

ABSTRACT

BARNARD, EMILY SARAH. The Canonical Join Representation in Algebraic Combinatorics. (Under the direction of Nathan Reading.)

We study the combinatorics of a certain minimal factorization of the elements in a finite lattice L called the canonical join representation. The join $\vee A = w$ is the canonical join representation of w if A is the unique lowest subset of L satisfying $\vee A = w$ (where “lowest” is made precise by comparing order ideals under containment). When each element in L has a canonical join representation, we define the canonical join complex to be the abstract simplicial complex of subsets A such that $\vee A$ is a canonical join representation. In the first chapter, we characterize the class of finite lattices whose canonical join complex is flag, and show how the canonical join complex is related to the topology of L .

Next, we study the canonical join complex of the Tamari lattice in types A and B. We realize the canonical join complex of the Tamari lattice as a complex of noncrossing arc diagrams, give a shelling order on its facets, and show that it is homotopy equivalent to a wedge of Catalan-many spheres. We extend these results to the c -Cambrian lattices of type A, which we show to be vertex decomposable.

We close this document by considering a family of counting problems, analogous to the well-studied Coxeter-Catalan combinatorics. In our construction, each object to be counted is obtained by doubling a Coxeter-Catalan object. We show that, given a finite Coxeter group W , each of these new counting problems has the same solution, which we call the W -biCatalan number.

© Copyright 2017 by Emily Sarah Barnard

All Rights Reserved

The Canonical Join Representation in Algebraic Combinatorics

by
Emily Sarah Barnard

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Mathematics

Raleigh, North Carolina

2017

APPROVED BY:

Patricia Hersh

Ricky Liu

Seth Sullivant

Nathan Reading
Chair of Advisory Committee

DEDICATION

For Brian.

BIOGRAPHY

Emily was born in Evergreen Park, IL. She spent her childhood split between the southwestern suburbs of Chicago and the northwestern woods of Michigan. In June 2010, she graduated from Northwestern University, and the following fall she began a year long post-baccalaureate program in mathematics at North Carolina State University. She chose to complete a Ph.D. in mathematics, in no small part, because of Nathan Reading's encouragement. At the time of this writing, she looks forward to a postdoctoral appointment with the Department of Mathematics at Northeastern University.

ACKNOWLEDGEMENTS

The author thanks her advisor, Nathan Reading, for his editorial suggestions and helpful conversations. She thanks her committee members: Patricia Hersh (who suggested the connection to the crosscut complex in Chapter 2), Ricky Liu, Blair D. Sullivan, and Seth Sullivant.

The author thanks the many other mathematicians whose conversations have shaped this document for the better: Andrew Carroll, Al Garver, Emily Gunawan, Christophe Hohlweg, Alexander Martsinkovsky, Thomas McConville, Victor Reiner, David Speyer, Gordana Todorov, and Shijie Zhu. This includes the fellow and former members of the Reading research group: Erin Bancroft, Chetak Hossain, Shirley Law, Emily Meehan, Salvatore Stella, and Shira Viel. She thanks her delightful and encouraging officemates, and the insightful crew at the graduate algebra and combinatorics seminar.

Finally, the author thanks Springer International Publishing: A version of the final chapter, Coxeter-biCatalan combinatorics will appear in the Journal of Algebraic Combinatorics.

TABLE OF CONTENTS

List of Tables	vii
List of Figures	viii
Chapter 1 Introduction	1
1.1 The canonical join representation	1
1.2 The combinatorics of the canonical join representation	2
1.3 Finite Coxeter groups	4
1.4 The topology of the canonical join complex	6
1.5 Coxeter-biCatalan combinatorics	7
Chapter 2 The Canonical Join Complex	9
2.1 Introduction	9
2.2 Motivation and examples	12
2.3 Finite semidistributive lattices	14
2.3.1 Definitions	14
2.3.2 The flag property	18
2.3.3 Crosscut-simplicial lattices	26
2.4 Lattice-theoretic constructions	29
2.4.1 sublattices and quotient lattices	29
2.4.2 Products and sums	30
2.4.3 Day’s doubling construction	31
2.5 Discussion and open problems	36
Chapter 3 The Canonical Join Complex of the Tamari Lattice	40
3.1 Introduction	40
3.2 Background	42
3.2.1 Lattice-theoretic background	42
3.2.2 The noncrossing arc complex	44
3.2.3 The c -Cambrian congruence and the Tamari lattice	46
3.2.4 The type-B Tamari lattice	47
3.2.5 Noncrossing perfect matchings	49
3.3 Shellability of the Tamari lattices	51
3.3.1 The Tamari lattice in type A	51
3.3.2 The Tamari lattice in type B	53
3.4 Vertex Decomposability of the c -Cambrian lattices	63
3.4.1 The c -Cambrian lattices	63
3.4.2 Vertex decomposability	64
Chapter 4 Coxeter BiCatalan Combinatorics	76
4.1 Introduction	76
4.2 BiCatalan objects	80
4.2.1 Antichains in the doubled root poset and twin nonnesting partitions	80

4.2.2	BiCambrian fans	84
4.2.3	The biCambrian congruence, twin sortable elements, and bisortable elements	88
4.2.4	Twin clusters and bicluster fans	91
4.2.5	Twin noncrossing partitions	94
4.3	Bipartite c -bisortable elements and alternating arc diagrams	95
4.3.1	Pattern avoidance	95
4.3.2	Noncrossing arc diagrams	96
4.3.3	Alternating arc diagrams	97
4.3.4	Counting alternating arc diagrams	99
4.3.5	Enumerating bipartite c -bisortable elements in type B	103
4.3.6	Simpliciality of the bipartite biCambrian fan in types A and B	106
4.4	Double-positive Catalan numbers and biCatalan numbers	110
4.4.1	Double-positivity	110
4.4.2	Counting twin nonnesting partitions	112
4.4.3	Canonical join representations and lattice congruences	114
4.4.4	Canonical join representations of c -bisortable elements	118
4.4.5	Counting bipartite c -bisortable elements	121
4.4.6	BiCatalan and Catalan formulas	123
4.4.7	The double-positive Catalan numbers	128
4.4.8	The Type D biCatalan number	133
4.4.9	Type-D biNarayana numbers	135
	References	136

LIST OF TABLES

Table 4.1	The W -biCatalan numbers	78
Table 4.2	The biNarayana numbers	79
Table 4.4	The type-D biNarayana numbers	80
Table 4.6	Some double positive Catalan numbers	128
Table 4.7	$\frac{(2k)!}{2^k} \cdot \text{biNar}_k(D_n)$ for small k	135

LIST OF FIGURES

Figure 1.1	Any map ϕ of from the set X to a lattice L extends uniquely to a lattice homomorphism $\tilde{\phi} : F \rightarrow L$	2
Figure 1.2	The weak order on the symmetric group S_3	5
Figure 2.1	The canonical join complex is an empty triangle.	10
Figure 2.2	Some examples of noncrossing arc diagrams.	13
Figure 2.3	Two finite lattices whose top elements have no canonical join representation.	15
Figure 2.4	Dashed lines represent order relations in L while solid lines represent cover relations.	18
Figure 2.5	Some order relations in the join-refinement order for L	21
Figure 2.6	Some relations in the join-refinement order for L	22
Figure 2.7	A depiction of the argument for Lemma 2.3.13. Dashed lines represent order relations in L while solid lines represent cover relations.	24
Figure 2.8	An illustration of the argument for the proof of Corollary 2.1.5. Dashed gray lines represent relations in L while solid black lines represent cover relations.	25
Figure 2.9	A finite crosscut-simplicial lattice failing both SD_{\vee} and SD_{\wedge}	26
Figure 2.10	Dashed lines represent relations in L and solid lines represent cover relations.	27
Figure 2.11	Dashed lines represent relations in L while solid lines represent cover relations.	28
Figure 2.12	The Tamari lattice T_3 and its canonical join complex.	30
Figure 2.13	The canonical labeled join graph of three non-isomorphic congruence uniform lattices.	33
Figure 2.14	The two leftmost graphs are isomorphic Hasse diagrams for the distributive lattice L . Rightmost is the lattice obtained by doubling the interval $[a, e]$ in L	35
Figure 2.15	Doubling the interval $[a, e]$ in the leftmost congruence uniform lattice yields the left-middle lattice, whose canonical join graph is isomorphic to C_6 . Doubling the interval $[a, e]$ in the right-middle lattice yields the rightmost lattice, whose canonical join graph is isomorphic to C_7	36
Figure 2.16	Left: The weak order for the symmetric group S_3 . Right: The lattice of 2-multichains.	38
Figure 2.17	Left: The canonical join complex of weak order for the symmetric group S_3 . Right: The canonical join complex of the lattice of 2-multichains.	38
Figure 3.1	The Tamari lattice T_3 and its canonical join complex.	40
Figure 3.2	The faces in the canonical join complex of the weak order on S_3	41
Figure 3.3	From left to right: $\delta(4123)$, $\delta(2413)$, $\delta(2341)$, and $\delta(3412)$	45
Figure 3.4	Left: $\delta(2431)$. Right: The arc corresponding to the canonical join and 2314.	46
Figure 3.5	A demonstration of the map μ	53
Figure 3.6	Each diagram contains two symmetric arcs.	54
Figure 3.7	An illustration for the proof of Lemma 3.3.9.	56

