Essays on Coxeter groups

Bruhat closures

Bill Casselman
University of British Columbia
cass@math.ubc.ca

If G is a split reductive group and B a Borel subgroup then G is the disjoint union of the BwB as w ranges over the Weyl group of a maximal torus in B. The closure of BwB is the union of certain double cosets BxB, where the x that occur can be characterized in purely combinatorial terms. Somewhat surprisingly, this combinatorial definition may be extended to define the Bruhat closure operation in any Coxeter group.

This essay will sketch this material, in so far as it is important in representation theory. The standard sources are [Dixmier:1974], §§5.8–5.11 of [Humphreys:1990], and—particularly thorough—Chapter 2 of [Bjorner-Brenti:2005]. My approach is somewhat different, at least at the beginning. In the last section my treatment is novel only in so far as it is accompanied by pictures that I believe make the argument clearer. But here, too, I am largely following §7.7 of Dixmier's book.

Contents

1. The closure of an element

I recall first some basic facts I'll need about Coxeter groups. Suppose (W, S) to be a Coxeter system with Coxeter matrix $(m_{s,t})$ (possibly with $m_{s,t} = \infty$). It is the group defined by generators in S and relations

$$s^2 = 1, \quad (st)^{m_{s,t}} = 1 .$$

Let $V = \mathbb{R}^S$ with basis (α_s). Define on V the inner product

$$\alpha_s \bullet \alpha_t = -\cos(\pi/m_{s,t}) .$$

For each s let

$$\rho_s\colon\ v \longmapsto v - 2(\alpha_s \bullet v)\alpha_s .$$

be the orthogonal reflection in the hyperplane $\alpha \bullet v = 0$. The map $s \mapsto \rho_s$ extends to an embedding of W in $\mathrm{GL}(V)$. Let C be the open cone

$$C = \{v \in V \mid \alpha_s \bullet v > 0 \text{ for all } s \in S\}$$

and let $\mathcal{C}$ (called by some the **Tits cone** and by others the **Vinberg cone**) be the interior of the union of the closures $w\overline{C}$ as w ranges over W. The group W acts discretely on $\mathcal{C}$ and $\overline{C} \cap \mathcal{C}$ is a strict fundamental domain for this action.

The α_s make up the set Δ of simple roots of the system. The roots Σ of the system are the W-transforms of the α_s. A root λ is said to be positive if $\lambda > 0$ on C, negative if $\lambda < 0$ on C. No root intersects the interior of C, so all roots are one or the other. For w in W and $\lambda > 0$, $w\lambda < 0$ if and only if C and $w^{-1}C$ are on opposite sides of the boundary of λ. Let Σ^+ the set of positive roots.

For every root λ, let s_λ be the orthogonal reflection in the hyperplane $\lambda = 0$. If $\lambda = w\alpha_s$ with $\alpha = \alpha_s$ in Δ, then $s_\lambda = ws_\alpha w^{-1}$.

A **word** in the alphabet S is the concatenation $s_1 \bullet \ldots \bullet s_n$ of elements of S. For w in W, let $\ell(w)$ be the length of the shortest word representing it (which is called a **reduced** word). Thus for s in S, $\ell(sw) = \ell(w) \pm 1$.

The basic result relating the combinatorics and geometry of W is that if s lies in S then $\ell(sw) = \ell(w) + 1$ if and only if $w^{-1}\alpha_s > 0$. In this case, I'll write $sw > w$. This is the simplest case of a more general fact. For any w in W, let

$$L_w = \{\lambda > 0 \,|\, w^{-1}\lambda < 0\}\,.$$

Then $\ell(xy) = \ell(x) + \ell(y)$ if and only if L_{xy} is the disjoint union of L_x and xL_y. As one consequence of this, together with the fact that $L_s = \{\alpha_s\}$, we have $\ell(w) = |L_w|$.

For $T \subseteq S$ let W_T be the subgroup of W generated by T. It too is a Coxeter group, and the length functions in W_T and $W = W_S$ are the same.

> *Every w in W can be represented uniquely as xy with (a) $x \in W_T$, (b) $ty > y$ for all t in T, and (c) $\ell(w) = \ell(x) + \ell(y)$.*

If $T = \{s, t\}$ and $m_{s,t} < \infty$ then W_T contains $2m_{s,t}$ elements. All but one of them has a unique expression as a word. The exception is the longest element, which is

$$w_{\ell,T} = st\ldots = ts\ldots \quad (m_{s,t} \text{ terms on each side})\,.$$

This is called a **braid relation**. This longest element w is also singled out in W_T by the conditions $sw < w$, $tw < w$.

The following is a special case of a result of [Tits:1968].

