

Sortable elements and Cambrian lattices

NATHAN READING

ABSTRACT. We show that the Coxeter-sortable elements in a finite Coxeter group W are the minimal congruence-class representatives of a lattice congruence of the weak order on W . We identify this congruence as the Cambrian congruence on W , so that the Cambrian lattice is the weak order on Coxeter-sortable elements. These results exhibit W -Catalan combinatorics arising in the context of the lattice theory of the weak order on W .

1. Introduction

The weak order on a finite Coxeter group is a lattice [3] which encodes much of the combinatorics and geometry of the Coxeter group. The weak order has been studied in the special case of the permutation lattice and in the broader generality of the poset of regions of a simplicial hyperplane arrangement. With varying levels of generality, many lattice and order properties of this lattice have been determined. (See, for example, references in [7], [9], [14], [16] and [17].)

This paper continues a program, begun in [18], of applying lattice theory to gain new insights into the combinatorics and geometry of Coxeter groups. Specifically, we solidify the connection, first explored in [20], between the lattice theory of the weak order and the combinatorics of the W -Catalan numbers. These numbers count, among other things, the vertices of the (simple) generalized associahedron (a polytope which encodes the underlying structure of cluster algebras of finite type [10, 11]), the W -noncrossing partitions (which provide an approach [2, 6] to the geometric group theory of the Artin group associated to W) and the sortable elements [21] of W (which we discuss below).

In [20], the *Cambrian lattices* were defined as lattice quotients of the weak order on W modulo certain congruences, identified as the join (in the lattice of congruences of the weak order) of a small list of join-irreducible congruences. Any lattice quotient of the weak order defines [19] a complete fan which coarsens the fan defined by the reflecting hyperplanes of W , and in [20] it was conjectured that the fan associated to a Cambrian lattice is combinatorially isomorphic to the normal fan of the

Presented by P. Dehornoy.

Received December 22, 2005; accepted in final form March 19, 2006.

2000 *Mathematics Subject Classification*: 20F55, 06B10.

The author was partially supported by NSF grants DMS-0202430 and DMS-0502170.

corresponding generalized associahedron. In particular, each Cambrian lattice was conjectured to have cardinality equal to the W -Catalan number. These conjectures were proved for two infinite families of finite Coxeter groups (A_n and B_n).

The definition of *Coxeter-sortable elements* (or simply *sortable elements*) of W was inspired by the effort to better understand Cambrian lattices. Sortable elements were introduced in [21] and used to give a bijective proof that W -noncrossing partitions are equinumerous with vertices of the generalized associahedron. In this paper, we make explicit the essential connection between sortable elements and Cambrian lattices, proving in particular that the elements of the Cambrian lattice are counted by the W -Catalan number. The conjecture from [20] on the combinatorial isomorphism between Cambrian fans and cluster fans is proven in [22].

The construction (see Section 2) of sortable elements involves the choice of a Coxeter element c of W . For each c , the corresponding sortable elements are called *c-sortable*. The first main result of this paper is the following theorem.

Theorem 1.1. *Let c be a Coxeter element of a finite Coxeter group W . There exists a lattice congruence Θ_c of the weak order on W such that the bottom elements of the congruence classes of Θ_c are exactly the c -sortable elements. In particular, the weak order on c -sortable elements is a lattice quotient of the weak order on all of W .*

The phrase “bottom elements” in Theorem 1.1 refers to the fact that for any congruence on a finite lattice, each congruence class has a unique minimal element. (See Section 3). The corresponding quotient lattice is isomorphic to the subposet induced by the set of elements which are minimal in their congruence class. In the course of proving Theorem 1.1, we also establish the following results. (The map $w \mapsto ww_0$ in Proposition 1.3 is explained in Section 2.)

Theorem 1.2. *Let c be a Coxeter element of a finite Coxeter group W . The c -sortable elements constitute a sublattice of the weak order on W .*

Proposition 1.3. *The map $w \mapsto ww_0$ maps the congruence Θ_c to the congruence $\Theta_{c^{-1}}$. In particular, the weak order on c -sortable elements is anti-isomorphic to the weak order on c^{-1} -sortable elements.*

The second main result of the paper concerns the connection between sortable elements and the Cambrian lattices of [20]. These lattices were originally proposed as a simple construction valid for arbitrary Coxeter groups which, in the cases of types A and B, was known to reproduce the combinatorics and geometry of the generalized associahedra. However, the lattice-theoretic definition does not immediately shed light on the combinatorics of the lattice, so that in particular the definition was not proven, outside of types A and B and small examples, to relate to the combinatorics of W -Catalan numbers. The content of the following

theorem is that sortable elements provide a concrete combinatorial realization of the Cambrian lattices for any W .

Theorem 1.4. *The congruence Θ_c is the Cambrian congruence associated to c . In particular, the Cambrian lattice associated to c is the weak order on c -sortable elements.*

The proof of Theorem 1.4 is accomplished using the geometric model for lattice congruences of the weak order laid out in [18]. The results of this paper prove some of the conjectures of [20]. Many of the remaining conjectures are proved in [22]. In light of Theorem 1.4, Proposition 1.3 is equivalent to [20, Theorem 3.5].

We conclude the introduction with two examples which illustrate the results of the paper. The definitions underlying these examples appear in later sections.

Example 1.5. Consider the case $W = B_2$ with $S = \{s_0, s_1\}$, $m(s_0, s_1) = 4$ and $c = s_0s_1$. Figure 1.a shows the congruence Θ_c on the weak order on W . The shaded 3-element chain is a congruence class and each other congruence class is a singleton. Figure 1.b shows the subposet of the weak order induced by the sortable elements, or equivalently the bottom elements of congruence classes. The map $w \mapsto ww_0$ acts on the Hasse diagram in Figure 1.a by rotating through a half-turn. One easily verifies Theorem 1.2 and Proposition 1.3 in this example.




FIGURE 1. An example of Θ_c and the associated Cambrian lattice.

Example 1.6. A more substantive example is provided by $W = A_3$ with $S = \{s_1, s_2, s_3\}$, $m(s_1, s_2) = m(s_2, s_3) = 3$, $m(s_1, s_3) = 2$ and $c = s_2s_1s_3$. This Coxeter group is isomorphic to the symmetric group S_4 as generated by the simple transpositions $s_i = (i \ i+1)$. Figure 2 shows the congruence Θ_c on the weak order on A_3 . Each element is represented by its c -sorting word (see Section 2) including the inert dividers “|”, with the symbols s_i replaced by i throughout. (The identity

element is represented by the empty word.) Non-trivial Θ_c -classes are indicated by shading, and each unshaded element is the unique element in its congruence class. The antiautomorphism $w \mapsto ww_0$ corresponds to rotating the Hasse diagram through a half-turn.




FIGURE 2. Another example of Θ_c .

The content of Theorem 1.4 in this example is that Θ_c is the smallest congruence of the weak order on A_3 which sets $s_1 \equiv s_1s_2$ and $s_3 \equiv s_3s_2$. Combining Theorem 1.4 with Proposition 1.3, we obtain the assertion that Θ_c is the smallest congruence which sets $s_1s_3s_2s_3 \equiv s_1s_3s_2s_1s_3 \equiv s_1s_3s_2s_1$.

2. Sortable elements

In this section we quickly review the definition and first properties of Coxeter groups. (More detail, including proofs of assertions not proven here, can be found in the standard references [4, 5, 13].) We then review the definition of sortable elements and quote and prove results which are used in later sections.

A *Coxeter group* W is a group with a presentation as the group generated by a set S subject to the relations $(st)^{m(s,t)} = 1$ for $s, t \in S$. Here $m(s, s) = 1$ and $m(s, t) \geq 2$ for $s \neq t$, with $m(s, t) = \infty$ meaning that no relation of the form $(st)^m$ holds. It can be shown that $m(s, t)$ is the order of st in W . (*A priori*, we only know that the order of st divides $m(s, t)$.) The *rank* of W is $|S|$. A group can have more than one non-equivalent presentation of this form but, as usual, we take the term “Coxeter group” to imply a distinguished choice of S , called the *simple generators* of W . The presentation of W is encoded in the *Coxeter diagram* of W . This is a graph whose vertices are the simple generators, with edges $s - t$ whenever $m(s, t) > 2$. Edges are labeled by the number $m(s, t)$ except that, by convention, if $m(s, t) = 3$ then the edge is left unlabeled.

Each $w \in W$ can be written, in many different ways, as a word in the alphabet S . The smallest length of a word for w is called the *length* of w , denoted $l(w)$. A word of length $l(w)$ representing w is called *reduced*. For any $s \in S$, $l(sw) = l(w) \pm 1$ and $l(sw) < l(w)$ if and only if there is some reduced word for w starting with the letter s . The analogous statement holds for $l(ws)$ and $l(w)$.

Every Coxeter group has a *reflection representation*: a representation as a group generated by orthogonal reflections of a real vector space. A Coxeter group is finite if and only if it has such a representation in a vector space with a Euclidean inner product. An element $t \in W$ acts as an orthogonal reflection in such a representation if and only if it is conjugate to a simple generator. Thus the set $T = \{wsw^{-1} : w \in W, s \in S\}$ is called the set of *reflections* in W . The (*left*) *inversion set* of $w \in W$ is $\{t \in T : l(tw) < l(w)\}$. The length $l(w)$ is equal to $|I(w)|$. An element of W is uniquely determined by its inversion set.

