Noncrossing Arc Diagrams of Type B

Emily Barnard[a] Nathan Reading[b] Ashley M. Tharp[c]

Submitted: May 2, 2025; Accepted: Sep 20, 2026; Published: Nov 6, 2026
© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

Noncrossing arc diagrams are combinatorial models for permutations that encode information about lattice congruences of the weak order and about the associated discrete geometry. In this paper, we consider two related, analogous models for signed permutations. One model features centrally symmetric noncrossing arc diagrams, while the other features their quotients modulo the central symmetry. We demonstrate the utility of the models by applying them to various questions about lattice quotients of the weak order.

Mathematics Subject Classifications: 05A05, 06B10, 05E16

1 Introduction

Noncrossing arc diagrams [18] are combinatorial objects, in bijection with permutations, that bring to light lattice-theoretic information about the weak order on permutations and related discrete-geometric information. The results of [18] include:

- An explicit bijection from permutations to noncrossing arc diagrams and an explicit inverse;
- a bijection from arcs to join-irreducible elements of the weak order;
- a bijection from arcs to shards;
- a characterization of compatibility of join-irreducible elements (in the sense of canonical join complexes);
- a construction of the canonical join complex of the weak order in a way that proves that this complex is flag; and

[a]Department of Mathematical Sciences, DePaul University, Chicago, Illinois U.S.A. (e.barnard@depaul.edu).
[b]Department of Mathematics, North Carolina State University, Raleigh, North Carolina, U.S.A. (nreadin@ncsu.edu).
[c]Department of Mathematics, University of Southern California, U.S.A. (tharpa@usc.edu).

- a characterization of lattice congruences on the weak order and their quotients.

In this paper, we consider two models for signed permutations and prove results analogous to the results listed above. (The theorem on flagness of the canonical join complex was already vastly generalized by the first author in [1], so we don't state it separately here.) In the sequel [3], we construct an analogous model for finite Coxeter groups of type D (even-signed permutations) and prove analogous results.

The models provide a tool for characterizing lattice congruences and quotients in type B (Corollaries 4.10 and 4.11 or Corollaries 5.10 and 5.11). This tool is a great improvement over the earlier state of the art [14, Theorem 7.3], which is usable (e.g. in [16, 19]) but unwieldy. The improvement is in two directions. First, the characterization of forcing relations among join-irreducible congruences is a one-step criterion, rather than a description of arrows whose transitive closure is the forcing relation. Second, the description of the entire forcing relation is much simpler even than the earlier description of single arrows.

The type-B model comes in two equivalent versions, analogous to the description of signed permutations as sequences $\pi_{-n}\pi_{-n+1}\cdots\pi_{-1}\pi_1\pi_2\cdots\pi_n$ or as sequences $\pi_1\pi_2\cdots\pi_n$. In the first version, signed permutations of $\{\pm 1,\ldots,\pm n\}$ are in bijection with centrally symmetric noncrossing arc diagrams on $2n$ points. See Figure 3. This version has the advantage that the comparison with type-A noncrossing arc diagrams is transparent. Centrally symmetric noncrossing arc diagrams appeared in [2] and much more extensively in Julian Ritter's thesis [24] and in a related paper [11]. In particular, the thesis proves the one-step characterization of forcing by a detailed look at the geometry of shards [24, Proposition 4.2.42]. (See also [24, Observation 4.2.43] and [11, Theorem 114].) Here, we shortcut some of that process using Theorem 2.10, a general reformulation of forcing among shards that is more easily translated into the language of arcs. The thesis [24] and paper [11] go on to construct quotientopes of types A and B (polytopes realizing lattice quotients of the weak order in the same way the permutohedron realizes the weak order [12]).

The second version of the type-B model is an orbifold model, obtained by passing to the quotient of centrally symmetric noncrossing diagrams modulo the symmetry. See Figure 9. In addition to the obvious advantage of spatial compactness, the orbifold model brings to light some more subtle and surprising connections with type-A diagrams: In the third author's thesis [26] (see also [23]), the meet and join operations in the shard intersection orders of types A and B are characterized in terms of the usual noncrossing arc diagrams and in terms of the orbifold model, and these characterizations are used to show that the shard intersection order of type A_n (permutations in S_{n+1}) embeds as a sublattice of the shard intersection order of type B_n.

The paper concludes with examples showing the utility of the orbifold model. In particular, we give a much simpler proof of a characterization from [19] of the generators of congruences of the weak order on B_n whose quotients are the weak order on A_n. We also greatly simplify the proof of a characterization of Cambrian lattices of type B from [16], and give the first characterization of the generators of the bipartite and linear biCambrian congruences of type B.

2 Preliminaries

In this section, we establish background on lattice theory, finite Coxeter groups, posets of regions, and shards. We assume the most basic notions, particularly for lattice theory and Coxeter groups. In Section 2.3, we prove a new characterization of shard arrows, which provides a convenient tool to study forcing among congruences in the case of lattices of regions of a simplicial arrangement.

2.1 Lattices and congruences

An element j of a finite lattice L is called ***join-irreducible*** if and only if it covers exactly one element. We write j_* for the unique element covered by j.

The ***canonical join representation*** (CJR) of an element x in a finite lattice L is the unique antichain X in L such that $x = \bigvee X$ and such that, if X' is any subset of L with $x = \bigvee X'$, the order ideal generated by X' contains the order ideal generated by X. In particular, X consists of join-irreducible elements called the ***canonical joinands*** of x. An element x may or may not have a canonical join representation.

The property that every element of a finite lattice L has a canonical join representation is equivalent to a property called ***join-semidistributivity*** [8, Theorem 2.24]. We will not need the usual definition of join-semidistributivity here, but for finite lattices, one can profitably take the existence of canonical join representations as the definition. If both join-semidistributivity and the dual condition hold, then L is called ***semidistributive***.

A natural question, given a finite join-semidistributive lattice L, is to characterize which subsets of L are canonical join representations of elements of L. Thus the question is to characterize the collection $\{X \subseteq L : X \text{ is the CJR of } \bigvee X\}$. This collection is called the ***canonical join complex***, because it is a simplicial complex (with vertex set the set of join-irreducible elements).

A simplicial complex is ***flag*** if each minimal non-face is a 2-element set. (Equivalently, a set X of vertices forms a face if and only if every 2-element subset of X is an edge.) The following theorem is part of [1, Theorem 2].

Theorem 2.1. *Suppose L is a finite join-semidistributive lattice. The canonical join complex of L is flag if and only if L is semidistributive.*

Theorem 2.1 is extremely important for understanding the combinatorics of finite semidistributive lattices. Given a finite semidistributive lattice L, we say that two join-irreducible elements j and j' of L are ***compatible*** if and only if $\{j, j'\}$ is a face of the canonical join complex. (That is, if and only if there is an element whose canonical join representation is $\{j, j'\}$. Equivalently, if and only if there is an element whose canonical join representation *contains* $\{j, j'\}$.) Since, in particular, L is join-semidistributive (so that every element has a canonical join representation), the elements of L are in bijection with the faces of the canonical join complex. But then since L is semidistributive, Theorem 2.1 implies that the faces of the canonical join complex are precisely the pairwise compatible sets of join-irreducible elements of L. In later sections, we will give a bijection between join-irreducible elements of the weak order and various kinds of "arcs", define

notions of compatibility of arcs that correspond to compatibility of join-irreducible elements, and thus show that elements of the Coxeter group are in bijection with pairwise compatible sets of arcs.

A ***congruence*** on a lattice L is an equivalence relation Θ such that if $x_1 \equiv x_2$ and $y_1 \equiv y_2$ modulo Θ, then $(x_1 \wedge y_1) \equiv (x_2 \wedge y_2)$ and $(x_1 \vee y_1) \equiv (x_2 \vee y_2)$. The ***quotient*** L/Θ of L modulo Θ is the lattice whose elements are the Θ-classes $[x]_\Theta$ and whose meet and join are given by $[x]_\Theta \wedge [y]_\Theta = [x \wedge y]_\Theta$ and $[x]_\Theta \vee [y]_\Theta = [x \vee y]_\Theta$.

Congruences and quotients have a nice order-theoretic description as well. An equivalence relation Θ on a finite lattice L is a lattice congruence if and only if (1) each equivalence class is an interval, (2) the map sending an element to the bottom of its equivalence class is order-preserving, and (3) the map sending an element to the top of its equivalence class is order-preserving. The quotient L/Θ is isomorphic to the subposet induced by the set of bottom elements of congruence classes. Said another way: An element $x \in L$ is said to be ***contracted*** by Θ if x is congruent to some element strictly less than x. The quotient L/Θ is isomorphic to the subposet of L induced by uncontracted elements.

A join-irreducible element j is contracted by Θ if and only if $j \equiv j_*$ modulo Θ. It is known that a congruence is completely determined by the set of join-irreducible elements it contracts. (For a precise statement, see [20, Theorem 9-5.12].) Thus, to completely characterize congruences on a given lattice, it suffices to characterize which sets of join-irreducible elements can be contracted by congruences. In this paper, we carry out this characterization in the case of the weak order on Coxeter groups of type B.

Canonical join representations behave well when passing to quotients. The following proposition is [20, Proposition 10-5.29].

Proposition 2.2. *Suppose L is a finite join-semidistributive lattice and Θ is a congruence on L. Then an element $x \in L$ is contracted by Θ if and only if one or more of its canonical joinands is contracted by Θ. If x is not contracted by Θ, then its canonical join representation in the quotient L/Θ coincides with its canonical join representation in L.*

In the proposition, the quotient L/Θ is realized, as before, as the subposet of L induced by the elements of L that are not contracted by Θ. The second assertion of the proposition implies, in particular, that the join-irreducible elements of the quotient are precisely the join-irreducible elements of L that are not contracted by Θ.

The following corollary is an immediate consequence of Proposition 2.2. (Recall that the vertices of the canonical join complex of L are the join-irreducible elements of L.)

Corollary 2.3. *Suppose L is a finite join-semidistributive lattice and Θ is a congruence on L. Then the canonical join complex of L/Θ is the subcomplex of the canonical join complex of L induced by the join-irreducible elements of L not contracted by Θ.*

Join-irreducible elements cannot be contracted independently. Instead, there is a partial (pre-)order that mediates which sets of join-irreducibles can be contracted by a congruence. We say that a join-irreducible element j_1 ***forces*** another join-irreducible element j_2 if every congruence that contracts j_1 also contracts j_2. In general, the forcing

relation is a pre-order (a reflexive, transitive, but not necessarily antisymmetric relation) on the set of join-irreducible elements.

A set is closed under forcing if, whenever a join-irreducible element j is contracted, then all join-irreducible elements forced by j are also contracted. Sets of join-irreducible elements closed under forcing correspond to lattice-congruences: Each such set is precisely the set of join-irreducible elements contracted by the corresponding congruence.

2.2 Coxeter groups and the weak order

Let (W, S) be a Coxeter system. We write $\ell(w)$ for the usual length function, the length of w in the alphabet S. A ***(left) inversion*** of an element $w \in W$ is a reflection t such that $\ell(tw) < \ell(w)$. Let $\mathrm{inv}(w)$ denote the set of inversions of w. We write "$\leqslant$" for the ***(right) weak order*** on W, which is the partial order on W with $u \leqslant w$ if and only if $\mathrm{inv}(u) \subseteq \mathrm{inv}(w)$. The cover relations in the weak order are precisely the relations $ws \lessdot w$ such that $w \in W$, $s \in S$, and $\ell(ws) < \ell(w)$.

Given an element w of a Coxeter group W, a ***cover reflection*** of w is a reflection t such that $tw \lessdot w$. Cover reflections of w are in bijection with elements covered by w: If $ws \lessdot w$, then wsw^{-1} is a cover reflection. Let $\mathrm{cov}(w)$ denote the set of cover reflections of w.

The weak order on a finite Coxeter group is a semidistributive lattice [10]. The following is [21, Theorem 10-3.9].

Theorem 2.4. *Suppose W is a finite Coxeter group and $w \in W$. For each $t \in \mathrm{cov}(w)$, there is a unique minimal element j_t in $\{v : v \leqslant w, t \in \mathrm{inv}(v)\}$. The canonical join representation of w is $w = \bigvee \{j_t : t \in \mathrm{cov}(w)\}$.*

The weak order on a finite Coxeter group has a property called congruence uniformity [7], which in particular implies that the forcing relation is a partial order. In later sections, we describe this forcing order on join-irreducible elements of the weak order in types A and B in terms of subarc relationships on arcs. In the sequel [3], we describe the forcing order in type D by an analogous but more complicated notion.

2.3 Shards, congruences, and forcing

We now briefly recall some discrete-geometric notions that can be used to understand lattice congruences of the weak order, and prove a characterization of forcing that is useful in describing the forcing order on join-irreducible elements. In fact, we prove the characterization in the more general setting of lattices of regions of simplicial arrangements. For more details on the background material, see [20, 21].

A ***(real, central) hyperplane arrangement*** is a finite collection $\mathcal{A}$ of linear hyperplanes in $\mathbb{R}^d$. (We use the adjective "real" because hyperplane arrangements in vector spaces over other fields are often of interest. The adjective "central" emphasizes that our hyperplanes are linear, as opposed to affine hyperplanes which might not contain the origin.) The ***regions*** of $\mathcal{A}$ are the closures of the connected components of $\mathbb{R}^d \setminus \cup\mathcal{A}$. Choosing a ***base region*** B, we define the ***separating set*** $S(R)$ of a region R to be the

set of hyperplanes in $\mathcal{A}$ that separate the interior of R from the interior of B. The ***poset of regions*** $\mathcal{P}(\mathcal{A}, B)$ is the set of regions, partially ordered by setting $Q \leqslant R$ if and only if $S(Q) \subseteq S(R)$.

A hyperplane arrangement is ***simplicial*** if every region is a simplicial cone (the nonnegative linear span of a linearly independent set of vectors). Semidistributivity in the following theorem is [13, Theorem 3]. The lattice property was proven earlier in [4, Theorem 3.4].

Theorem 2.5. *If $\mathcal{A}$ is a simplicial hyperplane arrangement, then $\mathcal{P}(\mathcal{A}, B)$ is a semidistributive lattice.*

The set $\mathcal{A}_W$ of reflecting hyperplanes of a Coxeter group W (in the usual reflection representation) constitutes the ***Coxeter arrangement*** associated to W. Choosing B to be a region bounded by the reflecting hyperplanes for the simple reflections S, the map $w \mapsto wB$ is an isomorphism from the weak order on W to the poset of regions $\mathcal{P}(\mathcal{A}_W, B)$. It is well known that the Coxeter arrangements are simplicial. (See [21, Theorem 10-2.1] and the citation notes at the end of that chapter.) Thus Theorem 2.5 in particular implies the semidistributivity of the weak order, already mentioned in Section 2.2.

A ***rank-two subarrangement*** of an arrangement $\mathcal{A}$ is a subset $\mathcal{A}'$ of $\mathcal{A}$, with $|\mathcal{A}'| > 1$ that can be described as $\mathcal{A}' = \{H \in \mathcal{A} : H \supset U\}$ for some *codimension*-2 linear subspace U. Any two distinct hyperplanes $H_1, H_2 \in \mathcal{A}$ are contained in a unique rank-two subarrangement, namely $\{H \in \mathcal{A} : H \supset (H_1 \cap H_2)\}$.

Suppose that we have fixed a base region B as in the definition of $\mathcal{P}(\mathcal{A}, B)$. Then a rank-two subarrangement has two distinguished hyperplanes called its ***basic hyperplanes***. The subarrangement cuts $\mathbb{R}^d$ into regions, and B is contained in one of these regions, call it B'. The basic hyperplanes of $\mathcal{A}'$ are the two hyperplanes that bound B'.

A hyperplane $H_1 \in \mathcal{A}$ ***cuts*** a hyperplane $H_2 \in \mathcal{A}$ if H_1 is basic in the rank-two subarrangement $\mathcal{A}'$ containing H_1 and H_2, and H_2 is not basic in $\mathcal{A}'$. Given $H \in \mathcal{A}$, the ***shards*** in H are the closures of connected components of $H \setminus \bigcup H'$, where the union ranges over all hyperplanes H' that cut H. The set of shards of $\mathcal{A}$ is the set of all shards in all hyperplanes in $\mathcal{A}$. Thus to make the shards of $\mathcal{A}$, we "slice" each hyperplane along all of its intersections with hyperplanes that cut it. The construction of shards depends on the choice of base region B, but when $\mathcal{A}$ is a Coxeter arrangement, all choices of B are symmetric.

Given a shard Σ, we write H_Σ for the hyperplane containing Σ. Suppose R is a region of $\mathcal{A}$ and Σ is a shard of $\mathcal{A}$. If $R \cap \Sigma$ is a facet of R and $H_\Sigma \in S(R)$, then R is an ***upper region*** of Σ and Σ is a ***lower shard*** of R. The following proposition is [14, Propositions 3.2, 3.5].

Proposition 2.6. *If $\mathcal{A}$ is a simplicial arrangement, then each shard Σ has a unique minimal upper region J_Σ. The region J_Σ is join-irreducible in $\mathcal{P}(\mathcal{A}, B)$ and is the unique join-irreducible upper region of Σ. The map $\Sigma \mapsto J_\Sigma$ is a bijection from the set of shards of $\mathcal{A}$ to the set of join-irreducible elements of $\mathcal{P}(\mathcal{A}, B)$. The inverse map takes a join-irreducible region J to its unique lower shard.*

The following theorem is [17, Theorem 3.6].

Theorem 2.7. *If $\mathcal{A}$ is a simplicial arrangement and R is a region, then the canonical join representation of R in $\mathcal{P}(\mathcal{A}, B)$ is $\{J_\Sigma : \Sigma$ is a lower shard of $R\}$.*

Recall from Section 2.1 that two join-irreducible elements j_1, j_2 of a semidistributive lattice L are compatible if and only if there exists an element of L whose canonical join representation is $\{j_1, j_2\}$ if and only if there exists an element of L (not necessarily the same element) whose canonical join representation contains $\{j_1, j_2\}$. We will say that two shards are ***compatible*** if and only if the intersection of their relative interiors is nonempty.

