then $\beta - \alpha'$ is a negative root.

9.2 Linear reflections

Let λ be an element of the coroot lattice, and fix a positive root β. Note that $s_\beta\lambda = \lambda \Leftrightarrow \beta(\lambda) = 0$. We say that the *opposite sign condition* is satisfied on (β, λ) if for every β-negative pair α, α' the values $\alpha(\lambda), \alpha'(\lambda)$ have opposite sign. In particular, both are nonzero.

Proposition 9.1 *Let $\beta \in \Phi^+$ and $\lambda \in Q^\vee$.*

(a) $\ell^S(s_\beta\lambda) < \ell^S(\lambda) \Leftrightarrow \beta(\lambda) < 0$.
(b) $\lambda \downarrow s_\beta\lambda \Leftrightarrow \beta(\lambda) < 0$ *and the opposite sign condition is satisfied.*

Proof: Note that application of a linear reflection s_β to λ can only affect the q term in the S-length formula (1). Furthermore the β-null positive roots and the β-positive pairs contribute zero to the change in the q term, henceforth denoted as Δq. Since we always have $\beta(s_\beta\lambda) = -\beta(\lambda)$, the change in length will be determined by what happens on the β-negative pairs α, α'.

Now suppose $\beta(\lambda) < 0$ and α, α' are a β-negative pair. Since $\alpha + \alpha'$ is a positive multiple of β, the values $\alpha(\lambda), \alpha'(\lambda)$ either have opposite sign, are both negative, or a negative and a zero. The pairs with opposite sign contribute zero to Δq, while in the other two cases we have respectively $\Delta q = 2$, $\Delta q = 1$. This proves $\Leftarrow$ in (a), and also (b).

If $\beta(\lambda) > 0$, then $\ell^S(s_\beta\lambda) > \ell^S(\lambda)$ (substitute $s_\beta\lambda$ for λ and apply the previous case). This yields $\Rightarrow$ in (a), completing the proof of the theorem.

The following application will be particularly useful. Given a fixed λ and nodes $s, t \in D$ with opposite sign, let I denote the unique minimal path between them, regarded as a subgraph of D. If all interior nodes of I vanish on λ, we say that s and t are *linked* by I.

Lemma 9.2 (*Linear ABC-moves*) *Suppose that nodes s, t have opposite sign, with s the negative node, and they are linked by a subgraph I of type A, B, or C. Then if any one of the following conditions holds, λ covers a pair.*

(a) *I has type A, B or C, and $\alpha_s(\lambda)+\alpha_t(\lambda)<0$. Furthermore, if I has type B or C, the positive node t is required to be the minuscule node of I.*
(b) *I has type BC, the positive node t is the minuscule node of I, $\alpha_s(\lambda)+\alpha_t(\lambda)>0$, and $2\alpha_s(\lambda)+\alpha_t(\lambda)<0$.*
(c) *I has type BC, the negative node s is the minuscule node of I, $\alpha_s(\lambda)+\alpha_t(\lambda)<0$ and $\alpha_s(\lambda)+2\alpha_t(\lambda)\neq 0$.*

Proof: We will find a non-simple $\beta\in\Phi_I^+$ such that (i) $\lambda\downarrow s_\beta\lambda$. It then follows that λ covers a pair. Since the β-negative pairs all lie in Φ_I^+, we may as well assume $I=D$; i.e., that Φ itself has type A_n, B_n, C_n and $\{s,t\}=\{s_1,s_n\}$. We will prove part (a) in detail and sketch the rest.

Case (a): Let $\beta=\alpha_1+\alpha_2+\ldots+\alpha_n$. Then by assumption $\beta(\lambda)<0$, and we claim that the opposite sign condition is satisfied. Let α,α' be a β-negative pair. If β is long, or β is short and the pair is also short, then $\alpha+\alpha'=\beta$ and it is clear that one of the two contains α_1 and the other contains α_n, proving our claim. In particular, this settles type A.

In type B_n we have $\beta=e_1$, which is a short root, and there are long β-negative pairs e_1-e_i, e_1+e_i. In that case we have $\alpha+\alpha'=2\beta$. Thus each root of the pair contains α_1. On the other hand $\alpha'=\beta+(\beta-\alpha)$, where $\beta-\alpha$ is a positive root. It follows that α' contains α_n twice, and since we are assuming $\alpha_n(\lambda)<0$, the opposite sign condition is satisfied as required.

In type C_n we have $\beta=e_1+e_n$, which is again short. There is one long β-negative pair $2e_1, 2e_n$, and again we have $\alpha+\alpha'=2\beta$. Here $\alpha'=2e_1=\alpha_0$, and so contains α_1 twice. The opposite sign condition follows as before.

Case (b): Take as β the smallest root containing $\alpha_s(\lambda)$ twice. In type B_n, $\beta=\alpha_1+\ldots\alpha_{n-1}+2\alpha_n=e_1+e_n$. In type C_n, $\beta=\alpha_0$. In each case it is easy to check the opposite sign condition.

Case (c): If $\alpha_s(\lambda)+2\alpha_t(\lambda)>0$, use the β of part (a). If $\alpha_s(\lambda)+2\alpha_t(\lambda)<0$, use the β of part (b).

9.3 Affine reflections

In this section we give necessary and sufficient conditions for $\lambda\downarrow r_\beta\lambda$ and $\lambda\uparrow r_\beta\lambda$. Note that $r_\beta\lambda=\lambda\Leftrightarrow 1-\beta(\lambda)=0$.

9.3.1 The positive and negative pair conditions

Suppose $\lambda\in Q^\vee$. We say that $\beta\in\Phi^+$ satisfies the *positive pair condition* if either:

(i) $1-\beta(\lambda)<0$ and whenever α,α' is a β-positive pair, either $\alpha(\lambda)>0$ or $\alpha'(\lambda)\leq 0$; or
(i)* $1-\beta(\lambda)>0$ and whenever α,α' is a β-positive pair, either $\alpha(\lambda)<0$ or $\alpha'(\lambda)>1$.

We say that $\beta \in \Phi^+$ satisfies the *negative pair condition* if either:

(ii) $1 - \beta(\lambda) < 0$ and either β is long, or β is short and whenever α, α' is a long β-negative pair, either $\alpha(\lambda) \leq 0$ or $\alpha'(\lambda) \leq 0$; or

(ii)* $1 - \beta(\lambda) > 0$ and either β is long, or β is short and whenever α, α' is a long β-negative pair, either $\alpha(\lambda) \geq 2$ or $\alpha'(\lambda) \geq 2$.

If we wish to refer only to a specific β-positive or β-negative pair we say that α, α' satisfies the positive pair condition or negative pair condition. In fact we will be concerned almost exclusively with the case $1 - \beta(\lambda) < 0$; the case $1 - \beta(\lambda) > 0$ is included for completeness.

Proposition 9.3 *Let $\beta \in \Phi^+$ and $\lambda \in Q^\vee$.*

a) $\ell^S(r_\beta(\lambda)) < \ell^S(\lambda) \Leftrightarrow 1 - \beta(\lambda) < 0$.
b) *Suppose $1 - \beta(\lambda) < 0$. Then $\lambda \downarrow r_\beta\lambda \Leftrightarrow$ the positive pair condition and the negative pair condition are satisfied.*
c) *The analogous statements hold for $1 - \beta(\lambda) > 0$.*

If there exists a β such that $\lambda \downarrow r_\beta\lambda$, we say λ *has an affine move*. In particular, if $\beta \neq \alpha_0$, then λ cannot be palindromic.

Proof: For $m \in \mathbb{Z}$ let

$$f(m) = \begin{cases} -m & \text{if } m \leq 0 \\ m - 1 & \text{if } m > 0 \end{cases}$$

Then

$$\ell^S(\lambda) = \sum_{\alpha \in \Phi^+} f(\alpha(\lambda)).$$

We will analyze the effect of r_β on ℓ^S by considering the four types of roots listed above separately. For any subset $\Gamma \subset \Phi^+$, we write Δ_Γ for the contribution of Γ to the change in ℓ^S; more precisely:

$$\sum_{\alpha \in \Gamma} f(\alpha(r_\beta\lambda)) = \sum_{\alpha \in \Gamma} f(\alpha(\lambda)) + \Delta_\Gamma.$$

First observe that

$$\beta(r_\beta\lambda) = 2 - \beta(\lambda),$$

while

$$f(2 - m) = \begin{cases} f(m) - 1 & \text{if } 1 - m < 0 \\ f(m) & \text{if } 1 - m = 0 \\ f(m) + 1 & \text{if } 1 - m > 0 \end{cases}$$

Hence

$$\Delta_{\{\beta\}} = \begin{cases} -1 & \text{if } 1 - \beta(\lambda) < 0 \\ 1 & \text{if } 1 - \beta(\lambda) > 0 \end{cases}$$

Second, if Γ is the set of positive roots α such that $s_\beta(\alpha) = \alpha$, then $\Delta_\Gamma = 0$. Therefore, the proposition then follows from the lemma below:

Lemma 9.4 *If α, α' is a β-positive (β-negative) pair, then*

$$1 - \beta(\lambda) < 0 \Rightarrow \Delta_{\{\alpha,\alpha'\}} \leq 0$$
$$1 - \beta(\lambda) > 0 \Rightarrow \Delta_{\{\alpha,\alpha'\}} \geq 0.$$

In each case $\Delta_{\{\alpha,\alpha'\}} = 0 \Leftrightarrow$ the positive (negative) pair condition holds for α, α'.