Figure 3.8	The nodes $[k_1, k_2 - 1]$ are filled in with a maximal collection of ordinary right arcs.	57
Figure 3.9	An illustration of the map μ_s when $i_l \neq -1$	58
Figure 3.10	An illustration of the map μ_s when $i_l = -1$	58
Figure 3.11	An illustration of the first two steps for the map ν . We curve some of the edges in the matching M to make them more suggestive of the arcs they will become in $\nu(M)$	59
Figure 3.12	An illustration of the map ν	60
Figure 3.13	The arcs $\alpha_{1,7}$ and $\alpha_{2,8}$ in $\Delta(9, \{4, 5, 8\}, \{2, 3, 6, 7\})$	67
Figure 3.14	The arc $\alpha_{1,7}$ in $\Delta(9, \{1, 2, 3, 4, 5, 7\}, \{6, 8\})$	69
Figure 3.15	The set $\{\alpha_{3,5}, \alpha_{2,3}\} \in \text{lk}(\alpha_{1,4})$, and it is a facet in $\Delta^{\setminus 2}(5, \{4\}, \{2, 3\}) \setminus \{\alpha_{1,5}, \alpha_{1,4}\}$	72
Figure 3.16	In $\Delta(9, \{4, 5, 8\}, \{2, 3, 6, 7\})$, arcs $\alpha_{2,8}$ and $\alpha_{1,3}$ intersect.	73
Figure 4.1	Some doubled root posets	81
Figure 4.2	Some posets of join-irreducibles of doubled root posets	83
Figure 4.3	Cambrian fans and the biCambrian fan in type B_2	85
Figure 4.4	The linear biCambrian fan in type A_3	86
Figure 4.5	The bipartite biCambrian fan in type A_3	86
Figure 4.6	The linear bicluster fan in type A_3	93
Figure 4.7	The bipartite bicluster fan in type A_3	93
Figure 4.8	Some alternating noncrossing arc diagrams	98

Chapter 1

Introduction

1.1 The canonical join representation

Throughout mathematics, and algebra in particular, one sees the unique decomposition of an object into irreducible or indecomposable components. In number theory, this is the prime factorization of an integer; in commutative algebra, this is the primary decomposition of an ideal; and in representation theory, this is the direct sum decomposition of a representation into indecomposable representations. The organizing principle of this thesis is the lattice-theoretic analogue: the *canonical join representation*.

Recall that a lattice L is a partially ordered set such that each pair of elements w and v has a smallest upper bound, called the *join*, and a greatest lower bound, called the *meet*. We write $w \vee v$ or $\vee\{w, v\}$ for the join of w and v , and $w \wedge v$ or $\wedge\{w, v\}$ for the meet. Alternatively, one can think of a lattice as a universal algebra, with operations (\vee, \wedge) that are associative, idempotent, and satisfy the absorption laws: $w \vee (w \wedge v) = w$ and $w \wedge (w \vee v) = w$. A *lattice homomorphism* is a map $\phi : L \rightarrow L'$ between lattices L and L' that respects the meet and join operations. We say that the image of ϕ is a *lattice quotient* of L .

Informally, the canonical join representation of an element w is its unique minimal “factorization” in terms of the join operation. (There is an analogous notion in terms of the meet operation that is called the *canonical meet representation*.) In the literature, the canonical join representation is sometimes called the *canonical form* [44, Section IV.2], and first appeared in the study of free lattices. The lattice F is *generated freely* by the set X if it is generated by X , and satisfies the usual universal property for free objects, as shown in Figure 1.1. Alternatively, one constructs F by writing down all possible polynomial equations in the operations (\vee, \wedge) with the set X as variables. We define an equivalence relation on these polynomial expressions, so that $a \equiv b$ if and only if a and b represent the same element in F . Determining which expressions belong to the same class is called the *word problem*, and its solution by Whitman is the basis




Figure 1.1: Any map ϕ of from the set X to a lattice L extends uniquely to a lattice homomorphism $\tilde{\phi}: F \rightarrow L$.

for much of the research on free lattices [39, Section 2]. As a part of his solution, Whitman constructed an algorithm that transforms any given polynomial expression into a “shortest” form [44, Appendix G, Section 1]. It turns out that this expression is also the canonical join representation of the corresponding element, and it is minimal in an order-theoretic sense: In a general lattice L , the expression $\bigvee A$ is the canonical join representation of an element w if it is the unique “lowest” irredundant expression for w as a join. One makes the notion of “lowest” precise by comparing order ideals. In this case, we also say that the set A is a canonical join representation (although, more precisely, we mean that $\bigvee A$ is a canonical join representation).

Later, Jónsson noticed a further connection between the algebraic structure of a lattice and this canonical form [53]. Certain elements in a lattice may not admit a canonical join representation. (For example, Figure 2.3 depicts two lattices, and in each the top element does not have a unique minimal expression $\bigvee A$.) When L is finite, each element admits a canonical join representation if and only if the lattice also satisfies a certain weakening of the distributive law called *join-semidistributivity*:

$$\text{If } x \vee y = x \vee z, \text{ then } x \vee (y \wedge z) = x \vee y. \quad (SD_{\vee})$$

We say that L is *join-semidistributive* if it satisfies SD_{\vee} for each x, y , and z . If L also satisfies the dual condition (where we replace \vee with \wedge) then it is *semidistributive*. For the remainder of this introduction, we assume that L is finite and join-semidistributive.

1.2 The combinatorics of the canonical join representation

We focus our attention on the discrete structure of the collection $\Delta(L)$ of subsets $A \in \mathbf{2}^L$ such that A is a canonical join representation. Recall that an *abstract simplicial complex* Δ on a set of vertices V is a collection of subsets of V satisfying: First, $\{v\} \in \Delta$ for each $v \in V$. Second, if $A \in \Delta$ then each subset $A' \subset A$ also belongs to Δ . We call the collection of edges and vertices in Δ its *one-skeleton*.

We will see that $\Delta(L)$ has the structure of an abstract simplicial complex. We call this

complex the *canonical join complex* of L . Its vertex set is the set of elements that cannot be written as a nontrivial join of lower elements. These elements are called *join-irreducible*. (That is, j is join-irreducible if $j = \bigvee A$ implies that $j \in A$.)

As combinatorialists, we ask questions like:

- How many faces does the canonical join complex have?
- What is the facial structure of the canonical join complex?

Because L is finite and join-semidistributive, the answer to the first question is immediate: Each element admits a canonical join representation, so the number of faces is just the size of L . (The empty face is the canonical join representation for the smallest element.)

Answering the second question will be the main focus of Chapter 2, where we consider a certain combinatorial property called the flag property. See Theorem 2.1.1. A complex Δ is *flag* if its minimal non-faces have size equal to 2. Informally, we can think of this condition as saying: There are no “hollow” simplices in Δ . More precisely, complex Δ is flag if and only if it is determined by its underlying one-skeleton as follows: Given a set A of vertices, A is a face in Δ if and only if A is a clique in the one-skeleton for Δ .

The flag property appears at the intersection of combinatorics with graph theory, differential geometry, and topology. In particular, its connection to the Charney-Davis conjecture(s) [21] has received much attention. The Charney-Davis conjecture is essentially the polyhedral analogue to a classical conjecture of Hopf. Hopf’s conjecture relates the geometry and topology of a Riemannian manifold M , and states: If M has dimension $2n$, and its sectional curvature is nonpositive, then $(-1)^n \chi(M) \geq 0$, where $\chi(M)$ is the Euler characteristic of M . (Informally, the *sectional curvature* is the Gaussian curvature of the surface we obtain by taking two-dimensional slices of M . See [38, Conjecture 54].) When we further restrict to polyhedral *flag* complexes, the Charney-Davis conjecture has a purely combinatorial reformulation. See [38, Conjecture 72] or [63, Conjecture 1]. Next, we discuss a few familiar examples of flag complexes.

Example 1.2.1. Suppose that Δ is a cell complex, and write $BCS(\Delta)$ for the barycentric subdivision of Δ . Geometrically, we construct $BCS(\Delta)$ by adding a vertex v_F at the barycenter of each face F in Δ . A subset of vertices $\{v_{F_1}, \dots, v_{F_k}\}$ is a face if and only if it corresponds with a *flag* of faces $F_1 \subset \dots \subset F_k$ in Δ . Clearly, each collection of vertices in $BCS(\Delta)$ satisfies: if each pair is a face, then the entire collection is a face. Thus, $BCS(\Delta)$ is flag.

Example 1.2.2. Let \mathcal{P} be a partially ordered set. The *order complex* for \mathcal{P} is the simplicial complex whose k -dimensional faces are the chains $x_0 < \dots < x_k$. Suppose that each pair of elements in a subset A of \mathcal{P} is comparable. Since A is totally ordered if and only if each pair is comparable, we conclude that the order complex for \mathcal{P} is flag.

Example 1.2.3. Fix a convex polygon P , and consider the simplicial complex $\Delta(P)$ whose faces correspond to partial tilings of P by triangles, so that its facets correspond to triangulations of P and its vertices correspond to the diagonals in P . It is well-known that this complex can be realized as the boundary of a convex polytope, called the *simplicial associahedron*. Observe that a collection of diagonals belongs to a (partial) triangulation if and only if each pair in the collection does not cross. Thus, $\Delta(P)$ is flag.

1.3 Finite Coxeter groups

Many of the most interesting join-semidistributive lattices are closely related to the *weak order* on a finite Coxeter group. We now turn our attention to the combinatorics of the canonical join representation in this context. In preparation for our results, we will give a gentle introduction to finite Coxeter groups and the weak order. (To find a complete discussion of finite Coxeter groups, with precise statements and proofs, see [12, 51].)