The (*right*) *weak order* is the partial order on W whose cover relations are $w < ws$ whenever $l(w) < l(ws)$. Equivalently, $v \leq w$ if and only if $I(v) \subseteq I(w)$. Throughout this paper, “ \leq ” denotes the weak order and “ W ” denotes both the group W and the partially ordered set W . This partial order is a meet-semilattice in general, and a lattice exactly when W is finite [3]. For a simple generator s and an interval $[u, v]$ in the weak order, if $l(su) < l(u)$ then the involution $w \mapsto sw$ is an isomorphism between the intervals $[su, sv]$ and $[u, v]$. The minimal element of W is the identity 1 and when W is finite, the longest element w_0 is the unique maximal element of W . In this case, the map $w \mapsto ww_0$ is an antiautomorphism of weak order, because $I(ww_0) = T - I(w)$.

A *Coxeter element* c of W is an element represented by a (necessarily reduced) word $a_1 a_2 \cdots a_n$ where $S = \{a_1, \dots, a_n\}$ and $n = |S|$. Fix c and a particular word $a_1 a_2 \cdots a_n$ for c and write a half-infinite word

$$c^\infty = a_1 a_2 \cdots a_n | a_1 a_2 \cdots a_n | a_1 a_2 \cdots a_n | \cdots$$

The symbols “|” are inert “dividers” which facilitate the definition of sortable elements. When subwords of c^∞ are interpreted as expressions for elements of W , the dividers are ignored. The c -*sorting word* for $w \in W$ is the lexicographically first (as a sequence of positions in c^∞) subword of c^∞ which is a reduced word for w . The c -sorting word can be interpreted as a sequence of subsets of S : Each subset in the sequence is the set of letters of the c -sorting word which occur between two adjacent dividers.

An element $w \in W$ is c -*sortable* if its c -sorting word defines a sequence of subsets which is weakly decreasing under inclusion. Since any two reduced words for c are related by commutation of letters, the c -sorting words for w arising from different reduced words for c are related by commutations of letters, with no commutations across dividers. In particular, the set of c -sortable elements does not depend on the choice of reduced word for c . Examples 1.5 and 1.6 in Section 1 illustrate the definition of sortable elements. (In Example 1.6, the c -sortable elements are the elements at the bottom of their congruence class, including elements which are unique in their class.)

For any $J \subseteq S$, let W_J be the subgroup of W generated by J . Such subgroups are called *standard parabolic subgroups*. For any $w \in W$, there is a unique factorization $w = w_J \cdot {}^J w$ such that $w_J \in W_J$ and ${}^J w$ has $l(s \cdot {}^J w) > l({}^J w)$ for all $s \in J$. The element w_J has inversion set $I(w_J) = I(w) \cap W_J$. The element $w_J \cdot {}^J(w_0)$ has inversion set $I(w) \cup \{t \in T : t \notin W_J\}$. The map $w \mapsto w_J$ is a lattice homomorphism [18, Section 6] for any $J \subseteq S$. (That is, $(x \vee y)_J = (x_J) \vee (y_J)$ and similarly for meets.) This map is also compatible with the antiautomorphism $w \mapsto ww_0$ in the sense that $(ww_0)_J = w_J(w_0)_J$. The set W_J is a lower interval in the poset W with maximal element $(w_0)_J$. Most often, the subset J is $\langle s \rangle := S - \{s\}$ for some element s of S .

An *initial letter* of a Coxeter element c is a simple generator which is the first letter of some reduced word for c . Similarly a *final letter* of c is a simple generator which occurs as the last letter of some reduced word for c . If s is initial in c then scs is a Coxeter element and s is final in scs . The next two lemmas [21, Lemmas 2.4 and 2.5] constitute an inductive characterization of sortable elements. As a base for the induction, 1 is c -sortable for any c .

Lemma 2.1. *Let s be an initial letter of c and let $w \in W$ have $l(sw) > l(w)$. Then w is c -sortable if and only if it is an sc -sortable element of $W_{\langle s \rangle}$.*

Lemma 2.2. *Let s be an initial letter of c and let $w \in W$ have $l(sw) < l(w)$. Then w is c -sortable if and only if sw is scs -sortable.*

For the rest of the paper we confine our attention to the case where W is a finite Coxeter group. In particular, all proofs should be assumed to apply only to the finite case.

The remainder of the section is devoted to quoting or proving preliminary results. For any $J \subseteq S$, the *restriction* of a Coxeter element c to W_J is the Coxeter element for W_J obtained by deleting the letters $S - J$ from any reduced word for c . The next lemma is immediate from the definition of sortable elements and the proposition following it is [21, Corollary 4.5].

Lemma 2.3. *Let c be a Coxeter element of W , let $J \subseteq S$ and let c' be the Coxeter element of W_J obtained by restriction. If $w \in W_J$ is c' -sortable then w is c -sortable as an element of W .*

Proposition 2.4. *Let c be a Coxeter element of W , let $J \subseteq S$ and let c' be the Coxeter element of W_J obtained by restriction. If w is c -sortable then w_J is c' -sortable.*

A *cover reflection* of $w \in W$ is a reflection $t \in T$ such that $tw = ws$ with $l(ws) < l(w)$. The term “cover reflection” refers to the fact that $w > ws$ is a cover relation in the weak order. Equivalently, $t \in T$ is a cover reflection of w if and only if $I(w) - \{t\}$ is the inversion set of some $w' \in W$. In this case $w' = tw < w$. Let $\text{cov}(w)$ denote the set of cover reflections of w . The following proposition is one direction of the rephrasing of [21, Theorem 6.1] described in [21, Remark 6.9].

Proposition 2.5. *If x and y are c -sortable and $\text{cov}(x) = \text{cov}(y)$ then $x = y$.*

The next lemma follows from [21, Lemma 5.2] and [21, Theorem 6.1] as described in [21, Remark 6.9].

Lemma 2.6. *If s is initial in c and $x \in W_{\langle s \rangle}$ is sc -sortable then there exists a c -sortable element w with $\text{cov}(w) = \{s\} \cup \text{cov}(x)$.*

We conclude with four lemmas on the join operation as it relates to cover reflections and sortable elements. The first two lemmas are used only in the proof of the last two, which are applied in Section 3.

Lemma 2.7. *For $w \in W$, if s is a cover reflection of w and every other cover reflection of w is in $W_{\langle s \rangle}$ then $w = s \vee w_{\langle s \rangle}$.*

Proof. Since s is a cover reflection of w , it is in particular an inversion of w , and since $I(s) = \{s\}$, we have $s \leq w$. Any element w has $w_{\langle s \rangle} \leq w$, so w is an upper bound for s and $w_{\langle s \rangle}$. Since s is a cover reflection of w and every other cover reflection is in $W_{\langle s \rangle}$, any element x covered by w has either $s \not\leq x$ or $w_{\langle s \rangle} \not\leq x$ (the latter because $I(w_{\langle s \rangle}) = I(w) \cap W_{\langle s \rangle}$). Thus $w = s \vee w_{\langle s \rangle}$. \square

Lemma 2.8. *For $x \in W_{\langle s \rangle}$, $\text{cov}(s \vee x) = \text{cov}(x) \cup \{s\}$.*

Proof. Let $w = s \vee x$. First, we show that the reflection s is a cover reflection of w . If not, then let w' be any element covered by w and weakly above x . Since s is not

a cover reflection of w we have $s \in I(w')$, so w' is above both s and x , contradicting the fact that $w = s \vee x$.

Now let $t \neq s$ be a cover reflection of w , so that $I(w) - I(tw) = \{t\}$. If $t \notin W_{\langle s \rangle}$ then since $I(x) \subset I(w)$ and $I(x) \subseteq W_{\langle s \rangle}$ we also have $I(x) \subseteq I(tw)$. Since $s \neq t$ we have $s \leq tw$, and this contradicts the fact that $w = s \vee x$. This contradiction proves that $t \in W_{\langle s \rangle}$. Now, $I(w_{\langle s \rangle}) = I(w) \cap W_{\langle s \rangle}$ and $I((tw)_{\langle s \rangle}) = I(w) \cap W_{\langle s \rangle}$ and since $t \in W_{\langle s \rangle}$, $I(w_{\langle s \rangle}) - I((tw)_{\langle s \rangle}) = \{t\}$. Thus t is a cover reflection of $w_{\langle s \rangle}$. Applying the lattice homomorphism $y \mapsto y_{\langle s \rangle}$ to the equality $w = s \vee x$, we have $w_{\langle s \rangle} = x$.