Proof: We have

$$\alpha(r_\beta \lambda) = \alpha'(\lambda) + \alpha(\beta^\vee)$$

$$\alpha'(r_\beta \lambda) = \alpha(\lambda) - \alpha(\beta^\vee).$$

Set $a = \alpha(\lambda)$, $a' = \alpha'(\lambda)$, and $k = -\alpha(\beta^\vee) = 1, 2, 3$. Thus we need only compute the effect of the transformation $(a, a') \mapsto (a - k, a' + k)$ on $f(a) + f(a')$. Taking $k = 1$ for simplicity, we find that

$\Delta_{\{\alpha,\alpha'\}} = 0$ if $a > 0$ and $a' > 1$, or $a < 0$ and $a' \leq 0$, or $a = 0$ and $a' = 1$;
$\Delta_{\{\alpha,\alpha'\}} < 0$ if $a = 0$ and $a' > 1$, or $a < 0$ and $a' \geq 1$;
$\Delta_{\{\alpha,\alpha'\}} > 0$ if $a > 0$ and $a' \leq 1$, or $a = 0$ and $a' < 0$.

If α, α' are a β-positive pair, the lemma now follows easily on inspection, making use of the fact that $\beta + \alpha = \alpha'$. Note, for example, that if $1 - \beta(\lambda) < 0$ and $a > 0$ then automatically $a' > 1$ and hence $\Delta_{\{\alpha,\alpha'\}} = 0$. If $k > 1$ a similar argument applies, making use of the fact that $k\beta + \alpha = \alpha'$.

If α, α' is a β-negative pair, then $\alpha(\beta^\vee) = \alpha'(\beta^\vee)$ and

$$\alpha(r_\beta \lambda) = \alpha(\beta^\vee) - \alpha'(\lambda)$$

$$\alpha'(r_\beta \lambda) = \alpha(\beta^\vee) - \alpha(\lambda).$$

Let $k = \alpha(\beta^\vee) = \alpha'(\beta^\vee) = 1, 2, 3$. Then we need only compute the effect of the transformation $(a, a') \mapsto (k - a, k - a')$ on $f(a) + f(a')$. Since $f(1 - a) = f(a)$ for all a, when $k = 1$ we find $\Delta_{\{\alpha,\alpha'\}} = 0$. If $k > 1$ and $1 - \beta(\lambda) < 0$ we find

$$\Delta_{\{\alpha,\alpha'\}} \begin{cases} = 0 & \text{if } \alpha(\lambda) \leq 0 \text{ or } \alpha'(\lambda) \leq 0 \\ < 0 & \text{otherwise} \end{cases}$$

Note that if, say, $\alpha(\lambda) \leq 0$, then $\alpha'(\lambda) \geq k\beta(\lambda) \geq 2k$. The case $k > 1$ and $1 - \beta(\lambda) > 0$ is similar. Since $k > 1 \Leftrightarrow \beta$ is short and α, α' are long, this completes the proof of the lemma.

9.3.2 Graph-splitting and affine ABC-moves

Two types of affine moves will be particularly useful. We continue to fix $\lambda \in \mathcal{Q}^\vee$.

Lemma 9.5 (*Graph-splitting moves*) *Let I be a proper connected subgraph of D, and let α_I denote the highest root of Φ_I. If $\alpha_I(\lambda) \geq 2$ and α_I satisfies the positive pair condition, then $\lambda \downarrow \alpha_I(\lambda)$ and λ covers a pair.*

Proof: Note that α_I is long in the sub-root system Φ_I, whether it is long in Φ or not. Since the negative pair condition involves only roots in Φ_I, it will automatically be satisfied in this case. If the positive pair condition also holds, then $\lambda \downarrow r_{\alpha_I}\lambda$ by Proposition 9.3. Since there is also a left descent $\lambda \downarrow s\lambda$ for some $s \in \tilde{S}$, it follows that λ covers a pair. (Note that $\alpha_I \neq \alpha_0$, since I is a proper subgraph.)

We call descents $\lambda \downarrow r_{\alpha_I}\lambda$ as above *graph-splitting moves*, since we are splitting the Dynkin graph D into I and the components of its complement.

Lemma 9.6 (*Affine A-moves*) *Suppose that s_0 is linked to a nonzero node t by a proper type A subgraph I, and that $1 - \alpha_0(\lambda)$, $\alpha_t(\lambda)$ have opposite sign. If $1 - \alpha_0(\lambda) + \alpha_t(\lambda) < 0$, then λ covers a pair.*

Proof: Let $\beta = \alpha_0 - \sum_{s\in I: s\neq s_0} \alpha_s$. Then $\beta(\lambda) > 1$ and β is a long root. Hence the negative pair condition holds. Moreover, the positive pair condition is also automatically satisfied: For suppose α, α' is a β-positive pair, so that $\beta + \alpha = \alpha'$. Then $\alpha(\lambda) = \alpha_t(\lambda)$ and $\alpha'(\lambda) = \alpha_0(\lambda)$. Hence either $\alpha(\lambda) > 0$ or $\alpha'(\lambda) \leq 0$. Then $\lambda \downarrow r_\beta\lambda$ and λ covers a pair.

We call the descent $\lambda \downarrow r_\beta\lambda$ constructed above an *affine A-move*. There are similar but less productive moves in the B,C cases; the following lemma will suffice for our purposes:

Lemma 9.7 (*Affine BC-moves*) *Suppose s_0 is a negative node of λ, linked to a positive node t by a subgraph I of type BC. If $1 - \alpha_0(\lambda) + \alpha_t(\lambda) < 0$ and $1 - \alpha_0(\lambda) + 2\alpha_t(\lambda) \neq 0$, then λ covers a pair.*

Proof: We may assume λ is dominant. There are three cases:

(1) Φ has type B_n and $t = s_n$;
(2) Φ has type C_n and $t = s_i$, with $i < n$;
(3) Φ has type F_4 and $t = s_3$.

Suppose $1 - \alpha_0(\lambda) + 2\alpha_t(\lambda) > 0$ (this rules out Case 3). Then λ has no other positive nodes. Let $\beta = \alpha_0 - \sum_{s\in I: s\neq s_0} \alpha_s$, as in the previous lemma. Then $\beta(\lambda) > 1$, and the positive pair condition is satisfied because β is the highest root containing α_t once. We also have (i) β is a short root; and (ii) if α, α' is a long β-negative pair, then α does not contain α_t (compare the proof of Lemma 9.2), and hence $\alpha(\lambda) = 0$. Hence the negative pair condition is satisfied, $\lambda \downarrow r_\beta\lambda$, and λ covers a pair.

Now suppose $1 - \alpha_0(\lambda) + 2\alpha_t(\lambda) < 0$. Let $J = \Phi - \{s_n\}$ in Case 1, $J = \{s_{i+1}, \ldots s_n\}$ in Case 2. Then $\alpha_J(\lambda) = \alpha_0(\lambda) - 2\alpha_t(\lambda) \geq 2$ (note that $m_s(\alpha_J) = m_s(\alpha_0)$ for all $s \in J$). Since the unique neighbor node of J in S is positive, λ covers a pair by Lemma 9.5.

In Case 3 let β be the highest root containing α_3 twice. Then $\beta(\lambda) \geq 2$ and the positive pair condition is satisfied. Moreover $\beta = 1222 = e_1 - e_4$ is a long root,[1] so the negative pair condition is satisfied. Hence $\lambda \downarrow r_\beta \lambda$ and λ covers a pair.

10 Proof of the Palindromy Theorem 1.4

The proof considers the anti-dominant, dominant, and "mixed" (i.e., neither dominant nor anti-dominant) cases separately, proceeding by a process of elimination based on the Palindromy Game II. The reader should keep at hand the list of chains (§7). The spiral varieties were shown to be palindromic in [16]; see §12. The exceptional case $\lambda = (3, 0, -1)$ in type B_3 was shown to be palindromic in Corollary 4.8, therefore, it remains to prove the "only if" part of the theorem. We will make frequent use of the graph-splitting moves and linear/affine ABC-moves introduced in §9.

10.1 Anti-dominant case

Theorem 10.1 *If λ is nonzero and anti-dominant, then X_λ is palindromic in precisely the following cases*:

(i) *Φ has type A_n and λ is a spiral class of the form $-k(n+1)\omega_i^\vee$ for $i = 1, n$ and $k \geq 1$;*
(ii) *Φ has type C_n or G_2 and $\lambda = -\alpha_0^\vee$. In these cases λ is a chain.*

Proof: Suppose λ is palindromic. There is a unique $s \in S$ such that $\alpha_s(\lambda) < 0$, since otherwise λ covers a pair. Thus $\lambda = -m\omega_s^\vee$ for some $s \in S$ and $m > 0$. If s is not a leaf node, firing it shows that λ covers a fork. This contradicts palindromy except in type A. In type A we have $\lambda = (\underline{0}, -m, \underline{0})$ with $-m$ in the i-th position, where $1 < i < n$ and $m \geq 2$ (otherwise $\lambda \notin \mathcal{Q}^\vee$). Then $\lambda \downarrow \mu = (\underline{0}, -m, m, -m, \underline{0})$. Then there is a graph-splitting move $\mu \downarrow r_{\alpha_i}\mu$. Hence λ covers a trident, a contradiction. Thus s is a leaf node.