1.1. Lemma. *If*

$$w = s_1 \ldots s_n = t_1 \ldots t_n$$

are two reduced expressions for w then one may be obtained from the other by a sequence of braid relations.

Proof. The proof is by induction on n. The cases $n = 1$ or 2 are trivial. So assume $n > 1$, and that

$$s_1 \ldots s_n = t_1 \ldots t_n$$

are reduced. If $s_1 = t_1$ we can cancel the common left factor and apply induction. Otherwise suppose $s = s_1 \neq t = t_1$, and in particular $n > 1$. Let

$$x = s_2 \ldots s_n, \quad y = t_2 \ldots t_n\,.$$

Then $sw < w$ so $w\alpha_s < 0$, and $tw < w$ so $w\alpha_t < 0$. If we write $w = uv$ with $u \in W_{s,t}$ and v such that $v\alpha_s > 0$, $v\alpha_t > 0$ then $su < u$, $tu < u$ hence $u = w_{\ell,s,t}$. We can write

$$w = sw_{s,t}z = tw_{t,s}z$$

where $sw_{s,t} = tw_{t,s} = w_{\ell,s,t}$. Since $sw_{s,t}z = ss_2 \ldots s_n$, we may cancel s and by induction obtain $w_{s,t}x$ from $s_2 \ldots s_n$ by a sequence of braid relations. Similarly for $w_{t,s}z$ and $t_2 \ldots t_n$. But then we can also obtain $sw_{s,t}z$ from $tw_{t,s}z$ by a single braid relation, so the Lemma is proved. $\square$.

If $\omega = s_1 \bullet \ldots \bullet s_n$ is a word in S, then I define the **closure** of ω to be the set of all words in the alphabet S that can be expressed as concatenations of ordered subwords of ω—i.e. the words determined by deleting some of its letters.

1.2. Proposition. *If w in W is represented by the reduced word ω, then the image in W of the words in the closure of ω depends only on w.*

This is called the **Bruhat closure** $\overline{w}$ of w. Another way to phrase it is that *the set of elements represented by subexpressions of a given reduced expression in W does not depend on the particular reduced expression.*

Proof. First suppose W to be generated by two elements s, t. Any reduced expression ω for an element w of W is an alternating product of s and t. By deleting successively elements at one end or another, one obtains all x with $\ell(x) < \ell(w)$. Since all elements in the closure of ω other than w have smaller length, the closure may be identified with all such x. Hence the Proposition is true in this case.

In general, by Lemma 1.1, it suffices to prove that the Proposition is true for two reduced expressions interchanged by a braid relation. But this will follow from the simple case of the Proposition in which S has two elements, which we have just seen.

What is not clear at the moment is that if x lies in $\overline{y}$ and y lies in $\overline{z}$ then x lies in $\overline{z}$, since deleting a generator might not leave behind a reduced expression.

Of course $x \le y$ if and only if $x^{-1} \le y^{-1}$. If W is finits and w_ℓ its longest element, then $x \to w_\ell x$ is an involution reversing closures.

2. Root reflections

If r is a root reflection, so is wrw^{-1}. If $w = urv$ then uv is what we get by deleting r. But $uv = uru^{-1} \cdot urv = uru^{-1}w$. Hence:

2.1. Proposition. *If $w = s_1 \ldots s_n$ is an expression for w as product of elements in S, then*

$$u = (s_1 \ldots s_{i-1}) \cdot (s_{i+1} \ldots s_n) = (s_1 \ldots s_{i-1}) \cdot s_i \cdot (s_{i-1} \ldots s_1) \cdot (s_1 \ldots s_n)$$

is of the form rw where r is a reflection in W.

Consequently, an element of W obtained from w by deleting terms in one of its reduced expressions may be expressed as a product $r_\ell \ldots r_1 w$ where each r_i is a reflection. As a partial converse:

2.2. Proposition. (Strong Exchange) *Let w be in W, $r = r_\lambda$ a root reflection with $\lambda > 0$. Then $\ell(rw) < \ell(w)$ if and only if $w^{-1}\lambda < 0$, and if $w = s_1 \ldots s_n$ then*

$$rw = s_1 \ldots s_{i-1} \cdot s_{i+1} \ldots s_n$$

for some intermediate s_i. If the expression for w is reduced, then s_i is unique.

Proof. Suppose the gallery $C, s_1C, \ldots, wC$ crosses the hyperplane $\lambda = 0$ in a wall labeled s_i. Then

$$rw = s_1 \ldots \widehat{s}_i \ldots s_n$$

for the usual geometric reasons, and $\ell(rw) < \ell(w)$.