A Coxeter group W is a group of transformations on Euclidean space, generated by orthogonal reflections. The collection of reflecting hyperplanes is called the *Coxeter arrangement* for W , and it is fixed by the action of the group. Each finite Coxeter group is equipped with a special set of generators called *simple generators*, and these are typically a proper subset of all of its reflections. The group W has the following presentation in terms of its simple generators s_1, \dots, s_n :

$$W = \langle s_1, \dots, s_n : (s_i s_k)^{m_{i,k}} = e \rangle$$

The numbers $m_{i,k}$ are symmetric in i and k , and satisfy: $m_{i,k} \in \mathbb{Z}^+ \cup \{\infty\}$, $m_{i,k} \geq 2$ when $i \neq k$, and $m_{i,i} = 1$.

We encode this data with a graph called the *Coxeter diagram* that is defined as follows: Take the simple generators s_1, \dots, s_n as nodes, and connect s_i to s_k whenever the number $m_{i,k} \geq 3$. We typically label the edge $\{s_i, s_k\}$ by the number $m_{i,k}$ whenever $m_{i,k} > 3$. A Coxeter group is *irreducible* if its associated Coxeter diagram is connected. Finite irreducible Coxeter groups have been classified by their Coxeter diagrams. There are four infinite families—called A_n , B_n , D_n and $I_2(n)$ —and six exceptional types. Examples include the symmetry groups for regular polytopes and the Weyl groups which appear in the study of semisimple Lie algebras. Below, we give two familiar examples.

Example 1.3.1. Consider the symmetry group W of an equilateral triangle drawn in \mathbb{R}^3 , so that its vertices are the standard basis vectors e_1, e_2 , and e_3 . Each reflecting hyperplane corresponds to an edge of the triangle as follows: The edge connecting e_i and e_k determines an orthogonal plane, $H_{i,k} = \{\mathbf{x} \in \mathbb{R}^3 : x_i = x_k\}$, that cuts the edge in half and contains the third vertex. The reflection through $H_{i,k}$ interchanges the i^{th} and k^{th} coordinates in \mathbb{R}^3 . The

reader will recognize that W is the symmetric group S_3 (or A_2 in the Weyl group notation). In general, we identify S_n with the symmetry group of the standard $(n - 1)$ -simplex. The simple generators are the set of reflections corresponding to the adjacent transpositions. We usually write s_i for the transposition $(i, i + 1)$, where $i \in \{1, 2, \dots, n - 1\}$. Thus, we have $m_{i,i+1} = 3$ and otherwise $m_{i,k} = 2$.

Example 1.3.2. Consider the symmetry group W of the regular n -cube, whose vertices in \mathbb{R}^n are $\{(e_1 \pm \dots \pm e_n)\}$. The reflecting hyperplanes for W have normal vectors $e_i, e_k \pm e_i$. We often realize W as the group of *signed permutations*. These are permutations on the set $\pm\{1, \dots, n\}$ that satisfy the symmetry condition $w(i) = -w(-i)$. In this permutation representation, each reflection corresponds either to a pair of transpositions $(i, k)(-i, -k)$ or a “symmetric” transposition $(-i, i)$. The simple generators correspond to the transpositions $(i, i + 1)(-i, -i - 1)$ and $(-1, 1)$, where $i \in \{1, 2, \dots, n - 1\}$. We usually write s_0 for the transposition $(-1, 1)$, and s_i for $(i, i + 1)(-i, -i - 1)$. Thus, $m_{0,1} = 4$, $m_{i,i+1} = 3$ for $i > 0$, and $m_{i,k} = 2$ otherwise. In the Weyl group notation, this is B_n .

We represent the elements of W as words in the simple generators S , although there are typically many such expressions for each element. The *length* $l(w)$ is the size of a reduced, or shortest possible, expression for w . The *weak order* on W is defined by the cover relations $w \prec ws$ whenever $l(w) < l(ws)$ and $s \in S$. Thus, the Hasse diagram for the weak order is just the Cayley graph for W (with generating set S). For each finite W , the weak order is a semidistributive lattice (see [26, Lemma 9]). In particular, each element has a canonical join representation.




Figure 1.2: The weak order on the symmetric group S_3

Example 1.3.3. Returning to the symmetric group S_n , from Example 1.3.1, we can describe the weak order as follows: Write each permutation in S_n in its one-line notation as $w_1 w_2 \dots w_n$, where $w(i) = w_i$. Acting on the right by the transposition $(i, i + 1)$ corresponds to swapping the entries w_i and w_{i+1} . Thus, one moves up in the weak order by swapping adjacent entries that are in order (that is $w_i < w_{i+1}$) and leaving all other entries fixed. See Figure 1.2.

1.4 The topology of the canonical join complex

In Chapter 3, we study the canonical join representation in certain lattice quotients of the weak order. These lattice quotients inherit semidistributivity, so each element admits a canonical join representation.

We begin by considering the Tamari lattice. The Tamari lattice is named for Dov Tamari, who proved that it is a lattice [40, 50], and defined it as follows: Consider a fixed word $a_1 a_2 \dots a_{n+1}$ and all of the possible ways to properly distribute brackets among its letters. We think of the bracketing as defining a binary operation, where each (rightward) application of the associative law corresponds to a cover relation. For example, the Tamari lattice T_2 consists of the single cover relation $(a_1 a_2) a_3 < a_1 (a_2 a_3)$.

Since its original definition, the Tamari lattice has made surprising appearances in algebra, topology, category theory, and even physics (not to mention combinatorics) [59]. Appropriately, it has many realizations. It is convenient for us to realize T_n as the lattice quotient of the weak order on S_n consisting of the permutations that avoid the pattern 312. We say that a permutation *avoids the 312-pattern* if, for each pair of numbers $i < k$ that appear out of order in $w_1 \dots w_n$, we have that each $j \in \{i + 1, \dots, k - 1\}$ precedes i and k . The Hasse diagram for T_n is an orientation of the one-skeleton of a polytope. This polytope is the *simple* associahedron—the dual (or polar) polytope of $\Delta(P)$ from Example 1.2.3.

Remark 1.4.1. The canonical join representation “sees” the geometry of the Hasse diagram for T_n . More precisely, for any finite join-semidistributive lattice, there is a bijection from the set $\{y : w \succ y\}$ to the canonical join representation of w . (This is Proposition 2.2.2.) Similar constructions were used to study the cover relations in free lattices. (See [39, Theorem 3.5].) When the Hasse diagram for L is the dual graph for a simplicial sphere Δ , as it is for the weak order and the Tamari lattice, the f -vector for the canonical join complex is equal to the h -vector for Δ .

Informally, a complex Δ is *shellable* if we can linearly order its facets F_1, \dots, F_m so that when we glue F_i into the complex $\bigcup_{r=1}^{i-1} F_r$ of earlier facets, one of two possible events occurs: Either the topology of the resulting complex does not change; or we close off a sphere of dimension $|F_i| - 1$. In Chapter 3, we show that the canonical join complex of T_n is shellable. See Theorem 3.1.1. We extend these results in two directions: We show that the type-B analogue to the Tamari lattice is also shellable (Theorem 3.1.2), and we show that each c -Cambrian lattice in type A is shellable (Theorem 3.1.3). The c -Cambrian lattices are a fundamental object of Coxeter-Catalan combinatorics, and are closely studied because of their connection with cluster algebras of finite type [68, 75]. See Example 1.5.2 below.

Remark 1.4.2. The canonical join complexes that we study here are all non-pure, meaning that their facets have different dimensions. We use the notion of non-pure shellability developed by Björner and Wachs in [13] and [14]. As an immediate consequence, we also obtain a direct-sum decomposition of the associated Stanley-Reisner ring that generalizes the Cohen-Macaulay property of pure complexes. See [14, Theorem 12.3]. Historically, the connection to the Cohen-Macaulay property was a major impetus behind the study of shellable complexes [13].

1.5 Coxeter-biCatalan combinatorics

In Chapter 4, we use the canonical join representation to solve an enumerative problem at the heart of *Coxeter-biCatalan combinatorics*. Before we outline our results, we make a very brief introduction to Coxeter-Catalan combinatorics. A more complete history and discussion of examples can be found in [2] and [33].

Our story begins with the classical *Catalan numbers*

$$C_n = \frac{1}{n+1} \binom{2n}{n}.$$

The study of the Catalan numbers dates back at least to Euler, who considered the problem of enumerating the triangulations of a fixed convex polygon [60]. Since that time, the Catalan numbers have appeared throughout algebraic combinatorics and enumerate more than 200 different combinatorial objects [83, Introduction]. Below, we call such an object a *Catalan object*.

Example 1.5.1. Our touchstone example is the Tamari lattice T_n . Recall that the Hasse diagram for the Tamari lattice can be realized as the one-skeleton for the simple associahedron. Since the vertices for the simple associahedron are parametrized by the triangulations of a fixed convex polygon, we conclude that there are Catalan many elements in T_n .

In Coxeter-Catalan combinatorics, many of the traditional Catalan objects are seen as a special case of a general construction. This general construction yields a Coxeter group analogue for each of the traditional Catalan objects. So, for example, each Coxeter group has its own version of the Tamari lattice. (More precisely, each Coxeter group has a family of Tamari-like lattices.) See Example 1.5.2. In general, this construction depends on a choice of a Coxeter group W , a set of simple generators S , an orientation c of the Coxeter diagram (c is also sometimes called a *Coxeter element*), and a collection of vectors Φ related to the combinatorics of the Coxeter arrangement.