In type A we then have $s = s_1, s_n$ and m divisible by $n+1$, so that λ is spiral as claimed. For the remainder of the proof we assume G is not of type A.

We claim $m = 1$; in other words, $\lambda = -\omega_s^\vee$ for some leaf node s. In particular, s is not minuscule. To prove the claim, suppose $m > 1$, and let $a \in S$ denote the unique node adjacent to s. If s is a long node, then firing it yields $\mu = s\lambda = m\omega_s^\vee - m\omega_a^\vee$. Then there is a graph-splitting move $\mu \downarrow r_{\alpha_s}\mu$ and hence λ covers a fork. To see this we need only check the positive pair condition. But if $\alpha_s + \alpha = \alpha'$, then since s is a leaf node we have $m_s(\alpha') < m_a(\alpha')$, and hence $\alpha'(\mu) \leq 0$. A similar argument works if s is short (note this only happens in type B_n with $s = s_n$, and in type G_2 with $s = s_1$). This proves the claim.

The theorem now follows immediately in types $CDEF$: In type DEF every leaf node is either minuscule or forks too soon, while in type C we can only have $s = s_1$. It remains to consider type B and G_2.

[1] Here we are following Bourbaki notation so $1222 = \alpha_1 + 2\alpha_2 + 2\alpha_3 + 2\alpha_4$.

In type B_n there is one non-minuscule extreme node $s = s_n$. Note that n is necessarily even, since otherwise $\omega_n^\vee$ is not in the coroot lattice. In particular $n \geq 4$. Then

$$-\omega_n^\vee = (\underline{0}, -1) \downarrow (\underline{0}, -2, 1) \downarrow (\underline{0}, -2, 2, -1)$$

so that now $(\underline{0}, -2, 2, -1)$ covers the pair $\mu = (\underline{0}, -2, 2, 0, -1)$ and $\eta = (\underline{0}, -2, 0, 1)$. Furthermore μ in turn covers the pair $\mu_1 = (\underline{0}, -2, 2, 0, 0, -1)$ (or $(2, 0, 0, -1)$ if $n = 4$) and $\mu_2 = (\underline{0}, -2, 2, -2, 1)$. Now let $\beta = \alpha_{n-1} + 2\alpha_n$. Then β is a long root with $\beta(\eta) = 2$. The positive pair condition for is immediately verified, so $\eta \downarrow r_\beta \eta$. Furthermore, η differs from μ_2 by a simple reflection. It follows that λ covers a scepter, a contradiction.

In type G_2, we have $-\omega_2^\vee = -\alpha_0^\vee$, which is a chain. On the other hand, $-\omega_1^\vee$ forks too soon (see the Hasse diagram in §13).

10.2 Dominant case

Throughout this section we assume $\lambda \in \mathcal{Q}^\vee$ is nontrivial and dominant. In particular, $\lambda \downarrow s_0\lambda$.

Theorem 10.2 *Suppose λ is dominant and nonzero. Then X_λ is palindromic in precisely the following cases*:

a) $\alpha_0(\lambda) = 2$, *in which case X_λ is a closed parabolic orbit.*
b) Φ *has type A_n and λ is a spiral class of the form $k(n+1)\omega_i^\vee$ for $i = 1, n$ and $k \geq 1$.*
c) Φ *has type G_2 and $\lambda = \omega_1^\vee$. In this case λ is a chain.*

Proof: The theorem is trivial in type A_1, so from now on we exclude that case. If $\alpha_0(\lambda) = 2$ then X_λ is a closed parabolic orbit by Proposition 4.4, while spiral classes are palindromic as discussed in § 11. So fix a dominant nonzero palindromic λ; we must show that one of the three conditions holds.

Let t be a positive node linked to s_0 by I.

Case 1: Suppose we can take I of type A. Then we must have $1 - \alpha_0(\lambda) + \alpha_t(\lambda) \geq 0$ (otherwise there is an affine A-move showing λ covers a pair). Hence $(m_t - 1)\alpha_t(\lambda) \leq 1$, and if equality holds then there are no other positive nodes.

1a: Suppose $m_t > 1$. Then $m_t = 2$ and $\lambda = \omega_t^\vee$; in particular $\alpha_0(\lambda) = 2$.

1b: Suppose $m_t = 1$ (i.e., t is minuscule) and there is one other positive node u. Then u must also be minuscule, forcing Φ of type A, D, or E_6; in particular, Φ is simply-laced. Then u is also connected to s_0 by a type A subgraph with all interior nodes vanishing, so by Lemma 9.6 we must have $1 - \alpha_0(\lambda) + \alpha_u(\lambda) \geq 0$. It follows that $\lambda = \omega_t^\vee + \omega_u^\vee$ (with t, u minuscule); in particular $\alpha_0(\lambda) = 2$.

1c: Suppose $\lambda = m\omega_t^\vee$ with t minuscule. Then $m \neq 1$ (otherwise $(\lambda \notin \mathcal{Q}^\vee)$. If $m = 2$ then $\alpha_0(\lambda) = 2$, so suppose $m > 2$. If Φ is not of type A, then t is not adjacent to s_0, and the node a adjacent to s_0 has $m_a = 2$. Hence firing s_0 yields $\mu = s_0\lambda$ with $\alpha_0(\mu) = 2(1-m) + m = 2 - m < 0$. Note that α_0 satisfies the opposite sign condition on μ: For if $\alpha_0 = \alpha + \alpha'$, then each of α, α' must contain α_a once. Thus if α, say,

contains α_t, we have $\alpha(\lambda) = 1$ and $\alpha'(\lambda) = 1 - m < 0$. Hence there is a linear move $\mu \downarrow s_{\alpha_0}\mu$ by Proposition 9.1, showing that λ covers a fork.

Now suppose Φ has type A. If t is not adjacent to s_0, a similar argument shows that λ covers a trident, again contradicting palindromy. Finally, if $t = s_1, s_n$ then $n+1$ divides m (otherwise $\lambda \notin \mathcal{Q}^\vee$), and hence $\lambda = k(n+1)\omega_i^\vee$ for $i = 1, n$ and $k \geq 1$.

This completes the proof of Theorem 10.2 in Case 1. In particular, the theorem is now proved in the simply-laced case.

Case 2: I has type C. This only happens when Φ itself has type C and $t = s_i$, $i < n$. Then $1 - \alpha_0(\lambda) + 2\alpha_t(\lambda) \neq 0$, since α_0 takes only even values in type C. By Lemma 9.7 we must have $1 - \alpha_0(\lambda) + \alpha_t(\lambda) \geq 0$ (otherwise λ covers a pair). Then $\lambda = \omega_t^\vee$ and $\alpha_0(\lambda) = 2$ as in Case 1a.

Case 3: I has type B. This can only happen in two ways:

3a: Φ itself has type B_n and $t = s_n$. Then $\alpha_1(\lambda) = 0$, because otherwise we can take $t = s_1$ in case 1b, a contradiction. Hence $\alpha_0(\lambda) = 2\alpha_n(\lambda)$ and $1 - \alpha_0(\lambda) + 2\alpha_n(\lambda) = 1 \neq 0$. Thus, as in Case 2, we conclude $\lambda = \omega_n^\vee$ and $\alpha_0(\lambda) = 2$.

3b: Φ has type F_4 and $t = s_3$. Then λ covers a pair by Lemma 9.7, a contradiction.

Case 4: I does not have type ABC. This can happen in three ways:

3a: Φ has type C_n and $t = s_n$. As in Case 1c, we conclude that $\lambda = 2\omega_n^\vee$ and $\alpha_0(\lambda) = 2$.

3b: Φ has type F_4 and $t = s_4$. Thus $\lambda = m\omega_4^\vee$ for some $m > 0$; we will show that $m = 1$ and hence $\alpha_0(\lambda) = 2$. We use only the following facts: (i) $m_4 = 2$; and (ii) if $m_4(\alpha) = 1$, then α is short. Now suppose $m > 1$, and let β denote the maximal positive root with $m_4(\beta) = 1$. Then β is short, and $1 - \beta(\lambda) < 0$. If α, α' is a short β-positive pair, then $\beta + \alpha = \alpha'$. Hence α contains α_4 (by the maximality of β) and $\alpha(\lambda) > 0$. Furthermore, there are no long β-positive pairs α, α'. For in that case $2\beta + \alpha = \alpha'$, and hence α does not contain α_4. But $\beta + \alpha$ is a root, so this contradicts the maximality of β. Hence the positive pair condition is satisfied. If α, α' is a long β-negative pair, then $\alpha + \alpha' = 2\beta$. It follows from (i) and (ii) that one of $\alpha(\lambda), \alpha'(\lambda)$ is zero. Hence the negative pair condition is also satisfied. Then $\lambda \downarrow r_\beta\lambda$ and λ covers a pair by Proposition 9.3, a contradiction.

Remark: Although it is not necessary to do so, one can easily write out the roots used above explicitly, using the tables in [5]: We have $\beta = 1231$. There are four β-positive pairs, with the α of the pair given by the graph roots containing α_4: 0001, 0011, 0111, 1111. There are two long β-negative pairs: (1342, 1120) and (1242, 1220).