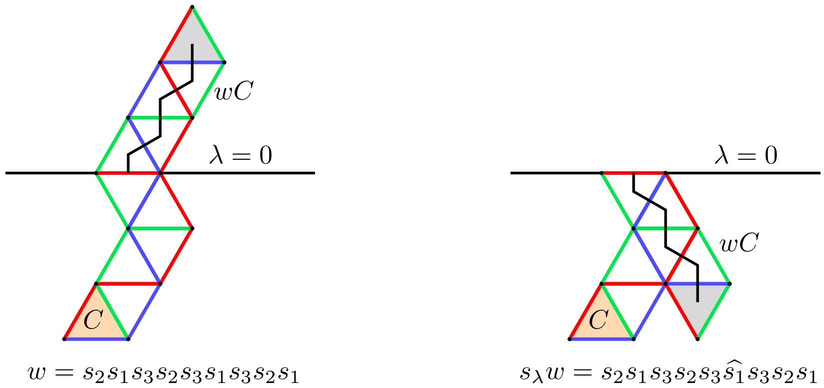

If we start with a reduced expression, the gallery crosses $\lambda = 0$ exactly once, guaranteeing uniqueness. If $w^{-1}\lambda > 0$, then $w^{-1}r^{-1}\lambda < 0$, so we can apply this argument to rwC.

If r is not in S, the reduced expression $s_1 \ldots \widehat{s}_i \ldots s_n$ may collapse further, as it does in the diagrams above.

Set $x \Leftarrow y$ if $\ell(x) < \ell(y)$ and $rx = y$ for some r in R, and define $x \le y$ to mean we can reach y from x by 0 or more such reflections. Since $wr = wrw^{-1} \cdot w$, it doesn't matter whether we use left or right multiplications by reflections in this definition. This order is called the **strong** or **Bruhat order**. I define the Bruhat graph to be that with elements of W as nodes and oriented edges $x \Leftarrow y$. The **closure** of y is the set of all $x \le y$, and if $x \le y$ the **interval** $[x, y]$ is the set of w with $x \le w \le y$.

Example. Let (W, S) be the dihedral group of order 8, with generators s, t. The following figure exhibits the Bruhat graph, with the orientation of every edge $x \Rightarrow y$ pointing down.

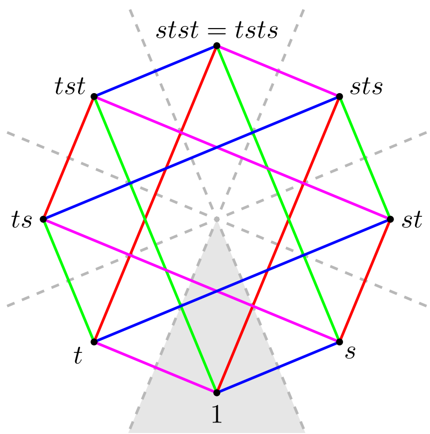

Here, $x < y$ if and only if $\ell(x) < \ell(y)$. All dihedral groups exhibit the same behaviour.

Strong Exchange implies that if $x = ry \Leftarrow y$ then a product expression for x may be obtained from a reduced expression for y by a single deletion. Repeating:

2.3. Proposition. *If $x < y$ then a product expression for x may be obtained by making one or more deletions in a reduced expression for y as a product of elements of S.*

That expression may not be reduced. We shall see later that the converse is also true.

3. The symmetric group

Let W be the symmetric group $\mathfrak{S}_n$, S the subset of $n-1$ elementary transpositions interchanging i and $i+1$. A permutation σ in $\mathfrak{S}_n$ is often expressed by the array $(\sigma(i))$. The reflections in W are the involutions $\langle i \,|\, j\rangle$ that swap two integers i and j. Multiplying σ on the right by $\langle i \,|\, j\rangle$ swaps the elements in the i-th and j-th places in the array of σ (some call it a **place permutation**), while multiplying by it on the left swaps the entries i and j in that array. I'll represent permutations by permutation matrices. In this scheme, a permutation is the horizontal array of colums in a matrix, an array of basis vectors. In column i is the basis element $\sigma_i = e_{\sigma(i)}$. Thus

$$[e_2, e_4, e_1, e_5, e_3] \cdot \langle 3 \,|\, 4\rangle = [e_2, e_4, e_5, e_1, e_3], \quad \langle 3 \,|\, 4\rangle \cdot [e_2, e_4, e_5, e_1, e_3] = [e_2, e_3, e_5, e_1, e_4] \,.$$

This description is consistent with the fact that right multiplication affects columns, left multiplication affects rows.

The length $\ell(\sigma)$ is the number of **inversions** in σ—the number of pairs (k, ℓ) with $k < \ell$ and $\sigma(k) > \sigma(\ell)$. It can be calculated by the well known `BubbleSort` algorithm.