Conversely suppose x covers tx . Let y be any element covered by w and weakly above $s \vee tx$. Then $y = t'w$ for some reflection t' . Since $I(x) \subset I(w)$, $I(x) \not\subseteq I(y)$ and $I(w) - I(y) = \{t'\}$, we must have $t' \in I(x)$. Since $I(tx) \subset I(s \vee tx) \subset I(y)$ we must have $t' \notin I(tx)$. But $I(x) - I(tx) = \{t\}$ so $t = t'$. Thus t is a cover reflection of w . \square

Lemma 2.9. *If s is initial in c and $x \in W_{\langle s \rangle}$ is sc -sortable then $s \vee x$ is c -sortable.*

Proof. Let s be initial in c , let $x \in W_{\langle s \rangle}$ be sc -sortable and let w be the c -sortable element, as in Lemma 2.6, such that $\text{cov}(w) = \{s\} \cup \text{cov}(x)$. By Lemma 2.7, $w = s \vee w_{\langle s \rangle}$. By Proposition 2.4, $w_{\langle s \rangle}$ is sc -sortable and by Lemma 2.8, $\text{cov}(w_{\langle s \rangle}) = \text{cov}(w) - \{s\} = \text{cov}(x)$. But Proposition 2.5 says that $w_{\langle s \rangle} = x$ and thus in particular $s \vee x = w$ is c -sortable. \square

Lemma 2.10. *If s is final in c and $w \in W$ is c -sortable with $l(sw) < l(w)$ then $w = w_{\langle s \rangle} \vee s$.*

Proof. By [21, Lemmas 6.6 and 6.7], s is a cover reflection of w and every other cover reflection is in $W_{\langle s \rangle}$. Thus Lemma 2.7 says that $w = w_{\langle s \rangle} \vee s$. \square

3. Sortable elements and lattice congruences

In this section we review the definition and an order-theoretic characterization of lattice congruences and prove Theorem 1.1, Theorem 1.2 and Proposition 1.3. We remind the reader that all proofs given here are valid only in the finite case.

A *congruence* on a lattice L is an equivalence relation Θ on L such that whenever $a_1 \equiv b_1$ and $a_2 \equiv b_2$ then $a_1 \wedge a_2 \equiv b_1 \wedge b_2$ and $a_1 \vee a_2 \equiv b_1 \vee b_2$. The quotient of L mod Θ is a lattice defined on the Θ -congruence classes. Denote the Θ -congruence class of a by $[a]_{\Theta}$ and set $[a_1]_{\Theta} \wedge [a_2]_{\Theta} = [a_1 \wedge a_2]_{\Theta}$ and $[a_1]_{\Theta} \vee [a_2]_{\Theta} = [a_1 \vee a_2]_{\Theta}$.

The following order-theoretic characterization of congruences of a finite lattice was introduced in [15]. It is a straightforward exercise to prove the characterization, which also follows from a characterization of congruences of general lattices due to Chajda and Snášel [8].

Proposition 3.1. *When L is finite, an equivalence relation on L is a lattice congruence if and only if it has the following three properties.*

- (i) *Every equivalence class is an interval.*
- (ii) *The downward projection $\pi_{\downarrow} : L \rightarrow L$, mapping each element to the minimal element in its equivalence class, is order-preserving.*
- (iii) *The upward projection $\pi_{\uparrow} : L \rightarrow L$, mapping each element to the maximal element in its equivalence class, is order-preserving.*

Furthermore the quotient of L mod Θ is isomorphic to the subposet of L induced by the set $\pi_{\downarrow}L$ of bottom elements of congruence classes and π_{\downarrow} is a homomorphism from L to $\pi_{\downarrow}L$.

The proof of Theorem 1.1 uses the characterization of congruences in Proposition 3.1. Specifically, we construct order-preserving maps π_{\downarrow}^c and π_c^{\uparrow} such that first, $\pi_{\downarrow}^c(x) = x$ if and only if x is c -sortable and second, the equivalence setting $x \equiv y$ if and only if $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^c(y)$ has equivalence classes of the form $[\pi_{\downarrow}^c(x), \pi_c^{\uparrow}(x)]$.

Some lemmas which contribute to the proof make use of induction on length and rank in a way that is similar to the proofs presented in [21]. In the course of these inductive proofs, we repeatedly apply the following fact: For $s \in S$ and $x, y \in W$, if $l(sx) < l(x)$ and $l(sy) < l(y)$ then $x \leq y$ if and only if $sx \leq sy$.

We also make use of the following notational convention: The explicit reference to c in π_{\downarrow}^c and π_c^{\uparrow} is used as a way of specifying a standard parabolic subgroup. For example, if s is initial in c then sc is a Coxeter element in the standard parabolic subgroup $W_{\langle s \rangle}$. Thus the notation π_{\downarrow}^{sc} refers to a map on $W_{\langle s \rangle}$.

We now construct a projection π_{\downarrow}^c . This is done by induction on the rank of W and on the length of the element to which π_{\downarrow}^c is applied. Define $\pi_{\downarrow}^c(1) = 1$ and for any initial letter s of c , define

$$\pi_{\downarrow}^c(w) = \begin{cases} s \cdot \pi_{\downarrow}^{scs}(sw) & \text{if } l(sw) < l(w), \\ \pi_{\downarrow}^{sc}(w_{\langle s \rangle}) & \text{if } l(sw) > l(w). \end{cases} \tag{3.1}$$

As written, each step of this inductive definition depends on a choice of an initial letter of a Coxeter element. The following proposition implies that π_{\downarrow}^c is well-defined and that $\pi_{\downarrow}^c \circ \pi_{\downarrow}^c = \pi_{\downarrow}^c$.

Proposition 3.2. *For any $w \in W$, $\pi_{\downarrow}^c(w)$ is the unique maximal c -sortable element weakly below w .*

Proof. Let $w \in W$ and let s be initial in c . If $l(sw) < l(w)$ then by induction $\pi_{\downarrow}^{scs}(sw)$ is the unique maximal scs -sortable element weakly below sw . Then $s \cdot \pi_{\downarrow}^{scs}(sw)$ is weakly below w , and by Lemma 2.2, $s \cdot \pi_{\downarrow}^{scs}(sw)$ is c -sortable. Let x be any c -sortable element below w . If $l(sx) < l(x)$ then sx is an scs -sortable element below sw , and therefore sx is below $\pi_{\downarrow}^{scs}(sw)$, so that x is below $s \cdot \pi_{\downarrow}^{scs}(sw) = \pi_{\downarrow}^c(w)$.

If $l(sx) > l(x)$ then by Lemmas 2.1 and 2.9, $x \vee s$ is c -sortable. Now $s \leq w$, so $x \vee s \leq w$, and since $x \vee s$ is shortened on the left by s , by the previous case $x \vee s$ is below $\pi_{\downarrow}^c(w)$, and therefore x is below $\pi_{\downarrow}^c(w)$.

If $l(sw) > l(w)$ then by Lemma 2.1, any c -sortable element x below w is an sc -sortable element of $W_{\langle s \rangle}$. In particular, $x \leq w_{\langle s \rangle}$. By induction on the rank of W , $\pi_{\downarrow}^c(w) = \pi_{\downarrow}^{sc}(w_{\langle s \rangle})$ is the unique maximal such. \square

Corollary 3.3. *The map π_{\downarrow}^c is order-preserving on the weak order.*

Proof. Suppose $x \leq y$. By Proposition 3.2, $\pi_{\downarrow}^c(x)$ is a c -sortable element below x and therefore below y . Thus Proposition 3.2 says that $\pi_{\downarrow}^c(x) \leq \pi_{\downarrow}^c(y)$. \square

We now prove Theorem 1.2, which asserts that c -sortable elements constitute a sublattice of the weak order.

Proof of Theorem 1.2. Let x and y be c -sortable. We first show that $x \wedge y$ is also c -sortable. Choose an initial letter s of c . If $l(sx) > l(x)$ and $l(sy) > l(y)$ then x and y are both in $W_{\langle s \rangle}$. By induction on the rank of W , $x \wedge y$ is an sc -sortable element of $W_{\langle s \rangle}$, so it is also c -sortable by Lemma 2.3.

If exactly one of x and y are shortened on the left by s , we may as well take $l(sx) < l(x)$ and $l(sy) > l(y)$. Then $y \in W_{\langle s \rangle}$. Since $W_{\langle s \rangle}$ is a lower interval in W and $w \mapsto w_{\langle s \rangle}$ is a lattice homomorphism,

$$x \wedge y = (x \wedge y)_{\langle s \rangle} = x_{\langle s \rangle} \wedge y_{\langle s \rangle} = x_{\langle s \rangle} \wedge y.$$

By Proposition 2.4, $x_{\langle s \rangle}$ is sc -sortable, so by the previous case $x \wedge y$ is c -sortable.

If $l(sx) < l(x)$ and $l(sy) < l(y)$ then by Lemma 2.2, sx and sy are both scs -sortable. By induction on length, $sx \wedge sy$ is scs -sortable. Since left multiplication by s is an isomorphism from the interval $[s, w_0]$ to the interval $[1, sw_0]$, $sx \wedge sy$ is lengthened on the left by s and $x \wedge y = s(sx \wedge sy)$. Now Lemma 2.2 says that $s(sx \wedge sy)$ is c -sortable.

We now show that $x \vee y$ is c -sortable. Since $x \vee y \geq x$ we have $\pi_{\downarrow}^c(x \vee y) \geq \pi_{\downarrow}^c(x)$ by Corollary 3.3. By Proposition 3.2 $\pi_{\downarrow}^c(x) = x$. Similarly $\pi_{\downarrow}^c(x \vee y) \geq y$, so $\pi_{\downarrow}^c(x \vee y)$ is an upper bound for x and y . By definition of join, $x \vee y \leq \pi_{\downarrow}^c(x \vee y)$, but Proposition 3.2 says that $\pi_{\downarrow}^c(x \vee y) \leq x \vee y$. Thus $x \vee y = \pi_{\downarrow}^c(x \vee y)$, which is c -sortable. \square

We continue towards a proof of Theorem 1.1 by defining the upward projection corresponding to π_{\downarrow}^c .