Proposition 2.8. *Suppose $\mathcal{A}$ is a simplicial arrangement. Two shards are compatible if and only if the two corresponding join-irreducible regions in $\mathcal{P}(\mathcal{A}, B)$ are compatible.*

Proof. By Theorem 2.7, two join-irreducible regions J_1 and J_2 in a simplicial poset of regions are compatible if and only if Σ_{J_1} and Σ_{J_2} are the two lower shards of some region R. In that case, the hyperplanes containing Σ_{J_1} and Σ_{J_2} are the two basic hyperplanes in the rank-two subarrangement containing them, so they don't cut each other. Their intersection has codimension 2 because it contains the intersection of two facets of the simplicial region R. For these two reasons, we see that the relative interiors of Σ_{J_1} and Σ_{J_2} have nonempty intersection. Conversely, if there are two shards Σ_1 and Σ_2 whose relative interiors have nonempty intersection, then there is a region Q having both as lower shards. (Such a region Q can be found by starting at a generic point in the intersection of the two relative interiors and moving a small distance in the direction away from the interior of B.) Then Theorem 2.7 says that the corresponding join-irreducible regions are in the canonical join representation of Q, and thus are compatible. □

We have seen that shards and their incidences encode the canonical join representations of regions in $\mathcal{P}(\mathcal{A}, B)$. We will also see that they encode the forcing (pre-)order on join-irreducible elements in $\mathcal{P}(\mathcal{A}, B)$. Define the ***shard digraph*** to be the directed graph whose vertices are the shards, with $\Sigma_1 \to \Sigma_2$ if and only if H_{Σ_1} cuts H_{Σ_2} and $\Sigma_1 \cap \Sigma_2$ has codimension 2. The following is [20, Theorem 9-7.17].

Theorem 2.9. *If $\mathcal{A}$ is simplicial and Σ_1 and Σ_2 are shards, then J_{Σ_1} forces J_{Σ_2} if and only if there is a directed path in the shard digraph from Σ_1 to Σ_2.*

Said another way, the map $\Sigma \mapsto J_\Sigma$ is an isomorphism from the reflexive-transitive closure of the shard digraph to the forcing (pre-)order defined in Section 2.1. We emphasize that, even when the forcing pre-order is a partial order, shard arrows are not necessarily cover relations in the forcing order. Instead, there may be pairs of shards that are related by a shard arrow and also by a longer path in the shard digraph.

We now state and prove the main theorem of this section, a technical result that rephrases the definition of the shard digraph in a way that is useful (in later sections) for describing forcing in terms of "subarc" relationships between arcs.

Theorem 2.10. *Suppose $\mathcal{A}$ is a simplicial hyperplane arrangement and suppose Σ_1 and Σ_2 are shards. Then $\Sigma_1 \to \Sigma_2$ in the shard digraph if and only if there exists a shard Σ_1' satisfying the following conditions:*

(i) *Σ_1 and Σ_1' are compatible,*

(ii) *H_{Σ_2} is in the rank-two subarrangement containing H_{Σ_1} and $H_{\Sigma_1'}$ but is not basic in that subarrangement, and*

(iii) *$\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$.*

Note that the compatibility of Σ_1 and Σ_1' in (i) implies that there is a region having both Σ_1 and Σ_1' as lower shards. In particular, H_{Σ_1} and $H_{\Sigma_1'}$ are basic in the rank-two subarrangement containing them.

A ***shard intersection*** is an intersection of a set of shards. The ***shard intersection order*** is the reverse containment order on the set of all shard intersections. The proof of Theorem 2.10 uses the fact that the reverse containment order on the set of all intersections of shards is graded by codimension [17, Proposition 5.1].

Proof of Theorem 2.10. Suppose $\Sigma_1 \to \Sigma_2$. Write H_1 for H_{Σ_1} and H_2 for H_{Σ_2}. Then in particular H_1 is basic in the rank-two subarrangement $\mathcal{A}'$ containing H_1 and H_2, but H_2 is not basic in $\mathcal{A}'$. Let H_1' be the other basic hyperplane in $\mathcal{A}'$. Since $\Sigma_1 \cap \Sigma_2$ has codimension 2 and is contained in the intersection of the hyperplanes of $\mathcal{A}'$, we can find a point $\mathbf{x} \in \Sigma_1 \cap \Sigma_2$ that is not contained in any hyperplane in $\mathcal{A} \setminus \mathcal{A}'$. Thus, since H_2 does not cut H_1, the point $\mathbf{x}$ is in the relative interior of Σ_1. Since H_2 also does not cut H_1', $\mathbf{x}$ is also in the relative interior of some shard Σ_1' in H_1'. By definition, Σ_1 and Σ_1' are compatible. Since $\mathbf{x}$ is in the relative interior of Σ_1', there is some ball about $\mathbf{x}$ whose intersection with Σ_1' is the same as its intersection with H_1'. Since H_1' contains $\Sigma_1 \cap \Sigma_2$, we see that $\Sigma_1 \cap \Sigma_1' \cap \Sigma_2$ also has codimension 2. Since H_1 and H_1' are not the same hyperplane, $\Sigma_1 \cap \Sigma_1'$ has codimension 2 as well. Since the shard intersection order is graded by codimension and $\Sigma_1 \cap \Sigma_1' \cap \Sigma_2 \subseteq \Sigma_1 \cap \Sigma_1'$, these two intersections must in fact be the same. Thus $\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$.

Conversely, suppose there exists Σ_1' such that the three conditions of the theorem hold. Then certainly H_{Σ_1} cuts H_{Σ_2}, so we must show that $\Sigma_1 \cap \Sigma_2$ has codimension 2. Since $H_{\Sigma_1} \neq H_{\Sigma_2}$, certainly the codimension of $\Sigma_1 \cap \Sigma_2$ is at least 2. But the fact that Σ_1 and Σ_1' are compatible means that they intersect in their relative interiors, so $\Sigma_1 \cap \Sigma_1'$ has codimension 2. But since $\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$, we have $\Sigma_1 \cap \Sigma_1' = \Sigma_1 \cap \Sigma_1' \cap \Sigma_2$. In particular, the codimension of $\Sigma_1 \cap \Sigma_2$ cannot be more than 2. □

The third condition in Theorem 2.10 can be restated as the existence of a certain order relation in the shard intersection order (in the sense of [17]). In this paper, we will use the third condition directly, without explicitly working with the shard intersection order.

3 Noncrossing arc diagrams of type A

In this section, we review results of [18] on noncrossing arc diagrams for permutations in S_n.

The Coxeter group of type A_{n-1} can be realized as the group of permutations of $\{1, \ldots, n\}$. The ***one-line notation*** of $\pi \in S_n$ is the sequence $\pi_1\pi_2\cdots\pi_n$, where $\pi_i = \pi(i)$. One can also realize S_n as a reflection group in $\mathbb{R}^n$ in the usual way, with each simple reflection $s_i = (i\ \ i+1)$ acting as a reflection orthogonal to $\mathbf{e}_{i+1} - \mathbf{e}_i$. (The $\mathbf{e}_i$ are the standard basis vectors.) The reflections in S_n are the transpositions $(i\ j)$ for $1 \leqslant i < j \leqslant n$.

The weak order on S_n has cover relations given by $\pi \lessdot \sigma$ if σ is obtained from π by exchanging the entries π_i and π_{i+1} for some $i \in \{1, \ldots, n-1\}$ such that $\pi_i < \pi_{i+1}$. The cover reflection of σ associated to this cover is $(\pi_i\ \ \pi_{i+1})$; multiplying π on the left by this reflection swaps the entries π_i and π_{i+1}. A join-irreducible element of S_n is a permutation whose one-line notation has exactly one descent $\pi_i > \pi_{i+1}$.

We now define noncrossing arc diagrams. We place n distinct points on a vertical line, identified with the numbers $1, \ldots, n$ from bottom to top. An ***arc*** is a curve connecting a point $q \in \{1, \ldots, n\}$ to a strictly lower point $p \in \{1, \ldots, n\}$, moving monotone downwards from q to p without touching any other numbered point, but rather passing to the left of some points and to the right of others. A ***noncrossing arc diagram*** is a collection of arcs that don't intersect, except possibly at their endpoints, such that no two arcs share the same upper endpoint or the same lower endpoint. The combinatorial data determining an arc consists of which pair of points it connects and which points in between are left or right of the arc. Two arcs are combinatorially equivalent if they have the same combinatorial data. We consider arcs and noncrossing arc diagrams up to combinatorial equivalence. When we need to distinguish these diagrams from the objects defined later (for Coxeter groups of type B), we will refer to them as ***type-A arcs*** and ***type-A noncrossing arc diagrams***.

Noncrossing arc diagrams can also be understood in terms of a compatibility relation on arcs. We say two arcs are ***compatible*** if they don't intersect except possibly at one common endpoint, and if they don't share the same upper endpoint or the same lower endpoint. Certainly, the arcs in a noncrossing arc diagram are pairwise compatible. Conceivably, there might be a set of pairwise compatible arcs that can't all be drawn together without crossing, but [18, Proposition 3.2] rules out that possibility. Thus every set of pairwise compatible arcs is a noncrossing arc diagram.

We now describe the bijection δ from S_n to the set of noncrossing arc diagrams on n points. Given $\pi = \pi_1 \cdots \pi_n \in S_n$, write each entry π_i at the point (i, π_i) in the plane. For every i such that $\pi_i > \pi_{i+1}$, draw a straight line segment from π_i to π_{i+1}. These line segments become arcs: We move the numbers $1, \ldots, n$ horizontally to put them into a single vertical line, allowing the line segments to curve, so that they avoid passing through any numbers and one another. We define $\delta(\pi)$ to be the resulting noncrossing arc diagram.

We can alternatively describe $\delta(\pi)$ by listing its arcs: For every descent $\pi_i > \pi_{i+1}$, there is an arc with endpoints π_i and π_{i+1}. This arc goes right of every entry π_j with

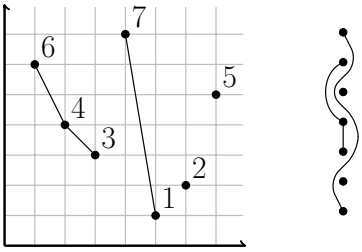


Figure 1: The map δ applied to 6437125

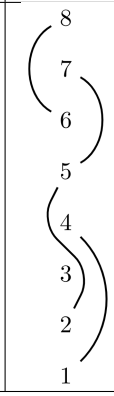
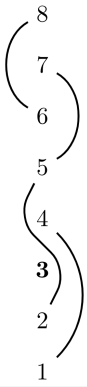
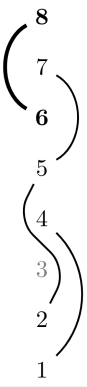
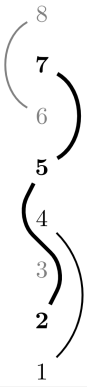
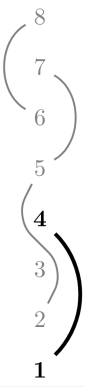
Step	Start	1	2	3	4
Permutation so far		**3**	3**86**	386**752**	38675241
Diagram remaining					

Figure 2: The map from type-A noncrossing arc diagrams to permutations

$\pi_{i+1} < \pi_j < \pi_i$ and $j < i$ and left of every entry π_j with $\pi_{i+1} < \pi_j < \pi_i$ and $j > i$. The following is [18, Theorem 3.1].

Theorem 3.1. *The map δ is a bijection from S_n to the set of noncrossing arc diagrams on n points.*

The proof of [18, Theorem 3.1] includes an explicit description of the inverse map, which we quote here. (An example is shown in Figure 2.) A noncrossing arc diagram has one or more ***blocks*** (the connected components of the diagram, viewed as an embedded graph). Each block is a single numbered point or a sequence of arcs sharing endpoints. Reading each block from top to bottom, we recover the maximal descending runs of the permutation. The noncrossing arc diagram has at least one ***left block***, meaning a block "with nothing to its left". More formally, no other arc or numbered point can be reached from that block by moving horizontally to the left. The left blocks can be totally ordered from lowest to highest. One can obtain π from $\delta(\pi)$ recursively by taking the lowest left block, writing its numbered points in decreasing order and then deleting it from the noncrossing arc diagram. Recursively, we then remove the lowest left block of what remains and continue writing the one-line notation for π from left to right.

Since δ maps a permutation in S_n with k descents to a noncrossing arc diagram with k arcs, in particular δ restricts to a bijection from permutations that are join-irreducible

in the weak order to single arcs. Suppose α is an arc connecting point q and p with $q > p$. Write $L(\alpha)$ for the set of ***left points of α*** (points in the interval (p, q) that are to the left of α) and write $R(\alpha)$ for the set of ***right points of α*** (points in the interval (p, q) that are to the right of α). Then the join-irreducible element corresponding to α is the permutation that starts with $1 \cdots (p-1)$, followed by the elements of $L(\alpha)$ in increasing order, then q, then p, then the elements of $R(\alpha)$ in increasing order, and finally $(q+1) \cdots n$.

Join-irreducible elements of the weak order are in bijection with shards, and thus shards are in bijection with arcs. Combining [21, Proposition 10-5.8] with the bijection described above between arcs and join-irreducible permutations, we can write the bijection between arcs and shards as follows.

Proposition 3.2. *Suppose α is an arc with endpoints $q > p$. The shard associated to α is*

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = x_q \text{ and } x_p \leqslant x_i \, \forall i \in R(\alpha) \text{ and } x_p \geqslant x_i \, \forall i \in L(\alpha)\}.$$

Noncrossing arc diagrams record canonical join representations of permutations, as explained in [18, Section 3] and summarized in the following theorem.

Theorem 3.3. *Given a permutation $\pi \in S_n$, the canonical join representation of π is the set of join-irreducible elements corresponding to the set of arcs in $\delta(\pi)$.*

Recall that two join-irreducible elements in a semidistributive lattice are called compatible if and only if they can appear together in a canonical join representation. As a consequence of Theorem 3.3, two join-irreducible elements of S_n are compatible if and only if the corresponding arcs are compatible.

Definition 3.4. An arc α' with endpoints $q' > p'$ is a ***subarc*** of an arc α with endpoints $q > p$ if and only if $q \geqslant q' > p' \geqslant p$ and $R(\alpha') = R(\alpha) \cap (p', q')$. If α' is a subarc of α, then we say α is a ***superarc*** of α'. Thus, given α', a superarc α is obtained by pushing α' right or left of the top and/or bottom endpoint, independently, and then extending upward and/or downward to make a longer arc.

Remark 3.5. It will be convenient in later generalizations to construct subarcs explicitly. Define a function h from the plane to $\mathbb{R}$ that returns the vertical height of a point, calibrated so that each numbered point i has height i. Given an arc α, parametrized as a function from the interval $[0, 1]$ into the plane, we obtain a subarc α' of α as follows: Choose t_1 and t_2 with $0 \leqslant t_1 < t_2 \leqslant 1$ such that $h(\alpha(t_1)) \in \{1, \ldots, n\}$ and $h(\alpha(t_2)) \in \{1, \ldots, n\}$. Also choose $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ such that $|h(\alpha(t_1 + \varepsilon_1)) - h(\alpha(t_1))| < 1$, $|h(\alpha(t_2 - \varepsilon_2)) - h(\alpha(t_2))| < 1$, and $t_1 + \varepsilon_1 < t_2 - \varepsilon_2$. Define α' to be the arc obtained by concatenating three curves: First, the straight line segment from the point numbered $h(\alpha(t_1))$ to the point $\alpha(t_1 + \varepsilon_1)$; second, the restriction of α to the interval $[t_1 + \varepsilon_1, t_2 - \varepsilon_2]$; and third, the straight line segment from $\alpha(t_2 - \varepsilon_2)$ to the point numbered $h(\alpha(t_2))$.

The following theorem is [18, Theorem 4.4]. A new proof of the theorem using Theorem 2.10 is found in [26, Section 2.3].

Theorem 3.6. *Let j_1 and j_2 be join-irreducible permutations, corresponding to arcs α_1 and α_2 respectively. Then j_1 forces j_2 if and only if α_1 is a subarc of α_2.*

We state two important corollaries of Theorem 3.6, both of which are stated as part of [18, Corollary 4.5]. The first is simply a rephrasing of Theorem 3.6. The second is obtained by combining Theorems 3.1 and 3.3 with Corollary 2.3.

Corollary 3.7. *A set U of arcs corresponds to the set of* uncontracted join-irreducible *permutations of some congruence Θ on S_n if and only if U is closed under passing to subarcs.*

Corollary 3.8. *If Θ is a congruence on S_n and U is the set of arcs corresponding to join-irreducible permutations* not *contracted by Θ, then δ restricts to a bijection from the quotient S_n/Θ (the set of permutations not contracted by Θ) to the set of noncrossing arc diagrams consisting only of arcs in U.*

We conclude this section with a result that will be helpful in Sections 4 and 5. A finite Coxeter group has an element w_0 that is longer than every other element (in the usual sense of length in terms of number of letters in a reduced word or number of inversions). The element w_0 is an involution.

In S_n, this element w_0 is $n(n-1)\cdots 321$. Given $\pi = \pi_1\pi_2\cdots\pi_n \in S_n$, πw_0 is the permutation with one-line notation $\pi_n\pi_{n-1}\cdots\pi_1$ and $w_0\pi$ is the permutation with one-line notation $(n+1-\pi_1)(n+1-\pi_2)\cdots(n+1-\pi_n)$. The one-line notation of $w_0\pi w_0$ is $(n+1-\pi_n)(n+1-\pi_{n-1})\cdots(n+1-\pi_1)$.

In the following proposition, we assume that the points are placed so that a rotation by a half turn sends the point labeled i to the point labeled $n+1-i$.

Proposition 3.9. *If $\pi \in S_n$, then the arc diagrams $\delta(\pi)$ and $\delta(w_0\pi w_0)$ are related by a half turn that sends the point labeled i to the point labeled $n+1-i$.*

Proof. There is a descent $\pi_i > \pi_{i+1}$ in π if and only if there is a descent $(n+1-\pi_{i+1}) > (n+1-\pi_i)$ in $w_0\pi w_0$, and an entry a with $\pi_{i+1} < a < \pi_i$ occurs left of π_i in π if and only if $(n+1-a)$ is right of $(n+1-\pi_i)$ in $w_0\pi w_0$. Thus $\delta(w_0\pi w_0)$ is obtained from $\delta(\pi)$ by a *reflection* taking each point labeled i to the point labeled $n+1-i$, followed by a reflection in the vertical line containing the points. The composition of these two reflections is a half turn that sends each point labeled i to the point labeled $n+1-i$. □

4 Centrally symmetric noncrossing arc diagrams

In this section and the next, we establish two notions of noncrossing arc diagrams for Coxeter groups of type B. While these notions are in some sense equivalent, they look quite different. We begin, in this section, with a centrally symmetric model, after briefly recalling some background about Coxeter groups of type B_n and establishing notation.

The Coxeter group of type B_n can be realized as the group of signed permutations of $\{\pm 1, \ldots, \pm n\}$. These are the permutations π of $\{\pm 1, \ldots, \pm n\}$ with the property that

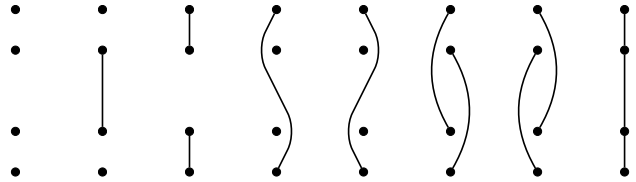


Figure 3: Centrally symmetric noncrossing arc diagrams for B_2

$\pi(-i) = -\pi(i)$ for $i = 1, \ldots, n$. The ***long one-line notation*** of $\pi \in B_n$ is the sequence $\pi_{-n}\pi_{-n+1}\cdots\pi_{-1}\pi_1\pi_2\cdots\pi_n$. But $\pi \in B_n$ is completely determined by its ***short one-line notation*** (or simply ***one-line notation***) $\pi_1\pi_2\cdots\pi_n$.