Thus $\sigma \Leftarrow \tau$ if τ is obtained from σ by swapping σ_p and σ_q, where $p < q$ and $\sigma(p) < \sigma(q)$. For example, $[e_2, e_4, e_{\mathbf{1}}, e_{\mathbf{5}}, e_3] \prec [e_2, e_4, e_{\mathbf{5}}, e_{\mathbf{1}}, e_3]$.

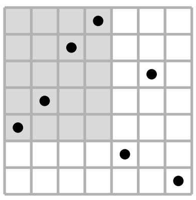

There is a simple criterion for determining whether $\sigma \le \tau$ or not. I imagine it is explained in many places in the literature, but I expand upon the short discussion in §3 of [Zhao:2007]. First of all, suppose σ to be in $\mathfrak{S}_n$. Associate to it the matrix $n(\sigma)$ with

$$n_{r,c} = \text{the number of non-zero entries in rows} \le r, \text{columns} \le c.$$

Thus in the figure at left $\sigma = (e_5, e_4, e_2, e_1, e_6, e_3, e_7)$ and $n_{5,4} = 4$.

Multiplying σ by $\langle k \,|\, \ell\rangle$ amounts to swapping columns k and ℓ in its matrix. How does that affect $n(\sigma)$?

3.1. Lemma. *Suppose*

$$\tau = \sigma \cdot \langle k \,|\, \ell\rangle$$

with $k < \ell$, $\sigma(k) < \sigma(\ell)$. Only the entries $n_{i,j}(\tau)$ for $k \le j < \ell$, $\sigma(k) \le i < \sigma(\ell)$ are different from $n_{i,j}(\sigma)$, and for these $n_{i,j}(\tau) = n_{i,j}(\sigma) - 1$.

In other words, going up in the Bruhat order decrements by 1 the entries in a certain rectangle in $n(\sigma)$.

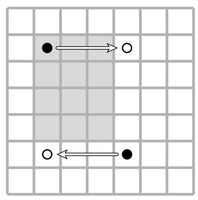

Proof. The figure at right should sufficiently explain why this is true.

If σ and τ are two permutations, define $\sigma \le_* \tau$ to mean that $n_{i,j}(\tau) \le n_{i,j}(\sigma)$ for all i, j.

3.2. Proposition. *If σ, τ are two permutations in $\mathfrak{S}_n$, then $\sigma \le \tau$ if and only if $\sigma \le_* \tau$.*

Proof. Lemma 3.1 implies immediately that if $\sigma < \tau$ then $\sigma \le_* \tau$.

Proving the other half of Proposition 3.2 is by induction on a certain measure of the difference between σ and τ, which I'll now define. Suppose σ and τ given in $\mathfrak{S}_n$. Let k be such that $\sigma(i) = \tau(i)$ for $i < k$, and then set

$$|\tau - \sigma| = \begin{cases} |\tau_k - \sigma_k| & \text{if } k \le n \\ 0 & \text{otherwise.} \end{cases}$$

Thus $|\sigma - \tau| = 0$ if and only if $\sigma = \tau$.

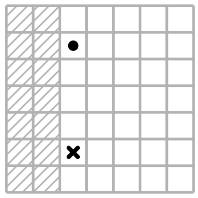

Now suppose that $\sigma \le_* \tau$. If $|\tau - \sigma| = 0$ there is nothing to prove. Otherwise, let k be smallest with $\sigma_k \ne \tau_k$. The assumption that $\sigma \le_* \tau$ implies that $\sigma_k < \tau_k$.

In the following figures, the region ▨ is where $\sigma = \tau$, σ is marked by $\boxed{\bullet}$, ρ by $\boxed{\circ}$, and τ by $\boxtimes$. In the figurea t right, $k = 3$.

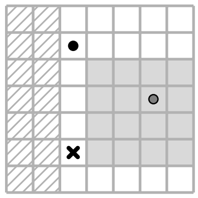

We shall find a permutation ρ with $\sigma \Leftarrow \rho \le_* \tau$ and apply the induction hypothesis. For this, we must find $\ell > k$ such that (a) $\sigma(k) < \sigma(\ell)$ and set $\rho = \sigma \cdot \langle k \,|\, \ell \rangle$ so as to have (b) $\rho \le_* \tau$.

Since $\rho(i) = \sigma(i) = \tau(i)$ for $i < k$, one thing condition (b) require is that $\rho(k) = \sigma(\ell) \le \tau(k)$. That is to say, we want to choose ℓ so that $(\ell, \sigma(\ell))$ lies in the grey region in the figure at right.

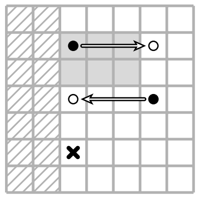

But in addition, Lemma 3.1 and condition (b) requires that $\pi_{i,j} = n_{i,j}(\sigma) - n_{i,j}(\tau) > 0$ in the region $k \le j < \ell$, $\sigma(k) \le i < \sigma(\ell)$, as at right. Why is it possible to find ℓ satisfying this?