Example 1.5.2. Recall that the Tamari lattice T_n is the lattice quotient of the weak order on S_n consisting of the 312-avoiding permutations. This pattern avoidance condition is a special

case of a more general lattice quotient construction that can be applied to the weak order on *any* finite Coxeter group. This general construction yields the so-called *c-Cambrian lattices*, a family of lattice quotients parametrized by an orientation c of the associated Coxeter diagram. In particular, when c is an orientation of the type-A Coxeter diagram in which all of the arrows point in the same direction, we recover a Tamari lattice. (In this case, we say that c is a *linear* orientation.) Like the Tamari lattice, each c -Cambrian lattice can be realized as an orientation of the one-skeleton for a simple polytope called the W -associahedron.

The next theorem is the cornerstone of Coxeter-Catalan combinatorics. In the statement, $\text{Cat}(W)$ is the *Coxeter-Catalan number*. When W is the symmetric group, we obtain the classical Catalan number. The numbers e_1, \dots, e_n are the *exponents* of W , certain numbers that originate in the study of the invariant theory for W . The number h is the *Coxeter number* for W . See [51, Section 3.20].

Theorem 1.5.3. *For each finite Coxeter group W , the enumeration of each Coxeter-Catalan object has the same solution:*

$$\text{Cat}(W) = \prod_{i=1}^n \frac{e_i + h + 1}{e_i + 1}.$$

In *Coxeter-biCatalan combinatorics*, we carry out a “doubling” or “twinning” process on each Coxeter-Catalan object, and obtain a new family of enumerative problems. In Example 1.5.4, we describe the doubled version of the Tamari lattice. In general, the doubled version of each c -Cambrian lattice is a lattice quotient of the weak order called the *c-biCambrian lattice*. When c is bipartite, we call the c -biCambrian lattice the *bipartite biCambrian lattice*. (We say that c is *bipartite* if each pair of adjacent arrows point in opposite directions.)

Example 1.5.4. Like the c -Cambrian lattices, each c -biCambrian lattice is a certain lattice quotient of the weak order on W . When W is the symmetric group, each c -biCambrian lattice is determined by pattern-avoidance conditions. For example, when c is a linear orientation for the type-A Coxeter diagram, the c -biCambrian lattice is the lattice quotient of S_n consisting of the permutations that avoid *both* the 41-2-3-pattern and the 2-3-41-pattern, and whose enumeration is given by the *Baxter numbers*. See [10, 24].

We closely study the canonical join representation in the bipartite biCambrian lattice of type-A. Our work here motivates a restrictive characterization of the canonical join representation in the bipartite biCambrian lattice *in general*. This characterization is a key ingredient in the proof of our main result: We show that, when c is bipartite, the analogue to Theorem 1.5.3 holds: For each finite Coxeter group, each of the “doubled” Coxeter-Catalan objects has the same enumeration, the *Coxeter-biCatalan number*. See Theorem 4.1.1.

Chapter 2

The Canonical Join Complex

2.1 Introduction

In this chapter, we consider the facial structure of the canonical join complex. The following theorem is our main result.

Theorem 2.1.1. *Suppose L is a finite join-semidistributive lattice. Then the canonical join complex of L is flag if and only if L is semidistributive.*

In light of Theorem 2.1.1, we define the *canonical join graph* of L to be the one-skeleton of its canonical join complex. Canonical join representations and canonical join graphs appear in many familiar guises. See Section 2.2 for connections to comparability graphs and noncrossing partitions.

Recall that the canonical join representation of an element w is the unique “lowest” irredundant expression for w in terms of the join operation. There is an analogous factorization in terms of the meet operation called the canonical meet representation that is defined dually (by replacing “lowest” with “highest” and “join” with “meet” in the sentence above). A finite lattice L is semidistributive if and only if each of its elements admits both a canonical join representation and a canonical meet representation [39, Theorem 2.24]. Suppose that L is a finite join-semidistributive lattice. Theorem 2.1.1 implies that the canonical join complex of L is flag if and only if each element admits a canonical meet representation.

It is not hard to find examples of finite join-semidistributive lattices whose canonical join complex is not flag. A key example is shown below in Figure 2.1. Observe that each pair of atoms in this lattice is a face in the canonical join complex. Since the join of all three atoms is redundant (because we can remove b and obtain the same join), the canonical complex is an empty triangle. Also, note that the bottom element $\hat{0}$ of this lattice does not have a canonical meet representation: Both $a \wedge e$ and $c \wedge d$ are meet-representations for $\hat{0}$ that are “as high as




Figure 2.1: The canonical join complex is an empty triangle.

possible”. This lattice also exhibits some unpleasant topological properties. We will see below that the combinatorics of the canonical join complex are closely related to the topology of its lattice.

The **crosscut complex** of L is the abstract simplicial complex whose faces are the subsets A' of atoms in L such that $\bigvee A' < \hat{1}$. A lattice is **crosscut-simplicial** if the crosscut complex for each interval is either a simplex or the boundary of a simplex. The *Crosscut Theorem* says that the order complex of a finite poset \mathcal{P} is homotopy equivalent to its crosscut complex ([11, Theorem 10.8]). Therefore, if L is crosscut-simplicial then each interval $[x, y]$ in L is either contractible or homotopy equivalent to a sphere with dimension two less than the number of atoms in $[x, y]$. (See also [47, Theorem 3.7].) In particular, $\mu(x, y) \in \{-1, 0, 1\}$.

Observe that the facets of the crosscut complex for the lattice L in Figure 2.1 are $\{a, b\}$ and $\{b, c\}$. Therefore, L is not crosscut-simplicial. By contrast, Hersh and Mészáros recently showed that a large class of finite semidistributive lattices—including the class of finite distributive lattices, the weak order on a finite Coxeter group, and the Tamari lattice ([47, Theorems 5.1, 5.3 and 5.5])—are crosscut-simplicial. Building on their work, McConville proved that if L is semidistributive, then it is crosscut-simplicial ([57, Theorem 3.1]). When each element in L has a canonical join representation, we prove that the converse is true.

Theorem 2.1.2. *Suppose that L is a finite join-semidistributive lattice. The following are equivalent:*

1. *The canonical join complex of L is flag.*
2. *L is crosscut-simplicial.*
3. *L is semidistributive.*

As an immediate corollary, we obtain the following topological obstruction to the flag property of the canonical join complex.

Corollary 2.1.3. *Suppose that L is a finite join-semidistributive lattice and its canonical join complex is flag. Then:*

1. Each interval $[x, y]$ in L is either contractible or homotopy equivalent to \mathbb{S}^{d-2} , where d is the number of atoms in $[x, y]$;
2. The Möbius function takes only the values $\{-1, 0, 1\}$ on the intervals of L .

McConville showed in [57, Corollary 5.4] that if L is crosscut-simplicial then so is each of its lattice quotients. Because semidistributivity is preserved under taking sublattices and quotients *when L is finite* (see Section 2.4.1), we immediately obtain the following extension of McConville’s result for finite join-semidistributive lattices.

Corollary 2.1.4. *Suppose that L is a finite join-semidistributive lattice that is crosscut-simplicial. Then each sublattice and quotient lattice of L is also crosscut-simplicial.*

Theorem 2.1.1 is surprising in part because its proof does not explicitly use the canonical meet representation of the elements in L . Instead, we make use of a local characterization of the canonical join representation in terms of cover relations, and a bijection κ from the join-irreducible to the meet-irreducible elements in L . As an easy consequence of this approach, we obtain the following nice result. In the statement, the **canonical meet complex** is the complex of subsets A in L such that the meet $\bigwedge A$ is a canonical meet representation.

Corollary 2.1.5. *Suppose that L is a finite semidistributive lattice. Then the bijection κ induces an isomorphism from the canonical join complex to the canonical meet complex of L .*

Using the isomorphism from Corollary 2.1.5, one obtains an operation on the canonical join complex that generalizes the operation of rowmotion (on the set of antichains in a poset) and the operation of Kreweras complementation (on the set of noncrossing partitions). See Remark 2.3.15.

The canonical join complex was first introduced in [72], in which Reading showed that it is flag for the special case of the weak order on the symmetric group (see Example 2.2.6). Recently, canonical join representations have played a role in the study of functorially finite torsion classes for the preprojective algebra of Dynkin-type W , when W is a simply laced Weyl group (see for example [41, 52]). In the forthcoming [9], the authors study the canonical join complex of *any* finite dimensional associative algebra Λ of finite representation type. Since the weak order on any finite Coxeter group W and the lattice of torsion classes for Λ of finite representation type are both examples of finite semidistributive lattices (see [26, Lemma 9] and [41, Theorem 4.5]), we obtain the following two applications of Theorem 2.1.1:

Corollary 2.1.6. *Suppose that W is a finite Coxeter group. Then the canonical join complex of the weak order on W is flag.*

Corollary 2.1.7. *Suppose that Λ is an associative algebra of finite representation type, and $\text{tors}(\Lambda)$ is its lattice of torsion classes ordered by containment. Then the canonical join complex of $\text{tors}(\Lambda)$ is flag.*

2.2 Motivation and examples

Before we give the technical background for our main results, we describe several familiar examples in which the combinatorics of canonical join representations appear. We begin with an example from number theory and commutative algebra.