We realize B_n as usual as a reflection group in $\mathbb{R}^n$, with s_0 acting as a reflection orthogonal to the standard basis vector $\mathbf{e}_1$ and with each s_i acting as a reflection orthogonal to $\mathbf{e}_{i+1} - \mathbf{e}_i$ for $i = 1, \ldots, n-1$. For convenience, we write $\mathbf{e}_{-i}$ to mean $-\mathbf{e}_i$ for each $i = 1, \ldots, n$. Given a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$, we write x_{-i} to mean $-x_i$.

The reflections in B_n are the signed permutations with cycle notation $(i\ -i)$ for $i \in \{\pm 1, \ldots, \pm n\}$ or $(i\ j)(-i\ -j)$ for $i, j \in \{\pm 1, \ldots, \pm n\}$, and $|j| \neq |i|$.

The weak order on B_n has cover relations given by $\pi \lessdot \sigma$ if σ is obtained in one of two ways from π. One way is that $\pi_{-1} < \pi_1$ and σ is obtained by exchanging π_{-1} and π_1. The cover reflection of σ associated to this cover is the involution $(\pi_1\ \pi_{-1})$. The other way is that $\pi_i < \pi_{i+1}$ for some $i \in \{1, \ldots, n-1\}$ (and equivalently $\pi_{-i-1} < \pi_{-i}$) and σ is obtained from π by exchanging π_i and π_{i+1} and exchanging π_{-i-1} and π_{-i}. The cover reflection of σ associated to this cover is $(\pi_i\ \pi_{i+1})(\pi_{-i}\ \pi_{-i-1})$.

A join-irreducible element of B_n is a signed permutation whose long one-line notation either has exactly one descent $\pi_{-1} > \pi_1$ or has exactly two descents that are symmetric to each other: $\pi_i > \pi_{i+1}$ and $\pi_{-i-1} > \pi_{-i}$ for some $i \in \{1, \ldots, n-1\}$. Equivalently, either $\pi_1 < 0$ and the (short) one-line notation is increasing or $\pi_1 > 0$ and there is exactly one $i \in \{1, \ldots, n-1\}$ such that $\pi_i > \pi_{i+1}$.

In the centrally symmetric model, we begin with $2n$ distinct points on a vertical line, and the points identified with the numbers $-n, \ldots, -2, -1, 1, 2, \ldots, n$, in order, with $-n$ at the bottom. The points are placed so that the antipodal map $x \mapsto -x$ (a half turn about the origin) takes i to $-i$ for all $i \in \{\pm 1, \ldots, \pm n\}$. A ***centrally symmetric noncrossing arc diagram*** is a collection of arcs on these points, satisfying the same requirements as in type A, with the additional requirement that the entire diagram is symmetric with respect to the half-turn symmetry. All centrally symmetric noncrossing arc diagrams for B_2 are shown in Figure 3.

A centrally symmetric noncrossing arc diagram consists of ***symmetric arcs*** (arcs that are fixed by the central symmetry) and ***symmetric pairs of arcs*** (pairs of arcs that are compatible in the type-A sense of Section 3 and are mapped to each other by the central symmetry). Symmetric pairs of arcs come in two types: A ***non-overlapping symmetric pair*** is a symmetric pair $\{\alpha, -\alpha\}$ such that α connects positive points and $-\alpha$ connects negative points, so that neither arc is left or right of the other. An ***overlapping symmetric pair*** is a symmetric pair $\{\alpha, -\alpha\}$ in which each arc connects a positive

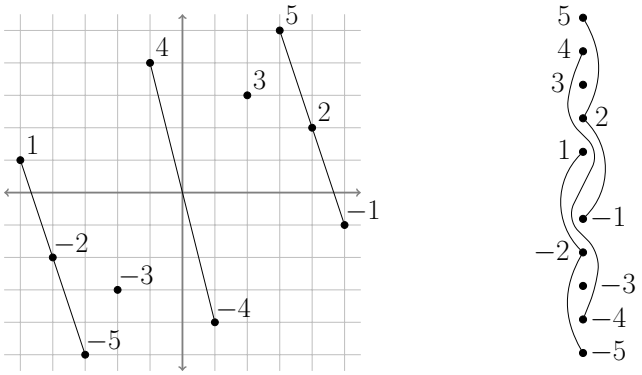


Figure 4: The map δ applied to $\pi = (-4)352(-1)$

point to a negative point. In an overlapping pair, since the two arcs are compatible, one is to the right of the other.

For the purposes of this section, we can harmlessly recast the results of Section 3 in terms of a group S_{2n} of permutations of $\{\pm 1, \ldots, \pm n\}$ and in terms of noncrossing arc diagrams on points labeled $-n, \ldots, -1, 1, \ldots, n$. As an immediate consequence of Theorem 3.1 and Proposition 3.9, we have the following theorem.

Theorem 4.1. *The map δ restricts to a bijection from B_n to the set of centrally symmetric noncrossing arc diagrams on $2n$ points.*

An example of δ applied to a signed permutation in B_5 is shown in Figure 4.

A join-irreducible signed permutation is mapped by δ to an arc diagram with only one arc (a symmetric arc), or only two arcs (a symmetric pair), and this is a bijection from join-irreducible signed permutations to symmetric arcs/pairs. We describe this bijection explicitly below as a map from symmetric arcs/pairs to join-irreducible signed permutations in (short) one-line notation. In this description, when we write a set as part of the one-line notation, we mean the elements of that set in increasing order. Also in this description, $i \cdots j$ will always stand for the sequence of elements increasing by 1 from i to j, or the empty sequence if $j = i - 1$.

Recall from Section 3 that the set $L(\alpha)$ of ***left points*** of an arc α is the set of numbers that are left of α, and the set $R(\alpha)$ of ***right points*** of α is the set of numbers that are right of α. (Applying this definition to arc diagrams on points $\{\pm 1, \ldots, \pm n\}$, of course 0 is never a right point or a left point.) We write $-L(\alpha)$ for $\{-i : i \in L(\alpha)\}$ and similarly $-R(\alpha)$. For a symmetric pair $\{\alpha, -\alpha\}$, $-L(\alpha) = R(-\alpha)$ and $-R(\alpha) = L(-\alpha)$.

A join-irreducible signed permutation π with one descent $\pi_{-1} > \pi_1$ corresponds to a symmetric arc α. If α has endpoints $-p < p$, then π is

$$(-p)\, R(\alpha)\, (p+1) \cdots n.$$

A join-irreducible signed permutation π that has two descents $\pi_i > \pi_{i+1}$ and $\pi_{-i-1} > \pi_{-i}$ corresponds to a symmetric pair $\{\alpha, -\alpha\}$ of arcs. Suppose α has endpoints $p < q$.

If $\{\alpha, -\alpha\}$ is nonoverlapping, then we assume that α is above $-\alpha$, so that $0 < p < q$. Then π is

$$1 \cdots (p-1)\, L(\alpha)\, q\, p\, R(\alpha)\, (q+1) \cdots n.$$

If $\{\alpha, -\alpha\}$ is overlapping, then we assume that α is right of $-\alpha$. If $p < 0 < -p < q$, then π is

$$[(0, -p) \cap L(\alpha) \cap -L(\alpha)]\, [(-p, q) \cap L(\alpha)]\, q\, p\, R(\alpha)\, (q+1) \cdots n.$$

If $p < 0 < q < -p$, then π is

$$[(0, q) \cap L(\alpha) \cap -L(\alpha)]\, q\, p\, R(\alpha)\, [(q, -p) \cap -L(\alpha)]\, (-p+1) \cdots n.$$

Given a signed permutation $\pi \in B_n$, each symmetric arc or symmetric pair of arcs in $\delta(\pi)$ is associated to a cover reflection t of π. Each symmetric arc or symmetric pair also specifies a join-irreducible element j as described above. By inspection of the three cases described above, we see that $j \leqslant \pi$ and that j is minimal with respect to the property that $t \in \mathrm{inv}(j)$. Thus Theorem 2.4 implies the following theorem.

Theorem 4.2. *Given $\pi \in B_n$, the canonical join representation of π is the set of join-irreducible elements corresponding to the symmetric arcs and symmetric pairs of arcs in $\delta(\pi)$.*

Recall from Section 3 that two arcs are called compatible if and only if they don't intersect, except possibly at endpoints, and don't share the same top endpoint or the same bottom endpoint. We now define one symmetric arc/pair A to be ***compatible*** with another symmetric arc/pair A' if and only if every arc in A is compatible with every arc in A'. As a consequence of Theorem 4.2, two symmetric arcs/pairs are compatible if and only if the corresponding join-irreducible elements are compatible.

The following two propositions are obtained by combining the correspondence between arcs and join-irreducible signed permutations with simple observations in [14, Sections 3,5]. In those sections, join-irreducible elements are shown to be in bijection with certain "signed subsets", and the signed subsets are used to write inequalities for shards.

Proposition 4.3. *Suppose α is a symmetric arc having endpoints $-p$ and p. The shard associated to α is*

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = 0 \text{ and } 0 \leqslant x_i \, \forall i \in R(\alpha)\}.$$

Proposition 4.4. *Suppose $\{\alpha, -\alpha\}$ is a symmetric pair of arcs such that α has endpoints p and q. The shard associated to $\{\alpha, -\alpha\}$ is*

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = x_q \text{ and } x_p \leqslant x_i \, \forall i \in R(\alpha) \text{ and } x_p \geqslant x_i \, \forall i \in L(\alpha)\}.$$

Remark 4.5. There is a global change in the direction of inequalities between Propositions 4.3 and 4.4 and [14], because in [14], hyperplanes were identified with their normal vectors pointing *away* from a "base region" B. A standard convention for root systems has the positive roots pointing *towards* a "fundamental chamber". This convention was used in [21], from which we quoted Proposition 3.2, so to keep conventions consistent in this paper, we apply the antipodal map relative to [14].

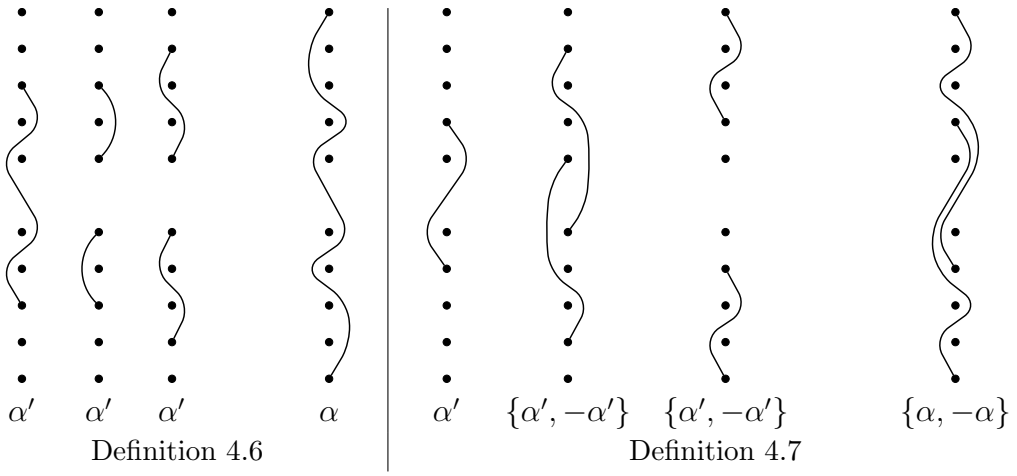


Figure 5: Subarcs and subarc pairs

We now use Theorem 2.10 to characterize forcing of join-irreducible signed permutations in terms of a centrally symmetric version of subarcs, which we now define. Since join-irreducible signed permutations correspond either to centrally symmetric arcs or to centrally symmetric pairs of arcs, we will use the terminology of a "subarc pair" of a symmetric arc or symmetric pair, or a "subarc" of a symmetric arc or symmetric pair. Figure 5 shows subarcs and subarc pairs of a symmetric arc and an overlapping symmetric pair.

Definition 4.6 (Subarcs/subarc pairs of a symmetric arc)**.** Suppose α is a symmetric arc with endpoints $-p < p$. A ***subarc*** of α is a symmetric arc α' with endpoints $-p' < p'$ with $p' \leqslant p$ and $R(\alpha') = R(\alpha) \cap (-p', p')$. A ***subarc pair*** of α is a non-overlapping pair $\{\alpha', -\alpha'\}$ of arcs such that α' has endpoints p' and q' with $0 < p' < q' \leqslant p$ and $R(\alpha') = R(\alpha) \cap (p', q')$.

Definition 4.7 (Subarcs/subarc pairs of a symmetric pair of arcs)**.** Suppose $\{\alpha, -\alpha\}$ is a symmetric pair of arcs and α has endpoints $p < q$. A ***subarc*** of $\{\alpha, -\alpha\}$ is a symmetric arc α' with endpoints $-p'$ and p' having $p \leqslant -p' < p' \leqslant q$ and $R(\alpha') = R(\alpha) \cap (-p', p') = R(-\alpha) \cap (-p', p')$. This can only happen if α and $-\alpha$ are overlapping and there are no numbered points between them on the interval $(-p', p')$. Now, if $\{\alpha, -\alpha\}$ is overlapping, suppose further that α is right of $-\alpha$. A ***subarc pair*** of $\{\alpha, -\alpha\}$ is a symmetric pair $\{\alpha', -\alpha'\}$ with endpoints p' and q' with $p \leqslant p' < q' \leqslant q$ and $R(\alpha') = R(\alpha) \cap (p', q')$, satisfying an additional requirement: If $\{\alpha', -\alpha'\}$ is overlapping then α' is also right of $-\alpha'$. The left two pictures of Figure 6 show a failure of this additional requirement.

Remark 4.8. To understand this additional requirement, it is useful to construct subarcs of $\{\alpha, -\alpha\}$ explicitly. Similarly to the type-A construction, we take h to be a function from the plane to $\mathbb{R}$ that returns the vertical height of a point, calibrated so that each numbered point $i \in \{\pm 1, \ldots, \pm n\}$ has height i. Choose t_1 and t_2 with $0 \leqslant t_1 < t_2 \leqslant 1$ such that $h(\alpha(t_1)), h(\alpha(t_2)) \in \{\pm 1, \ldots, \pm n\}$ and choose $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ such that

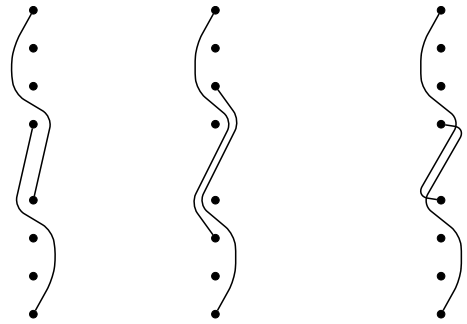


Figure 6: A pair not satisfying the additional requirement of Definition 4.7

$|h(\alpha(t_1+\varepsilon_1)) - h(\alpha(t_1))| < 1$, $|h(\alpha(t_2-\varepsilon_2)) - h(\alpha(t_2))| < 1$, and $t_1+\varepsilon_1 < t_2-\varepsilon_2$. Then α' is obtained by concatenating three curves: the segment from the point numbered $h(\alpha(t_1))$ to the point $\alpha(t_1+\varepsilon_1)$, the restriction of α to $[t_1+\varepsilon_1, t_2-\varepsilon_2]$ and the segment from $\alpha(t_2-\varepsilon_2)$ to the point numbered $h(\alpha(t_2))$. The additional requirement on $\{\alpha', -\alpha'\}$ is that (up to choosing another symmetric arc pair combinatorially equivalent to $\{\alpha, -\alpha\}$) the specific curves α' and $-\alpha'$ constructed here do not intersect each other. The right picture of Figure 6 shows an example of α' and $-\alpha'$ crossing.

We will prove the following theorem.

Theorem 4.9. *Let j_1 and j_2 be join-irreducible signed permutations. Then j_1 forces j_2 if and only if the arc or pair of arcs corresponding to j_1 is a subarc or subarc pair of the arc or pair of arcs corresponding to j_2.*

Theorem 4.9 lets us prove the type-B analogues of Corollaries 3.7 and 3.8. The following result is a rephrasing of Theorem 4.9, in light of the correspondence between congruences and sets of join-irreducible elements closed under forcing.

Corollary 4.10. *A set U of symmetric arcs/symmetric pairs corresponds to the set of* un*contracted join-irreducible signed permutations of some congruence Θ on B_n if and only if U is closed under passing to subarcs/subarc pairs.*

Combining Corollary 4.10 with Theorems 4.1 and 4.2 and Corollary 2.3, we obtain the following result.

Corollary 4.11. *If Θ is a congruence on B_n and U is the set of symmetric arcs/pairs corresponding to join-irreducible signed permutations* not *contracted by Θ, then δ restricts to a bijection from the quotient B_n/Θ (the set of signed permutations not contracted by Θ) to the set of centrally symmetric noncrossing arc diagrams consisting only of arcs in U.*

It is convenient to understand the opposite of the subarc relation. We say that a symmetric arc/pair A_2 is a ***superarc*** or ***superarc pair*** of a symmetric arc/pair A_1 if and only if A_1 is a subarc/subarc pair of A_2. We now describe how to construct

a superarc/superarc pair of a given symmetric arc/pair. The construction is a direct rephrasing of Definitions 4.6 and 4.7.

Superarcs/superarc pairs of a symmetric arc α. We construct a superarc by pushing α right or left of its top endpoint, symmetrically pushing it left or right of its bottom endpoint, and extending α symmetrically upwards and downwards. We construct a superarc pair by first replacing α with two arcs that are combinatorially equivalent to α, antipodal images of each other, and disjoint except at their endpoints. We then push one or both copies right or left of the top endpoint, independently but without making the curves cross each other, make the symmetric change at the bottom endpoint, and extend the curves upward and downward symmetrically.

Superarcs/superarc pairs of a non-overlapping symmetric pair of arcs $\{\alpha, -\alpha\}$. To construct a superarc or superarc pair, we first push α and $-\alpha$ left or right of their inner endpoints and/or outer endpoints in a way that preserves the symmetry. We construct a superarc by extending the curves inward and connecting them, and also possibly extending both outwards. We construct a superarc pair by extending the curves symmetrically to create a symmetric pair of compatible arcs, either overlapping or not.