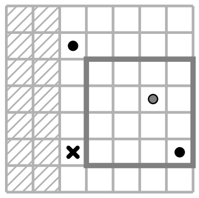

Since $\sigma \le_* \tau$, we must have the σ-entry in row τ_k to the right of column k. Therefore the set of (i, j) in the region $\sigma(k) < i \le \tau_k$, $k < j$ with $\sigma_{i,j} \ne 0$ is not empty. Choose ℓ the smallest value of j occurring.

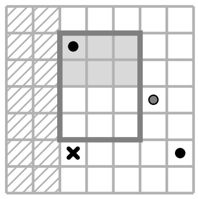

By definition of ℓ, the only non-zero entry $\sigma_{i,j}$ with $\sigma(k) \le i < \tau(k)$, $k \le j < \ell$ is at the upper left corner. I claim that $\pi_{i,j} > 0$ there, also. For if there were some $\pi_{i,j} = 0$ in that region, then $\pi_{\tau(k),j}$ would be < 0, a contradiction.

But since this region includes the region $\sigma(k) \le i < \sigma(\ell)$, $k \le j < \ell$, we are through.

This concludes the proof of Proposition 3.2.

Remark. Verifying the condition in this Proposition involves checking n^2 items. I do not know if there is a significantly more efficient test. There is, however, another test of roughly the same theoretical efficiency that is in practice a bit faster, at least for small n. For σ in $\mathfrak{S}_n$ and each $k \le n$ let $((\sigma_1, \dots, \sigma_k))$ be the sorted initial array of σ of length k. Then $\sigma \le \tau$ if and only if

$$((\sigma_1, \dots, \sigma_k)) \le ((\tau_1, \dots, \tau_k))$$

for each k, in the sense of term-by-term inequality. As explained in §2.6 of [Bjorner-Brenti:2005], there is a generalization of this to arbitrary Coxeter groups due to Deodhar.

4. Structure of the graph

Multiplication by s is an involution of the group. *How does this involution relate to the closure graph?* Very nicely. Multiplication takes takes edges to edges, if we neglect orientation, and in a very simple way:

4.1. Lemma. *Suppose s in S, $x \Leftarrow y$. Then exactly one of the following occurs:*

(a) $sx = y$, so that s reverses the edge in the strong Bruhat graph between them;
(b) s maps the edge $x \Leftarrow y$ to the edge $sx \Leftarrow sy$.

In other words, applying s to the edge doesn't reverse the orientation of the edge, unless it just exchanges its endpoints.

Proof. Suppose $x \Leftarrow y$, say $x = r_\lambda y$ with $\ell(x) < \ell(y)$, $\lambda > 0$. If $s = r$ then multiplication by s clearly reverses this edge, so suppose $s \neq r$.

Since $r_\lambda y < y$, Proposition 2.2 implies that $y^{-1}\lambda < 0$. But then

$$sx = sr_\lambda y = ss_\lambda s \cdot sy = s_{s\lambda} sy\,.$$

Since $r \neq s$, $s\lambda > 0$, so that $sx < sy$ if and only if $(sy)^{-1}s\lambda < 0$. But

$$(sy)^{-1}s\lambda = y^{-1}\lambda < 0\,.$$

Let me analyze the situation a bit more closely. Since $sx = sry = srs \cdot sy$, it suffices to show that if $s \neq r$ then $\ell(sx) < \ell(sy)$. (a) If $sy > y$ then $\ell(sy) > \ell(y) \geq \ell(sx)$, so this case is trivial. (b) If $\ell(sx) < \ell(x)$ the conclusion is also trivial. (c) If $\ell(x) < \ell(y) - 1$ then since $\ell(y) - \ell(x)$ has to be odd, $\ell(x) \leq \ell(y) - 3$, and again $\ell(sx) < \ell(sy)$ is easy. So it is the case $sy < y$, $sx > x$, $\ell(x) = \ell(y) - 1$, in other words the configuration in the following diagram, that has to be ruled out.

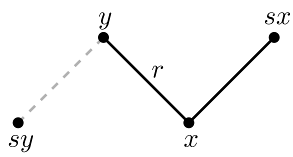

This is useful to keep in mind.