Example 2.2.1 (The divisibility poset). It is often useful to give a canonical factorization of the elements in a set equipped with some algebraic structure. A familiar example is the primary decomposition of an ideal in a Noetherian ring. The canonical join representation is the natural lattice-theoretic analogue. Indeed, when L is the divisibility poset (whose elements are the positive integers ordered $r \leq s$ if and only if $r|s$), the canonical join representation of $x \in L$ coincides with the primary decomposition of the ideal generated by x :

$$x = \bigvee \{p^d : p \text{ is prime and } p^d \text{ is the largest power of } p \text{ dividing } x\}.$$

Suppose that L is a finite lattice, such that each element in L admits a canonical join representation. One pleasant property of the canonical join representation (and its dual, the canonical meet representation) is that it “sees” the geometry the Hasse diagram for L . Suppose that $w \in L$ has the canonical join representation $\bigvee A$. We will shortly prove that the factors that appear in A are naturally in bijection with the elements covered by w . So, the down-degree of w is equal to the size of A . Specifically, we will prove the following proposition. (See Lemma 2.3.3 and Proposition 2.3.4. Similar constructions appear in the literature, for example see [39, Theorem 3.5] which gives essentially the same statement for free lattices.)

Proposition 2.2.2. *Suppose that $\bigvee A = w$ is a face in the canonical join complex of L . Then, for each element y that is covered by w there is a corresponding element $j \in A$ such that $j \vee y = w$, and j is the unique minimal element in L with this property. The correspondence $y \mapsto j$ is a bijection.*

With this proposition in mind, we consider the class of finite distributive lattices.

Example 2.2.3 (Finite distributive lattices). Suppose that L is a finite distributive lattice. Recall that the fundamental theorem of finite distributive lattices (see for example [81, Theorem 3.4.1]) says that L is the lattice $J(\mathcal{P})$ of order ideals of some finite poset \mathcal{P} . Suppose that A is an antichain in \mathcal{P} . We write I_A for the order ideal generated by A (that is, the elements of

A are the maximal elements of I_A). Dually, we write I^A for the order ideal satisfying: A is the set of minimal elements in $\mathcal{P} \setminus I^A$. Observe that the order ideals covered by I_A are exactly of the form $I_{A \setminus \{y\}} = I_A \setminus \{y\}$, where $y \in A$. Since I_y is the smallest order ideal in $J(\mathcal{P})$ containing y , it follows immediately from Proposition 2.2.2 that the canonical join representation of I_A is $\bigcup\{I_y : y \in A\}$. (Dually, the canonical meet representation for the ideal I^A is $\bigcap\{I^y : y \in A\}$.) It follows that the canonical join graph of $J(\mathcal{P})$ is the incomparability graph of \mathcal{P} .

Comparability graphs were classified by a theorem of Gallai which we quote from [88, Theorem 2.1] below.

Theorem 2.2.4. *A graph G is a comparability graph for a finite poset if and only if it does not contain as an induced subgraph any graph from [88, Table 1] or the complement of any graph appearing in [88, Table 2].*

As an immediate corollary we have the following characterization of the canonical join graphs for finite distributive lattices.

Proposition 2.2.5. *The graph G is the canonical join graph for a finite distributive lattice if and only if G does not contain, as an induced subgraph, the complement of any graph forbidden by Theorem 2.2.4.*

Example 2.2.6 (The Symmetric group and noncrossing arc diagrams). Reading gave an explicit combinatorial model for the canonical join complex of the weak order on the symmetric group S_n in terms of certain noncrossing arc diagrams. A **noncrossing arc diagram** is a diagram consisting of n nodes arranged vertically, together with a collection of curves called arcs that satisfy certain compatibility conditions. In particular, the arcs in a noncrossing arc diagram do not intersect in their interiors. (See [72] or Section 3.2.2 for details.) Each diagram is determined by its combinatorial data: the endpoints of its arcs and on which side (either left or right) each arc passes the nodes in the diagram.




Figure 2.2: Some examples of noncrossing arc diagrams.

We say that two arcs are **compatible** if there is a noncrossing arc diagram that contains them. The following is a combination of [72, Corollary 3.4 and Corollary 3.6]. (In the statement

of the Theorem, we take “a collection of arcs” to also mean a collection of noncrossing arc diagrams, each containing a single arc.)

Theorem 2.2.7. *There is a bijection δ from the set of join-irreducible permutations in S_n to the set of noncrossing arc diagrams on n nodes that contain precisely one arc. Moreover, a collection of arcs \mathcal{E} corresponds to a face in the canonical join complex of S_n if and only if the arcs in \mathcal{E} are pairwise compatible.*

Example 2.2.8 (The Tamari lattice and noncrossing partitions). Recall from Example 1.5.2, the Tamari lattice T_n is a finite semidistributive lattice (see for example [42, Theorem 3.5]), which can be realized as an ordering on the set of triangulations for a fixed convex polygon P . The simple associahedron is a convex polytope, whose faces are in bijection with the collections of pairwise noncrossing diagonals of P (see [33, Figure 3.5]). The Hasse for T_n is an orientation for the one-skeleton of the associahedron. Since the number of factors in a canonical join representation (called the *canonical joinands*) for $w \in T_n$ is equal to the down-degree of w , we obtain the following result:

Proposition 2.2.9. *The f -vector for the canonical join complex of the Tamari lattice T_n is equal to the h -vector of the rank $n - 1$ associahedron. Specifically, the number of size- k faces in the canonical join complex is equal to the Narayana number*

$$N(n, k) = \frac{1}{n} \binom{n}{k+1} \binom{n}{k}.$$

Indeed, the canonical join representation of $w \in T_n$ is essentially a noncrossing partition. Recall that the Tamari lattice T_n may be realized as the set of permutations avoiding the 312-pattern. It is a fact that a permutation avoids the 312-pattern if and only if its image under the bijection δ (from Theorem 2.2.7) is a noncrossing arc diagram consisting of only right arcs. (A *right arc* is an arc that does not pass to the left of any node. See the leftmost noncrossing diagram in Figure 2.2.) Rotating such a diagram by a quarter-turn gives the familiar representation of a noncrossing partition as a bump diagram. (See [72, Example 4.5] for details, and [75, Theorem 2.7] and the discussion following [75, Proposition 8.8] for a type-free discussion.)

2.3 Finite semidistributive lattices

2.3.1 Definitions

In this chapter, we study only finite lattices. We write $\hat{0}$ for the unique smallest element in L and $\hat{1}$ for the unique largest element. A *join-representation* of w is an expression $\bigvee A$ which evaluates to w in L . At times we will also refer to the set A as a join-representation. We write

$\text{cov}_\downarrow(w)$ for the set $\{y \in L : w \succ y\}$. Similarly, we write $\text{cov}^\uparrow(w)$ for the set of upper covers of w . Recall that w is **join-irreducible** if $w = \bigvee A$ implies that $w \in A$. (In particular, the bottom element $\hat{0}$ is not join-irreducible, because it is equal to the empty join.) Since L is finite, w is join-irreducible when $\text{cov}_\downarrow(w)$ has exactly one element. **Meet-irreducible** elements satisfy the dual condition. We write $\text{Irr}(L)$ for the set of join-irreducible elements in L .

A join-representation $\bigvee A$ of w is **irredundant** if $\bigvee A' < \bigvee A$ for each proper subset $A' \subset A$. Each irredundant join-representation is an antichain in L . We say that the subset A of L **join-refines** a subset B if, for each element $a \in A$, there exists some element $b \in B$ such that $a \leq b$. Join-refinement defines a preorder on the subsets of L that is a partial order (corresponding to the containment of order ideals) when restricted to the set of antichains in L .

We write $\text{ijr}(w)$ for the set of irredundant join-representations of w . The **canonical join representation** of w in L is the unique minimal element, in the sense of join-refinement, of $\text{ijr}(w)$, when such an element exists. We write $\text{can}(w)$ for the canonical join representation of w . An element $j \in \text{can}(w)$ is a **canonical joinand** for w . If $A = \text{can}(w)$, we say that A **joins canonically**, or A is a canonical join representation in L (although, more precisely, we mean that the expression $\bigvee A$ is a canonical join representation). It follows immediately from the definition that each canonical joinand of w is join-irreducible. Moreover, the canonical join representation of each join-irreducible element j exists and is equal to $\{j\}$. The **canonical meet representation** of w is defined dually (when it exists).

In Figure 2.3, we give two examples in which the canonical join representation of $\hat{1}$ does not exist. In the modular lattice on the left each pair of atoms is a lowest-possible, irredundant




Figure 2.3: Two finite lattices whose top elements have no canonical join representation.

join-representation for the top element. Since there is no *unique* such join-representation, the canonical join representation for $\hat{1}$ does not exist. Arguing dually, we see that the canonical meet representation for the bottom element $\hat{0}$ does not exist either. In the lattice on right, each element has a canonical meet representation. However, both $a \vee d$ and $b \vee c$ are minimal elements of $\text{ijr}(\hat{1})$. Again, the canonical join representation of $\hat{1}$ does not exist.

In the lattice on the right, we observe the following failure of the distributive law: both $e \vee a$

and $e \vee b$ are expressions for $\hat{1}$, but $e \vee (a \wedge b)$ is equal to e . (A similar failure is easily verified among the atoms of the modular lattice.) We will see that correcting for precisely this kind of failure of distributivity guarantees the existence of canonical join representations, when L is finite.

A lattice L is **join-semidistributive** if L satisfies the following implication for every x, y and z :

$$\text{If } x \vee y = x \vee z, \text{ then } x \vee (y \wedge z) = x \vee y \quad (SD_{\vee})$$

L is **meet-semidistributive** if it satisfies the dual condition:

$$\text{If } x \wedge y = x \wedge z, \text{ then } x \wedge (y \vee z) = x \wedge y \quad (SD_{\wedge})$$

A lattice is **semidistributive** if it is join-semidistributive and meet-semidistributive. The following result is the finite case of [39, Theorem 2.24].