Superarc pairs of an overlapping symmetric pair of arcs $\{\alpha, -\alpha\}$. (There are no superarcs of $\{\alpha, -\alpha\}$ and no non-overlapping superarc pairs, only overlapping superarc pairs.) We construct a superarc pair by pushing α right or left of one or both endpoints independently, making the symmetric change to $-\alpha$, and then extending the curves symmetrically to create a symmetric pair of compatible overlapping arcs. The additional requirement in Definition 4.7 is implied by this description of superarcs: If $\{\alpha_1, -\alpha_1\}$ fails to be a subarc pair of $\{\alpha_2, -\alpha_2\}$ because it fails the additional requirement, then the attempt to extend α_1 and $-\alpha_1$ to obtain α_2 and $-\alpha_2$ will result in a pair of arcs that cross each other. (See Figure 6.)

A first step towards the proof of Theorem 4.9 is the following observation.

Proposition 4.12. *The subarc/subarc pair relation on symmetric arcs/pairs is transitive.*

The proposition is verified by checking the various cases (e.g. a nonoverlapping subarc pair of a symmetric subarc of an overlapping arc pair). There are many cases to check, but each individual case is easy.

As a next step towards proving Theorem 4.9, we use Theorem 2.10 to characterize arrows in the shard digraph in terms of arcs and arc pairs. Given symmetric arcs/pairs A_1 and A_2, we write $A_1 \to A_2$ if and only if the corresponding shards Σ_1 and Σ_2 have $\Sigma_1 \to \Sigma_2$ in the shard digraph.

Proposition 4.13. *Suppose A_1 is a symmetric arc or pair and A_2 is another symmetric arc or pair. Then $A_1 \to A_2$ if and only if A_1 is a subarc/subarc pair of A_2 and one of the following conditions holds.*

(i) *A_2 is a symmetric pair $\{\alpha_2, -\alpha_2\}$ such that α_2 has endpoints p and q with $p < q$ and A_1 is a symmetric pair $\{\alpha_1, -\alpha_1\}$ such that α_1 has endpoints p' and q with $p' \neq -p$ or α_1 has endpoints p and q' with $q' \neq -q$.*

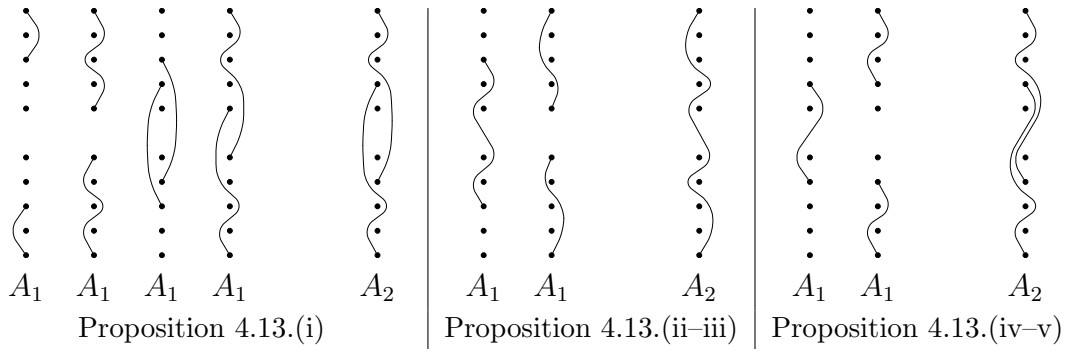


Figure 7: Arrows $A_1 \to A_2$ among symmetric arcs/pairs

(ii) *A_1 and A_2 are both symmetric arcs with no endpoints in common.*

(iii) *A_2 is a symmetric arc α_2 with endpoints p and $-p$ and A_1 is a non-overlapping symmetric pair $\{\alpha_1, -\alpha_1\}$ such that α_1 has an endpoint p.*

(iv) *A_2 is a symmetric pair $\{\alpha_2, -\alpha_2\}$ such that α_2 has endpoints p and q with $-q < p < 0 < -p < q$ and such that each point in $\{\pm 1, \ldots, \pm(-p-1)\}$ is left of α_2 if and only if it is left of $-\alpha_2$, and A_1 is a symmetric arc with endpoints p and $-p$.*

(v) *A_2 is a symmetric pair $\{\alpha_2, -\alpha_2\}$ such that α_2 has endpoints p and q with $-q < p < 0 < -p < q$ and such that each point in $\{\pm 1, \ldots, \pm(-p-1)\}$ is left of α_2 if and only if it is left of $-\alpha_2$, and A_1 is a non-overlapping symmetric pair $\{\alpha_1, -\alpha_1\}$ such that α_1 has endpoints $-p$ and q.*

Examples of arrows $A_1 \to A_2$ are shown in Figure 7. We emphasize that the conditions for an arrow in Proposition 4.13 include the condition that A_1 is a subarc/subarc pair of A_2. This requirement includes the additional requirement in Definition 4.7, which is crucial for the arrows satisfying condition (i) in Proposition 4.13, as explained in Case 1 of the proof below.

Proof. Let Σ_1 and Σ_2 be the shards corresponding to A_1 and A_2. Theorem 2.10 says that $A_1 \to A_2$ if and only if there exists a symmetric arc/pair A_1' (corresponding to a shard Σ_1') satisfying certain properties. In particular, A_1 and A_1' must be compatible and there needs to exist a non-basic hyperplane in the rank-two subarrangement containing H_{Σ_1} and $H_{\Sigma_1'}$. In other words, that subarrangement must have more than two hyperplanes. Each rank-two subarrangement looks like a rank-two parabolic subgroup, so in type B it can have 2, 3, or 4 hyperplanes.

A symmetric arc with endpoints $\pm p$ determines a shard whose hyperplane is orthogonal to $\mathbf{e}_p$. A symmetric pair such that one arc has endpoints p and q determines a shard whose hyperplane is orthogonal to $\mathbf{e}_q - \mathbf{e}_p$ (recalling the convention that $\mathbf{e}_{-i}$ means $-\mathbf{e}_i$). If A_1 and A_1' have no endpoints in common, then H_{Σ_1} and $H_{\Sigma_1'}$ are orthogonal, and thus, since these hyperplanes are basic in the rank-two subarrangement they determine, that

subarrangement has only two hyperplanes. Thus we need only consider cases where A_1 and A_1' have endpoints in common. We break into two cases, where A_1 and A_1' determine a rank-two subarrangement with 3 or 4 hyperplanes.

Case 1. H_{Σ_1} and $H_{\Sigma_1'}$ are the basic hyperplanes in a rank-two subarrangement with 3 hyperplanes. Then A_1 is a symmetric pair $\{\alpha_1, -\alpha_1\}$ such that α_1 has endpoints $p < q$ and A_1' is a symmetric pair $\{\alpha_1', -\alpha_1'\}$ such that α_1' has endpoints $q < r$ with $q \notin \{-p, -r\}$. At most one of the pairs A_1 and A_1' is overlapping, and if one is, we can assume (up to renaming arcs and points) that α_1 or α_1' is to the right of its negative. The basic hyperplane H_{Σ_1} is orthogonal to $\mathbf{e}_q - \mathbf{e}_p$ and the other basic hyperplane $H_{\Sigma_1'}$ is orthogonal to $\mathbf{e}_r - \mathbf{e}_q$. The unique non-basic hyperplane in the subarrangement is orthogonal to $\mathbf{e}_r - \mathbf{e}_p$. Thus an arc/arc pair specifies a non-basic hyperplane in the subarrangement if and only if it is a symmetric arc pair $\{\alpha_2, -\alpha_2\}$ such that α_2 has endpoints p and r. Proposition 4.4 implies that

$$\begin{aligned}\Sigma_1 \cap \Sigma_1' = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = x_r,& \\ x_p \leqslant x_i \,\forall i \in (R(\alpha_1) \cup R(\alpha_1')),& \\ x_p \geqslant x_i \,\forall i \in (L(\alpha_1) \cup L(\alpha_1'))\}.&\end{aligned}$$

On the other hand,

$$\Sigma_2 = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_r,\, x_p \leqslant x_i \,\forall i \in R(\alpha_2),\, x_p \geqslant x_i \,\forall i \in L(\alpha_2)\}.$$

Since $\Sigma_1 \cap \Sigma_1'$ is in the subspace where $x_p = x_q$, it is contained in Σ_2 if and only if it is contained in $\{\mathbf{x} \in \Sigma_2 : x_p = x_q\}$, which equals

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = x_r,\, x_p \leqslant x_i \,\forall i \in R(\alpha_2) \setminus \{q\},\, x_p \geqslant x_i \,\forall i \in L(\alpha_2) \setminus \{q\}\}.$$

Thus $\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$ if and only if $R(\alpha_2) \setminus \{q\} = R(\alpha_1) \cup R(\alpha_1')$ (equivalently, if and only if $L(\alpha_2) \setminus \{q\} = L(\alpha_1) \cup L(\alpha_1')$). This is in turn equivalent to the condition that α_2 is obtained by pushing the curve $\alpha_1 \cup \alpha_1'$ left or right of the shared endpoint q.

Now Theorem 2.10 implies that every arrow among shards that determine rank-two subarrangements with 3 hyperplanes is of the form $A_1 \to A_2$ or $A_1' \to A_2$ for A_1, A_1', and A_2 as above such that α_2 is obtained by pushing the curve $\alpha_1 \cup \alpha_1'$ left or right of q. Given A_1 and A_2, the existence of A_1' such that α_2 is obtained by pushing $\alpha_1 \cup \alpha_1'$ left or right of q is exactly the condition that A_1 is a subarc pair of A_2, *except for the additional requirement on subarc pairs of symmetric arc pairs* in Definition 4.7. Similarly, given A_1' and A_2, the existence of an appropriate A_1 is exactly that A_1' is a subarc pair, without the additional requirement.

Furthermore, Theorem 2.10 says that if α_2 is obtained by pushing $\alpha_1 \cup \alpha_1'$ left or right of q, then these arrows $A_1 \to A_2$ or $A_1' \to A_2$ exist if and only if A_1 and A_1' are compatible. We will show that, under the condition that α_2 is obtained by pushing $\alpha_1 \cup \alpha_1'$ left or right of q, the compatibility of A_1 and A_1' is equivalent to the additional requirement.

Compatibility of A_1 and A_1' means that all four arcs α_1, $-\alpha_1$, α_1', and $-\alpha_1'$ can be in the same noncrossing arc diagram. Recall also that if one of the pairs $\{\alpha_1, -\alpha_1\}$ or $\{\alpha_1', -\alpha_1'\}$

is overlapping, then α_1 or α_1' is on the right. For each pair which is nonoverlapping, α_1 or α_1' is above its opposite.

Suppose A_1 and A_1' are compatible. Then the union $\alpha_1 \cup \alpha_1'$ is either above or right of its opposite $(-\alpha_1) \cup (-\alpha_1')$. Since α_2 is obtained by pushing $\alpha_1 \cup \alpha_1'$ left or right of q, either $\{\alpha_2, -\alpha_2\}$ is nonoverlapping and α_2 is above $-\alpha_2$ or $\{\alpha_2, -\alpha_2\}$ is overlapping and α_2 is right of $-\alpha_2$.

Now, suppose A_1 and A_1' are not compatible. The individual arcs α_1 and α_1' are compatible in any case, so the failure of compatibility of A_1 and A_1' implies that α_1 and $-\alpha_1'$ cross each other. Thus the embedding of α_2 obtained by pushing $\alpha_1 \cup \alpha_1'$ left or right of q also crosses its antipodal opposite. Since by supposition $\{\alpha_2, -\alpha_2\}$ is a symmetric arc pair, there is an embedding of α_2 that does not cross its opposite. This embedding must therefore have α_2 left of $-\alpha_2$, as in Figure 6.

We have shown that arrows of the form $A_1 \to A_2$ and $A_1' \to A_2$ arising from Case 1 are precisely the arrows described in condition (i).

Case 2. H_{Σ_1} and $H_{\Sigma_1'}$ are the basic hyperplanes in a rank-two subarrangement with 4 hyperplanes. Then (up to swapping A_1 and A_1') A_1 is a symmetric arc α_1 with endpoints $\pm p$ for $p > 0$ and A_1' is a non-overlapping arc pair $\{\alpha_1', -\alpha_1'\}$ such that α_1' has endpoints p and q with $p < q$. Any such A_1 and A_1' are compatible, with no additional requirements needed.

The basic hyperplanes in the subarrangement are orthogonal to $\mathbf{e}_p$ and $\mathbf{e}_q - \mathbf{e}_p$. The two non-basic hyperplanes are $\mathbf{e}_q$ and $\mathbf{e}_q + \mathbf{e}_p$. Thus an arc/arc-pair A_2 specifies a non-basic hyperplane in the subarrangement if and only if it is a symmetric arc with endpoints $\pm q$ or an overlapping symmetric pair one of whose arcs has endpoints $-p$ and q. To determine all possible arrows arising from this compatible pair, it remains to determine necessary and sufficient conditions on A_2 so that the corresponding shard has $\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$.

Propositions 4.3 and 4.4 combine to say that

$$\Sigma_1 \cap \Sigma_1' = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in (R(\alpha_1) \cup R(\alpha_1')),\, 0 \geqslant x_i \forall i \in L(\alpha_1')\}.$$

Since $x_{-i} = -x_i$, we can rewrite this as

$$\Sigma_1 \cap \Sigma_1' = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in (R(\alpha_1) \cup R(\alpha_1') \cup (-L(\alpha_1')))\}. \tag{4.1}$$

Since $\Sigma_1 \cap \Sigma_1'$ is in the subspace where $x_p = 0$, it is contained in Σ_2 if and only if it is contained in $\{\mathbf{x} \in \Sigma_2 : x_p = 0\}$.

Case 2a. If A_2 is a symmetric arc α_2 with endpoints $\pm q$, then Proposition 4.3 says that

$$\{\mathbf{x} \in \Sigma_2 : x_p = 0\} = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in R(\alpha_2)\}.$$

If A_1 is a subarc of A_2 as in condition (ii) and A_1' is a subarc pair of A_2 as in condition (iii), then in particular $R(\alpha_2) \setminus \{\pm p\} = R(\alpha_1) \cup R(\alpha_1') \cup (-L(\alpha_1'))$, so $\Sigma_1 \cap \Sigma_1'$ is $\{\mathbf{x} \in \Sigma_2 : x_p = 0\}$.

Conversely, if $\Sigma_1 \cap \Sigma_1' \subseteq \Sigma_2$, then for every $i \in R(\alpha_2)$, the inequality $0 \leqslant x_i$ holds in $\Sigma_1 \cap \Sigma_1'$. Thus $R(\alpha_2) \setminus \{\pm p\} = R(\alpha_1) \cup R(\alpha_1') \cup (-L(\alpha_1'))$, so A_1 is a subarc of A_2 and A_1' is a subarc pair of A_2.

We have shown that arrows of the form $A_1 \to A_2$ and $A_1' \to A_2$ in Subcase 2a are precisely the arrows described in conditions (ii) and (iii).

Case 2b. If A_2 is an overlapping symmetric pair $\{\alpha_2, -\alpha_2\}$ such that α_2 has endpoints $-p$ and q, then we consider Proposition 4.4 in two further cases, given by whether α_2 is to the right or left of $-\alpha_2$. If α_2 is to the right of $-\alpha_2$, then Proposition 4.4 implies that

$$\{\mathbf{x} \in \Sigma_2 : x_p = 0\} = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in R(\alpha_2),\, 0 \geqslant x_i \forall i \in L(\alpha_2)\}.$$

If α_2 is to the left of $-\alpha_2$, then since $-L(\alpha_2)$ is the set of points to the *right* of $-\alpha_2$ and $-R(\alpha_2)$ is the set of points to the *left* of $-\alpha_2$, Proposition 4.4 implies that

$$\{\mathbf{x} \in \Sigma_2 : x_p = 0\} = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in -L(\alpha_2),\, 0 \geqslant x_i \forall i \in -R(\alpha_2)\}.$$

In either case,

$$\{\mathbf{x} \in \Sigma_2 : x_p = 0\} \\ = \{\mathbf{x} \in \mathbb{R}^n : x_p = x_q = 0,\, 0 \leqslant x_i \forall i \in (R(\alpha_2) \cup (-L(\alpha_2)))\}. \quad (4.2)$$

If A_1 is a subarc of A_2 as in condition (iv) and A_1' is a subarc pair of A_2 as in condition (v), then no points are between α_2 and $-\alpha_2$, so $R(\alpha_2) \cup (-L(\alpha_2))$ contains no pairs $\pm i$. Furthermore, $R(\alpha_2) \cap (-p, p) = R(\alpha_1)$, and $R(\alpha_2) \cap (p, q) = R(\alpha_1')$, and $L(\alpha_2) \cap (p, q) = L(\alpha_1')$. Comparing (4.2) with (4.1) yields $\Sigma_1 \cap \Sigma_1' = \{\mathbf{x} \in \Sigma_2 : x_p = 0\}$.

Conversely, if $\Sigma_1 \cap \Sigma_1' \subseteq \{\mathbf{x} \in \Sigma_2 : x_p = 0\}$, then by (4.1) and (4.2), we know that $(R(\alpha_2) \cup (-L(\alpha_2))) \subseteq (R(\alpha_1) \cup R(\alpha_1') \cup (-L(\alpha_1')))$. In this case, no point can be between α_2 and $-\alpha_2$, because $R(\alpha_1)$ contains no pairs $\pm i$. Thus $R(\alpha_2)$ and $-L(\alpha_2)$ coincide along the interval $(-p, p)$. Therefore $R(\alpha_2) \setminus \{p\} \subseteq R(\alpha_1) \cup R(\alpha_1')$ and $L(\alpha_2) \setminus \{p\} \subseteq (-R(\alpha_1)) \cup L(\alpha_1') = L(\alpha_1) \cup L(\alpha_1')$, so A_1 is a subarc of A_2 and A_1' is a subarc pair of A_2.

We have shown that arrows of the form $A_1 \to A_2$ and $A_1' \to A_2$ in Subcase 2b are precisely the arrows described in conditions (iv) and (v). □

Proof of Theorem 4.9. Both directions of the theorem use Proposition 4.13. Let A_1 be the symmetric arc/pair corresponding to j_1 and A_2 be the symmetric arc/pair corresponding to j_2. We continue to use the notation $A_1 \to A_2$ to mean $j_1 \to j_2$.

If j_1 forces j_2, then there is a sequence of arrows from j_1 to j_2. Thus Propositions 4.12 and 4.13 imply that A_1 is a subarc/subarc pair of A_2.

Conversely, suppose that A_1 is a subarc/subarc pair of A_2, with $A_1 \neq A_2$. We will show that, in every case, there is a sequence of one, two, or three arrows from j_1 to j_2. For the purposes of this proof, we will refer to the kinds of arrows described in Proposition 4.13 as "type (i)", etc.

Case 1. A_2 is a symmetric arc.

Case 1a. A_1 is also a symmetric arc. In this case, $A_1 \to A_2$ by an arrow of type (ii).