The content of the Proposition is that multiplication by s preserves the links in the strong Bruhat graph, but will reverse orientation in exactly one case, that of an edge between x and sx. There are thus essentially three kinds of edge-swaps, given and edge $x < y$: (a) an edge reverses itself; (b) $sx < y$ and $sx < x$; or (c) $sx < x$, $sy > y$. All three cases occur already for dihedral groups:

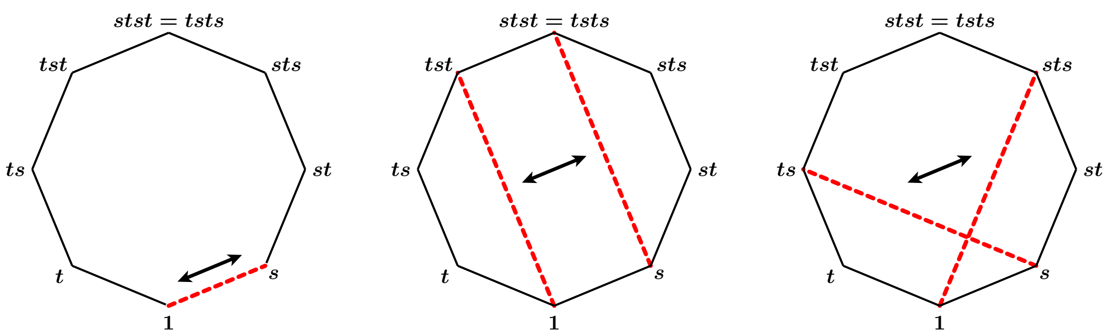

4.2. Theorem. *If $sy < y$ then the set of $x \leq y$ is stable under multiplication by s.*

Proof. By induction on the length n of a chain

$$x = x_0 \Leftarrow x_1 \Leftarrow \ldots \Leftarrow x_n = y$$

of minimal length. The case $n = 1$ follows immediately from the Lemma. Otherwise, say $n > 1$. The cases $sx < x$ and $sx = x_1$ are also immediate. Otherwise, $sx > x$, and we may apply induction to see that $sx_1 < y$. But then the Lemma implies that $sx < sx_1$, and we are again through.

4.3. Corollary. *Suppose $y = sx > x$, and let*

$$X = \{z \mid z \le x\}, \quad Y = \{z \mid z \le y\}.$$

Then $Y = X \cup sX$.

Of course these may overlap.

4.4. Corollary. *Suppose $x < y$, with $\ell(y) - \ell(x) = 2$, $sy < y$. Either $sx > x$ and $[x, y] = \{x, sx, sy, y\}$ or $sx < x$ and the interval $[x, y]$ is isomorphic to $[sx, sy]$.*

The diagrams above show that both possibilities can occur.

Proof. Since $x < y$, parity considerations require that the interval between x and y be filled with edges of length 1. If $[x, y] \ne \{x, sx, sy, y\}$ then there exists $x < z < y$ with $z \ne sx$, $z \ne sy$. In this case the Proposition implies that $sx < sz < sy$, and since $sy < y$ we must have $sx < x$. In particular $sx \notin [x, y]$.

Now there is a further dichotomy: either $sy \in [x, y]$ or not. In the second case, s is an isomorphism of $[x, y]$ with $[sx, sy]$. In the first case, the map $z \mapsto sz$, $sy \mapsto x$ is an isomorphism of $[x, y]$ with $[sx, sy]$.

4.5. Corollary. *Suppose $x < y$ and $\ell(x) = \ell(y) - 2$. Then there exist exactly two w with $x < w < y$.*

That is to say, the Bruhat interval $[x, y]$ in this case is very simple.

Proof. By induction on $\ell(y)$. The minimum this can be is 2, in which case $x = 1$, $y = st$, and $[x, y] = \{1, s, t, st\}$.

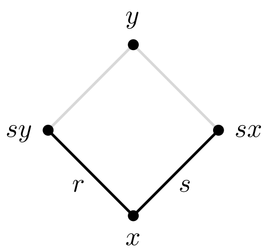

Otherwise, choose s with $sy < y$. If $sx < x$, then Corollary 4.4 tells us that $[x, y]$ is isomorphic to $[sx, sy]$, and we apply induction. If $sx > x$ the same result tells us $[x, y] = \{x, sx, sy, y\}$.

4.6. Corollary. *If x is obtained from y by one or more deletions in a reduced word for y, then $x < y$.*

Proof. What makes this not quite obvious is that one deletion might lead to a number of collapses of terms in the product.

The proof is by induction on $\ell(y)$. Suppose y has the reduced word $s_1 \ldots s_n$. Suppose first that s_1 is not one of the deleted items, so $x = s_1 w$ and w is obtained from $z = s_2 \ldots s_n$ by one or more deletions. By the induction hypothesis, $w < z$. Now we have $y = s_1 z > z$, $x = s_1 w$, with $w < z$. But then implies that $x < y$.

Otherwise, let $z = s_2 \ldots s_n$, so $y = s_1 z$ and x is obtained from z by deletions. By induction $x < z$ and $x < z < s_1 z = y$.