Theorem 2.3.1. *Suppose that L is a finite lattice. Then L satisfies SD_{\vee} if and only if each element in L has a canonical join representation. Dually, L satisfies SD_{\wedge} if and only if each element in L has a canonical meet representation.*

Assume that L is a finite join-semidistributive lattice, and let $j \in \text{Irr}(L)$. We write j_* for the unique element covered by j , and $\mathcal{K}(j)$ for the set of elements $a \in L$ such that $a \geq j_*$ and $a \not\leq j$. When it exists, we write $\kappa(j)$ for the unique maximal element of $\mathcal{K}(j)$. It is immediate that $\kappa(j)$ is meet-irreducible. Below, we quote [39, Theorem 2.56]:

Proposition 2.3.2. *A finite lattice L is meet-semidistributive if and only if $\kappa(j)$ exists for each join-irreducible element j in L .*

Below we establish a bijection from the set $\text{cov}_{\downarrow}(w)$ to $\text{can}(w)$. A similar construction also appears in [39, Theorem 3.5]. Suppose that $w \in L$. For each $y \in \text{cov}_{\downarrow}(w)$, there is some element $j \in \text{can}(w)$ such that $y \vee j = w$ (because there is some element $j \in \text{can}(w)$ such that $j \not\leq y$). For this j , the set $\text{can}(w)$ join-refines $\{j, y\}$. Because $\text{can}(w)$ is an antichain, each $j' \in \text{can}(w) \setminus \{j\}$ satisfies $j' \leq y$. Therefore, j is the unique canonical joinand of w such that $y \vee j = w$. We define a map $\eta : \text{cov}_{\downarrow}(w) \rightarrow \text{can}(w)$ which sends y to the unique canonical joinand j such that $y \vee j = w$.

Lemma 2.3.3. *Suppose that L is a finite join-semidistributive lattice, and $w \in L$. Then the map $\eta : \text{cov}_{\downarrow}(w) \rightarrow \text{can}(w)$ is a bijection such that $y \geq \bigvee \text{can}(w) \setminus \{\eta(y)\}$ and $y \in \mathcal{K}(\eta(y))$ for each $y \in \text{cov}_{\downarrow}(w)$.*

Proof. Suppose there exist distinct y and y' in $\text{cov}_{\downarrow}(w)$ satisfying $\eta(y) = \eta(y')$. Then, $y \vee y' = w$, and $\text{can}(w)$ does not join-refine $\{y, y'\}$ (because $\eta(y)$ is below neither y nor y'). We have a contradiction, because $\text{can}(w)$ is the unique minimal element (in join-refinement) of $\text{ijr}(w)$. By

this contradiction, we conclude that η is injective. Suppose that $j \in \text{can}(w)$. Since $\bigvee \text{can}(w)$ is irredundant, $\bigvee(\text{can}(w) \setminus \{j\}) < w$. Thus, there is some $y \in \text{cov}_\downarrow(w)$ such that $y \geq \bigvee(\text{can}(w) \setminus \{j\})$. If $y \geq j$ then $y = w$, and that is absurd. We conclude that $j = \eta(y)$, and that η is a bijection.

We have already argued, in the paragraph above the statement of the proposition, that $y \geq \bigvee \text{can}(w) \setminus \{\eta(y)\}$. To complete the proof, suppose that $y \vee \eta(y)_* = w$. Since, $\text{can}(w)$ does not join-refine $\{y, \eta(y)_*\}$ (because $\eta(y) \not\leq \eta(y)_*$ and $\eta(y) \not\leq y$), we obtain a contradiction as above. We conclude that $y \vee \eta(y)_* < w$. Since y is covered by w , we have $y \vee \eta(y)_* = y$. Thus, $y \in \mathcal{K}(\eta(y))$, for each $y \in \text{cov}_\downarrow(w)$. \square

As a consequence of Lemma 2.3.3, we obtain a proof of Proposition 2.2.2, which we restate here with the notation from of Lemma 2.3.3.

Proposition 2.3.4. *Suppose that L is a finite join-semidistributive lattice, and y is covered by w in L . Then, $\eta(y)$ is the unique minimal element of L such that $\eta(y) \vee y = w$.*

Proof. Suppose that $x \in L$ has $x \vee y = w$. Since $\text{can}(w)$ join-refines $\{x, y\}$ and $\eta(y)$ and y are incomparable, we conclude that $\eta(y) \leq x$. \square

In fact, the previous proposition characterizes of finite join-semidistributive lattices. (Similar constructions exist; for example, see the proof of [1, Theorem 3-1.4].) Because the proof is similar to the proof of Lemma 2.3.3, we leave the details to the reader.

Proposition 2.3.5. *Suppose that L is a finite lattice. The following conditions are equivalent:*

1. *For each $w \in L$ and each $y \in \text{cov}_\downarrow(w)$, there is a unique minimal element $\eta(y) \in L$ satisfying $y \vee \eta(y) = w$.*
2. *L is join-semidistributive.*

Suppose that L is a finite join-semidistributive lattice, $j \in \text{Irr}(L)$, and F is a canonical join representation. The next lemma, in particular, implies that $F \cup \{j\}$ is a canonical join representation if and only if $\bigvee F \vee j > \bigvee F \vee j_*$.

Lemma 2.3.6. *Suppose that L is a finite join-semidistributive lattice and $j \in \text{Irr}(L)$. Then:*

1. *j is a canonical joinand of $y \vee j$, for each $y \in \mathcal{K}(j)$;*
2. *j is a canonical joinand of $\bigvee F \vee j$ if and only if $\bigvee F \vee j > \bigvee F \vee j_*$, for each subset F of $L \setminus \{j\}$.*

Proof. If $y = j_*$, then the first statement is obvious (because $\{j\}$ is the canonical join representation), so we assume that y and j are incomparable. We write w for the join $j \vee y$, and we write $A = \{j' \in \text{can}(w) : j' \leq j\}$ and $A' = \{j' \in \text{can}(w) : j' \leq y\}$. Because $\text{can}(w)$ join-refines $\{j, y\}$,

we have $A \cup A' = \text{can}(w)$. Also, the set A is not empty because the join $y \vee j$ is irredundant. We want to show that $A = \{j\}$. Since j is join-irreducible, it is enough to show that $j = \bigvee A$. Since $y \geq \bigvee A'$, we see that $\bigvee A \vee y = j \vee y$. If $\bigvee A < j$, then $j_* \vee y = j \vee y$, and that is impossible because $y \in \mathcal{K}(j)$. We conclude that j is a canonical joinand of $y \vee j$.

If $\bigvee F \vee j > \bigvee F \vee j_*$, then $\bigvee F \vee j_* \in \mathcal{K}(j)$. We conclude that j is a canonical joinand of $\bigvee F \vee j$. The remaining direction of the second item is straightforward to verify. \square

We close this subsection by quoting the following easy proposition (for example see [72, Proposition 2.2]), which says that the canonical join complex is indeed a simplicial complex.

Proposition 2.3.7. *Suppose L is a finite lattice, and the join $\bigvee A$ is a canonical join representation in L . Then each proper subset of A also joins canonically.*

2.3.2 The flag property

In this section we prove Theorem 2.1.1. We begin by presenting the key arguments in one direction the proof: If L is a finite semidistributive lattice, then its canonical join complex is flag. Most of the work is done in the following two lemmas.

Lemma 2.3.8. *Suppose that L is a finite semidistributive lattice, and F is a subset of $\text{Irr}(L)$ such that $|F| \geq 3$ and each proper subset of F is a face in the canonical join complex. Then the joins $\bigvee(F \setminus \{j\})$ and $\bigvee(F \setminus \{j'\})$ are incomparable for each distinct j and j' in F .*

Proof. Without loss of generality we assume that $\bigvee F = \hat{1}$. Suppose there exists distinct $j, j' \in F$ such that $\bigvee(F \setminus \{j\}) \geq \bigvee(F \setminus \{j'\})$. On the one hand, we have $(\bigvee(F \setminus \{j\})) \vee (\bigvee(F \setminus \{j'\}))$ is equal to $\bigvee F = \hat{1}$. On the other hand, $(\bigvee(F \setminus \{j\})) \vee (\bigvee(F \setminus \{j'\})) = \bigvee(F \setminus \{j\})$. Thus, $\bigvee(F \setminus \{j\}) = \hat{1}$. Since F has at least three elements, there exists $j'' \in F \setminus \{j, j'\}$. We write w' for




Figure 2.4: Dashed lines represent order relations in L while solid lines represent cover relations.

$\bigvee(F \setminus \{j'\})$ and w'' for $\bigvee(F \setminus \{j''\})$. See Figure 2.4. Because both $F \setminus \{j'\}$ and $F \setminus \{j''\}$ are

faces in the canonical join complex, j is a canonical joinand for both w' and w'' . Lemma 2.3.3 implies that there exists $y' \in \text{cov}_\downarrow(w')$ and $y'' \in \text{cov}_\downarrow(w'')$ such that $y', y'' \in \mathcal{K}(j)$. Moreover, $y' \geq \vee(F \setminus \{j, j'\})$ and similarly $y'' \geq \vee(F \setminus \{j, j''\})$.