Call $w \in W$ c -antisortable if ww_0 is c^{-1} -sortable. Define an upward projection map π_c^{\uparrow} by setting $\pi_c^{\uparrow}(w) = (\pi_{\downarrow}^{(c^{-1})}(ww_0))w_0$. Since $w \mapsto ww_0$ is an antiautomorphism of the right weak order, it is immediate from Proposition 3.2 and Corollary 3.3 that $\pi_c^{\uparrow}(w)$ is the unique minimal element among c -antisortable elements above w , that $\pi_c^{\uparrow} \circ \pi_c^{\uparrow} = \pi_c^{\uparrow}$ and that π_c^{\uparrow} is order-preserving.

Lemma 3.4. *If s is a final letter of c then*

$$\pi_c^\uparrow(w) = \begin{cases} s \cdot \pi_{scs}^\uparrow(sw) & \text{if } l(sw) > l(w), \\ \pi_{cs}^\uparrow(w_{\langle s \rangle}) \cdot \langle s \rangle w_0 & \text{if } l(sw) < l(w). \end{cases}$$

Proof. If $l(sw) > l(w)$ then $l(sww_0) < l(ww_0)$, so

$$\pi_c^\uparrow(w) = s \cdot (\pi_\downarrow^{(sc^{-1})}(sww_0))w_0 = s \cdot \pi_{scs}^\uparrow(sw).$$

If $l(sw) < l(w)$ then $l(sww_0) > l(ww_0)$, so $\pi_c^\uparrow(w) = \pi_\downarrow^{(sc^{-1})}((ww_0)_{\langle s \rangle})w_0$. But $(ww_0)_{\langle s \rangle} = w_{\langle s \rangle}(w_0)_{\langle s \rangle}$, so

$$\pi_\downarrow^{(sc^{-1})}((ww_0)_{\langle s \rangle})w_0 = \pi_\downarrow^{(sc^{-1})}(w_{\langle s \rangle}(w_0)_{\langle s \rangle})(w_0)_{\langle s \rangle} \cdot \langle s \rangle w_0 = \pi_{cs}^\uparrow(w_{\langle s \rangle}) \cdot \langle s \rangle w_0.$$

□

In order to construct the congruence Θ_c , it is necessary to relate the fibers of π_\downarrow^c to those of π_c^\uparrow . This is done by induction on rank and length as in the proofs of Proposition 3.2 and Theorem 1.2. However, the recursive definition of π_\downarrow^c requires an initial letter of c and the corresponding property of π_c^\uparrow (Lemma 3.4) requires a final letter of c . Thus an additional tool is needed, and this tool is provided by the dual (Lemma 3.6) of the following lemma.

Lemma 3.5. *If s is a final letter of c and $l(sw) < l(w)$ then $\pi_\downarrow^c(w) = s \vee \pi_\downarrow^{cs}(w_{\langle s \rangle})$.*

Proof. Let $x = s \vee \pi_\downarrow^{cs}(w_{\langle s \rangle})$. By Lemma 2.3 and Proposition 3.2, $\pi_\downarrow^{cs}(w_{\langle s \rangle})$ is a c -sortable element. Thus both s and $\pi_\downarrow^{cs}(w_{\langle s \rangle})$ are c -sortable elements below w and therefore x is a c -sortable element below w by Theorem 1.2. (Note that the weaker result, Lemma 2.9, applies to an initial letter, and thus cannot be used to show that x is c -sortable.) By Lemma 3.2, $x \leq \pi_\downarrow^c(w)$. Since $l(sx) < l(x)$, $\pi_\downarrow^c(w)$ is also shortened on the left by s . Now Lemma 2.10 says that $\pi_\downarrow^c(w) = s \vee (\pi_\downarrow^c(w)_{\langle s \rangle})$. Because $\pi_\downarrow^c(w) \leq w$, we have $\pi_\downarrow^c(w)_{\langle s \rangle} \leq w_{\langle s \rangle}$. Since $\pi_\downarrow^c(w)_{\langle s \rangle}$ is a cs -sortable element below $w_{\langle s \rangle}$, it is below $\pi_\downarrow^{cs}(w_{\langle s \rangle})$ by Proposition 3.2. Therefore $x = s \vee \pi_\downarrow^{cs}(w_{\langle s \rangle}) \geq s \vee (\pi_\downarrow^c(w)_{\langle s \rangle}) = \pi_\downarrow^c(w)$. □

Lemma 3.6. *If s is an initial letter of c and $l(sw) > l(w)$ then $\pi_c^\uparrow(w) = sw_0 \wedge (\pi_{sc}^\uparrow(w_{\langle s \rangle}) \cdot \langle s \rangle w_0)$.*

Proof. Starting with Lemma 3.5, replace w by ww_0 and c by c^{-1} , multiply both sides on the right by w_0 and apply the fact that $w \mapsto ww_0$ is an antiautomorphism. □

We now relate π_\downarrow^c to π_c^\uparrow .

Proposition 3.7. *The maps π_\downarrow^c and π_c^\uparrow are compatible in the following senses.*

- (i) *For any $x, y \in W$, $\pi_\downarrow^c(x) = \pi_\downarrow^c(y)$ if and only if $\pi_c^\uparrow(x) = \pi_c^\uparrow(y)$.*
- (ii) *$\pi_c^\uparrow \circ \pi_\downarrow^c = \pi_c^\uparrow$ and $\pi_\downarrow^c \circ \pi_c^\uparrow = \pi_\downarrow^c$.*

Proof. By the antisymmetry $w \mapsto ww_0$, to prove (i) it suffices to prove the “only if” direction. We treat first the special case of (i) where $x \leq y$. Suppose $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^c(y)$ and $x \leq y$. Let s be initial in c . If $l(sx) < l(x)$ then since $x \leq y$, $l(sy) < l(y)$. Thus $\pi_{\downarrow}^c(x) = s \cdot \pi_{\downarrow}^{scs}(sx)$ and $\pi_{\downarrow}^c(y) = s \cdot \pi_{\downarrow}^{scs}(sy)$, so $\pi_{\downarrow}^{scs}(sx) = \pi_{\downarrow}^{scs}(sy)$. Since $sx \leq sy$, by induction on $l(x)$, $\pi_{\uparrow}^{scs}(sx) = \pi_{\uparrow}^{scs}(sy)$. By Lemma 3.4 (with c replaced by scs and w replaced by x or y) $\pi_{\uparrow}^{scs}(sx) = s \cdot \pi_c^{\uparrow}(x)$ and $\pi_{\uparrow}^{scs}(sy) = s \cdot \pi_c^{\uparrow}(y)$ so $\pi_c^{\uparrow}(x) = \pi_c^{\uparrow}(y)$.

If $l(sx) > l(x)$, we claim that $l(sy) > l(y)$. If not then $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^{sc}(x_{(s)})$ and $\pi_{\downarrow}^c(y) = s \cdot \pi_{\downarrow}^{scs}(sy)$. In particular, $\pi_{\downarrow}^c(x)$ is lengthened on the left by s but $\pi_{\downarrow}^c(y)$ is shortened on the left by s , contradicting the supposition that $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^c(y)$. This contradiction proves the claim. Thus $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^{sc}(x_{(s)})$ and $\pi_{\downarrow}^c(y) = \pi_{\downarrow}^{sc}(y_{(s)})$, so that $\pi_{\downarrow}^{sc}(x_{(s)}) = \pi_{\downarrow}^{sc}(y_{(s)})$. Since $x_{(s)} \leq y_{(s)}$, by induction on the rank of W $\pi_{\uparrow}^{sc}(x_{(s)}) = \pi_{\uparrow}^{sc}(y_{(s)})$. By Lemma 3.6, $\pi_c^{\uparrow}(x) = \pi_c^{\uparrow}(y)$.

Having established (i) in the case $x \leq y$, we now prove (ii). Any $y \in W$ has $\pi_{\downarrow}^c(y) \leq y$ and $\pi_{\downarrow}^c(\pi_{\downarrow}^c(y)) = \pi_{\downarrow}^c(y)$, so setting $x = \pi_{\downarrow}^c(y)$ in the special case of (i) already proved, $\pi_c^{\uparrow}(\pi_{\downarrow}^c(y)) = \pi_c^{\uparrow}(y)$. Thus $\pi_c^{\uparrow} \circ \pi_{\downarrow}^c = \pi_c^{\uparrow}$. The antisymmetry $w \mapsto ww_0$ now implies that $\pi_{\downarrow}^c \circ \pi_c^{\uparrow} = \pi_{\downarrow}^c$ as well.

Finally, we prove the general case of (i). If x and y are unrelated in the weak order and $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^c(y)$ then $\pi_c^{\uparrow}(x) = \pi_c^{\uparrow}(\pi_{\downarrow}^c(x)) = \pi_c^{\uparrow}(\pi_{\downarrow}^c(y)) = \pi_c^{\uparrow}(y)$. \square

We now prove the main theorem of the section.

Proof of Theorem 1.1. For each $w \in W$, let $D(w) = \{y \in W : \pi_{\downarrow}^c(y) = \pi_{\downarrow}^c(w)\}$ and let $U(w) = \{y \in W : \pi_c^{\uparrow}(y) = \pi_c^{\uparrow}(w)\}$. Since π_{\downarrow}^c is order-preserving, $\pi_{\downarrow}^c(w)$ is the unique minimal element of $D(w)$, and similarly, $\pi_c^{\uparrow}(w)$ is the unique maximal element of $U(w)$. Proposition 3.7 states that $D(w) = U(w)$, and since π_{\downarrow}^c is order preserving, $D(w)$ is the entire interval $[\pi_{\downarrow}^c(w), \pi_c^{\uparrow}(w)]$ in the right weak order.