Case 1b. A_1 is a non-overlapping symmetric pair. In this case, suppose that the endpoints of A_2 are $\pm p$ with $p > 0$ and the endpoints of one arc of A_1 are p' and q' with

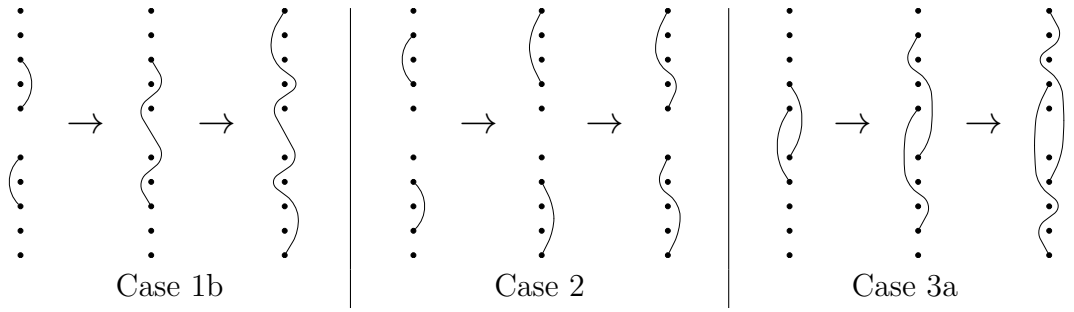


Figure 8: Some illustrations of the proof of Theorem 4.9

$0 < p' < q' \leqslant p$. Then there is an arrow of type (iii) from A_1 to a symmetric arc A' with endpoints $\pm q'$ that is a subarc of A_2, and (unless $A' = A_2$) an arrow of type (ii) from A' to A_2, as illustrated in the left picture of Figure 8.

Case 2. A_2 is a nonoverlapping symmetric pair. In this case, A_1 is also a nonoverlapping symmetric pair. If the outer endpoints of A_1 and A_2 are not the same, then there is an arrow of type (i) from A_1 to an arc pair A' whose outer endpoints agree with those of A_2; otherwise, let $A' = A_1$. If the inner endpoints of A' and A_2 agree then $A' = A_2$. If not, then there is another arrow of type (i) from A' to A_2, as illustrated in the center picture of Figure 8.

Case 3. $A_2 = \{\alpha_2, -\alpha_2\}$ is an overlapping symmetric pair. Write p and q with $p < q$ for the endpoints of α_2.

Case 3a. $A_1 = \{\alpha_1, -\alpha_1\}$ is also an overlapping symmetric pair. In this case, assume that α_2 is right of $-\alpha_2$, that α_1 is right of $-\alpha_1$, and that α_1 has endpoints p' and q' with $p \leqslant p' < 0 < q' \leqslant q$. If $p < p' < 0 < q' < q$, then we find a sequence of two arrows of type (i) from A_1 to A_2. If $q' = q$ and/or $p = p'$, then one or both of these arrows is replaced by equality. If $-q > p$, then also $-q' > p$, so there is a type (i) arrow from A_1 to a subarc pair $A' = \{\alpha', -\alpha'\}$ of A_2 such that α' has endpoints p and q' and an arrow of type (i) from A' to A_2, as illustrated in the right picture of Figure 8. If $p > -q$, then also $p' > -q$, so there is a type (i) arrow from A_1 to a subarc pair $A' = \{\alpha', -\alpha'\}$ of A_2 such that α' has endpoints p' and q and an arrow of type (i) from A' to A_2. (Separating into two cases $-q > p$ and $p > -q$ is necessary. For example, when $-q > p$, it is possible that $p' = -q$, so that there is no subarc pair $A' = \{\alpha', -\alpha'\}$ of A_2 such that α' has endpoints p' and q. This is the case illustrated in the right picture of Figure 8.)

Case 3b. $A_1 = \{\alpha_1, -\alpha_1\}$ is a non-overlapping symmetric pair. In this case, assume that α_1 is a subarc of α_2 (making no assumption about which of α_2 or $-\alpha_2$ is to the right). Write p' and q' for the endpoints of α_1, this time with $p < 0 < p' < q' \leqslant q$. If $p' \neq -p$, then there is a subarc pair A' of A_2 with endpoints p' and q and an arrow $A' \to A_2$ of type (i). If $p' = -p$, then there exists a subarc pair $A' = \{\alpha', -\alpha'\}$ of A_2 such that α' has endpoints 1 and q and an arrow from A' to A_2. The arrow is of type (i) if $p < -1$ or (v) if $p = -1$. Either way, A_1 is a subarc of A' and by Case 2, there is a (possibly empty) sequence of arrows from A_1 to A'.

Case 3c. A_1 is a symmetric arc α_1. Write p' and $-p'$ with $p \leqslant p' < 0 < -p' \leqslant q$ for the endpoints of α_1. If $p' = p$ or $-p' = q$, then there is an arrow of type (iv) from A_1 to A_2, so it remains to consider the case where $p < p' < 0 < -p' < q$. Without loss of generality (up to swapping $\pm\alpha_2$), we may as well assume that $p > -q$. Therefore also $p' > -q$, so there is a type (iv) arrow from A_1 to a subarc pair $A' = \{\alpha', -\alpha'\}$ of A_2 such that α' has endpoints p' and q and an arrow from A' to A_2 of type (i). □

5 Orbifold noncrossing arc diagrams

We now take the centrally symmetric model for Coxeter groups of type B_n and pass to a quotient modulo the central symmetry, obtaining what we call an orbifold model. Most simply—but not very "drawably"—the orbifold model lives in the quotient space where each point of the plane is identified with its antipodal opposite. The quotient map takes a point x in the plane to $\{\pm x\}$.

To make a more "drawable" model, we consider the same space as a different quotient of the plane: We cut the plane in half with a horizontal line through the origin. Points strictly above the horizontal line are not identified with any other points. Each point on or below the horizontal line is identified with all other points on or below the line at the same distance from the origin. Thus, each "point" in the quotient is either a point strictly above the horizontal line (a ***point in the upper halfplane***), the origin, or a semicircle below the line with endpoints on the line (or a degenerate semicircle consisting only of the origin, called a ***semicircle point***). We will refer to this quotient as the ***orbifold plane***.

We are not interested in the natural quotient map associated to the orbifold plane. Rather, we are interested in the map ϕ defined as follows: If x is not on the horizontal line, then $\phi(x)$ is the point in $\{\pm x\}$ that is above the horizontal line. If x is on the horizontal line, then $\phi(x)$ is the semicircle of points of length $|x|$ on or below the horizontal line (or $\phi(0)$ is the degenerate semicircle at the origin).

A symmetric arc α or pair $\{\alpha, -\alpha\}$ in the centrally symmetric model can be uniquely recovered from $\phi(\alpha)$. Thus also symmetric arc diagrams and the various results and constructions in Section 4 can be recovered from their images under ϕ. The goal now is to define arcs, arc diagrams, a bijection to B_n, etc. in the orbifold plane so that the results of Section 4 translate to results in the orbifold plane via the map ϕ. With the right definitions, these new results are simply "translations" into a new setting, and will not require new proofs. Collectively, we will refer to these constructions and results as the ***orbifold model*** for B_n.

The orbifold model for B_n starts with n distinct points on a vertical line containing the origin, with each point strictly above the origin. The origin itself is called the ***orbifold point***, and is marked with an "×". (We will sometimes also refer, in prose, to the orbifold point as "×".) The n points above the origin are identified with the numbers $1, 2, \ldots, n$, in order, with n at the top, and are called ***numbered points***.

A ***(type-B) orbifold arc*** (or in context simply an ***arc***) is a curve in the orbifold plane with each endpoint at a numbered point or at × (the origin), satisfying one of the

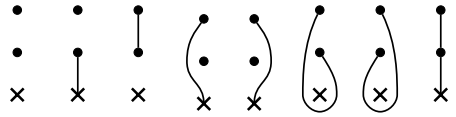


Figure 9: Orbifold noncrossing arc diagrams for $n = 2$

following three descriptions:

- An ***ordinary arc*** is an arc on the points $1, \ldots, n$ satisfying the same rules as an arc in type A (Section 3).
- A ***central arc*** is an arc with one endpoint at $\times$ and the other at a numbered point p, moving monotone downwards from p to $\times$ without touching any other numbered point, passing to the left or right of any numbered points below p.
- A ***long arc*** is an arc α containing exactly one semicircle point, which is not degenerate (i.e. not equal to the orbifold point). The ***left piece*** of α moves monotone downward from the ***left endpoint*** of α and hits the semicircle point left of $\times$, passing left or right of numbered points between. The ***right piece*** moves monotone downward from the ***right endpoint*** and hits the semicircle point right of $\times$, passing left or right of numbered points between. The left and right pieces do not intersect and their endpoints do not coincide.

An ordinary arc or central arc α is specified combinatorially by its endpoints and the set $R(\alpha)$ of numbered points to its right, or equivalently by its endpoints and the set $L(\alpha)$ of numbered points to its left. A long arc α is specified combinatorially by its endpoints (right and left), the set $R(\alpha)$ of points right of its right piece *and* the set $L(\alpha)$ of points left of its left piece. We emphasize that for long arcs (in contrast to the situation for ordinary and central arcs), the sets $R(\alpha)$ and $L(\alpha)$ do not completely determine each other, because points not in $R(\alpha)$ could be in $L(\alpha)$ or between the left and right pieces. However, $R(\alpha)$ and $L(\alpha)$ must be disjoint, and both sets are necessary to determine α combinatorially. We also emphasize that to determine α, it is necessary to specify which endpoint is right and which is left. (For example, for each pair of distinct endpoints, there are two long arcs α with $R(\alpha) = \emptyset$ and $L(\alpha) = \emptyset$, depending on which endpoint is right and which is left.)

An ***orbifold noncrossing arc diagram*** is a collection of orbifold arcs on points $\times$ and $1, \ldots, n$ that don't intersect, except possibly at their endpoints, with no two arcs sharing an endpoint from which they both go down or both go up. Again, we consider arcs and noncrossing arc diagrams up to combinatorial equivalence. All type-B_2 and type-B_3 orbifold noncrossing arc diagrams are shown in Figures 9 and 10. For easy comparison, the diagrams in Figures 3 and 9 appear in the same order.

The map ϕ is a bijection from the set of centrally symmetric arcs and centrally symmetric pairs of arcs to the set of orbifold arcs, and also induces a bijection on combinatorial equivalence classes. Nonoverlapping symmetric pairs, symmetric arcs, and overlapping

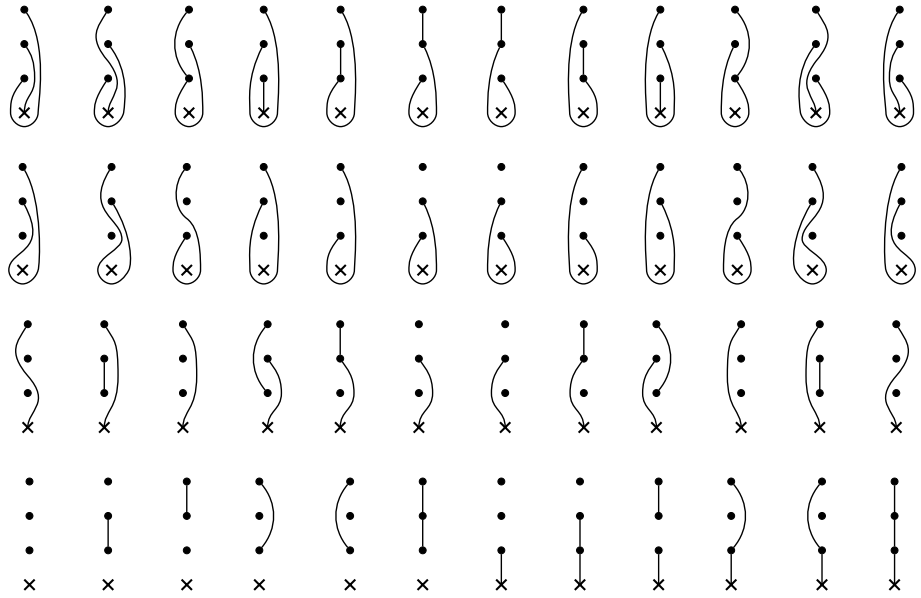


Figure 10: Orbifold noncrossing arc diagrams for $n = 3$

symmetric pairs map respectively to ordinary, central, and long arcs. Furthermore, ϕ induces a bijection from centrally symmetric noncrossing arc diagrams to orbifold noncrossing arc diagrams. We write δ° for the map $\phi \circ \delta$, where δ is understood as in Section 4 to map signed permutations to centrally symmetric noncrossing arc diagrams. The following theorem is a restatement of Theorem 4.1.

Theorem 5.1. *The map δ° is a bijection from B_n to the set of orbifold noncrossing arc diagrams.*

We now give a direct description of the bijection δ° from B_n to the set of orbifold noncrossing arc diagrams on n points, as illustrated in Figure 11. Given $\pi \in B_n$ with one-line notation $\pi_1 \cdots \pi_n$, write each entry π_i at the point (i, π_i) in the plane, for $i = 1, \ldots, n$. For every i such that $\pi_i > \pi_{i+1}$, draw a straight line segment from π_i to π_{i+1}. Additionally, if $\pi_1 < 0$, draw a line segment from the origin to π_1. These line segments become arcs: First, we move the numbers $\pi_1, \ldots, \pi_n$ horizontally to put them into a single vertical line, allowing the line segments to curve, so that they don't pass through any of the numbers or one another. Then, we rotate the negative numbers 180 degrees clockwise, allowing the (already curved) line segments connecting positive numbers to negative numbers to stretch around the origin. We remove the negative signs on the numbers. Since $\{|\pi_1|, \ldots, |\pi_n|\} = \{1, \ldots, n\}$, the numbers are now $1, \ldots, n$ from bottom to top. The resulting orbifold noncrossing arc diagram is $\delta^\circ(\pi)$.

We continue to translate the results of Section 4 into the language of the orbifold model. We next describe the bijection from arcs to join-irreducible elements of B_n. We continue the conventions from Section 4 for describing join-irreducible signed permutations.

A join-irreducible signed permutation π with a single descent $\pi_{-1} > \pi_1$ corresponds

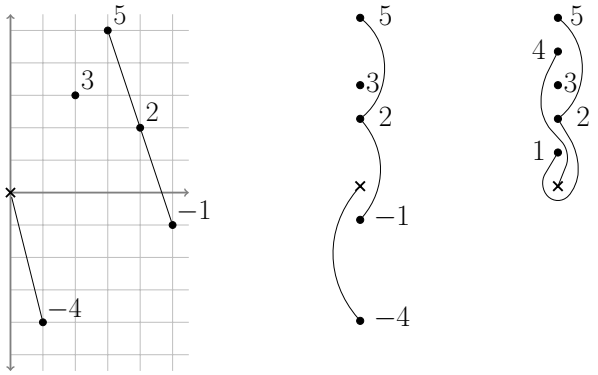


Figure 11: The map δ° applied to $(-4)352(-1)$

to a central arc α. If α has upper endpoint p, then π is

$$(-p)\,(-L(\alpha))\,R(\alpha)\,(p+1)\cdots n.$$

Every other join-irreducible signed permutation π has a symmetric pair of descents $\pi_i > \pi_{i+1}$ and $\pi_{-i-1} > \pi_{-i}$, with $i > 0$. In this case, π_i must be positive, because if π_i is negative, then π must have additional descents. (If π_i is negative, then to avoid additional descents in the short one-line notation, the entries $\pi_{i-1}, \ldots, \pi_1$ are also negative. But then π has the additional descent $\pi_{-1} > \pi_1$.)

If $\pi_i > \pi_{i+1} > 0$, then π corresponds to an ordinary arc α. If α has endpoints p and q with $0 < p < q$, then π is

$$1\cdots(p-1)\,L(\alpha)\,q\,p\,R(\alpha)\,(q+1)\cdots n.$$

If $\pi_i > 0 > \pi_{i+1}$, then π corresponds to a long arc α. Suppose α has left endpoint p and right endpoint q (noting that as we translate this description from Section 4, the p here is the $-p$ there). Necessarily, $p \neq q$. Write $B(\alpha)$ (suggesting "between") to denote the set of points that are left of the right piece of α and right of the left piece. Thus $B(\alpha)$ is a subset of $(0, \min(p,q))$. If $p < q$, then π is

$$B(\alpha)\,[(p,q) \setminus R(\alpha)]\,q\,(-p)\,(-L(\alpha))\,R(\alpha)(q+1)\cdots n.$$

If $q < p$, then π is

$$B(\alpha)\,q\,(-p)\,(-L(\alpha))\,R(\alpha)\,[(q,p) \setminus L(\alpha)]\,(p+1)\cdots n.$$

The following is a restatement of Theorem 4.2 in the orbifold model.

Theorem 5.2. *Given $\pi \in B_n$, the canonical join representation of π is the set of join-irreducible elements corresponding to the arcs in $\delta^\circ(\pi)$.*

Two orbifold arcs are ***compatible*** if and only if they do not intersect, except possibly at numbered endpoints, and don't share the same top endpoint or the same bottom

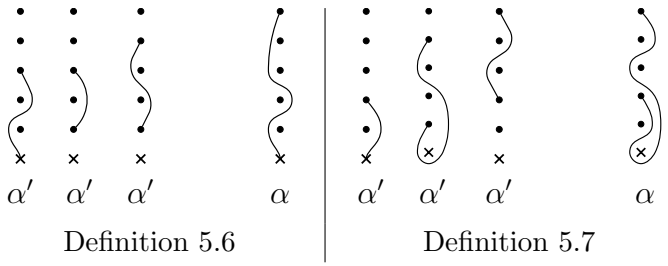


Figure 12: Subarcs

endpoint. In other words, they are compatible if and only if they can appear together in an orbifold noncrossing arc diagram. As a consequence of Theorems 5.1 and 5.2, two join-irreducible elements of B_n are compatible (can appear together in a canonical join representation) if and only if the corresponding orbifold arcs are compatible.

The following three propositions are restatements of Propositions 4.3 and 4.4. (Propositions 5.4 and 5.5 break Proposition 4.4 into two cases.)