So, we have: $y' \vee y'' \geq (\vee(F \setminus \{j, j'\})) \vee (\vee(F \setminus \{j, j''\})) = \vee(F \setminus \{j\})$. Since $\vee(F \setminus \{j\}) = \hat{1}$, we conclude that $\vee \mathcal{K}(j) = \hat{1}$, contradicting Proposition 2.3.2. \square

Lemma 2.3.9. *Suppose that L is a finite join-semidistributive lattice, and F is a subset of $\text{Irr}(L)$ satisfying the following conditions: First, $|F| \geq 3$; second, each proper subset of F is a face in the canonical join complex of L ; third, $\vee F$ is irredundant; fourth F is not a face of the canonical join complex. Then there exists $j \in F$ such that $\kappa(j)$ does not exist.*

Proof. Without loss of generality, we assume that $\vee F = \hat{1}$. Since the join $\vee F$ is irredundant, there exists some $j \in F$ such that $j \notin \text{can}(\hat{1})$. Lemma 2.3.6 implies that $(\vee(F \setminus \{j\})) \vee j_* = \hat{1}$. Let j' and j'' be distinct elements in $F \setminus \{j\}$. As in the proof of Lemma 2.3.8, let y' and y'' be the unique elements covered by $\vee F \setminus \{j'\}$ and $\vee F \setminus \{j''\}$, respectively, with $y', y'' \in \mathcal{K}(j)$. Thus, $y', y'' \geq j_*$. Also $y' \geq \vee(F \setminus \{j, j'\})$ and $y'' \geq \vee(F \setminus \{j, j''\})$. Therefore, $y' \vee y'' \geq \vee(F \setminus \{j\}) \vee j_* = \hat{1}$. The statement follows. \square

Proof of one direction of Theorem 2.1.1. We show that if L is semidistributive, then its canonical join complex is flag. Suppose that $F \subset \text{Irr}(L)$ such that $|F| \geq 3$ and each proper subset of F is a face of the canonical join complex. Without loss of generality, assume that $\vee F = \hat{1}$. Lemma 2.3.8 says that for each distinct j and j' in F , the joins $\vee(F \setminus \{j\})$ and $\vee(F \setminus \{j'\})$ are incomparable. Thus, we have

$$\vee(F \setminus \{j\}) < (\vee(F \setminus \{j\})) \vee (\vee(F \setminus \{j'\})) = \vee F.$$

We conclude that $\vee F$ is irredundant. Lemma 2.3.9 implies that F is a face of the canonical join complex. \square

We now turn to the other direction of Theorem 2.1.1. In the following lemmas we will assume that L is a finite join-semidistributive lattice in which fails SD_\wedge . By Proposition 2.3.2, there is some $j \in \text{Irr}(L)$ such that $\kappa(j)$ does not exist. Our goal is to construct a set $A \subset \text{Irr}(L)$ such that $A \cup \{j\}$ is a “hollow face” in the canonical join complex. More precisely, the set A must satisfy the following conditions. (NF stands for “not-flag”.)

(NF1) $A \cup \{j\}$ is *not* a face in the canonical join complex of L .

(NF2) Each pair of elements in $A \cup \{j\}$ is a face in the canonical join complex.

The essential idea is that among all of the subsets of $\text{Irr}(L)$ satisfying (NF1), a set A chosen as low as possible in L will also satisfy (NF2). For us, “as low as possible” means that A is chosen

to be minimal in join-refinement. The argument is somewhat delicate because join-refinement is a preorder, not a partial order, on subsets of L . So, we must take extra care to compare only antichains $Y \subset \text{Irr}(L)$ satisfying (NF1). To further emphasize this point, we write $A \ll B$ when A join-refines B , for *antichains* A and B . We write $A(j)$ for the collection of antichains $Y \subseteq L \setminus \{j\}$ satisfying $Y \cup \{j\}$ is an antichain. We write $E(j)$ for the set of $j' \in \text{Irr}(L) \setminus \{j\}$ such that $j' \vee j$ is a canonical join representation. When it is possible, we suppress j , and simply write E .

Lemma 2.3.10. *Suppose that L is a finite join-semidistributive lattice and $j \in \text{Irr}(L)$ such that $\kappa(j)$ does not exist. Let E be the set of $j' \in \text{Irr}(L) \setminus \{j\}$ such that $j' \vee j$ is a canonical join representation. Then:*

1. $\bigvee E \vee j = \bigvee E \vee j_*$;
2. *There exists a nonempty antichain Y in $A(j)$ such that $\bigvee Y \vee j = \bigvee Y \vee j_*$.*

Proof. Assume that $\bigvee E \vee j > \bigvee E \vee j_*$. Lemma 2.3.6 says that j is a canonical joinand of $\bigvee E \vee j$. Also, for each element a in $\mathcal{K}(j)$, j is a canonical joinand of $a \vee j$. That is, $a \vee j$ has the canonical join representation $\bigvee E' \vee j$ for some subset $E' \subset E$. Thus $a \vee j \leq \bigvee E \vee j$, and in particular $a \leq \bigvee E \vee j$. Lemma 2.3.3 implies that there is a unique element $y \in \mathcal{K}(j)$ covered by $\bigvee E \vee j$. If $a \not\leq y$, then $y \vee a = \bigvee E \vee j$. Proposition 2.3.4 says that j is the unique minimal element of L whose join with y is equal to $\bigvee E \vee j$. Therefore, $j \leq a$, contradicting the fact that $a \in \mathcal{K}(j)$. We conclude that $a \leq y$. We have proved that $y = \kappa(j)$, contradicting our hypothesis. Thus, $\bigvee E \vee j = \bigvee E \vee j_*$.

For the second statement, observe that if E is empty, then Lemma 2.3.6 implies that $\mathcal{K}(j) = \{j_*\}$, contradicting the assumption that $\kappa(j)$ does not exist. We conclude that E is nonempty. Since the antichain of maximal elements $Y \subseteq E$ satisfies $\bigvee Y = \bigvee E$, we have the desired result. \square

Lemma 2.3.10 says that the collection of antichains Y in $A(j)$ satisfying

$$\bigvee Y \vee j = \bigvee Y \vee j_* \tag{NC}$$

is nonempty. (Actually, we have shown something stronger: The collection of antichains $Y \subseteq E(j)$ that satisfy (NC) is nonempty.) We write (NC) for “not-canonical” because Lemma 2.3.6 implies that $\bigvee Y \vee j$ is not a canonical join representation. In particular, j is *not* a canonical joinand of $\bigvee Y \vee j$.

We choose such an antichain in $A(j)$ so that it is minimal in join-refinement. Suppose that B is this antichain, taken “as low as possible”. The next lemma is the difficult part of the proof of the remaining direction of Theorem 2.1.1. We argue that $(B \setminus \{b\}) \cup \{j\}$ is a face in canonical

join complex for each $b \in B$. Thus, if B has at least three elements, then $B \cup \{j\}$ is the “hollow face” that we want to construct. We will deal with the case where $|B| \leq 2$ in Lemma 2.3.12 and Lemma 2.3.13.

Lemma 2.3.11. *Suppose that L is a finite join-semidistributive lattice and $j \in \text{Irr}(L)$ such that $\kappa(j)$ does not exist. Among all antichains in $A(j)$ that satisfy (NC), let B be minimal in join-refinement. Then $(B \setminus \{b\}) \cup \{j\}$ is a canonical join representation, for each $b \in B$.*

Proof. We begin by pointing out two easy observations about the join-refinement relation. (Note that the second observation, (JR2), may fail if $S \cup \{x\}$ and $T \cup \{x\}$ are not antichains.)

(JR1) For any pair of subsets S and T , if S join-refines T then each subset $S' \subseteq S$ also join-refines T .

(JR2) Suppose that $S \cup \{x\}$ and $T \cup \{x\}$ are antichains. Then, $S \cup \{x\} \leq T \cup \{x\}$ if and only if $S \leq T$.

In particular, (JR1) implies that $B \setminus \{b\} \leq B$. Thus, $\vee(B \setminus \{b\}) \vee j_* < \vee(B \setminus \{b\}) \vee j$. Lemma 2.3.6 says that j is a canonical joinand of $\vee(B \setminus \{b\}) \vee j$. We write $C \cup j$ for the canonical join representation of $\vee(B \setminus \{b\}) \vee j$, where $j \notin C$. We claim that $C \cup \{b\} = B$. In Figure 2.5, we depict the relationship between C , B , and $B \setminus \{b\}$ in the join-refinement order. In the figure, we have $C \cup \{j\} \leq (B \setminus \{b\}) \cup \{j\}$, because $C \cup \{j\}$ is the canonical join representation for $\vee(B \setminus \{b\}) \vee j$. By (JR2), we have $C \leq B \setminus \{b\}$.




Figure 2.5: Some order relations in the join-refinement order for L .

We make two observations that follow immediately from Lemma 2.3.6. First, we observe that j *not* a canonical joinand of $\vee B \vee j$. Thus,

$$\vee C \vee j = \vee(B \setminus \{b\}) \vee j < \vee B \vee j. \quad (2.3.1)$$

Second, we observe that:

$$\vee(C \cup \{b\}) \vee j = \vee(C \cup \{b\}) \vee j_*. \quad (2.3.2)$$

Indeed, if $\vee(C \cup \{b\}) \vee j_* < \vee(C \cup \{b\}) \vee j$ then Lemma 2.3.6 says that j is a canonical joinand of $\vee(C \cup \{b\}) \vee j = \vee B \vee j$. We have just noted that j is a *not* a canonical joinand for $\vee B \vee j$.

If $C \cup \{b\}$ is an antichain, then applying (JR2) to the relation $C \leq B \setminus \{b\}$, we have $C \cup \{b\} \leq B$. Thus, we have $C \cup \{b\}$ is an antichain in $A(j)$ that satisfies (NC) and join-refines B . By minimality of B , we conclude that $C \cup \{b\} = B$ as desired. So, we assume that $C \cup \{b\}$ is not an antichain. The inequality in (2.3.1) implies that there exists no $c \in C$ with $b \leq c$. We write C' for the set $\{c \in C : c < b\}$.