Thus the fibers of the map π_{\downarrow}^c form a decomposition of the weak order on W into intervals and π_{\downarrow}^c and π_c^{\uparrow} are the order-preserving maps required in Proposition 3.1. Therefore these intervals are the congruence classes of some congruence Θ_c . \square

The preceding considerations also prove Proposition 1.3, as we now explain. The congruence classes of Θ_c are of the form $[\pi_{\downarrow}^c(w), \pi_c^{\uparrow}(w)]$. Applying the anti-automorphism $w \mapsto ww_0$ we obtain intervals of the form $[\pi_{\downarrow}^{(c^{-1})}(ww_0), \pi_{(c^{-1})}^{\uparrow}(ww_0)]$. Thus the anti-automorphism $w \mapsto ww_0$ takes Θ_c to $\Theta_{c^{-1}}$.

Remark 3.8. The c -sortable elements of a finite Coxeter group W are counted by the W -Catalan number, as is proven bijectively in two different ways in [21, Theorem 6.1] and [21, Theorem 8.1]. In particular, the number of c -sortable elements is independent of c . Lemma 2.2 gives a bijection between c -sortable elements shortened on the left by s and scs -sortable elements lengthened on the left by s .

Thus there is a bijection between c -sortable elements w with $l(sw) > l(w)$ and scs -sortable elements x with $l(sx) < l(x)$. Using Theorem 1.2 and Lemma 2.10, one can show that $w \mapsto s \vee w$ is such a bijection with inverse $x \mapsto (x)_{(s)}$.

We conclude the section with an easy technical lemma which is useful in the proof of Theorem 1.4.

Lemma 3.9. *Let s be initial in c and let $y \in W$ have $l(sy) < l(y)$. If $x < y$ but $x \neq sy$ then $\pi_{\downarrow}^c(x) = \pi_{\downarrow}^c(y)$ if and only if $\pi_{\downarrow}^{scs}(sx) = \pi_{\downarrow}^{scs}(sy)$.*

Proof. Note that s is an inversion of y and the inversion sets of x and y differ by one reflection. This reflection is not s because if so, we would have $x = sy$. Thus $l(sx) < l(x)$. So $\pi_{\downarrow}^c(x) = s \cdot \pi_{\downarrow}^{scs}(sx)$ and $\pi_{\downarrow}^c(y) = s \cdot \pi_{\downarrow}^{scs}(sy)$. \square

4. Lattice congruences of the weak order

In preparation for the proof (in Section 6) of Theorem 1.4, we review some well-known facts about congruences of finite lattices as well as some facts from [18] about the geometry underlying lattice congruences of the weak order on a finite Coxeter group. (For more information about lattice congruences in a general setting, see [12].) We also discuss an observation (Observation 4.7) which connects the geometry of lattice congruences to the action of a simple generator. We remind the reader that the weak order on a Coxeter group W is a lattice if and only if W is finite. Thus the results and observations of this section should be applied only in the context of finite W .

For a finite lattice L , let $\text{Con}(L)$ be the set of congruences on L , partially ordered by refinement. The poset $\text{Con}(L)$ is a distributive lattice, and thus is completely specified by the subposet $\text{Irr}(\text{Con}(L))$ induced by the join-irreducible congruences. (Recall that a join-irreducible element of a finite lattice is an element which covers exactly one element.) For a cover relation $x < y$ in L , a congruence is said to *contract* the edge $x < y$ if $x \equiv y$. For any cover relation $x < y$ there is a unique smallest (i.e., finest in refinement order) congruence $\text{Cg}(x < y)$ of L which contracts that edge, and this congruence is join-irreducible in $\text{Con}(L)$. Every join-irreducible in $\text{Con}(L)$ arises in this manner, and in fact, every join-irreducible congruence arises from a cover of the form $j_* < j$, where j is a join-irreducible element of L and j_* is the unique element of L covered by j . We say the congruence contracts j if it contracts the edge $j_* < j$ and write $\text{Cg}(j)$ for the join-irreducible congruence $\text{Cg}(j_* < j)$. A congruence is determined by the set of join-irreducible elements it contracts.

The weak order on a finite Coxeter group W has a property called congruence uniformity or boundedness [7], meaning in particular that the map Cg is a bijection from join-irreducibles of W to join-irreducibles of $\text{Con}(W)$. In what follows,

we tacitly use the map Cg to blur the distinction between join-irreducible congruences and join-irreducible elements of the weak order on W , so that $\text{Irr}(\text{Con}(W))$ is considered to be a partial order on the join-irreducibles of W . To distinguish this partial order from the weak order on W , we denote it “ \leq_{Con} ”. This partial order has the following description: $j_2 \leq_{\text{Con}} j_1$ if and only if every congruence contracting j_1 must also contract j_2 . Thus congruences of W are identified with order ideals of contracted join-irreducibles in $\text{Irr}(\text{Con}(W))$. (Order ideals in a poset are subsets I such that $x \in I$ and $y \leq x$ implies $y \in I$.)

We now review a geometric characterization of the partial order \leq_{Con} from [16]. Continuing under the assumption that W is finite, we fix a reflection representation of W as a group of orthogonal transformations of a real Euclidean vector space V of dimension d . Each reflection $t \in T$ acts on V as an orthogonal reflection and the *Coxeter arrangement* for W is the collection \mathcal{A} of reflecting hyperplanes for reflections $t \in T$. The hyperplanes in \mathcal{A} are permuted by the action of W . The *regions* of \mathcal{A} are the closures of the connected components of the complement $V - (\bigcup \mathcal{A})$ of \mathcal{A} . Choosing any region B to represent the identity element of W , the elements of W are in one-to-one correspondence $w \mapsto wB$ with the regions of \mathcal{A} . The inversion set of an element is the *separating set* of the corresponding region R (the set of reflections whose hyperplanes separate R from B). The weak order on W is thus containment order on separating sets.

Say a subset U of V is *above* a hyperplane H if every point in U is either contained in H or is separated from B by H . If U is above H and does not intersect H , then U is *strictly above* H . Similarly, U is *below* H if points in U are either contained in H or on the same side of H as B , and U is *strictly below* H if it is disjoint from and below H .

We say \mathcal{A}' is a *rank-two subarrangement* of \mathcal{A} if \mathcal{A}' consists of all the hyperplanes of \mathcal{A} containing some subspace of dimension $d - 2$ and $|\mathcal{A}'| \geq 2$. (The rank-two subarrangements of \mathcal{A} are exactly the collections of reflecting hyperplanes of the rank-two parabolic subgroups considered in [21].) For each rank-two subarrangement \mathcal{A}' , there is a unique region B' of \mathcal{A}' containing B . The two facet hyperplanes of B' are called *basic* hyperplanes in \mathcal{A}' . (The basic hyperplanes of \mathcal{A}' are the reflecting hyperplanes for the canonical generators of the corresponding rank-two parabolic subgroups, as defined in [21, Section 1].)

We cut the hyperplanes of \mathcal{A} into pieces called *shards* as follows. For each non-basic H in a rank-two subarrangement \mathcal{A}' , cut H into connected components by removing the subspace $\bigcap \mathcal{A}'$ from H . Equivalently, H is cut along its intersection with either of the basic hyperplanes of \mathcal{A}' . Do this cutting for each rank-two subarrangement, and call the closures of the resulting connected components of the hyperplanes *shards*. (In some earlier papers [16, 17] where shards were considered,

closures were not taken.) For illustrations of the shards for $W = A_3$ and $W = B_3$, see [18, Figures 1 and 3].

The shards of \mathcal{A} are important because of their connection to the join-irreducible elements of W . For any shard Σ , let $U(\Sigma)$ be the set of *upper elements* for Σ . That is, $U(\Sigma)$ is the set of regions of \mathcal{A} having a facet contained in Σ such that the region adjacent through that facet is lower (necessarily by a cover) in the weak order. We partially order $U(\Sigma)$ as an induced subposet of the weak order. By [17, Proposition 2.2], an element of W is join-irreducible in the weak order if and only if it is minimal in $U(\Sigma)$ for some shard Σ . Thus, to every join-irreducible j in W we associate a shard Σ_j . In fact, the map $j \mapsto \Sigma_j$ is a bijection [18, Proposition 3.5] between join-irreducibles of the weak order and shards in \mathcal{A} . The inverse map is written $\Sigma \mapsto j_\Sigma$. We now use this bijection to describe $\text{Irr}(\text{Con}(W))$ in terms of shards.

For each shard Σ , write H_Σ for the hyperplane in \mathcal{A} containing Σ . Define the *shard digraph* $Sh(W)$ to be the directed graph whose vertices are the shards, and whose arrows are as follows: Given shards Σ_1 and Σ_2 , let \mathcal{A}' be the rank-two subarrangement containing H_{Σ_1} and H_{Σ_2} . There is a directed arrow $\Sigma_1 \rightarrow \Sigma_2$ if and only if

- (i) H_{Σ_1} is basic in \mathcal{A}' but H_{Σ_2} is not, and
- (ii) $\Sigma_1 \cap \Sigma_2$ has dimension $d - 2$.

As explained in [16, Section 8], this directed graph is acyclic. The following is a special case of one of the assertions of [16, Theorem 25].