3c: Φ has type G_2 and $t = s_1$. If $\alpha_2(\lambda) > 0$, then $\lambda = \omega_2^\vee$ by Case 1a, and $\alpha_0(\lambda) = 2$. It remains to show that if $\lambda = m\omega_1^\vee$, then $m = 1$. Let $\beta = \alpha_1 + \alpha_2$ (the highest root containing α_1 once). If $m > 1$ then $\beta(\lambda) \geq 2$ and the positive pair condition is satisfied. There is one β-negative pair α_2, α_0 (a long pair). But $\alpha_2(\lambda) = 0$, so the negative pair condition is satisfied and λ covers a pair, a contradiction.

This completes the proof of the theorem.

10.3 Mixed case

We assume throughout this section that λ is palindromic of mixed type. Thus λ has a unique negative node $s \in S$ and at least one positive node $t \in S$. Most of the work is done in a series of preliminary lemmas, culminating in Corollary 10.8.

Recall, we say that an element λ of mixed type is *overweight* if $\alpha_0(\lambda) \geq 2$, or equivalently, s_0 is a negative node. An overweight λ of mixed type covers a pair, and hence is not palindromic.

10.3.1 Preliminary lemmas

Lemma 10.3 *Suppose that Φ is not of type G_2. Then either*
(*a*) *there exists a positive node t linked to s with $\alpha_s(\lambda) + \alpha_t(\lambda) \geq 0$, or*
(*b*) *Φ has type C_n and λ is the chain $(\underline{0}, 1, -2)$.*

Proof: Case 1: There is a positive node t linked to s by a subgraph I of type ABC, where in the BC case either (i) t is the minuscule node of I, or (ii) s is the minuscule node of I and $\alpha_s(\lambda) + 2\alpha_t(\lambda) \neq 0$. In this case $\alpha_s(\lambda) + \alpha_t(\lambda) \geq 0$ by Lemma 9.2 (otherwise λ covers a pair).

Case 2: There is a positive node t linked to s by a subgraph I of type BC, where s is the minuscule node of I and $\alpha_s(\lambda) + 2\alpha_t(\lambda) = 0$.

In type B_n we have $t = s_n$, and we may assume there are no other positive nodes, since such a node would be linked to s by a type A subgraph and we are back in Case 1. Then there is a graph-splitting move based on $I - \{s\}$; hence λ covers a pair, a contradiction.

In type C_n we have $s = s_n$, so the assumption $\alpha_s(\lambda) + 2\alpha_t(\lambda) = 0$ implies t is the only positive node (otherwise λ is overweight). If $\alpha_t(\lambda) \geq 2$ there is again a graph-splitting move and λ covers a pair, a contradiction. Hence $\alpha_t(\lambda) = 1$, with $t = s_i$ for some $i < n$. If $i < n - 1$ then after firing s_n there is a linear A-move showing that λ covers a fork. Hence $\lambda = (\underline{0}, 1, -2)$, which is a chain.

In type F_4 with $s = s_1$ and $t = s_3$, we have $\alpha_4(\lambda) = 0$ (otherwise λ is overweight). Since $\alpha_s(\lambda) + \alpha_t(\lambda) < 0$, firing down shows that λ forks too soon. If $s = s_2, s_3$ then we may assume that $\alpha_1(\lambda) = 0$, since otherwise we are back in Case 1. In both cases it follows that λ covers a fork, a contradiction.

Case 3: Φ has type F_4 and $\{s, t\} = \{s_1, s_4\}$. In both cases λ forks too soon. This completes the proof of Lemma 10.3.

Lemma 10.4 *Let t be any positive node (not necessarily linked to s). Suppose $m_t > m_s$ and $\alpha_s(\lambda) + \alpha_t(\lambda) \geq 0$. Then*
(1) *$m_t = m_s + 1$;*
(2) *$\alpha_s(\lambda) = -1$ and $\alpha_t(\lambda) = 1$;*
(3) *there are no other positive nodes.*

Proof: If any one of the three conditions is not satisfied, then λ is overweight, a contradiction.

Lemma 10.5 *Let t be a positive node linked to s. Suppose $m_s = m_t$ and $\alpha_s(\lambda) + \alpha_t(\lambda) > 0$. Then*

(1) $m_s = m_t = 1$;
(2) $\alpha_s(\lambda) + \alpha_t(\lambda) = 1$;
(3) *there are no other positive nodes.*

Furthermore, either $\alpha_s(\lambda) = -1$ or Φ has type A and λ is spiral

Proof: It is immediate that conditions (1)-(3) hold (otherwise λ is overweight). In particular Φ has type A, D or E_6, since there are two minuscule nodes. Now suppose $\alpha_s(\lambda) \leq -2$, and let $\mu = s\lambda$. Then $\alpha_s(\mu) = -\alpha_s(\lambda) \geq 2$. In types D and E_6 there is a unique node $a \in S$ adjacent to s, and $a \neq t$. Moreover, a is minuscule in $S - \{s\}$. Hence if α, α' is an α_s-positive pair, so that $\alpha_s + \alpha = \alpha'$, it follows that α_a occurs exactly once in α. Therefore, either $\alpha(\mu) = \alpha_a(\mu) + \alpha_t(\mu) = 1$, or $\alpha'(\mu) = \alpha_s(\mu) + \alpha_a(\mu) = 0$. Thus μ covers a pair by Lemma 9.5 (with $I = \{s\}$), and hence λ covers a fork, a contradiction.

Now suppose Φ has type A. Then if s and t are not adjacent, a similar argument shows that λ covers a trident, a contradiction. If s and t are adjacent, λ is spiral by Proposition 11.1.

Lemma 10.6 *Let t be a positive node linked to s. Suppose $m_s = m_t$ and $\alpha_s(\lambda) + \alpha_t(\lambda) = 0$. Then $\alpha_s(\lambda) = -1$.*

Proof: Case 1: Assume there are no other positive nodes. Suppose $\alpha_s(\lambda) \leq -2$, and let β denote the maximal root such that $m_t\beta = m_s(\beta) + 1$. Then $\beta(\lambda) \geq 2$. If β is a long root and α, α' is a β-positive pair, then since $\beta + \alpha = \alpha'$, we must have $m_t(\alpha) \neq m_s(\alpha)$ by the maximality of β. If $m_t(\alpha) > m_s(\alpha)$ then $\alpha(\lambda) > 0$. If $m_t(\alpha) < m_s(\alpha)$ then $\alpha'(\lambda) \leq 0$. Hence the positive pair condition is satisfied and we conclude that λ covers a pair by Proposition 9.3. In particular, this completes the proof of Case 1 in the simply-laced case.

In types B and F_4, β is always a long root. This is clear on inspection in type B. In F_4 we have $\{s, t\} = \{s_1, s_4\}$. If s_1 is the negative node then $\beta = 1342$, the second highest root. This is clearly long since it belongs to the type A_2 subgraph of $\tilde{D}$ on s_0, s_1. If s_1 is the positive node then $\beta = 1220$, the highest root of the B_3 subsystem. Hence β is long, completing the proof in B, F_4.

In type C_n, β will be a short root and more care is required in the case of long pairs. Let $s = s_i, t = s_j$, where $1 \leq i, j < n$. If $i < j$ then $\beta = e_1 + e_{i+1}$. There are no long β-positive pairs, so the positive pair condition follows as before. However, there is one β-negative pair $2e_1, 2e_{i+1}$. But $2e_1 = \alpha_0$ and $\alpha_0(\lambda) = 0$, so the negative pair condition holds as well and $\lambda \downarrow r_\beta\lambda$ by Proposition 9.3. If $i > j$ then $\beta = e_1 - e_i$. There are no long β-negative pairs, but there is one long β-positive pair $2e_1, 2e_i$. Thus to check the positive pair condition we have to consider $2\beta + 2e_i = 2e_1 = \alpha_0$. But $\alpha_0(\lambda) = 0$ so again $\lambda \downarrow r_\beta\lambda$. Hence, in either case, λ covers a pair, a contradiction.

Case 2: Assume there is more than one positive node. Then the following three conditions hold (otherwise λ is overweight):

(1) there is only one additional positive node u;

(2) u is minuscule;
(3) $\alpha_u(\lambda) = 1$.

This rules out E_8, F_4 and G_2 (since there is a minuscule node) and also C_n (since $\alpha_u(\lambda)$ is odd by§ 3.2). Now assume that $\alpha_s(\lambda) \leq -2$.

Suppose first that u, t, s lie on a type A subgraph. If s is linked to u, then λ covers a pair by Lemma 9.2. So suppose that t lies between u and s. If Φ has type A, then there is a graph-splitting move showing that λ covers a pair where I is the connected component of $S - \{u, s\}$ containing t. In type BD we must have $u = s_1$ (if $u = s_{n-1}, s_n$ in type D_n, then $\lambda \notin \mathcal{Q}^\vee$). Thus $t = s_i$, $s = s_j$, with $1 < i < j$. Now let $\beta = e_1 + e_i$, which is the smallest root containing α_u such that $m_t(\beta) = 1, m_s(\beta) = 2$. Then $\beta(\lambda) = 1 + \alpha_s(\lambda) < 0$, and one easily checks that the opposite sign condition is satisfied. Hence λ covers a pair in this case by Proposition 9.3.