Combining this and Proposition 2.3:

4.7. Theorem. *For x and y in W, $x \le y$ if and only if x lies in the Bruhat closure of y.*

5. Maximal chains

There is one more thing to notice about the strong Bruhat graph of the dihedral group seen earlier. There is some redundancy in it, in the sense that there are more links than necessary to describe the partial order. For example, the reflection sts takes t to $stst$, so $t \le stst$. But this can be seen also by the chain t-ts-sts-$stst$. With the redundant links removed, the graph of the order looks like this:

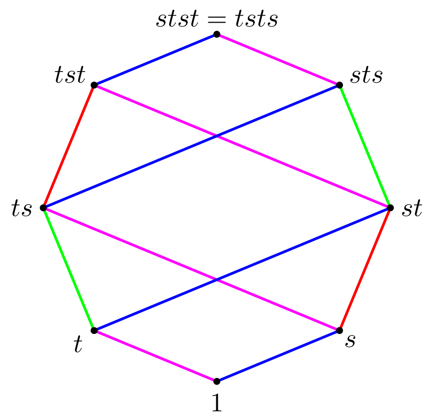

It is easy to see for all dihedral groups that the Bruhat order is generated by pairs $x = ry$ with $\ell(x) = \ell(y) - 1$. This is a general fact, and the second of the two most important results about Bruhat order.

Define $x \prec y$ to mean $x = ry < y$ and $\ell(y) - \ell(x) = 1$.

5.1. Theorem. *If $x < y$, then there exists a chain $x = x_0 \prec x_1 \prec \ldots \prec x_n = y$.*

This allows a very simple algorithmic description of closures. In the proof, I follow closely [Dixmier:1974], pp. 250–252.

Proof. We may assume that $x = ry < y$. We proceed by induction on $\ell(y) + \big(\ell(y) - \ell(x)\big)$. If $\ell(x) = \ell(y) - 1$, there is nothing to be proven. So we may assume $\ell(y) \ge \ell(x) + 3$.

Choose s with $sy < y$. Then $sx = sry = srs \cdot sy$ and

$$\ell(sx) < \ell(x) + 1 \le \ell(y) - 2 < \ell(y) - 1 = \ell(sy)\,.$$

So $sx < sy$. We may apply induction to get a chain from sx to sy:

$$sx = w_0 < w_1 < w_2 < \ldots < w_n = sy < w_{n+1} = y$$

with (say) $w_{i+1} = r_i w_i$. In particular, $r_n = s$.

• If $x < sx$, we can just extend the chain to include x:

$$x < sx = w_0 < w_1 < w_2 < \ldots < w_n = sy < w_{n+1} = y$$

• If $x > sx$ and $w_1 = x$, the chain we want is

$$x = w_1 < w_2 < \ldots < w_n = sy < w_{n+1} = y\,.$$

• Otherwise, $sx < x$ and $w_1 \ne x$. The situation is indicated by this diagram:

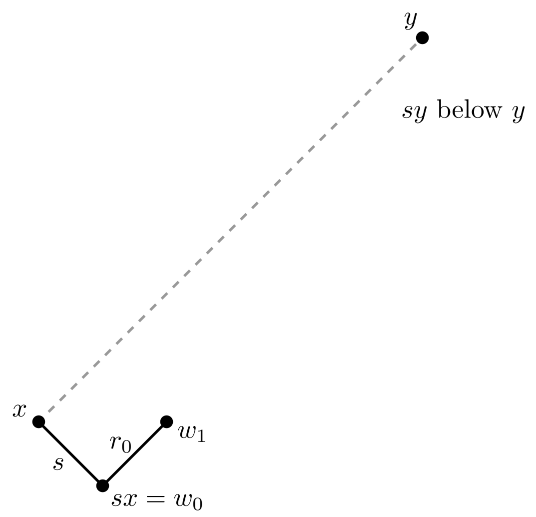

Let $t_0 = sr_0s$. Since $s \neq r_0$, we know that $sw_1 > w_1$ and that $t_0x = sw_1$, so we may fill in the diagram.

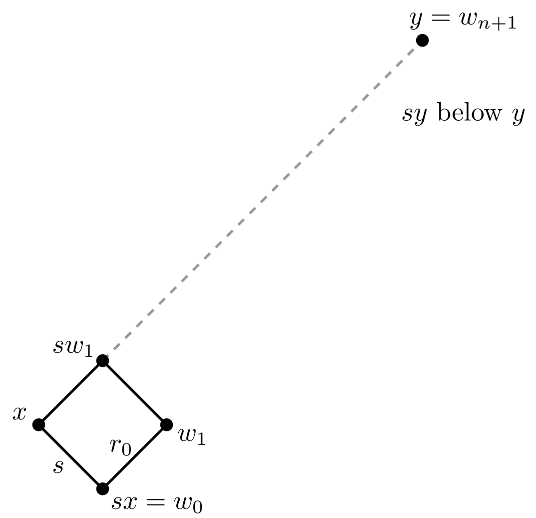

Since $\ell(y) - \ell(x) \geq 3$, $\ell(sw_1) = \ell(x) + 1 < \ell(y) - 1 = \ell(sy)$.