Theorem 4.1. *The poset $\text{Irr}(\text{Con}(W))$ is isomorphic to the transitive closure of $Sh(W)$ via the bijection $j \mapsto \Sigma_j$.*

We interpret the transitive closure of $Sh(W)$ as a poset by the rule that an arrow “ \rightarrow ” corresponds to an order relation “ \geq ”. In other words two join-irreducibles j_1 and j_2 in the weak order on W have $j_2 \leq_{\text{Con}} j_1$ if and only if there is a directed path in $Sh(W)$ from Σ_{j_1} to Σ_{j_2} .

Example 4.2. For W of rank two and $S = \{s, t\}$ the partial order \leq_{Con} on join-irreducibles is illustrated in Figure 3. (See Figure 1.a in Section 1, which shows the weak order on the rank-two Coxeter group B_2 .)

The following is a special case of [18, Lemma 3.9].

Lemma 4.3. *Let Σ be a shard. The following are equivalent:*

- (i) Σ is a source in $Sh(W)$.
- (ii) Σ consists of an entire facet hyperplane of B (the region representing the identity element of W).
- (iii) There is no facet of Σ intersecting the region for j_Σ in dimension $d - 2$.




FIGURE 3. $\text{Irr}(\text{Con}(W))$ for W of rank 2.

For $J \subseteq S$, let \mathcal{A}_J be the hyperplane arrangement associated to the standard parabolic subgroup W_J . We think of the arrangement \mathcal{A}_J as a subset of the arrangement \mathcal{A} by letting W_J inherit the reflection representation fixed for W . Recall the notation $\langle s \rangle = S - \{s\}$. For a join-irreducible j in W , write H_j for the hyperplane containing Σ_j . This is also the hyperplane separating j from j_* , the unique element covered by j . The next two lemmas are special cases of [18, Lemma 6.6] and [18, Lemma 6.8] respectively.

Lemma 4.4. *For $s \in S$, suppose $H_1 \in (\mathcal{A} - \mathcal{A}_{\langle s \rangle})$ and $H_2 \in \mathcal{A}_{\langle s \rangle}$. Let \mathcal{A}' be the rank-two subarrangement containing H_1 and H_2 . Then $(\mathcal{A}' \cap \mathcal{A}_{\langle s \rangle}) = \{H_2\}$ and H_2 is basic in \mathcal{A}' .*

Lemma 4.5. *Let $s \in S$ and let j be a join-irreducible. Then $H_j \in \mathcal{A}_{\langle s \rangle}$ if and only if $j \in W_{\langle s \rangle}$.*

In what follows, note that each $s \in S$ is a join-irreducible element of W and H_s (the hyperplane separating s from the unique element 1 covered by s) happens to be the reflecting hyperplane for the reflection s .

Lemma 4.6. *Let $s \in S$ and suppose H is a hyperplane in \mathcal{A} such that $H \notin \{H_r : r \in S\}$. Then the following are equivalent:*

- (i) *H is not basic in the rank-two subarrangement containing H and H_s but is basic in every other rank-two subarrangement containing H .*
- (ii) *There are exactly two shards in H , and their intersection is $H \cap H_s$.*
- (iii) *$H \in \mathcal{A}_{\{r,s\}}$ for some $r \in S$ with $r \neq s$.*

Proof. The equivalence of (i) and (ii) is immediate from the definition of shards.

Suppose H satisfies (i) and (ii), but not (iii). Let Σ be the shard in H below H_s , let j be the unique minimal element of $U(\Sigma)$ and let $r \in S$ have $r \neq s$. Since (iii) fails, H_r, H_s and H are not in the same rank-two subarrangement. Thus H_r intersects the interior of Σ . Since H is basic in the rank-two subarrangement it shares with H_r , there is an element x of $U(\Sigma)$ such that the region for x is not above H_r . Therefore j is not above r in the weak order, for each $r \in S$ with $r \neq s$.

But also j is not above s because Σ is below the hyperplane H_s . The only element of the weak order not above some element of S is the identity 1. But 1 is not join-irreducible, and this contradiction shows that (i) implies (iii).

Conversely, suppose H satisfies (iii). Then the rank-two subarrangement containing H and H_s is $\mathcal{A}_{\{r,s\}}$. By hypothesis, H is not basic in $\mathcal{A}_{\{r,s\}}$, because the basic hyperplanes of $\mathcal{A}_{\{r,s\}}$ are H_r and H_s . Every other rank-two parabolic \mathcal{A}' containing H also contains at least one hyperplane not in $\mathcal{A}_{\{r,s\}}$, so by Lemma 4.4, H is basic in \mathcal{A}' . Thus H satisfies (i). \square

We conclude this section by discussing an observation that is crucial to the proof of Theorem 1.4.

Observation 4.7. Let H be a hyperplane distinct from H_s , let $sH \in \mathcal{A}$ be the hyperplane obtained from H by reflecting by s , and let \mathcal{A}' be the rank-two subarrangement containing H , sH and H_s . Then the decomposition of sH into shards agrees (via the reflection s) with the decomposition of H into shards except that one of H and sH may be cut by H_s while the other may not. This exception occurs precisely when one of H and sH is basic in \mathcal{A}' but the other is not. (If both H and sH are basic in \mathcal{A}' then necessarily $H = sH$.)

Observation 4.7 is readily justified by the definition of shards. Consider any rank-two subarrangement \mathcal{A}' containing H . The case where \mathcal{A}' also contains H_s is depicted in the three illustrations in Figure 4. Each illustration shows some of the hyperplanes in \mathcal{A} and indicates the position, with respect to these hyperplanes, of the region for B . We insert space between two shards in the same hyperplane to show where the cut is. The hyperplanes H and sH are in black while the other hyperplanes are gray. In the first two illustrations, the case $|\mathcal{A}'| = 5$ is depicted.




FIGURE 4. Illustrations for Observation 4.7.

If neither H nor sH is basic in \mathcal{A}' then both H and sH are cut by H_s (possibly with $H = sH$) as exemplified in Figure 4.a. If exactly one of the two is basic in \mathcal{A}' then only the one that is non-basic is cut by H_s , as exemplified in Figure 4.b. If both are basic then since H_s is also basic and \mathcal{A}' has only two basic hyperplanes, Figure 4.c must apply. These considerations explain how and why the shard decompositions of H and sH may or may not differ where these hyperplanes intersect H_s .




FIGURE 5. Another illustration for Observation 4.7.

To see why the shard decompositions of H and sH must agree away from H_s , consider a rank-two subarrangement \mathcal{A}' containing H but not H_s . Let $s\mathcal{A}'$ be the rank-two subarrangement obtained from \mathcal{A}' by the reflection s . Since H_s is the only hyperplane separating the region B (representing 1) from the region sB (representing s) and $H_s \notin \mathcal{A}'$, B and sB are contained in the same \mathcal{A}' -region B' . The reflection s maps B to sB , so B is contained in the $(s\mathcal{A}')$ -region sB' . Thus the reflection s maps the basic hyperplanes of \mathcal{A}' to the basic hyperplanes of $s\mathcal{A}'$. (This is illustrated in Figure 5, where B' is shaded gray and sB' is shaded in stripes.) We conclude that H is non-basic in \mathcal{A}' (and is therefore cut by the basic hyperplanes of \mathcal{A}') if and only if sH is non-basic in $s\mathcal{A}'$ (and is therefore cut by the basic hyperplanes of $s\mathcal{A}'$).

5. Cambrian lattices

In this section we define Cambrian congruences and Cambrian lattices and prove two lemmas which are useful in the proof of Theorem 1.4. We continue to restrict our attention to the case of a finite Coxeter group W .

An *orientation* of the Coxeter diagram for W is obtained by replacing each edge of the diagram by a single directed edge, connecting the same pair of vertices in either direction. Orientations of the Coxeter diagram correspond to Coxeter elements (cf. [23]). Specifically, for each pair of noncommuting simple generators s and t , the edge $s - t$ is oriented $s \rightarrow t$ if and only if s precedes t in every reduced word for c . Each directed edge in the orientation corresponds to a pair of elements which are required to be congruent in the Cambrian congruence. Specifically, if $s \rightarrow t$ then the requirement is that the element t be congruent to the element with reduced word $tsts \cdots$ of length $m(s, t) - 1$. (Recall that $m(s, t)$ is the order of the product st in W .) The *Cambrian congruence* associated to c is the smallest (i.e., finest as a partition) lattice congruence satisfying this requirement for each directed edge. The Cambrian lattice associated to c is the quotient of the weak order on W modulo the Cambrian congruence. For brevity, we refer to these as the c -Cambrian congruence and the c -Cambrian lattice. Examples 1.5 and 1.6 illustrate these definitions.

The definition of the Cambrian congruence can be rephrased as follows: For each directed edge $s \rightarrow t$, we require that the join irreducibles $ts, tst, \dots, tsts, \dots$ of lengths 2 to $m(s, t) - 1$ be contracted. The join-irreducibles thus required to be contracted are called the *defining join-irreducibles* of the c -Cambrian congruence. Requiring that these join-irreducibles be contracted specifies an order ideal of join-irreducibles under the partial order \leq_{Con} , namely the smallest order ideal containing the defining join-irreducibles. This order ideal, interpreted as an order ideal in $\text{Irr}(\text{Con}(W))$, specifies the congruence.