In type E_6 we have $m_s = m_t = 2$. In all cases s has two adjacent zero nodes in $\tilde{D}$, and therefore λ covers a fork. In type E_7 we have $s = s_1, s_2, s_3$. In all cases λ forks too soon.

If u, t, s do not lie on a type A subgraph, then we are in type B_n with $u = s_1$ and $s = s_n$. Then after firing s there is a linear A-move showing that λ covers a fork. (Note that this works whether or not t is adjacent to s, bearing in mind that back-firing along the double bond adds $2\alpha_s(\lambda)$ to the value of its neighbor.)

Lemma 10.7 *Suppose Φ is not of type G_2 and $m_t < m_s$ for all positive nodes t. Then $\alpha_s(\lambda) = -1$.*

Proof: Suppose $\alpha_s(\lambda) \leq -2$.

Case 1: Φ has type BCD and there is only one positive node.

Note that the positive node is minuscule. We may assume that $2\alpha_s(\lambda) + \alpha_t(\lambda) \leq 1$ (otherwise λ is overweight). If $\alpha_s(\lambda) + \alpha_t(\lambda) < 0$ there is a linear ABC-move and λ covers a pair. If $\alpha_s(\lambda) + \alpha_t(\lambda) = 0$ and $\alpha_s(\lambda) \leq -2$ there is a graph-splitting move, and again λ covers a pair. So we assume $\alpha_s(\lambda) + \alpha_t(\lambda) > 0$ and consider three cases:

$2\alpha_s(\lambda) + \alpha_t(\lambda) < 0$: Let β be the smallest root containing $2\alpha_s(\lambda) + \alpha_t(\lambda)$. In all cases β is a long root. It is then clear that the opposite sign condition is satisfied, as desired. Hence $\lambda \downarrow s_\beta\lambda$ by Proposition 9.1 and λ covers a pair.

$2\alpha_s(\lambda) + \alpha_t(\lambda) = 0$: Let β be the highest root containing α_s once. Then $\beta(\lambda) = \alpha_s(\lambda) + \alpha_t(\lambda) \geq 2$ since $\alpha_s(\lambda) \leq 2$. If β is long then the positive pair condition is clearly satisfied and we are done. If β is short and there is a long β-negative pair α, α', then $\alpha + \alpha' = 2\beta$ and $\alpha' = \beta + (\beta - \alpha)$, where $\beta - \alpha$ is a positive root. By the maximality of β, it follows that α' must contain α_s twice, as well as α_t. Hence $\alpha'(\lambda) = 0$, and both the positive pair condition and the negative pair condition are satisfied. By Proposition 9.3, λ covers a pair.

$2\alpha_s(\lambda) + \alpha_t(\lambda) = 1$: Note this rules out type C, since $\alpha_t(\lambda)$ is odd. In type D there are two zero nodes adjacent to s, so λ covers a fork. Type B_n is similar, except in the case $s = s_n$. In that case, let $\lambda = (a, \underline{0}, b)$ with $b < 0$ and $a + 2b = 1$. Firing s_n yields $\mu = (a, \underline{0}, 2b, -b)$. If $b \leq -2$ there is an affine move $\mu \downarrow r_{\alpha_n}\mu$ by Proposition 9.5, showing that λ again covers a fork. Here we note that if α, α' is a β-positive pair, then $2\alpha_n + \alpha = \alpha'$, and hence $\alpha(\lambda) = 2b$ or $\alpha(\lambda) = 1$. In either case the positive pair condition is satisfied.

This completes the proof in case 1.

Case 2: Φ has type BCD and there is more than one positive node.

Since the positive nodes must be minuscule, this can only happen in type D. If t is any positive node then it is linked to s and $\alpha_s(\lambda)+\alpha_t(\lambda)\geq 0$ by Lemma 9.2. Hence there are two positive nodes t,u. If $\alpha_s(\lambda)+\alpha_t(\lambda)=0=\alpha_s(\lambda)+\alpha_u(\lambda)$, there is an affine move by Lemma 9.2 with β being the highest root containing α_s once. Thus, λ covers a pair. If (say) $\alpha_s(\lambda)+\alpha_t(\lambda)>0$, then $\alpha_s(\lambda)+\alpha_t(\lambda)=1$ and $\alpha_s(\lambda)+\alpha_u(\lambda)=0$ (otherwise λ is overweight). In other words, up to symmetry the diagram of λ has one of the following forms, with $a\geq 2$: $(a,\underline{0},-a,\underline{0},0,a+1)$ and $(\underline{0},-a,\underline{0},a+1,a)$.

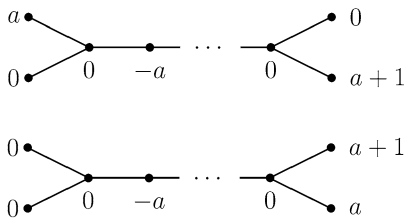

In the first case λ covers a fork. The second diagram is not even in the coroot lattice, as the reader can check.

Case 3: Φ has type E. It is a pleasant exercise in the palindromy game to check that λ forks too soon or is overweight. Details are left to the reader.

Case 4: Φ has type F_4. Then $s=s_2,s_3$. If $s=s_2$ then $\alpha_3(\lambda)=0$ by assumption. Then $\alpha_1(\lambda)+\alpha_2(\lambda)\geq 0$ by Lemma 9.2a, and $\alpha_2(\lambda)+\alpha_4(\lambda)\geq 0$ (otherwise λ forks too soon). But then λ is overweight.

If $s=s_3$ then for any positive node t linked to s we have $\alpha_s(\lambda)+\alpha_t(\lambda)\geq 0$, by Lemma 9.2a. If $\alpha_2(\lambda)>0$ then $\alpha_2(\lambda)+\alpha_3(\lambda)\geq 0$. By Proposition 9.2b we have either $\alpha_2(\lambda)+\alpha_3(\lambda)=0$ or $\alpha_2(\lambda)+2\alpha_3(\lambda)\geq 0$. In the first case we must have $\alpha_4(\lambda)>0$ (otherwise λ forks too soon) and λ is overweight. In the second case λ is again overweight. Finally if $\alpha_2(\lambda)=0$ then $\alpha_4(\lambda)>0$ and $\alpha_1(\lambda)+2\alpha_3(\lambda)\geq 0$ (otherwise λ forks too soon), and again λ is overweight. (This last argument doesn't use the assumption $\alpha_s(\lambda)\leq -2$.)

The main conclusions of the five lemmas above can be summarized as follows:

Corollary 10.8 *Suppose Φ is not of type G_2 and λ is palindromic of mixed type, with unique negative node $s\in S$. Then at least one of the following conditions holds*:

(1) $\alpha_s(\lambda)=-1$;

(2) λ *is a chain*;

(3) Φ *has type A and* λ *is spiral*.

10.3.2 Proof of the palindromy theorem in the mixed case

By Corollary 10.8 we may assume $\alpha_s(\lambda)=-1$ (except in type G_2). If s is a long node, then the palindromic λ is admissible, and hence is a closed parabolic orbit or a chain by Lemma 8.1. In particular, the proof of Theorem 1.4 is now complete in

the simply-laced case. It remains to consider types BCF when s is a short node with $\alpha_s(\lambda) = -1$, and type G_2.

C_n: Since s is short, $s = s_i$ for some $i < n$. Let s_j be a positive node linked to s_i. If $j < n$, by Lemma 10.3 we may assume $\alpha_s(\lambda) + \alpha_t(\lambda) = 0$ (otherwise λ is overweight) and that there are no other positive nodes by Lemma 10.5. If s and t are not adjacent, then firing s shows that λ covers a fork (note this works even when $i = 1$, since then back-firing along the double bond adds -2 to the initial value +1 on the node s_0). If s and t are adjacent then $\lambda = \pm(\underline{0}, 1, -1, \underline{0})$, which is a chain.

Now suppose $t = s_n$ is the unique positive node. Since $\alpha_n(\lambda)$ is necessarily even, we must have $\alpha_n(\lambda) = 2$ (otherwise λ is overweight). If $i < n - 1$, then firing s shows that λ covers a fork (note this works even when $i = 1$, since back-firing along the double bond puts -1 on the s_0 node). If $i = n - 1$ we have $\lambda = (\underline{0}, -1, 2)$, which is a chain.

B_n: Here $s = s_n$. Let s_j be the unique positive node linked to s_n. Then $\alpha_j(\lambda) = 1, 2, 3$, where the values 2, 3 can occur only when $j = 1$ (otherwise λ is overweight).

Case 1: $\alpha_j(\lambda) = 1$. If $j < n - 1$ then firing s_n yields $(\underline{0}, 1, \underline{0}, -2, 1)$ and there is a linear A-move, showing that λ covers a fork. If $j = n - 1$ then there must be another positive node (otherwise $\lambda \notin \mathcal{Q}^\vee$). This forces $\lambda = (1, \underline{0}, 1, -1)$, which is a singular chain.

Case 2: $\alpha_j(\lambda) = 2$. This forces $\lambda = (2, \underline{0}, -1)$. Thus n must be even (otherwise $\lambda \notin \mathcal{Q}^\vee$), and in particular $n \geq 4$. Then λ covers a fork: Firing s_n yields $(2, \underline{0}, -2, 1)$, which admits an affine move based on the long root $\beta = \alpha_1 + \ldots + \alpha_{n-1} + 2\alpha_n = e_1 + e_n$. This is very similar to the anti-dominant case $(\underline{0}, -1)$; details are left to the reader.