Proposition 5.3. *Suppose α is a central arc with upper endpoint p. The shard associated to α is*

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = 0 \text{ and } 0 \leqslant x_i \,\forall i \in R(\alpha) \text{ and } 0 \geqslant x_i \,\forall i \in L(\alpha)\}.$$

Proposition 5.4. *Suppose α is an ordinary arc with endpoints p and q. The shard associated to α is*

$$\{\mathbf{x} \in \mathbb{R}^n : x_p = x_q \text{ and } x_p \leqslant x_i \,\forall i \in R(\alpha) \text{ and } x_p \geqslant x_i \,\forall i \in L(\alpha)\}.$$

Proposition 5.5. *Suppose α is a long arc with left endpoint p and right endpoint q. The shard associated to α is*

$$\big\{\mathbf{x} \in \mathbb{R}^n : x_q = -x_p \text{ and } x_q \leqslant x_i \,\forall i \in (R(\alpha) \cup (-L(\alpha))) \\ \text{and } x_q \geqslant x_i \,\forall i \in ((-p, q) \setminus (R(\alpha) \cup (-L(\alpha)))\big\}.$$

We now turn to translating the definitions of subarcs and superarcs, as well as Theorem 4.9, to the orbifold model, using the map ϕ. There is a case-free way to define subarcs of an orbifold arc α: Informally, we can cut α shorter at one or both of its endpoints and reattach the shortened curve to a numbered point. If the result is an orbifold arc (or can be easily modified to become a central arc as made precise below), then it is a subarc of α. Together, Definitions 5.6 and 5.7, below, correspond to Definitions 4.6 and 4.7, but the split into two definitions is not the same here as in Section 4.

Definition 5.6 (Subarcs of ordinary and central arcs)**.** Suppose α is an ordinary or central arc with endpoints $p < q$, with $\times = 0$ for the purposes of this definition. A ***subarc*** of α is an arc α' with endpoints p' and q' having $p \leqslant p' < q' \leqslant q$ and $R(\alpha') = R(\alpha) \cap (p', q')$. The arc α' is ordinary unless $p' = 0$, in which case α' is a central arc.

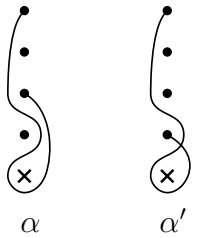


Figure 13: A failed construction of a subarc, per Definition 5.7 and Remark 5.8

Definition 5.7 (Subarcs of long arcs)**.** Suppose α is a long arc with left endpoint p and right endpoint q. A ***subarc*** of α can be ordinary, central, or long.

- An ordinary arc α' with endpoints $p' < q'$ is a subarc of α if and only if one of the following occurs: either $q' \leqslant p$ and $L(\alpha') = L(\alpha) \cap (p', q')$ or $q' \leqslant q$ and $R(\alpha') = R(\alpha) \cap (p', q')$.
- A central arc α' with upper endpoint p' is a subarc of α if and only if $p' \leqslant \min(p, q)$, $L(\alpha') = L(\alpha) \cap (0, p')$, and $R(\alpha') = R(\alpha) \cap (0, p')$.
- A long arc α' with left endpoint p' and right endpoint q' is a subarc of α if and only if $p' \leqslant p$, $q' \leqslant q$, $L(\alpha') = L(\alpha) \cap (0, p')$, $R(\alpha') = R(\alpha) \cap (0, q')$, and the following additional conditions hold: $p' \notin R(\alpha) \cap (0, q')$ and $q' \notin L(\alpha) \cap (0, p')$.

Remark 5.8. To better understand long subarcs of long arcs in Definition 5.7, we discuss an explicit construction of these subarcs. Compare similar constructions in Remarks 3.5 and 4.8. Take h to be a height function on the orbifold plane with each numbered point i having height i and with every semicircle point in the closed lower halfplane having height 0.

Suppose α is a long arc, parametrized as a function from the interval $[0, 1]$ into the orbifold plane. The construction of a long subarc α' begins as follows: Choose t_1 and t_2 with $0 \leqslant t_1 < t_2 \leqslant 1$ such that $h(\alpha(t_1))$ and $h(\alpha(t_2))$ are distinct elements of $\{1, \ldots, n\}$. Choose $\varepsilon_1 > 0$ such that $|h(\alpha(t_1 + \varepsilon_1)) - h(\alpha(t_1))| < 1$ and such that $h(\alpha(t)) > 0$ for all t in the interval $(t_1, t_1 + \varepsilon_1)$. Similarly, choose $\varepsilon_2 > 0$ such that $t_1 + \varepsilon_1 < t_2 - \varepsilon_2$, such that $|h(\alpha(t_2 - \varepsilon_2)) - h(\alpha(t_2))| < 1$, and such that $h(\alpha(t)) > 0$ for all t in the interval $(t_2 - \varepsilon_2, t_2)$. Define α' to be the *curve* obtained by concatenating three curves: First, the straight line segment from the point numbered $h(\alpha(t_1))$ to the point $\alpha(t_1 + \varepsilon_1)$; second, the restriction of α to the interval $[t_1 + \varepsilon_1, t_2 - \varepsilon_2]$; and third, the straight line segment from $\alpha(t_2 - \varepsilon_2)$ to the point numbered $h(\alpha(t_2))$.

If α' is a valid long arc, then it is a subarc of α. It fails to be an orbifold arc if and only if it crosses itself, in which case, this choice of t_1 and t_2 does not produce a subarc of α. See Figure 13.

The map ϕ translates the subarc relation on symmetric arcs/pairs exactly to the subarc relation on orbifold arcs. The following results are direct translations of Theorem 4.9 and Corollaries 4.10 and 4.11.

Theorem 5.9. *Let j_1 and j_2 be join-irreducible signed permutations. Then j_1 forces j_2 if and only if the orbifold arc corresponding to j_1 is a subarc of the orbifold arc corresponding to j_2.*

Corollary 5.10. *A set U of orbifold arcs corresponds to the set of* uncontracted *join-irreducible signed permutations of some congruence Θ on B_n if and only if U is closed under passing to subarcs.*

Corollary 5.11. *If Θ is a congruence on B_n and U is the set of orbifold arcs corresponding to join-irreducible permutations* not *contracted by Θ, then δ° restricts to a bijection from the quotient B_n/Θ (the set of signed permutations not contracted by Θ) to the set of orbifold noncrossing arc diagrams consisting only of arcs in U.*

An orbifold arc α' is a ***superarc*** of an orbifold arc α if and only if α is a subarc of α'. We now describe how to construct superarcs.

Superarcs of a central arc $\boldsymbol{\alpha}$. These can be central arcs or long arcs. We construct an orbifold superarc by pushing α left or right of its top endpoint and extending it upwards to a new upper endpoint. We construct a long superarc by first replacing α with a curve that goes around $\times$, with both endpoints at the original endpoint of α, and with its left piece and its right piece having exactly the same right points as α. We then push one or both ends off of the endpoint (one left and one staying, one right and one staying, both left, one on each side, or both right, but not creating a self-intersection), and extend upwards without creating self-intersections, to make a long superarc.

Superarcs of an ordinary arc $\boldsymbol{\alpha}$. These can be of any of the three kinds. To construct a superarc, we push the top or bottom of α independently to the left or right. We then extend upwards and/or downwards. The downward extension can end before reaching $\times$ (to make an ordinary superarc), can have an endpoint at $\times$ (to make an orbifold superarc), or can go around $\times$ before ending (to make a long superarc), provided that no self-intersections are created.

Superarcs of a long arc $\boldsymbol{\alpha}$. Every superarc of α is long. We construct a superarc by pushing one or both endpoints of α left or right and extending upwards, without creating self-intersections.

We close this section with a translation of Proposition 4.13, which will be useful in the sequel. As before, we write arrows $\to$ between orbifold arcs to indicate arrows between the corresponding shards.

Proposition 5.12. *Suppose α_1 and α_2 are orbifold arcs. Then $\alpha_1 \to \alpha_2$ if and only if α_1 is a subarc of α_2 and one of the following conditions holds.*

(i) *Neither α_1 nor α_2 is a central arc, and α_1 and α_2 have exactly one endpoint in common.*

(ii) *α_1 and α_2 are central arcs with different upper endpoints.*

(iii) *α_2 is a central arc and α_1 is ordinary, with the same upper endpoint.*

(iv) *α_2 is a long arc with endpoints $p < q$ and no numbered points between its left and right pieces, and α_1 is a central arc with upper endpoint p.*

(v) *α_2 is a long arc with endpoints $p < q$ and no numbered points between its left and right pieces, and α_1 is an ordinary arc with endpoints p and q.*

6 Congruence examples

In this section, we demonstrate the use of Corollaries 5.10 and 5.11 by characterizing some important lattice congruences of the weak order of type B. Throughout this section, it will be convenient to say that a congruence "contracts an arc" as shorthand for "contracts the join-irreducible element corresponding to that arc".

6.1 Parabolic congruences

We warm up with the simplest example, verifying easy known results. Given a Coxeter system (W, S), each element of S is in particular a join-irreducible element of the weak order on W. A ***parabolic congruence*** of the weak order on a Coxeter group W is a congruence generated by contracting some subset J of S. (See [14, Section 6] or [21, Example 10-7.1].) The quotient of the weak order on W modulo the parabolic congruence generated by J is isomorphic to the weak order on the parabolic subgroup $W_{S\setminus J}$ [14, Corollary 6.10]. To understand parabolic congruences on W, it is enough to understand the case where J is a singleton.

When (W, S) is of type B_n, the elements of S and the associated arcs are as follows: The simple reflection $s_0 = (-1)2\cdots n$ is the central arc with upper endpoint 1, and for $i = 1, \ldots, n-1$, the simple reflection $s_i = 1\cdots(i-1)(i+1)i(i+2)\cdots n$ is the ordinary arc connecting i to $i+1$.

The parabolic congruence generated by contracting s_i contracts all superarcs of the corresponding arc. For s_0, the set of these superarcs consists of all central arcs and all long arcs. Thus the elements of the quotient are the orbifold noncrossing arc diagrams that only have ordinary arcs. This is in accordance with [14, Corollary 6.10], which implies in this case that the quotient is isomorphic to S_n. For s_i with $i > 0$, these superarcs are the ordinary arcs connecting a point at or below i to a point strictly above i, the central arcs with upper endpoint strictly above i, and the long arcs such that one or both endpoints are strictly above i. Again, this accords with [14, Corollary 6.10], which implies that the quotient is isomorphic to $B_i \times S_{n-i}$.

6.2 Surjective homomorphisms from B_n to S_{n+1}

A surprising result of [19] is that surjective lattice homomorphisms exist between the weak orders on different finite Coxeter groups *of the same rank*. Such homomorphisms always restrict to bijections between the corresponding sets of simple reflections. Thus, we take

(W, S) and (W', S') to be finite Coxeter groups with the same rank $|S| = |S'|$ and fix a bijection $\phi : S \to S'$. According to [19, Theorem 1.6], there exists a surjective lattice homomorphism from W to W' restricting to ϕ if and only if $m(r, s) \geqslant m(\phi(r), \phi(s))$ for every $r, s \in S$.

It is not hard to see that a surjective lattice homomorphism from W to W' restricts, for any $r, s \in S$, to a surjective lattice homomorphism from the lower interval $[1, r \vee s]$ to the lower interval $[1, \phi(r) \vee \phi(s)]$. (For a more general statement, see [19, Proposition 2.6].) Less obviously a surjective homomorphism is uniquely determined by all of these restrictions [19, Corollary 1.8].

We consider the case where $W = B_n$ and $W' = S_{n+1}$ (see [19, Section 6]) and fix ϕ to be the bijection sending s_i to s_{i+1} for $i = 0, \ldots, n-1$. In this case, $m(r, s) = m(\phi(r), \phi(s))$ for almost every pair $r, s \in S$, and for those pairs the only possibility is that the restriction of the homomorphism to $[1, r \vee s]$ is an isomorphism. The one exception is that $m(s_0, s_1) = 4$ and $m(\phi(s_0), \phi(s_1)) = m(s_1, s_2) = 3$. The interval $[1, s_0 \vee s_1]$ is an octagon and the interval $[1, s_1 \vee s_2]$ is a hexagon. There are precisely 4 surjective lattice homomorphisms from the octagon to the hexagon. Specifically, the congruence associated to a surjective homomorphism from the octagon to the hexagon contracts exactly one of the join-irreducible elements $s_0 s_1$ or $s_0 s_1 s_0$ and exactly one of the join-irreducible elements $s_1 s_0$ or $s_1 s_0 s_1$. As a consequence, for each for $n \geqslant 2$, there are precisely 4 surjective lattice homomorphisms from B_n to S_{n+1} that restrict to ϕ, each contracting exactly one of $s_0 s_1$ or $s_0 s_1 s_0$ and exactly one of $s_1 s_0$ or $s_1 s_0 s_1$.

For each surjective homomorphism η from B_n to S_{n+1}, there is a natural injection from S_{n+1} into B_n, mapping each element $\pi \in S_{n+1}$ to the bottom element of $\eta^{-1}(\pi)$. In [26] (see also [23]), these injections for each of the 4 homomorphisms are characterized in terms of noncrossing arc diagrams of types A and B, and 3 of the injections are shown to embed the shard intersection order on S_{n+1} as a sublattice of the shard intersection order on B_n. Here, we show how the subarc/superarc tools developed in this paper recover results about these 4 homomorphisms much more easily than in [19].

Specifically, we construct, in terms of orbifold noncrossing arc diagrams, four congruences, each contracting exactly one of $s_0 s_1$ or $s_0 s_1 s_0$ and exactly one of $s_1 s_0$ or $s_1 s_0 s_1$. In three cases, contracting these two join-irreducible elements generates a congruence associated to surjective homomorphisms from B_n to S_{n+1}. The exception is the congruence generated by contracting $s_0 s_1 s_0$ and $s_1 s_0$. In this case, to get the congruence associated to a homomorphism from B_n to S_{n+1}, we must contract additional join-irreducible elements. Everything we need to know about these congruences follows from some easy facts about arcs and superarcs:

- $s_0 s_1 = 2(-1)3 \cdots n$ corresponds to the long arc with left endpoint 1 and right endpoint 2, with right piece passing right of 1. Its superarcs are the long arcs whose right piece passes right of 1.

- $s_0 s_1 s_0 = (-2)(-1)3 \cdots n$ corresponds to the central arc with upper endpoint 2, passing right of 1. Its superarcs are the central arcs that pass right of 1 and the long arcs both of whose pieces pass right of 1.

- $s_1s_0 = (-2)13\cdots n$ corresponds to the central arc with upper endpoint 2, passing left of 1. Its superarcs are the central arcs that pass left of 1 and the long arcs both of whose pieces pass left of 1.

- $s_1s_0s_1 = 1(-2)3\cdots n$ corresponds to the long arc with right endpoint 1 and left endpoint 2, with left piece passing left of 1. Its superarcs are the long arcs whose left piece passes left of 1.

- $s_1s_0s_1s_2 = 13(-2)4\cdots n$ corresponds to the long arc with left endpoint 2 and right endpoint 3, with left piece passing left of 1 and right piece passing right of 1 and 2. Its superarcs are the long arcs whose left piece passes left of 1 and whose right piece passes right of 1 and 2.

- $s_2s_1s_0s_1s_2 = 12(-3)4\cdots n$ corresponds to the long arc with right endpoint 2 and left endpoint 3, with right piece passing right of 1 and left piece passing left of 1 and 2. Its superarcs are the long arcs whose right piece passes right of 1 and whose left piece passes left of 1 and 2.

These easy facts, along with Corollary 5.10, lead almost immediately to the proofs of the following three propositions, which are [19, Proposition 6.3], [19, Proposition 6.6], and [19, Proposition 6.9], translated from the language of shards to the language of orbifold arcs. The missing fourth proposition (describing the congruence generated by contracting s_0s_1 and s_1s_0) is symmetric to the third by swapping left and right, and this symmetry is also exploited in [19].

Proposition 6.1. *The congruence generated by contracting s_0s_1 and $s_1s_0s_1$ contracts the long arcs and no other arcs.*

Proof. This congruence contracts precisely the arcs that are long, with right piece passing right of 1 or left piece passing left of 1, or in other words, all long arcs. □

Proposition 6.2. *The congruence generated by contracting $s_0s_1s_0$, s_1s_0, $s_1s_0s_1s_2$, and $s_2s_1s_0s_1s_2$ contracts an arc if and only if it is not ordinary and only has endpoints $\geqslant 2$.*

Proof. This congruence contracts no ordinary arcs. It contracts a central arc if and only if it passes left or right of 1, or in other words, if and only if its upper endpoint is $\geqslant 2$. It contracts a long arc if and only if both of its pieces pass right of 1 or both of its pieces pass left of 1 or its left piece passes left of 1 and its right piece passes right of 1. In other words, it contracts a long arc if and only if its endpoints are both $\geqslant 2$. □

Proposition 6.3. *The congruence generated by contracting $s_0s_1s_0$ and $s_1s_0s_1$ contracts an arc if and only if it is a central arc passing right of* 1 *or a long arc whose left endpoint is not* 1*.*

Proof. This congruence also contracts no ordinary arcs. It contracts a central arc if and only if it passes right of 1. It contracts a long arc if and only if both of its pieces pass right of 1 or its left piece passes left of 1. In other words, it contracts a long arc if and only if its left endpoint is not 1. □

6.3 Cambrian congruences

Let (W, S) be a finite Coxeter system. A Cambrian congruence on the weak order on W is defined by choosing a set of join-irreducible elements to contract in the following way: For each edge in its Coxeter diagram, connecting distinct simple reflections $r, s \in S$ and labeled $m(r, s)$, choose either to contract every join-irreducible element $rs, rsr, rsrsr, \ldots$ up to length $m(r, s) - 1$ or to contract every join-irreducible element $sr, srs, srsr, \ldots$ up to length $m(r, s) - 1$. (Since most edges of a Coxeter diagram of finite type are labeled 3, for most edges the choice is between one join-irreducible element and another.) The ***Cambrian congruence*** is the congruence generated by contracting the chosen join-irreducible elements. (Often, these choices are described as the choice of a Coxeter element of W or as the choice of an orientation of the Coxeter diagram, but we do not need to be so specific here.) The quotient of the weak order modulo a Cambrian congruence is called a ***Cambrian lattice***.

To give context to Cambrian lattices of type B, we review the description of Cambrian lattices of type A in terms of noncrossing arc diagrams, following [18, Example 4.9], and add some straightforward observations. Each numbered point from 2 to $n-1$ is designated as either a left point or a right point. Designating i as a right point corresponds to contracting the join-irreducible element $s_{i-1}s_i$, while designating i as a left point corresponds to contracting $s_i s_{i-1}$. The superarcs of these chosen contracted join-irreducible elements are precisely the arcs that pass left of a left point or right of a right point. Corollaries 3.7 and 3.8 imply that δ restricts to a bijection from the elements of the Cambrian lattice to the noncrossing arc diagrams containing arcs that only pass right of left points and only pass left of right points.

If there are no right points or if there are no left points, then the resulting Cambrian lattice is isomorphic to the ***Tamari lattice*** [16].

A convenient way to draw noncrossing arc diagrams for elements of the Cambrian lattice is to draw a horizontal "barrier" ray going right from each right point and going left from each left point. Elements of the Cambrian lattice are then noncrossing arc diagrams such that every arc is disjoint from the barriers, except at endpoints. For example, the left picture of Figure 14 shows the noncrossing arc diagram for the permutation 497862153, which is an element of the Cambrian lattice generated by contracting $s_{i-1}s_i$ for $i \in \{2, 3, 5, 8\}$ and $s_i s_{i-1}$ for $i \in \{4, 6, 7\}$.