The diagram is not deceptive. According to . since $sy < s \cdot sy$ implies that since $w_1 < sy$ we also have $sw_1 < y$. The proof can be concluded by induction. But I'll continue the proof in a way that will suggest an efficient algorithm.

We may keep on filling in as long as $r_i \neq s$:

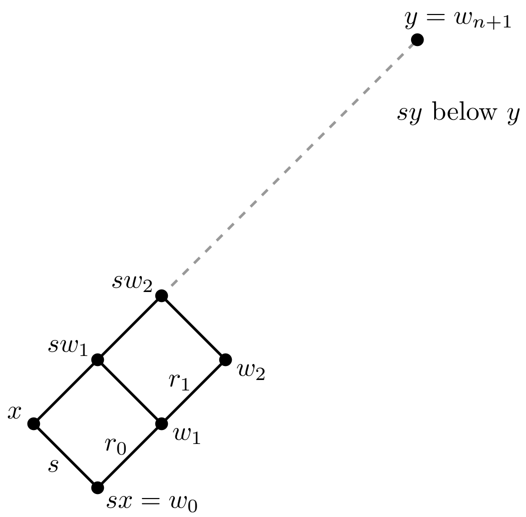

We have $r_n = s$; let i be least with $r_i = s$. So then we get a chain

$$x < sw_1 < sw_2 < \ldots < sw_i = w_{i+1} < w_{i+2} < \ldots sy < y$$

If $i = n$, the picture is this:

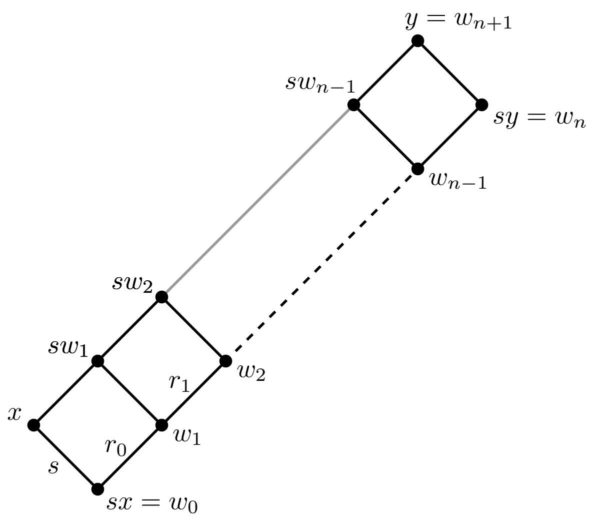

In this case, $sw_{n-1} < y$ by . But then $x < sw_1 < sw_2 < \ldots < sw_{n-1} < y$ is the chain we want. Otherwise $i < n$, and the picture is this:

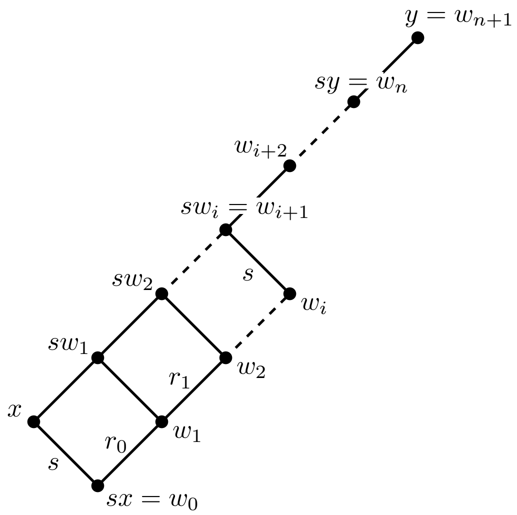

In this case, the chain is indicated in the diagram.

6. References

1. Anders Bjorner and Francesco Brenti, **Combinatorics of Coxeter groups**, Springer, 2005.
2. Jacques Dixmier, **Algèbres enveloppantes**, Gauthier-Villars, Paris, 1974.
3. James E. Humphreys, **Reflection groups and Coxeter groups**, Cambridge University Press, 1990.
4. Jacques Tits, 'Le problème de mots dans les groupes de Coxeter', *Symposia Math.* **1** (1968), 175–185.
5. Yufei Zhao, 'On the Bruhat order of the symmetric group and its shellability', preprint, 2007.