Case 3: $\alpha_j(\lambda) = 3$. This forces $\lambda = (3, \underline{0}, -1)$. Thus n must be odd (otherwise $\lambda \notin \mathcal{Q}^\vee$). If $n \geq 5$, we can proceed exactly as we did in the anti-dominant case with $(\underline{0}, -1)$. The key point is again that $(3, \underline{0}, -2, 0, 1)$ admits an affine move using $\beta = \alpha_{n-1} + 2\alpha_n$, and as a result λ covers a scepter, hence is not palindromic. This leaves the case $n = 3$, $\lambda = (3, 0, -1)$ which is known to be palindromic by Corollary 4.8.

This completes the proof of Theorem 1.4 in type B.

F_4: Here the short nodes are s_3, s_4. If $s = s_4$, then $\lambda = (0, 1, 0, -1)$ or $\lambda = (1, 0, 0, -1)$ (otherwise λ is overweight). Then $(0, 1, 0, -1)$ is a chain, while $(1, 0, 0, -1)$ forks too soon (barely!).

If $s = s_3$, the argument used in the proof of Lemma 10.7 shows that $\lambda = (0, 1, -1, 1)$, which is a chain.

G_2: Before starting the proof, we recall (§7.3) that there are four chains of mixed type, namely $\pm(1, -1)$ and $\pm(1, -2)$. Recall also that the initial chain coming down from $(-1, 0)$ (see the Hasse diagram in §13)

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

with $(2, -3)$ covering a pair. Hence all of the displayed elements fork too soon.

Now suppose λ is generic, meaning that $\alpha(\lambda) \neq 0$ for all $\alpha \in \Phi$. Let γ range over the four non-simple positive roots $\alpha_1 + \alpha_2$, $2\alpha_1 + \alpha_2$ $3\alpha_1 + \alpha_2$, $3\alpha_1 + 2\alpha_2$. If at least

one of the γ's is negative on λ, then the minimal such γ satisfies the opposite sign condition, and λ covers a pair by Proposition 9.3.

Now suppose all of the γ's are positive on λ. If s_2 is the positive node, then λ is overweight and again λ covers a pair. If s_2 is the negative node, firing it yields a generic dominant class μ: $\lambda = (a, b) \downarrow (a+b, -b) = \mu$. We saw earlier that any such class in type G_2 covers a pair, so λ covers a fork.

It remains to consider the case $\gamma(\lambda) = 0$ for some γ. Let a denote a positive integer.

$\gamma = \alpha_1 + \alpha_2$: If $\lambda = (a, -a)$ and $a \geq 2$ then λ is overweight. If $a = 1$ we have the chain $(1, -1)$. If $\lambda = (-a, a)$ and $a \geq 2$, λ has an affine move $\lambda \downarrow r_{\alpha_2}\lambda$ by Proposition 9.3. Here α_2 is long and there are two β-positive pairs α, α': $(\alpha_1, \alpha_1 + \alpha_2)$ and $(3\alpha_1 + \alpha_2, 3\alpha_1 + 2\alpha_2)$. In each case $\alpha'(\lambda) < 0$ and the positive pair condition is satisfied. If $a = 1$ we again have a chain.

$\gamma = 2\alpha_1 + \alpha_2$: If $\lambda = (a, -2a)$ and $a \geq 2$, there is an affine move r_β with $\beta = 3\alpha_1 + \alpha_2$. Note that β is long and there is just one β-positive pair α_2, α_0, so the positive pair condition is satisfied. If $a = 1$ we have the chain $(1, -2)$. If $\lambda = (-a, 2a)$ then λ is overweight unless $a = 1$, in which case we again have a chain.

$\gamma = 3\alpha_1 + \alpha_2$: If $\lambda = (-a, 3a)$ then λ is overweight. If $\lambda = (a, -3a)$ and $a \geq 2$, there is an affine move r_{α_1} and λ covers a pair. Finally $(1, -3)$ forks too soon as noted above.

$\gamma = 3\alpha_1 + 2\alpha_2$: Note that in this case $|\alpha_1(\lambda)|, |\alpha_2(\lambda)| \geq 2$. If s_2 is the negative node then there is an affine move r_β with $\beta = 3\alpha_1 + \alpha_2$. If s_1 is the negative node and $\alpha_1(\lambda) + \alpha_2(\lambda) \leq 0$ there is an affine move r_{α_2}. If $\alpha_1(\lambda) + \alpha_2(\lambda) > 0$ and $2\alpha_1(\lambda) + \alpha_2(\lambda) \leq 0$, there is an affine move r_β with $\beta = \alpha_1 + \alpha_2$, provided that $\beta(\lambda) \geq 2$. If $\beta(\lambda) = 1$ then $\lambda = (-2, 3)$, which forks too soon as noted above. Finally if $2\alpha_1(\lambda) + \alpha_2(\lambda) > 0$, there is an affine move r_β with $\beta = 2\alpha_1 + \alpha_2$.

11 The spiral varieties in type A

Most of the results in this section are from [16], to which the reader is referred for the missing proofs. The varieties $X_{n,k}$ are denoted $F_{n+1,k}$ in [16]. We include these results for the readers' convenience and to highlight some unsaid consequences of the previous work.

Let σ_d (resp. σ'_d) denote the word in $\tilde{W}$ obtained by starting at s_0 and—writing the word from right to left—proceeding clockwise (resp. counterclockwise) d steps around the Coxeter diagram. For example, if $n = 3$ then $\sigma_6 = s_3 s_0 s_1 s_2 s_3 s_0$. Note that these words are reduced, in $\tilde{W}^S$, and rigid. Note also that σ'_d is conjugate to σ_d under the involution of the Dynkin diagram fixing s_0.

Let $\sigma_{n,k} = \sigma_{kn}$, $\sigma'_{n,k} = \sigma'_{kn}$. We call these elements and the varieties associated to them *spiral*. The term is suggested by the manner in which $\sigma_{n,k}$, $\sigma'_{n,k}$ spiral around the affine Dynkin diagram and up in length as k increases.

Warning: Note that the spiral classes have length divisible by n, not $n + 1$. Hence they are out of phase with the natural period of the affine Dynkin diagram; the first (left-hand) factor s of $\sigma_{n,k}$ rotates around the diagram as k increases.

The corresponding coroot lattice representatives $\lambda_{n,k}, \lambda'_{n,k}$ are described as follows:

Proposition 11.1 *An element $\lambda \in \mathcal{Q}^\vee$ represents a spiral class if and only if*

(1) λ *has exactly two nonzero nodes* $s, t \in \tilde{S}$;
(2) s, t *are adjacent*; *and*
(3) $\alpha_s(\lambda) + \alpha_t(\lambda) = 1$.

More precisely, let $k = r(n+1) + i$, where $0 \leq i < n+1$. Then

$$\lambda_{n,k} = (\underline{0}, k+1, -k, \underline{0}),$$

where $k+1$ is in the i-th coordinate, or $\lambda_{n,k} = (-k, \underline{0})$, $\lambda_{n,k} = (\underline{0}, k+1)$. Similarly

$$\lambda'_{n,k} = (\underline{0}, -k, k+1, \underline{0})$$

or $\lambda'_{n,k} = (k+1, \underline{0})$, $\lambda'_{n,k} = (\underline{0}, -k)$.

Proof: The explicit formulas for $\lambda_{n,k}, \lambda'_{n,k}$ are easily obtained by induction on k, and show that the spiral classes satisfy properties (1)-(3). Conversely, suppose that (1)-(3) hold, and consider the case $s = s_i, t = s_{i+1}$ with $0 < i < n$ and $\alpha_s(\lambda) > 0$. Then since λ is in the coroot lattice, we have $i\alpha_s(\lambda) + (i+1)\alpha_t(\lambda) = 0 \, mod \, n+1$. Then $\lambda = \lambda_{n,k}$ with $k = -\alpha_t(\lambda)$. The remaining cases are similar.

Let $X_{n,k} = X_{\sigma_{n,k}}$, $X'_{n,k} = X'_{\sigma_{n,k}}$. Since $X'_{n,k}$ is canonically isomorphic to $X_{n,k}$ as a variety, in what follows we will state the results only for $X_{n,k}$.

Proposition 11.2 *$X_{n,k}$ is isomorphic to the variety of k-dimensional $\mathbb{C}[z]/z^k$-submodules in $\mathbb{C}[z]/z^k \otimes \mathbb{C}^{n+1}$. In particular, $X_{n,1} \cong \mathbb{C}P^n$.*

This description arises from the Quillen model, which identifies $\mathcal{L}_{SU(n+1)}$ with certain spaces of $\mathbb{C}[z]$-lattices in $\mathbb{C}[z, z^{-1}] \otimes \mathbb{C}^{n+1}$. In fact from this point of view, $X_{n,k}$ is precisely the intersection in BU of the Ind-varieties $\mathcal{L}_{SU(n+1)}$ and $BU(k)$.

Let $\begin{bmatrix} n+k \\ n \end{bmatrix}_q$ denote the "q-binomial coefficient" or equivalently the Poincaré polynomial for the classical Grassmannian $G_n\mathbb{C}^{n+k}$.