Another way to draw noncrossing arc diagrams for elements of the Cambrian lattice highlights the bijection between the Cambrian lattice and the set of noncrossing partitions of an appropriate cycle. The numbers 1 and n are placed at the bottom and top of a circle respectively, and the numbers $2, \ldots, n-1$ are placed on the circle, increasing in height as the numbers increase in value, with left points on the left of the circle and right points on the right of the circle. The noncrossing arc diagrams for elements of the Cambrian lattice can then be drawn with "arcs" that never leave the circle (and may as well be straight line segments). For example, the right picture of Figure 14 shows the same arc diagram as the left picture, drawn with points on a circle.

The numbered points define a cycle, starting with 1, passing through the left points in increasing order, then n, then the right points in decreasing order. The bijection

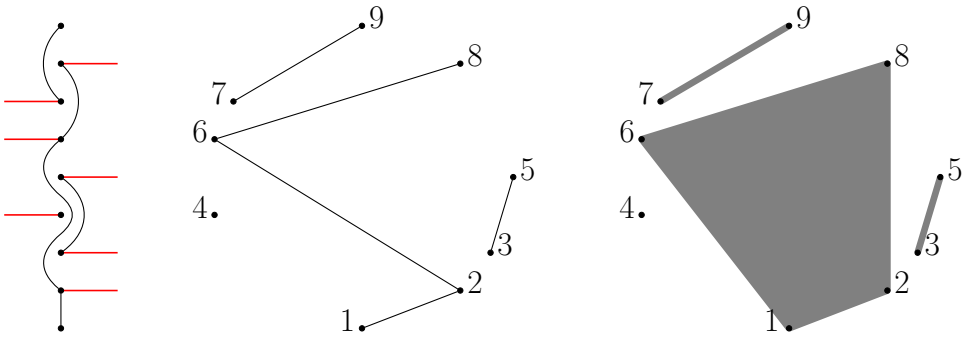


Figure 14: Cambrian lattice and noncrossing partitions, type A

between elements of the Cambrian lattice and noncrossing partitions of the cycle is thus clear: Each connected component of the noncrossing arc diagram contains some set of numbered points, and these sets of numbered points are the blocks of the noncrossing partition. Inversely, given a noncrossing partition $\mathcal{P}$ of the cycle, the noncrossing arc diagram corresponding to $\mathcal{P}$ consists of the following arcs: For each block in $\mathcal{P}$ with elements $b_1 < b_2 < \cdots < b_k$, there is an arc connecting b_{i+1} to b_i for each $i = 1, \ldots, k-1$, with no arc passing left of a left point or right of a right point.

There is another interesting characterization of the Cambrian congruences of type A due to Santocanale and Wehrung [25]. To state this characterization, we recall some facts about the congruence lattice. The congruences of a fixed finite lattice L constitute a distributive lattice $\mathrm{Con}(L)$ under refinement order (and in fact a sublattice of the partition lattice). The join-irreducible congruences are the congruences generated by contracting a single join-irreducible element of the lattice. The meet-irreducible congruences are also indexed by join-irreducible elements of the lattice: A meet-irreducible congruence is the largest (i.e. coarsest) congruence not contracting some join-irreducible element.

Although Cambrian congruences are defined as the join of join-irreducible congruences (the congruences generated by each contracted element), remarkably, when W is of type A, the Cambrian congruences are themselves meet-irreducible congruences. Specifically, [25, Corollary 6.10] says that the Cambrian congruences of the weak order in type A are precisely the *minimal* meet-irreducible congruences of the weak order. This statement also has a simple proof in terms of noncrossing arc diagrams: A meet-irreducible congruence is minimal if and only if it is associated to an arc that has no proper superarcs. After designating the points $2, \ldots, n-1$ as left and right points, there is a unique longest arc α that goes right of every left point and left of every right point. This arc α has no proper superarcs, and the arcs not contracted by the Cambrian congruence are precisely the subarcs of α. In other words, the Cambrian congruence is the meet-irreducible congruence associated to α.

In this section, we use Corollaries 5.10 and 5.11 to characterize Cambrian congruences of type B combinatorially and lattice-theoretically. The lattice-theoretic characterization is related to [25, Corollary 6.10] and characterizes Cambrian congruences in type B as meets of meet-irreducible congruences.

As part of determining a Cambrian congruence of the weak order on B_n, we choose to contract either $s_0s_1 = 2(-1)3\cdots n$ and $s_0s_1s_0 = (-2)(-1)3\cdots n$, or $s_1s_0 = (-2)13\cdots n$ and $s_1s_0s_1 = 1(-2)3\cdots n$. In both cases, the corresponding arcs are a central arc with upper endpoint 2 and a long arc with endpoints 1 and 2. In one case, both arcs pass right of 1 and an arc is a superarc of one or both of them if and only if some part of it passes right of 1. In the other case both arcs pass left of 1, and the superarcs are those that pass left of 1.

To determine a Cambrian congruence, it remains to choose, for each $i \in \{2, \ldots, n-1\}$, either to contract the join-irreducible $s_{i-1}s_i = 1\cdots(i-2)\,i\,(i+1)(i-1)(i+2)\cdots n$ or to contract the join-irreducible $s_is_{i-1} = 1\cdots(i-2)(i+1)(i-1)\,i\,(i+2)\cdots n$. The arc corresponding to $s_{i-1}s_i$ connects $i+1$ to $i-1$ and passes right of i. Its superarcs are precisely the arcs that have some part passing right of i. The arc corresponding to s_is_{i-1} connects $i+1$ to $i-1$ and passes left of i, and its superarcs are the arcs that pass left of i.

We see that choosing a Cambrian congruence on the weak order on B_n corresponds to designating each point from 1 to $n-1$ as a left point or right point. By Corollary 5.10, the Cambrian congruence contracts all arcs that pass right of a right point and/or left of a left point. Thus Corollary 5.11 implies the following theorem.

Theorem 6.4. *Given a Cambrian congruence, encoded by a designation of* $1, \ldots, n-1$ *as left and right points, the map* δ° *on signed permutations restricts to a bijection from the associated Cambrian lattice to the set of orbifold noncrossing arc diagrams consisting of arcs that never pass left of a left point and never pass right of a right point.*

For example, the left pictures of Figure 15 show the orbifold noncrossing arc diagram for two elements of the Cambrian lattice for a particular choice of left and right points. Other important examples arise when there are no right points or there are no left points. In these two cases, the resulting Cambrian lattices are ***type-B Tamari lattices*** in the sense of [16, 27]. The two type-B Tamari lattices are dual to each other.

If one works in the centrally symmetric model (and writes down a centrally symmetric version of Theorem 6.4 in the obvious way), one can connect elements of the Cambrian lattice with centrally symmetric noncrossing partitions by imposing a central symmetry condition throughout the procedure described above and exemplified in Figure 14. Here, we will stay in the orbifold model and use Theorem 6.4 to describe the bijection between elements of the Cambrian lattice and noncrossing partitions of a disk with one orbifold point. (See the discussion in [22, Section 4.2] or make the obvious modification of the discussion in [6, Section 4.2] of noncrossing partitions of a disk with two orbifold points.)

Consider a closed disk D with one distinguished point $\times$ in its interior and numbered points $1, \ldots, n$ on its boundary. Given a subset of the numbered points, an ***embedding*** of the subset is the choice of a closed disk $B \subset D$ containing the numbered points in that block but otherwise disjoint from the boundary of D. The disk B may be disjoint from $\times$ or may contain $\times$ in its interior. (The disk B may not contain $\times$ on its boundary.) A ***noncrossing partition of the disk with one orbifold point*** is a set partition of $\{1, \ldots, n\}$ together with a choice, for each block of the partition, of an embedding, such

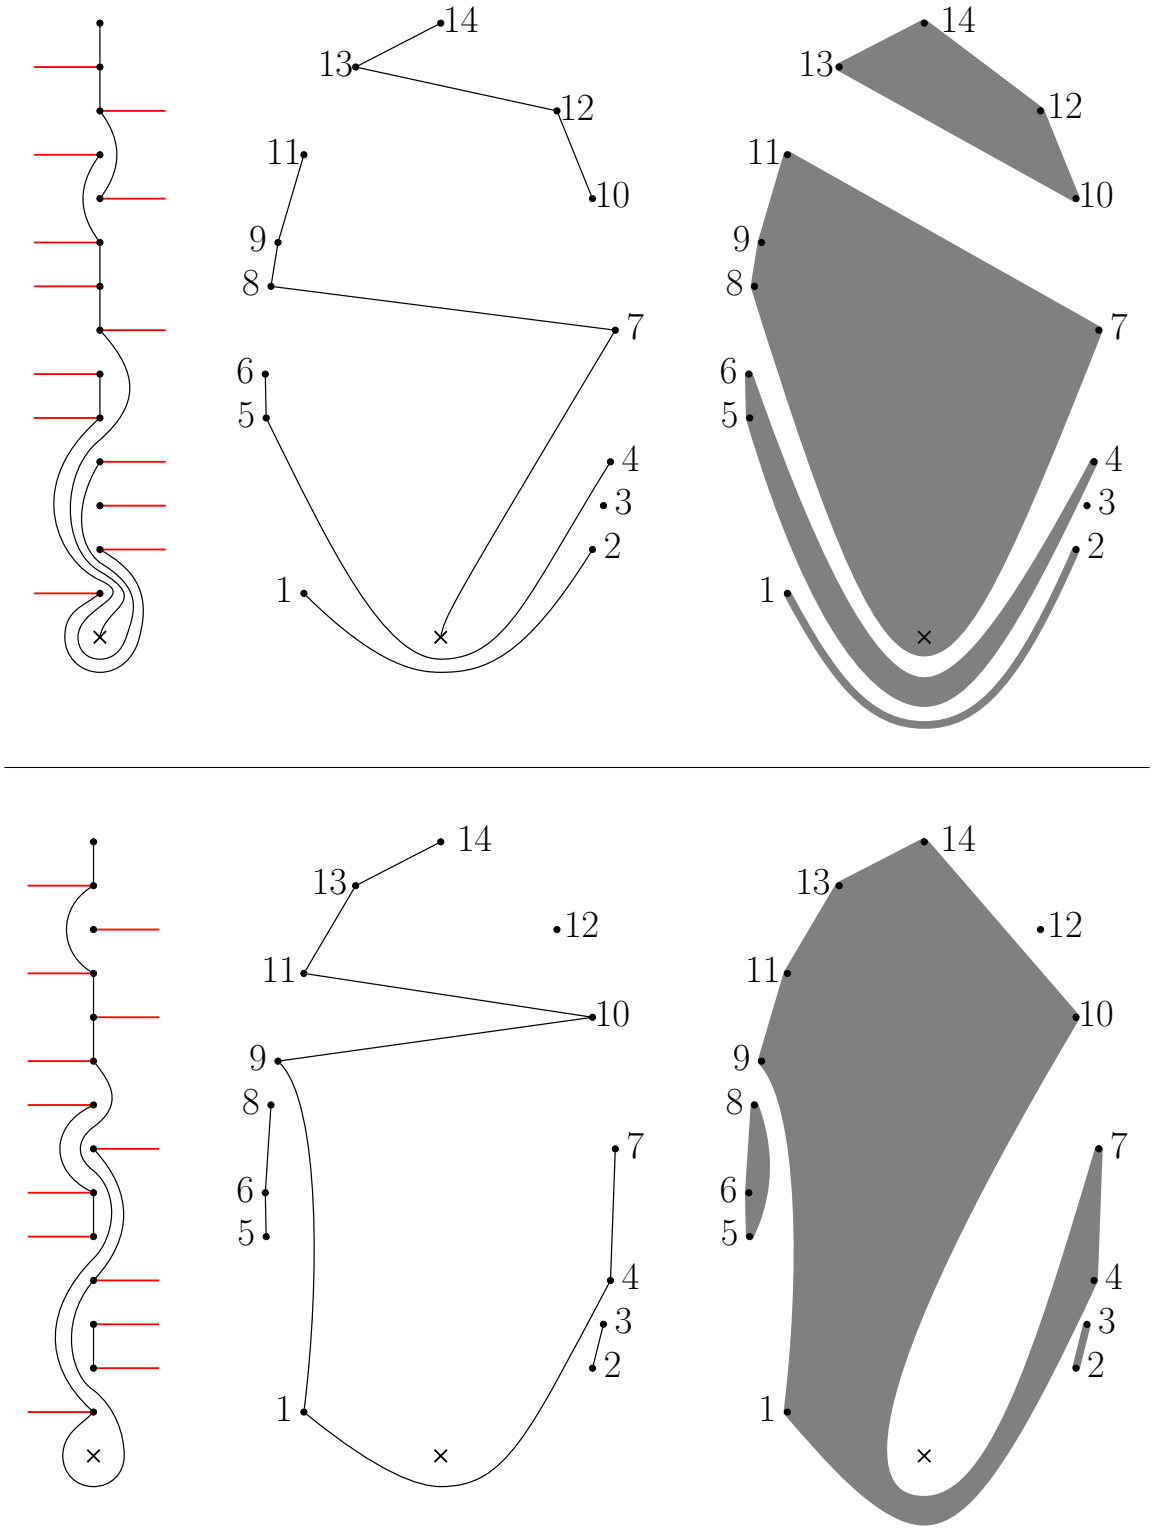


Figure 15: Cambrian lattice and noncrossing partitions, type B

that the embeddings of any two blocks are disjoint. These noncrossing partitions are considered up to isotopy, fixing the boundary of D and the point $\times$. The natural partial order on these noncrossing partitions is the usual type-B noncrossing partition lattice.

The designation of each point in $\{1, \ldots, n-1\}$ as a left point or a right point determines a placement of numbered points on the boundary of D as follows. The number n is placed at the top of D, and the numbers $1, \ldots, n-1$ are placed on the boundary of D, increasing in height as the numbers increase in value, with left points on the left and right points on the right, and $\times$ is placed lower than 1 in the interior of D.

The bijection between elements of the Cambrian lattice and noncrossing partitions of the one-orbifold disk is now apparent, and is illustrated in Figure 15. Each block of the noncrossing arc diagram contains some set of numbered points, and these sets of numbered points are the blocks of the noncrossing partition. The embedding of each block is obtained from the noncrossing arc diagram by moving left points to the left, onto the left boundary of D and right points onto the right boundary, moving the arcs correspondingly, and choosing a disc that contains the arcs in the block (or the unique point in the block if it has no arcs). Inversely, given a noncrossing partition of the one-orbifold disk, we can read off the corresponding orbifold noncrossing arc diagram. Specifically, we read from the noncrossing partition certain pairs of points that should be connected by arcs in the orbifold noncrossing arc diagram. The arcs that connect the pairs must, of course, satisfy the requirement of Theorem 6.4, in that they never pass left of a left point or right of a right point. If there is a block containing $\times$, then we connect the lowest numbered point in the block to $\times$ by a central arc and connect each other numbered point in the block to the next lower numbered point in the block by an ordinary arc. For each block that passes below $\times$, we first treat the block as if it had two pieces, separated at the part of the block that goes below $\times$. We connect each numbered point within each piece (except the lowest numbered point in the piece) to the next lower numbered point in that piece by an ordinary arc. We then connect the lowest point in one piece to the lowest point in the other piece by a long arc. For each block that passes above $\times$, we connect each numbered point in the block, except the lowest, to the next lower numbered point in the block, by an ordinary arc.

Theorem 6.4 also lets us recover (much more easily) the characterization of the type-B Cambrian lattices given in [16]: The Cambrian lattice consists of the signed permutations whose long one-line notation has no patterns 231 such that the "2" is a right point or the negative of a left point. More formally, we have the following theorem, which is [16, Theorem 7.5], restated in the language of this section.

Theorem 6.5. *Given a designation of* $1, \ldots, n-1$ *as left and right points, the associated Cambrian lattice is the subposet of the weak order induced by signed permutations whose long one-line notation has no subsequence bca with* $a < b < c$ *such that* b *is a right point or the negative of a left point.*

By the symmetry of the long one-line notation, the criterion in Theorem 6.5 is equivalent to requiring that the long one-line notation has no pattern 312 such that the "2" is a left point or the negative of a right point.

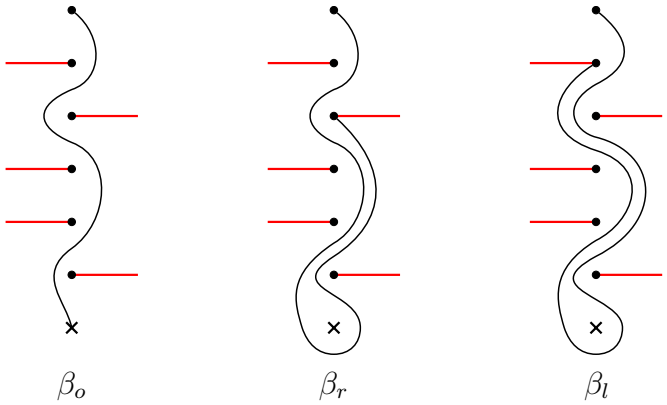


Figure 16: A meet representation of the Cambrian congruence

To write the Cambrian congruence as a meet of meet-irreducible congruences, we need to find the arcs that are not contracted by the Cambrian congruence but all of their proper superarcs are contracted. In every case, one such arc is the unique central arc β_o with upper endpoint n that passes left of every right point and right of every left point. Besides this central arc, there are one or two additional uncontracted arcs that have no uncontracted proper superarcs. Each is a long arc with nothing between its left and right pieces, passing left of every right point and right of every left point. One arc, β_r, has left endpoint at n and right endpoint at the highest right point and the other, β_l, has right endpoint at n and left endpoint at the highest left point. If there are no right points or no left points (the case where the Cambrian lattice is a type-B Tamari lattice), then only one of the arcs β_r or β_l exists. Figure 16 shows β_o, β_r, and β_l for a particular choice of right and left points, in the case $n = 6$.

Let Θ_o be the largest (i.e. coarsest) congruence that does not contract β_o. Define Θ_r and/or Θ_l similarly (if β_r and/or β_l exist). Each congruence Θ_o, Θ_r, and Θ_l is minimal, among meet-irreducible congruences in the congruence lattice, because each arc β_o, β_r, and β_l has no proper superarcs. We have proven the following theorem.

Theorem 6.6. *Let Θ be the Cambrian congruence on B_n associated to a designation of $1, \ldots, n-1$ as left and right points. Then Θ is the meet, in the lattice of congruences, of the minimal meet-irreducible congruences Θ_o, Θ_r, and Θ_l (but leaving out Θ_r if there are no right points or leaving out Θ_l if there are no left points).*

6.4 Bipartite biCambrian congruences

Recall from Section 6.3 that a Cambrian congruence on the weak order on W is defined by a binary choice for each edge in the Coxeter diagram. Two Cambrian congruences are ***opposite*** if they arise from making the opposite choice for each edge. A ***biCambrian congruence*** is the meet, in the congruence lattice of the weak order, of two opposite Cambrian congruences. (The meet of two congruences contracts an element if and only if both congruences contract that element.) More details on biCambrian congruences can be found in [2, Section 2.3].