Example 5.1. For W of rank two and $S = \{s, t\}$ the partial order \leq_{Con} on join-irreducibles is illustrated in Figure 3 in Section 4. The defining join-irreducibles for the Cambrian congruence for $c = st$ are the join-irreducibles of length greater than one whose unique reduced word starts with t . This set of join-irreducibles is an order ideal. Thus the Cambrian congruence has the interval $[t, tsts \cdots]$ as its only nontrivial congruence class, so that the bottom elements of the congruence classes are exactly the c -sortable elements (cf. Example 1.5).

We conclude this section with two lemmas which contribute to the proof of Theorem 1.4.

Lemma 5.2. *Let W_J be the standard parabolic subgroup generated by $J \subseteq S$ and let c' be the restriction of c to W_J . Then $x \in W_J$ is contracted by the c -Cambrian congruence if and only if it is contracted by the c' -Cambrian congruence.*

Proof. Recall that a lattice congruence is uniquely determined by the set of join-irreducibles it contracts. As a special case of [18, Lemma 6.12], the identity map embeds $\text{Irr}(\text{Con}(W_J))$ as an induced subposet of $\text{Irr}(\text{Con}(W))$, and the complement of this induced subposet is an order ideal. The defining join-irreducibles of the c' -Cambrian congruence are exactly those defining join-irreducibles of the c -Cambrian congruence which happen to be in W_J . Thus a join-irreducible in W_J is below (in \leq_{Con}) some defining join-irreducible of the c -Cambrian congruence if and only if it is below some defining join-irreducible of the c' -Cambrian congruence. \square

The *degree* of a join-irreducible j is the cardinality of the smallest $J \subseteq S$ such that $j \in W_J$. By [18, Lemma 6.12], if $j_2 \leq_{\text{Con}} j_1$ in $\text{Irr}(\text{Con}(W))$ then the degree of j_1 is less than or equal to the degree of j_2 .

Lemma 5.3. *Let s be initial in c . If j is join-irreducible with $l(sj) > l(j)$ but $j \notin W_{\langle s \rangle}$ then j is contracted by the c -Cambrian congruence.*

Proof. This is a modification of the proof of [18, Theorem 6.9]. We argue by induction on the dual of $\text{Irr}(\text{Con}(W))$, the base case being where j is of degree one or two. The case where j is of degree one is vacuous. If j is of degree two then j is a defining join-irreducible of the c -Cambrian lattice.

Let j be a join-irreducible of degree more than two satisfying the hypotheses of the lemma. To accomplish the inductive proof, we need only find a join-irreducible j' above j in $\text{Irr}(\text{Con}(W))$ such that $l(sj') > l(j')$ but $j' \notin W_{\langle s \rangle}$. Let H stand for H_j and let Σ be Σ_j . Since j is of degree more than two and each $s \in S$ has degree one, H is not a facet hyperplane of B (the region associated to 1).

First, consider the case where H is basic in the rank-two subarrangement \mathcal{A}' containing H and H_s . By Lemma 4.3 it has a facet, as a polyhedral subset of H , and moreover, there is a facet of Σ intersecting the region for j in dimension $d - 2$. There are two hyperplanes in \mathcal{A} which define this facet, and since H_j is basic in \mathcal{A}' , neither of these two hyperplanes is H_s . By Lemma 4.5, at least one of the two hyperplanes (call it H') is not in $\mathcal{A}_{\langle s \rangle}$. Some shard Σ' contained in H' arrows Σ in $\text{Sh}(W)$ and thus intersects the region for j in dimension $d - 2$. Since the region for j is below H_s , there is a region in $U(\Sigma')$ which is below H_s . In particular, $l(sj_{\Sigma'}) > l(j_{\Sigma'})$, and since $H' \notin \mathcal{A}_{\langle s \rangle}$, Lemma 4.5 says $j' \notin W_{\langle s \rangle}$. Thus $j_{\Sigma'}$ is the desired j' in this case.

Next, consider the case where H is not basic in the rank-two subarrangement containing H and H_s . In particular, Σ is below H_s . By Lemmas 4.3 and 4.6 (since j has degree greater than two), Σ has a facet which is not defined by H_s . As

in the previous case, this facet is defined by at least one hyperplane H' not in $\mathcal{A}_{(s)}$. Some shard Σ' in H' arrows Σ , and since every region in $U(\Sigma)$ is below H_s , the join-irreducible j' associated to Σ' has $l(sj') > l(j')$. Finally, j' is not in $W_{(s)}$ by Lemma 4.5. \square

6. Sortable elements and Cambrian congruences

In this section, we prove Theorem 1.4, which states that the Cambrian congruence associated to c coincides with the congruence Θ_c of Theorem 1.1 and that therefore the associated Cambrian lattice is the restriction of the weak order to c -sortable elements. Part of the proof is accomplished by Theorem 1.1. The strategy of the remainder of the proof is to relate the Cambrian congruences to the recursive structure of c -sortable elements as described in Lemmas 2.1 and 2.2.

The definition of the Cambrian congruence can be further restated as follows: The Cambrian congruence is the unique smallest congruence contracting all non- c -sortable join-irreducibles of degree 2. (In fact, for any $s, t \in S$, every non- c -sortable element of $W_{\{s,t\}}$ is join-irreducible.) The congruence Θ_c contracts all non- c -sortable join-irreducibles, and thus in particular contracts all non- c -sortable join-irreducibles of degree 2. Therefore, by the definition of the Cambrian congruence, Θ_c is a weakly coarser congruence than the c -Cambrian congruence. Since Θ_c does not contract any c -sortable join-irreducibles, we have the following corollary of Theorem 1.1:

Corollary 6.1. *For a finite Coxeter group W and any Coxeter element c , the c -Cambrian congruence does not contract any c -sortable join-irreducibles.*

In light of Corollary 6.1 and the fact that a congruence is determined by the join-irreducibles it contracts, Theorem 1.4 is equivalent to the following assertion.

Proposition 6.2. *For a finite Coxeter group W and any Coxeter element c , the c -Cambrian congruence contracts every non- c -sortable join-irreducible.*

Proof. Let j be a non- c -sortable join-irreducible in W and fix an initial letter s of c . We argue by induction on the length of j , on the rank of W and on the dual of the poset $\text{Irr}(\text{Con}(W))$ that j is contracted by the c -Cambrian congruence.

We first discuss the bases of the induction. The proposition is trivial when $\text{rank}(W)$ is 0 or 1, and is immediate from the definition for $\text{rank}(W) = 2$. The base of the induction on length is the case $l(j) = 2$. The proposition holds in this case because any non- c -sortable element of length 2 is a defining join-irreducible for the c -Cambrian congruence. The base of the induction on the dual of $\text{Irr}(\text{Con}(W))$ is the case where j is a defining join-irreducible (possibly of length greater than 2).

In this case the proposition holds by definition. Suppose now that j fits none of these base cases.

Consider the case where $l(sj) > l(j)$ and break into two subcases. If $j \in W_{\langle s \rangle}$ then by Lemma 2.3, j is not sc -sortable. By induction on rank, j is contracted by the sc -Cambrian congruence, and by Lemma 5.2, j is contracted by the c -Cambrian congruence. If $j \notin W_{\langle s \rangle}$ then by Lemma 5.3, j is contracted by the c -Cambrian congruence.

For the remaining case, $l(sj) < l(j)$, we apply Observation 4.7 and induction on length and on the dual of $\text{Irr}(\text{Con}(W))$. To avoid repetitions of the phrase “the region associated to,” we identify each group element w with its corresponding region wB .

Since s is c -sortable and j is not, $s \neq j$, so that $sj \neq 1$. Therefore sj is join-irreducible in light of the isomorphism between $[s, w_0]$ and $[1, sw_0]$. Lemma 2.2 implies that sj is not scs -sortable, and by induction on length, sj is contracted by the scs -Cambrian congruence. Let Σ and H be the shard and hyperplane for sj and let \mathcal{A}' be the rank-two subarrangement containing H and H_s . Since sj is below H_s and $sj \in U(\Sigma)$, the shard for Σ contains points strictly below H_s .

First consider the case where Σ also contains a point strictly above H_s . We first claim that \mathcal{A}' contains at least one additional hyperplane. To prove the claim, suppose to the contrary that $|\mathcal{A}'| = 2$, as illustrated in Figure 6. In light of Observation 4.7 (particularly as illustrated in Figure 4.c) Σ is fixed, as a set, by s . But then j is also in $U(\Sigma)$ so that by the uniqueness of join-irreducibles in $U(\Sigma)$, we must have $j = sj$, which is absurd. This proves the claim.




FIGURE 6. An impossible case in the proof of Proposition 6.2.

Let $s\Sigma$ be the image of Σ under the reflection s . Since Σ contains both points above H_s and points below H_s , the hyperplane H is basic in \mathcal{A}' . The other basic hyperplane is H_s , so in light of Observation 4.7 and the claim of the previous paragraph, $s\Sigma$ is the union of two shards, one above H_s and one below. Let Σ_1 be

the shard in $s\Sigma$ below H_s and let Σ_2 be the shard in $s\Sigma$ above H_s , as illustrated in Figure 7.




FIGURE 7. A case in the proof of Proposition 6.2.