Proposition 11.3 $|X_{n,k}| = \begin{bmatrix} n+k \\ n \end{bmatrix}_q$. *In particular, $X_{n,k}$ is palindromic.*

Undoubtedly a direct combinatorial proof of this could be given. Here we will mention two topological proofs. The first is based on:

Proposition 11.4 *$X_{n,k} - X_{n,k-1} = E(\gamma^k \downarrow X_{n-1,k})$, the total space of the canonical k-plane bundle defined by Proposition 11.2. Hence $X_{n,k}/X_{n,k-1} = T(\gamma^k)$.*

Thus $|X_{n,k}| = |X_{n,k-1}| + t^k|X_{n-1,k}|$, which is exactly the Pascalian recursion formula for the q-binomial coefficients. This yields one proof of Proposition 11.3. The second is based on:

Theorem 11.5 $[X_{n,j}] \cdot [X_{n,k}] = [X_{n,j+k}]$, *where the dot denotes Pontrjagin product in* $H_*\Omega G \cong H_*\mathcal{L}_G$. *Moreover the natural map* $Sym^k(H_*X_{n,1}) \longrightarrow H_*X_{n,k}$ *is an isomorphism.* (Sym^k *denotes the* k*-th symmetric power.*)

This theorem gives another proof of Proposition 11.3, since one can easily check that $|Sym^k(H_*\mathbb{C}P^n)| = \begin{bmatrix} n+k \\ n \end{bmatrix}_q$.

Theorem 11.6 $X_{n,k}$ *satisfies Poincaré duality integrally if and only if* $k = 1$. *It satisfies Poincaré duality rationally for all* k.

Remark: The first assertion has already been proved in Theorem 1.1. By Proposition 11.3 $X_{n,k}$ is palindromic, and hence is rationally smooth by the Carrell-Peterson theorem, proving the second assertion. Here we provide alternative proofs of both assertions.

Proof: First observe that Poincaré duality can be expressed in terms of homology as follows: If $[X]$ is the fundamental class, then there are bases e_i, e_i' for H_*X that are dual in the sense that $\Delta_*[X] = \sum e_i \otimes e_i'$, where Δ is the diagonal map. It follows that if X_λ is any Schubert variety and $[X_\lambda] = y^k$ for some $y \in H_*\mathcal{L}_G$ and $k > 1$, then X_λ does not satisfy Poincaré duality: For we may assume $k = p$ is a prime, and then $\Delta_*[X_\lambda] = (\Delta_* y)^p \, mod \, p$. Hence $\Delta_*[X_\lambda]$ is concentrated in bidegrees divisible by p, and no such dual bases exist. Since $[X_{n,k}] = [X_{n,1}]^k$ by Theorem 11.5, this proves that $X_{n,k}$ does not satisfy Poincaré duality for $k > 1$. For $k = 1$, we have $X_{n,1} = \mathbb{P}^n$.

Now let $b_1, \ldots, b_n$ denote the standard basis for $\tilde{H}_*X_{n,1} = \tilde{H}_*\mathbb{C}P^n$. Then $H_*X_{n,k}$ consists of polynomials of degree at most k in $b_1, \ldots, b_n$, with b_n^k the fundamental class. We set $b_0 = 1$. Now let $r = (r_0, \ldots, r_n)$ be a multi-index with $r_i \geq 0$ and $\sum r_i = k$. Let $b^r = b_0^{r_0} \ldots b_n^{r_n}$, and let $r^* = (r_n, \ldots, r_0)$. Then the b_r, b_{r^*} are bases, and if we take coefficients in $\mathbb{Q}$, then up to scalar multiples they are dual in the above sense:

$$\Delta_*(b_n^k) = (\sum_{i+j=n} b_i \otimes b_j)^k = \sum_r c_r (b_0 \otimes b_n)^{r_0} \ldots (b_n \otimes b_0)^{r_n} = \sum_r c_r b^r \otimes b^{r^*},$$

where the multinomial coefficients c_r are all nonzero. Thus, up to nonzero scalar multiples, b_r, b_{r^*} are dual bases over $\mathbb{Q}$. This completes the proof.

12 Hasse diagrams

In this section we give some examples of Hasse diagrams for the Bruhat order on $\tilde{W}^S$ in a range of dimensions, showing in particular the palindromics and the closed

parabolic orbits. These diagrams are easily generated by hand in the following way: First of all, the length generating function $|\tilde{W}^S|(t)$ is given by Bott's formula Theorem 3.1 where in the latter case variable t is assigned dimension 2. So we know in advance the number of nodes at each level. Then we begin firing up from $0 \in \mathcal{Q}^\vee$; this easily yields the weak order on the coroot lattice representatives λ. By recording the node fired at each step, we have reduced expressions for the minimal length representatives $\sigma \in \tilde{W}^S$ as well (incidentally, this also yields the cup product coefficients occurring in Chevalley's formula). Then we fill in the missing covering relations by using the standard criterion: $\sigma \downarrow \tau$ if and only if τ can be obtained from σ by omitting one generator from some reduced expression for σ. A further interesting exercise is to find and check the affine or linear moves that produce these descents.

The circled nodes are the non-trivial palindromic classes. A double circle indicates a closed parabolic orbit.

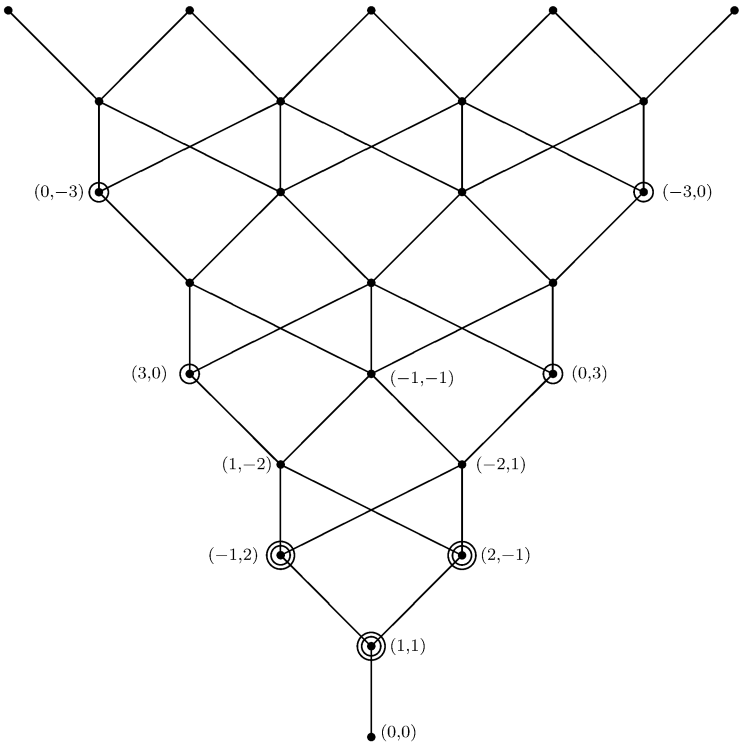

Type $\tilde{A}_2$

The spiral classes $\lambda_{2,k}, \lambda'_{2,k}$ are the elements of even length along the two edges of the diagram.

Note also that for $k = 2$ the Poincaré polynomial of $X_{(3,0)}$ or $X_{(0,3)}$ is the same as that of $G_2\mathbb{C}^4$. The ring structure in cohomology, however, is different; Poincaré duality fails. When $k = 3$ the Poincaré polynomial is the same as that of $G_2\mathbb{C}^5$. The order ideal below $(-3, 0)$ or $(-3, 0)$ is visibly not self-dual; in particular it is not isomorphic to Bruhat poset of $G_2\mathbb{C}^5$.

The bilateral symmetry in the diagram reflects the automorphism of the affine Dynkin diagram fixing the special node s_0.

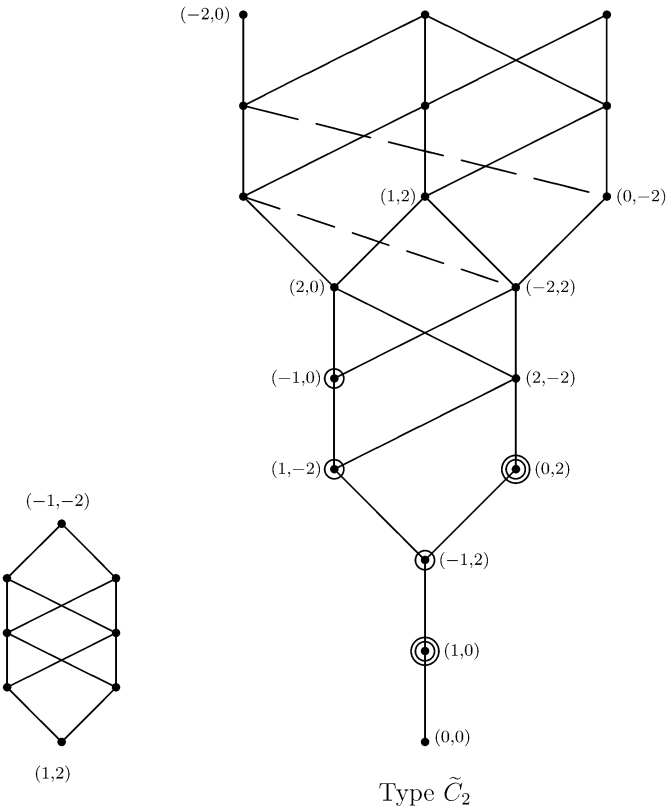

Type $\widetilde{C}_2$

Note the asymmetry in the diagram. No global symmetry is expected, because there are no automorphisms of the affine Dynkin diagram fixing the special node s_0. However, the range of the diagram is too low to show all the patterns present. In particular the smallest generic λ is the dominant class $(1, 2)$. It lies at the bottom of the first generic orbit, as shown in the small diagram on the left.