A Cambrian congruence is ***bipartite*** if, whenever r and s form an edge and s and t also form an edge, the congruence either contracts the elements $rs, rsr, rsrsr, \ldots$ and $ts, tst, tstst, \ldots$ or contracts the elements $sr, srs, srsr, \ldots$ and $st, sts, stst, \ldots$. Assuming W is *irreducible*, there are precisely two bipartite Cambrian congruences, and they are opposite each other. The ***bipartite biCambrian congruence*** is the biCambrian congruence that is the meet of these two bipartite Cambrian congruences.

We briefly review the combinatorics of bipartite biCambrian congruences in type A. (More details are in [2, Section 3.3].) One bipartite Cambrian congruence corresponds to setting all even numbers in $2, \ldots, n-1$ to be left points and all odd numbers to be right points. The other is the same, but with left and right switched. Thus the bipartite biCambrian congruence contracts an arc if and only if it goes right of both an even number and an odd number or goes left of both an even number and an odd number. The uncontracted arcs are the ***alternating arcs***, which don't pass two consecutive points on the same side (and thus alternate left-right or right-left). Every arc with endpoints differing by less than 3 is alternating, because it doesn't pass two consecutive points at all. Thus the smallest non-alternating arcs, in the sense of subarcs, are the arcs with endpoints i and $i+3$, passing to the same side of $i+1$ and $i+2$, for $i = 1, \ldots, n-3$. The bipartite biCambrian congruence is thus generated by contracting the join-irreducible elements $s_i s_{i+1} s_{i+2}$ and $s_{i+2} s_{i+1} s_i$ corresponding to these arcs. The permutations not contracted by the bipartite biCambrian congruence are described by a bivincular pattern avoidance condition in the sense of [5, Section 2]. (See [2, Proposition 3.7].)

We now turn to type B. The bipartite biCambrian congruence on signed permutations is described in [2, Section 3.5] in terms of symmetric alternating arc diagrams. Here, we give an account in the orbifold model and use the subarc relation on orbifold arcs to add a new description of the generators of the congruence.

Similarly to type A, one bipartite Cambrian congruence of type B corresponds to designating all even numbers in $1, \ldots, n-1$ to be left points and all odd numbers to be right points. The other is the same with right and left reversed. Thus the uncontracted arcs of the bipartite biCambrian congruence of type B are the ***alternating orbifold arcs***, meaning:

- ordinary arcs that don't pass consecutive points on the same side,
- central arcs that don't pass consecutive points on the same side, and
- long arcs with no numbered points between them, with left piece and right piece each not passing consecutive points on the same side.

The following theorem is easily verified by checking that the orbifold arcs listed are the minimal non-alternating orbifold arcs, in the sense of subarcs.

Theorem 6.7. *The type-B bipartite biCambrian congruence is generated by contracting the following join-irreducible elements.*

- *for each $i = 1, \ldots, n-3$, the join-irreducible elements $s_i s_{i+1} s_{i+2}$ and $s_{i+2} s_{i+1} s_i$ corresponding to ordinary arcs with endpoints i and $i+3$ passing to the same side of $i+1$ and $i+2$,*

- *the join-irreducible elements $s_0s_1s_0s_2s_1s_0$ and $s_2s_1s_0$ corresponding to central arcs with upper endpoint 3 passing to the same side of 1 and 2, and*

- *the join-irreducible elements $s_0s_1s_2$ and $s_2s_1s_0s_1$ corresponding to long arcs with endpoints at 1 and 3 passing to the same side of 1 and 2.*

6.5 Linear biCambrian congruences

When W is of type A or B (or F or H, which are not considered in this paper), there are two ***linear Cambrian congruences***. (The corresponding Cambrian lattices are the Tamari lattices and type-B Tamari lattices. See [16, 27].) One of them contracts all join-irreducible elements of the form $s_is_{i+1}, s_is_{i+1}s_i, \ldots$, and the other contracts all join-irreducible elements of the form $s_{i+1}s_i, s_{i+1}s_is_{i+1}, \ldots$. The ***linear biCambrian congruence*** is the meet of these two opposite linear Cambrian congruences.

In type A, the linear biCambrian congruence contracts an arc if and only if it goes left of some point and right of some other point. It is generated by contracting, for each $i = 1, \ldots, n-3$, the join-irreducible elements $s_is_{i+2}s_{i+1}$ and $s_{i+1}s_is_{i+2}s_{i+1}$, associated to arcs with endpoints i and $i+3$, passing to opposite sides of $i+1$ and $i+2$. The permutations not contracted by the linear biCambrian congruence, called the ***twisted Baxter permutations***, are defined by a vincular pattern and are counted by the Baxter number. (See [15, Section 10] and [9, Section 8].)

In type B, the linear biCambrian congruence also contracts an arc if and only if it goes left of some point and right of another.

Theorem 6.8. *The type-B linear biCambrian congruence is generated by contracting the following join-irreducible elements.*

- *for each $i = 1, \ldots, n-3$, the join-irreducible elements $s_is_{i+2}s_{i+1}$ and $s_{i+1}s_is_{i+2}s_{i+1}$ corresponding to ordinary arcs with endpoints i and $i+3$ passing to opposite sides of $i+1$ and $i+2$,*

- *the join-irreducible elements $s_0s_2s_1s_0$ and $s_1s_2s_0s_1s_0$ corresponding to central arcs with upper endpoint 3 passing to opposite sides of 1 and 2,*

- *the join-irreducible elements $s_0s_2s_1$ and $s_1s_0s_1s_2s_1s_0s_1$ corresponding to long arcs with endpoints at 1 and 3 passing to opposite sides of 1 and 2, and*

- *the join-irreducible elements $s_1s_0s_1s_2$ and $s_2s_1s_0s_1s_2$ corresponding to long arcs with endpoints at 2 and 3 with 1 between the sides.*

The theorem is easily verified by checking that the arcs listed are minimal, in the sense of subarcs, among arcs that pass left of some point and right of some other point.

6.6 Symmetric congruences of the weak order on S_{2n}

The weak order on signed permutations is a sublattice of the weak order on all permutations of $\{\pm 1, \ldots, \pm n\}$. (In light of Proposition 3.9 and as discussed in Section 4, this sublattice consists of all permutations whose corresponding noncrossing arc diagrams are centrally symmetric.)

Given a lattice and a sublattice, any congruence on the lattice restricts to a congruence on the sublattice. However, there may exist congruences on the sublattice that are not the restriction of congruences on the larger lattice. (See [11, Remark 119].) We write $\mathrm{Con}_A(B_n)$ for the set of congruences on the weak order on B_n that are the restriction of congruences on S_{2n}. The purpose of this section is to characterize $\mathrm{Con}_A(B_n)$.

Recall that conjugation by w_0 is an involutive automorphism of the weak order on S_{2n}, and when S_{2n} is represented as permutations of $\{\pm 1, \ldots, \pm n\}$, this map is

$$\pi_{-n} \cdots \pi_{-2}\pi_{-1}\pi_1\pi_2 \cdots \pi_n \mapsto (-\pi_n) \cdots (-\pi_2)(-\pi_1)(-\pi_{-1})(-\pi_{-2}) \cdots (-\pi_{-n}).$$

A ***symmetric congruence*** of S_{2n} is a congruence such that $\pi \equiv \tau$ if and only if $w_0\pi w_0 \equiv w_0\tau w_0$ for all π and τ.

Definition 6.9. An orbifold arc α' is a ***loose subarc*** of an orbifold arc α if and only if one of the following conditions holds:

- α' is a subarc of α in the sense of Definition 5.6 or Definition 5.7;
- α is a central arc with endpoint q and α' is a long arc with left endpoint p' and right endpoint q' with $p', q' \leqslant q$ such that $L(\alpha') = L(\alpha) \cap (0, p')$ and $R(\alpha') = R(\alpha) \cap (0, q')$. (In particular, there are no numbered points between the two sides of α'.)
- α is a long arc with left endpoint p and right endpoint q and α' is a long arc with left endpoint p' and right endpoint q' with $p' \leqslant q$ and $q' \leqslant p$, $L(\alpha') = (0, p') \setminus R(\alpha)$, $R(\alpha') = (0, q') \setminus L(\alpha)$, and there are no numbered points that are between the two pieces of α'.

We emphasize the comparison between Definition 6.9 and Definitions 5.6 and 5.7. First, Definition 6.9 allows long loose subarcs of central arcs, while Definition 5.6 disallows long subarcs of central arcs. Second, Definition 6.9 allows, loosely speaking, the *left* piece of α' to be obtained by cutting the *right* piece of α and the *right* piece of α' to be obtained by cutting the *left* piece of α as long as there are no numbered points between the pieces of α'. (The definition of loose subarcs is equivalent to removing the additional requirement in Definition 4.7 for subarc pairs of a symmetric pair as long as there are no numbered points between the two arcs in the subarc pair.) Figure 17 shows some pairs α' and α of orbifold arcs such that α' is a loose subarc of α but not a subarc of α.

The following theorem characterizes $\mathrm{Con}_A(B_n)$.

Theorem 6.10. *If Θ is a lattice congruence on the weak order on B_n, the following are equivalent.*

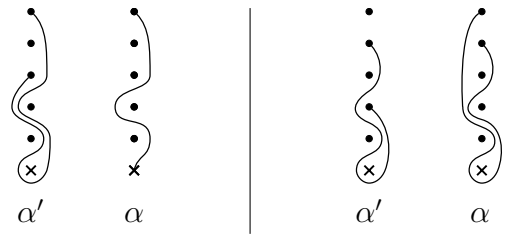


Figure 17: Loose subarcs that are not subarcs

(i) Θ *is the restriction of a congruence on* S_{2n}.

(ii) Θ *is the restriction of a symmetric congruence on* S_{2n}. *That is,* $\Theta \in \mathrm{Con}_A(B_n)$.

(iii) *The set of orbifold arcs corresponding to join-irreducible signed permutations* not *contracted by* Θ *is closed under passing to loose subarcs.*

Proof. The equivalence of (i) and (ii) is a special case of a standard lattice argument. Suppose Θ is the restriction of some congruence Θ' on S_{2n} and let Θ'' be the congruence on S_{2n} with $\pi \equiv \tau$ modulo Θ'' if and only if $w_0\pi w_0 \equiv w_0\tau w_0$ modulo Θ'. Then Θ'' also restricts to Θ. The meet in the lattice of congruences is coarsest common refinement, so $\Theta' \wedge \Theta''$ also restricts to Θ and is manifestly symmetric.

Suppose Θ is a congruence on B_n and let U be the set of orbifold arcs corresponding to join-irreducibles not contracted by Θ. Since Θ is a congruence, U is closed under passing to subarcs in the usual (orbifold) sense. We need to show that there exists a congruence Θ' of S_{2n} that restricts to Θ if and only if U is also closed under passing to loose subarcs that are not subarcs, as described in the second and third bullet points of Definition 6.9.

Given an orbifold arc α, write $\tilde{\alpha}$ for the symmetric arc/pair associated to α. (In the language of Section 5, $\alpha = \phi(\tilde{\alpha})$.) In both directions of the proof, the essential point is the following observation, which is easily verified, perhaps with the help of Figure 17: Given orbifold arcs α' and α, the case where α' is a loose subarc of α but not a subarc of α is precisely the case where there is a subarc relationship, in the type-A sense, between some arc in $\tilde{\alpha}'$ and some arc in $\tilde{\alpha}$.

If Θ is the restriction of a symmetric congruence Θ', then considering the symmetric arcs/pairs corresponding to α and α', closure of U under passing to loose subarcs follows immediately from the observation in the previous paragraph and the fact that Θ' is closed under passing to subarcs in the type-A sense.

Conversely, suppose U is closed under passing to loose subarcs. Let C be the complement of U in the set of all orbifold arcs. Each orbifold arc in C corresponds to a symmetric arc or symmetric pair on $2n$ points. Let C' be the set of all of these type-A arcs on $2n$ points. Let Θ' be the congruence on S_{2n} generated by contracting the arcs in C'. Because C' is a symmetric set of arcs, Θ' is a symmetric congruence. We will show that Θ' restricts to Θ.

If Θ' fails to restrict to Θ, then it restricts to a strictly coarser congruence $\tilde{\Theta}$. In particular, there is some join-irreducible element contracted by $\tilde{\Theta}$ but not by Θ, corresponding

to an orbifold arc β. Let β' be the symmetric arc associated to β or one of the arcs in the symmetric pair associated to β. Since β is contracted by $\tilde{\Theta}$, which is symmetric, there is some arc $\alpha' \in C'$ that is a subarc of β' in the type-A sense. Let α be the orbifold arc that arises from α' by passing to the orbifold model. Since β is not contracted by Θ, α is not a subarc of β in the orbifold sense. The observation says that α is a loose subarc of β, and this contradiction implies that Θ' restricts to Θ. □

The following corollary is immediate.

Corollary 6.11. $\mathrm{Con}_A(B_n)$ *induces a sublattice of the congruence lattice of* B_n.

To illustrate Theorem 6.10, we consider which of the congruences in Section 6 are in $\mathrm{Con}_A(B_n)$. The parabolic congruences of Section 6.1 are easily seen to be restrictions of parabolic congruences on S_{2n}. Similarly, each Cambrian congruence (Section 6.3) on B_n is the restriction of a Cambrian congruence on S_{2n}. Therefore also each biCambrian congruence (including those in Sections 6.4 and 6.5) is in $\mathrm{Con}_A(B_n)$ because $\mathrm{Con}_A(B_n)$ induces a sublattice of $\mathrm{Con}(B_n)$. The reader can also verify these statements using Theorem 6.10.

Theorem 6.10 is the simplest way to determine whether congruences from Section 6.2 are in $\mathrm{Con}_A(B_n)$. The congruence described in Proposition 6.1 is not in $\mathrm{Con}_A(B_n)$ because it contracts all long arcs, some of which are loose subarcs of uncontracted central arcs. The congruence in Proposition 6.2 is in $\mathrm{Con}_A(B_n)$: the only uncontracted central arc has upper endpoint 1 and the only uncontracted long arcs are those with lower endpoint 1, so no loose subarc relationship seen in Figure 17 may occur. The congruence in Proposition 6.3 is not in $\mathrm{Con}_A(B_n)$, since there exists a long arc with right endpoint at 1 which is contracted and is a loose subarc of an uncontracted central arc passing left of 1.

Acknowledgments

Nathan Reading and Ashley Tharp were partially supported by the National Science Foundation under award number DMS-2054489. We thank an anonymous referee for helpful comments and suggestions.

References

[1] E. Barnard, *The canonical join complex.* Electron. J. Combin. **26** (2019), no. 1, #P1.24.

[2] E. Barnard and N. Reading, *Coxeter-biCatalan combinatorics.* J. Algebraic Combin. **47(2)** (2018), 241–300.

[3] E. Barnard, N. Reading, and A. M. Tharp, *Noncrossing arc diagrams of type D.* In preparation, 2026.

[4] A. Björner, P. Edelman and G. Ziegler, *Hyperplane Arrangements with a Lattice of Regions.* Discrete Comput. Geom. **5** (1990), 263–288.

[5] M. Bousquet-Mélou, A. Claesson, M. Dukes, and S. Kitaev, (2+2)-*free posets, ascent sequences and pattern avoiding permutations.* J. Combin. Theory Ser. A **117** (2010), no. 7, 884–909.

[6] L. Brestensky and N. Reading, *Noncrossing partitions of an annulus.* Comb. Theory. **5(1)**, 2025.

[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] R. Freese, J. Ježek, and J. Nation, *Free lattices.* Mathematical Surveys and Monographs, **42**. American Mathematical Society, Providence, RI, 1995.

[9] S. E. Law and N. Reading, *The Hopf algebra of diagonal rectangulations.* J. Combin. Theory Ser. A **119** (2012), no. 3, 788–824.

[10] C. Le Conte de Poly-Barbut, *Sur les treillis de Coxeter finis.* Math. Inform. Sci. Humaines No. **125** (1994), 41–57.

[11] A. Padrol, V. Pilaud, and J. Ritter, *Shard Polytopes.* International Mathematics Research Notices **2023**, Issue 9, 7686–7796.

[12] V. Pilaud and F. Santos, *Quotientopes.* Bull. Lond. Math. Soc. **51** (2019), no. 3, 406–420.

[13] N. Reading, *Lattice and order properties of the poset of regions in a hyperplane arrangement*, Algebra Universalis, **50** (2003), 179–205.

[14] N. Reading, *Lattice congruences of the weak order*, Order **21** (2004) no. 4, 315–344.

[15] N. Reading, *Lattice congruences, fans and Hopf algebras.* J. Combin. Theory Ser. A **110** (2005) no. 2, 237–273.

[16] N. Reading, *Cambrian Lattices.* Adv. Math. **205** (2006) no. 2, 313–353.

[17] N. Reading, *Noncrossing partitions and the shard intersection order.* J. Algebraic Combin. **33** (2011), no. 4, 483–530.

[18] N. Reading, *Noncrossing arc diagrams and canonical join representations.* SIAM J. Discrete Math. **29** (2015), no. 2, 736–750.

[19] N. Reading, *Lattice homomorphisms between weak orders.* Electron. J. Combin. **26**, (2019), #P2.23.

[20] N. Reading, *Lattice theory of the poset of regions.* Lattice Theory: Special Topics and Applications, Volume **2**, ed. G. Grätzer and F. Wehrung, 399–487, Birkhäuser/Springer, Cham, 2016.

[21] N. Reading, *Finite Coxeter Groups and the weak order.* Lattice Theory: Special Topics and Applications, Volume **2**, ed. G. Grätzer and F. Wehrung, 489–561, Birkhäuser/Springer, Cham, 2016.

[22] N. Reading. *Noncrossing partitions of a marked surface.* SIAM J. Discrete Math. **39** (2025) 231-266.

[23] N. Reading and A. M. Tharp, *The shard intersection order on noncrossing arc diagrams.* In preparation, 2026.

[24] J. Ritter, *Shard polytopes and quotientopes for lattice congruences of the weak order.* Ph.D. Thesis, Institut Polytechnique de Paris, 2021.

[25] L. Santocanale and F. Wehrung, *Sublattices of associahedra and permutohedra.* Adv. in Appl. Math. **51** (2013), no. 3, 419-445.

[26] A. M. Tharp, *Arcs and shards.* Ph.D. Thesis, North Carolina State University, 2023.

[27] H. Thomas, *Tamari lattices and noncrossing partitions in type B.* Discrete Math. **306** (2006), no. 21, 2711–2723.