Now Σ arrows Σ_1 , so the join-irreducible y associated to Σ_1 is also contracted by the *scs*-Cambrian congruence. Thus by Corollary 6.1, y is non-*scs*-sortable. Since Σ_1 is below H_s , $l(sy) > l(y)$. Let y_* be the unique element covered by y , so that Σ_1 is the shard associated to the edge $y_* < y$. Since y is a non-*scs*-sortable join-irreducible, it is contracted by Θ_{scs} , or in other words $\pi_{\downarrow}^{scs}(y_*) = \pi_{\downarrow}^{scs}(y)$. By Lemma 3.9, we have $\pi_{\downarrow}^c(sy) = \pi_{\downarrow}^c(s(y_*))$, or in other words, the edge $s(y_*) < sy$ is contracted by Θ_c . But the shard associated to the edge $s(y_*) < sy$ is Σ , so sj is contracted by Θ_c . By Corollary 6.1, sj not *c*-sortable. The shard associated to j is Σ_2 , which is arrowed to by Σ . By induction on length, sj is contracted by the *c*-Cambrian congruence, so j is also contracted by the *c*-Cambrian congruence.

By virtue of the cases we have already argued, we may now proceed under the assumptions that $l(sj) < l(j)$ and that the shard Σ associated to sj is below the hyperplane H_s .

We claim that under these assumptions, sj is not a defining join-irreducible of the *scs*-Cambrian congruence. Suppose to the contrary that sj is defining. Then $sj \in W_{\langle s \rangle}$, because all defining join-irreducibles for the *scs*-Cambrian congruence which are not in $W_{\langle s \rangle}$ are shortened on the left by s . By Lemma 4.5, the hyperplane H containing Σ is in $\mathcal{A}_{\{r_1, r_2\}}$ with $s \notin \{r_1, r_2\}$. By Lemma 4.6, H is basic in the rank-two subarrangement containing H and H_s and Σ contains both points strictly above H_s and points strictly below H_s . This contradiction to the assumption proves the claim.

Thus our list of assumptions becomes: $l(sj) < l(j)$; Σ is below H_s ; and sj is contracted by the *scs*-Cambrian congruence but is not a defining join-irreducible for the *scs*-Cambrian congruence. By the last of these assumptions, there is a join-irreducible x which is contracted by the *scs*-Cambrian congruence and which arrows sj in $\mathcal{S}h(W)$. By Corollary 6.1, x is non-*scs*-sortable. Since x arrows sj in $\mathcal{S}h(W)$, the corresponding shards Σ_x and Σ intersect in codimension 2, and the corresponding hyperplanes H_x and H are contained in a rank-two subarrangement \mathcal{A}' in which H_x is basic and H is not. By the assumption that Σ is weakly below H_s , there are two cases: either $\Sigma_x \cap \Sigma$ is contained in H_s or $\Sigma_x \cap \Sigma$ contains a point strictly below H_s . We now complete the proof by considering these two cases.




FIGURE 8. Another case in the proof of Proposition 6.2.

The case where the intersection $\Sigma_x \cap \Sigma$ is contained in H_s is illustrated in Figure 8. In this case, \mathcal{A}' contains H_s , and since H_x is basic in \mathcal{A}' , some upper region of Σ_x is below H_s . Thus $l(sx) > l(x)$. Since x arrows sj , Σ has a facet defined by H_x , or equivalently by H_s . We claim that the hyperplane H is not the image of H_x under s . If it is, then by Observation 4.7, the image $s\Sigma$ of Σ under s is contained in a shard with points on both sides of H_s (namely, the shard Σ_x). However, $s\Sigma$ is Σ_j , which lies above H_s because $l(sj) < l(j)$. This contradiction proves the claim. Thus the hyperplane containing j is not basic in \mathcal{A}' , so that j is arrowed to by x , which is not *c*-sortable.

We next claim that x is not *c*-sortable. If it is, then Lemma 2.1 implies that x is an *sc*-sortable element of $W_{(s)}$. But then by Lemma 2.3, x is *scs*-sortable as an element of W , contradicting what was established previously. This contradiction establishes the claim. By induction on the dual of $\text{Irr}(\text{Con}(W))$, x is contracted by the *c*-Cambrian congruence, and therefore so is j .

Finally, we deal with the case where $\Sigma_x \cap \Sigma$ contains a point strictly below H_s , as illustrated in Figure 9. In this case some upper region of Σ_x is below H_s , so that the minimal upper region x has $l(sx) > l(x)$. The reflection s relates Σ_x to a shard Σ' having sx in its upper set. The reflection s carries Σ into the shard Σ_j . Possibly Σ_x and Σ' differ to the extent allowed by Observation 4.7, and similarly Σ and Σ_j may differ. However, those differences occur only at the intersection with H_s , in particular not affecting the following assertions.




FIGURE 9. One last case in the proof of Proposition 6.2.

The intersection $\Sigma' \cap \Sigma_j$ has dimension $d - 2$ because $\Sigma_x \cap \Sigma$ has dimension $d - 2$. As explained in connection with Observation 4.7, reflection by s maps the basic hyperplanes of \mathcal{A}' to the basic hyperplanes of $s\mathcal{A}'$. Thus Σ' arrows Σ_j . Let x_* be the unique element covered by x . By Lemma 3.9, and the fact that Θ_{scs} contracts x , the congruence Θ_c contracts the edge $sx_* < sx$. Thus Θ_c contracts the join-irreducible y associated to Σ' . By Corollary 6.1, the join-irreducible y is non- c -sortable and by construction y is above j in $\text{Irr}(\text{Con}(W))$. By induction

on the dual of $\text{Irr}(\text{Con}(W))$, y is contracted by the c -Cambrian congruence, and therefore so is j . \square

This concludes the proof of Theorem 1.4.

Acknowledgments. I am grateful to David Speyer, John Stembridge and Hugh Thomas for helpful conversations. I particularly thank David Speyer for pointing out, in the course of joint work on [22], that Lemma 2.8 was not immediately obvious, and for cooperating in writing down the proof.

REFERENCES

- [1] I. Bernštein, I. Gel'fand and V. Ponomarev, *Coxeter functors, and Gabriel's theorem*, Uspehi Mat. Nauk **28** (1973) no. 2(170), 19–33. English translation in Russian Math. Surveys **28** (1973), no. 2, 17–32.
- [2] D. Bessis, *The dual braid monoid*, Ann. Sci. Ecole Norm. Sup. **36** (2003) 647–683.
- [3] A. Björner, *Orderings of Coxeter groups*, Contemp. Math. **34** (1984) 175–195.
- [4] A. Björner and F. Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, **231**, Springer, New York, 2005.
- [5] N. Bourbaki, *Lie groups and Lie algebras, Chapters 4–6*, Springer-Verlag, Berlin, 2002.
- [6] T. Brady and C. Watt, *$K(\pi, 1)$'s for Artin groups of finite type*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000), Geom. Dedicata **94** (2002), 225–250.
- [7] N. Caspard, C. Le Conte de Poly-Barbut and M. Morvan, *Cayley lattices of finite Coxeter groups are bounded*, Adv. in Appl. Math. **33** (2004), no. 1, 71–94.
- [8] I. Chajda and V. Snásel, *Congruences in Ordered Sets*, Math. Bohem. **123** (1998) no. 1, 95–100.
- [9] V. Duquenne and A. Cherfouh, *On permutation lattices*, Math. Social Sci. **27** (1994), no. 1, 73–89.
- [10] S. Fomin and A. Zelevinsky, *Y -systems and generalized associahedra*, Ann. of Math. **158** (2003), 977–1018.
- [11] S. Fomin and A. Zelevinsky, *Cluster algebras. II. Finite type classification*, Invent. Math. **154** (2003), no. 1, 63–121.
- [12] G. Grätzer, *General lattice theory*, Second edition, Birkhäuser Verlag, Basel, 1998.
- [13] J. Humphreys, *Reflection Groups and Coxeter Groups*, Cambridge Studies in Advanced Mathematics **29**, Cambridge Univ. Press, 1990.
- [14] G. Markowsky, *Permutation lattices revisited*, Math. Social Sci. **27** (1994), no. 1, 59–72.
- [15] N. Reading, *Order dimension, strong Bruhat order and lattice properties for posets*, Order **19** (2002) no. 1, 73–100.
- [16] N. Reading, *Lattice and order properties of the poset of regions in a hyperplane arrangement*, Algebra Universalis, **50** (2003), 179–205.
- [17] N. Reading, *The order dimension of the poset of regions in a hyperplane arrangement*, J. Combin. Theory Ser. A, **104** (2003) no. 2, 265–285.
- [18] N. Reading, *Lattice congruences of the weak order*, Order **21** (2004) no. 4, 315–344.
- [19] N. Reading, *Lattice congruences, fans and Hopf algebras*, J. Combin. Theory Ser. A **110** (2005) no. 2, 237–273.
- [20] N. Reading, *Cambrian Lattices*, Adv. Math. **205** (2006) no. 2, 313–353.
- [21] N. Reading, *Clusters, Coxeter-sortable elements and noncrossing partitions*, (math.CO/0507186), Trans. Amer. Math. Soc., to appear.

- [22] N. Reading and D. Speyer, *Cambrian Fans*, ([math.CO/0606201](https://arxiv.org/abs/math/0606201)), preprint, 2006.
- [23] J. Shi, *The enumeration of Coxeter elements*, *J. Algebraic Combin.* **6** (1997), no. 2, 161–171.

NATHAN READING

Mathematics Department, University of Michigan, Ann Arbor, MI 48109-1043, USA
Current address: Department of Mathematics, North Carolina State University, Raleigh, NC
27695-8205, USA
e-mail: nathan_reading@ncsu.edu
URL: <http://www4.ncsu.edu/~nreadin>