On the other hand, the only non-trivial palindromic Schubert varieties are the five visible in the diagram, each of which happens to be a chain.

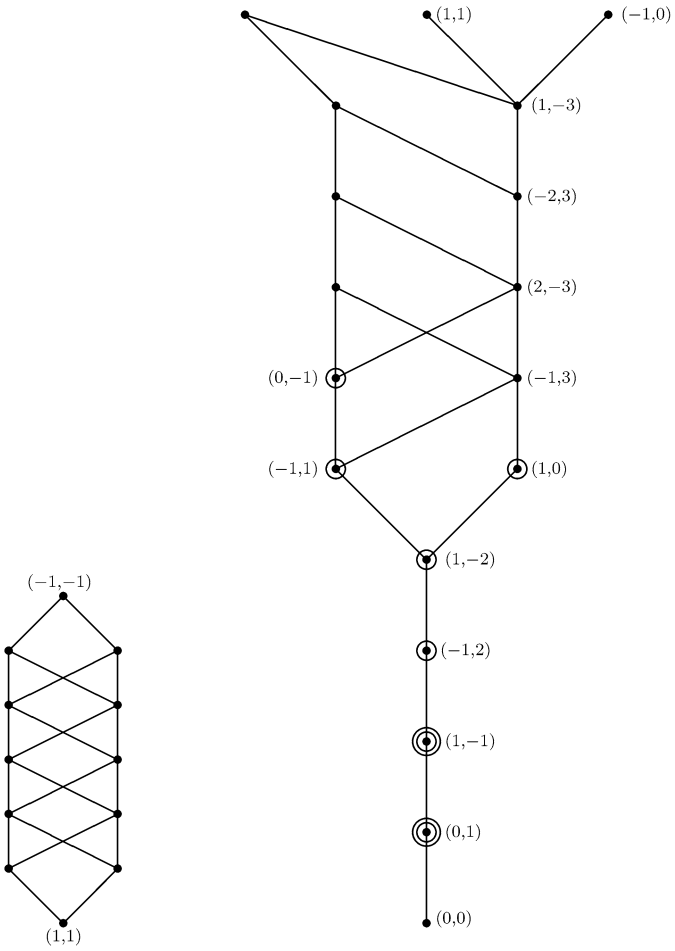

Type $\widetilde{G}_2$

The range of the diagram is too low to show all the patterns. The lowest generic orbit begins with the dominant class $\lambda = (1, 1)$ as shown in the diagram on the left. However, the only non-trivial palindromic Schubert varieties are the seven visible in the diagram, each of which happens to be a chain.

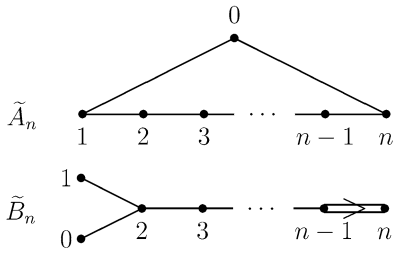

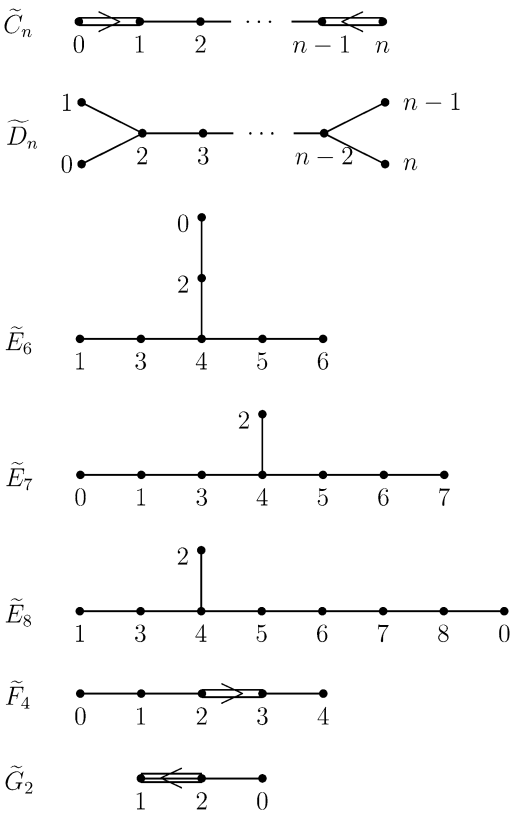

Table of coefficients m_i:

A_n: $m_i = 1$ for all i.

B_n: $m_1 = 1$; $m_i = 2$ otherwise.

C_n: $m_n = 1$; $m_i = 2$ otherwise.

D_n: $m_i = 1$ for $i = 1, n-1, n$; $m_i = 2$ otherwise.

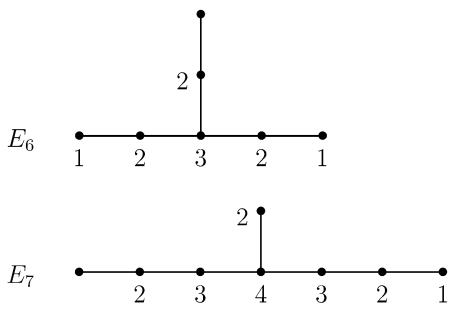

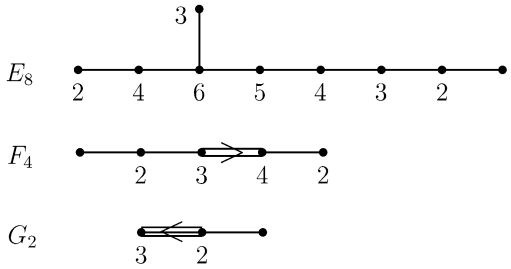

Acknowledgements The authors would like to thank Shrawan Kumar, Thomas Lam, Mark Shimozono, Pete Littig, and Tricia Hersh for helpful conversations. We thank Mary Sheetz for preparing the Hasse diagrams.

References

1. Billey, S., Lakshmibai, V.: Singular Loci of Schubert Varieties. Birkauser, Boston (2000)
2. Billey, S., Mitchell, S.: Affine partitions and the affine Grassmannians (in preparation)
3. Björner, A., Brenti, F.: Combinatorics of Coxeter groups. Graduate Texts in Mathematics, vol. 231. Springer, New York (2005)
4. Bott, R.: An application of the Morse theory to the topology of Lie-groups. Bull. Soc. Math. France **84**, 251–281 (1956)
5. Bourbaki, N.: Lie Groups and Lie Algebras. Springer, Berlin (2002)
6. Carrell, J.: The Bruhat graph of a Coxeter group, a conjecture of Deodhar, and rational smoothness of Schubert varieties. In: Algebraic groups and their generalizations: classical methods, University Park, PA, 1991. Proc. Sympos. Pure Math., vol. 56, pp. 53–61. Am. Math. Soc., Providence (1994), Part 1
7. Carrell, J., Kuttler, J.: Singular points of t-varieties in g/p and the Peterson map. Invent. Math. **151**, 353–379 (2003)
8. Cohen, R., Lupercio, E., Segal, G.: Holomorphic Spheres in Loop Groups and Bott Periodicity. Asian J. Math. **3**, 501–518 (1999)
9. Eriksson, K.: Strongly convergent games and Coxeter groups. Ph.D. Thesis, KTH, Stockholm (1993)
10. Humphreys, J.: Reflection Groups and Coxeter Groups, Cambridge Studies in Advanced Mathematics, vol. 29. Cambridge University Press, Cambridge (1990)
11. Iwahori, N., Matsumoto, M.: On some Bruhat decomposition and the structure of Hecke rings of p-adic Chevalley groups. Publ. Math. I.H.E.S. **25**, 5–48 (1965)
12. Kumar, S.: The nil Hecke ring and singularity of Schubert varieties. Inventiones Math. **123**, 471–506 (1996)
13. Kumar, S.: Kac-Moody Groups, their Flag Varieties and Representation Theory. Birkhauser, Boston (2002)
14. Littig, P.: University of Washington Ph.D. Thesis (2005)
15. McCrory, C.: A characterization of homology manifolds. J. London Math. Soc. **16**, 146–159 (1977)
16. Mitchell, S.A.: A filtration of the loops on SU(n) by Schubert varieties. Math. Z. **193**, 347–362 (1986)
17. Mitchell, S.A.: Parabolic orbits in flag varieties. Preprint (2006). http://www.math.washington.edu/~mitchell/papers/
18. Mozes, S.: Reflection processes on graphs and Weyl groups. J. Combin. Theory Ser. A **53**, 128–142 (1990)
19. Stembridge, J.R.: On the fully commutative elements of Coxeter groups. J. Algebraic Combin. **5**(4), 353–385 (1996)