Figure 2.6: Some relations in the join-refinement order for L .

We make three easy observations: First, $(C \setminus C') \cup \{b\}$ is member of $A(j)$. Second, applying (JR1) to the relation $C \leq B \setminus \{b\}$, we have that $C \setminus C' \leq B \setminus \{b\}$. By (JR2), we conclude that $(C \setminus C') \cup \{b\} \leq B$. We depict these relations in Figure 2.6. Third,

$$\vee((C \setminus C') \cup \{b\}) \vee j = \vee(C \cup \{b\}) \vee j = \vee(C \cup \{b\}) \vee j_* = \vee((C \setminus C') \cup \{b\}) \vee j_*,$$

where the first and third equalities follow from the fact that $\vee(C \cup \{b\})$ is equal to $\vee((C \setminus C') \cup \{b\})$, and the middle equality is (2.3.2).

Therefore, the set $(C \setminus C') \cup \{b\}$ is an antichain in $A(j)$ that satisfies (NC) and join-refines B . By the minimality of B , we have $B = (C \setminus C') \cup \{b\}$. Since $C \leq B \setminus \{b\}$, we have that C join-refines its proper subset $C \setminus C'$. That is a contradiction (because C is an antichain). Thus, C' is empty. We have proved the desired result. \square

Our candidate for a “hollow face” in the canonical join complex is the antichain $B \cup \{j\}$ from Lemma 2.3.11. As we have noted, if B has at least three elements then B satisfies both (NF1) and (NF2).

Suppose that $B = \{b_1, b_2\}$. By Lemma 2.3.11, $\{j, b_i\}$ is a canonical join representation, for $i = 1, 2$. (Thus, B is minimal in join-refinement among the antichains *in* $E(j)$ that satisfy (NC).) The next lemma, in particular, implies that $\{b_1, b_2\}$ is a canonical join representation.

Lemma 2.3.12. *Suppose that L is a finite join-semidistributive lattice and $j \in \text{Irr}(L)$ such that $\kappa(j)$ does not exist. Among all antichains in $A(j)$ that satisfy (NC) , let B be minimal in join-refinement. Suppose that B has at least two elements. Then each pair of elements in $B \cup \{j\}$ is a face in the canonical join complex.*

Proof. If B has three or more elements, then the statement follows from Lemma 2.3.11 and Proposition 3.2.1. Assume that B has two elements, b_1 and b_2 . By Lemma 2.3.11, we have $\{j, b_i\}$ is a canonical join representation, for $i = 1, 2$. Consider $\{b_1, b_2\}$. We will argue that b_1 is a canonical joinand of $b_1 \vee b_2$, and complete the proof by symmetry.

Assume that $(b_1)_* \vee b_2 = b_1 \vee b_2$. We observe that $(b_1)_* \vee b_2 \vee j = (b_1)_* \vee b_2 \vee j_*$. If $(b_1)_* \leq j_*$, then we have $b_2 \vee j = b_2 \vee j_*$, contradicting Lemma 2.3.6. By the same reasoning, $(b_1)_* \not\leq b_2$. Also, $j \not\leq (b_1)_*$ because b_1 and j are incomparable. Similarly, $b_2 \not\leq (b_1)_*$. Thus, $\{(b_1)_*, b_2\}$ is an antichain in $A(j)$ that satisfies (NC) and join-refines $\{b_1, b_2\}$. But this contradicts our hypothesis, which says that B is minimal in join-refinement among all such antichains. By this contradiction, we conclude that $(b_1)_* \vee b_2 < b_1 \vee b_2$. Lemma 2.3.6 says that b_1 is a canonical joinand of $b_1 \vee b_2$. \square

Finally, we turn to the case where B is a singleton. This turns out to be a non-issue. The next lemma says that we can always find such an antichain in $A(j)$ with at least two elements.

Lemma 2.3.13. *Suppose that L is a finite join-semidistributive lattice and $j \in \text{Irr}(L)$ such that $\kappa(j)$ does not exist. Then there exists an antichain $A \in A(j)$ satisfying:*

1. *A has at least two elements; and*
2. *A is minimal in join-refinement among all antichains in $A(j)$ that satisfy (NC) .*

Proof. Recall that $E(j)$ is the set of $j' \in \text{Irr}(L) \setminus \{j\}$ such that $j' \vee j$ is a canonical join representation. Take A to be a nonempty antichain that is minimal in join-refinement among all antichains $Y \subseteq E(j)$ that satisfy (NC) . Lemma 2.3.10 implies that such an antichain A exists. For each $a \in A$, we have $a \vee j$ is a canonical join representation. It follows immediately from Lemma 2.3.6 that A has at least two elements.

Now we prove that A is minimal in join-refinement among all antichains in $A(j)$ that satisfy (NC) . Suppose that $B \in A(j)$ satisfies (NC) , and $B \leq A$. Without loss of generality, assume that B is minimal in join-refinement with this property. If B has two or more elements, then Lemma 2.3.11 implies that B is a subset of $E(j)$. Therefore, $B = A$. Thus we can assume that $B = \{b\}$. Since B join-refines A , there is some $a \in A$ such that $b \leq a$.




Figure 2.7: A depiction of the argument for Lemma 2.3.13. Dashed lines represent order relations in L while solid lines represent cover relations.

Write w for $a \vee j$. Since $a \vee j$ is the canonical join representation of w , Lemma 2.3.3 implies that $\text{cov}_\downarrow(w)$ has precisely two elements, y and y' . Let $\eta(y) = j$ and $\eta(y') = a$, so that $y \in \mathcal{K}(j)$ and $y \geq a$. See Figure 2.7. Thus, we have $b \leq a \leq y$. On the one hand, $(b \vee j) \vee y = (b \vee j_*) \vee y = y$. On the other hand, $b \vee (j \vee y) = b \vee w = w$. By this contradiction, we have proved the result. \square

Finally, we complete the proof of the main result.

Proof of the remaining direction of Theorem 2.1.1. Now we argue that if L is a finite join-semidistributive lattice and the canonical join complex of L is flag, then L is semidistributive. By Proposition 2.3.2, it is enough to show that for each $j \in \text{Irr}(L)$ the element $\kappa(j)$ exists.

Suppose $j \in \text{Irr}(L)$ and $\kappa(j)$ does not exist. Among all nonempty antichains in $A(j)$ that satisfy (NC) , let A be minimal in join-refinement, and choose A so that it has at least two elements. Lemma 2.3.13 says that such an antichain A exists. Lemma 2.3.6 implies that $A \cup \{j\}$ is not face of the canonical join complex. Finally, Lemma 2.3.12 says that each pair of elements in $A \cup \{j\}$ is face in the canonical join complex. We have reached a contradiction to our hypothesis that the canonical join complex is flag. By this contradiction, we conclude that L is semidistributive. \square

Suppose that m is meet-irreducible and write m_* for the unique element covering m . When it exists, let $\kappa_*(m)$ be the smallest element $j \in L$ with $j \leq m_*$ and $j \not\leq m$. It is immediate that $\kappa_*(m)$ is join-irreducible. Proposition 2.3.2, applied to the dual lattice, says that L is meet-semidistributive if and only if $\kappa_*(m)$ exists for each meet-irreducible element m . In fact, L is semidistributive if and only if κ is a bijection, with inverse map κ_* ; this is the finite case of [39, Corollary 2.55]. Applying the dual argument for the canonical meet complex, we immediately obtain the following result. (Recall that Theorem 2.3.1 says that each element in L has a canonical meet representation if and only if L is meet-semidistributive.)

Corollary 2.3.14. *Suppose that L is a finite meet-semidistributive lattice. Then, the canonical meet complex for L is flag if and only if L is semidistributive.*

Next, we prove Corollary 2.1.5 by showing that the bijection κ taking a join-irreducible element j to $\kappa(j)$ induces an isomorphism from the canonical join complex of L to the canonical meet complex of L .

Proof of Corollary 2.1.5. Corollary 2.3.14 says that the canonical meet complex of L is flag, so it is enough to show that κ bijectively maps edges of the canonical join complex to edges of the canonical meet complex. Suppose that $\{j_1, j_2\}$ is a face of the canonical join complex, and write m_1 for $\kappa(j_1)$ and m_2 for $\kappa(j_2)$. Suppose that $m_1 \wedge m_2 = (m_1)_* \wedge m_2$. Lemma 2.3.3 implies that there exists some $y \in \text{cov}_\downarrow(j_1 \vee j_2)$ satisfying $j_1 \leq y \leq \kappa(j_2)$. (See Figure 2.8 for an illustration.) Since $j_1 \leq (m_1)_*$, we conclude that $j_1 \leq (m_1)_* \wedge m_2 = m_1 \wedge m_2$. We see that $j_1 \leq m_1$ and that is a contradiction. Therefore, $(m_1)_* \wedge m_2 > m_1 \wedge m_2$. By the dual statement of Lemma 2.3.6, we conclude that m_1 is a canonical meetand of $m_1 \wedge m_2$, and by symmetry m_2 is also a canonical meetand of $m_1 \wedge m_2$. The dual argument establishes the desired isomorphism. \square




Figure 2.8: An illustration of the argument for the proof of Corollary 2.1.5. Dashed gray lines represent relations in L while solid black lines represent cover relations.

We close this section by relating Corollary 2.1.5 to Example 2.2.3 and Example 2.2.8, from Section 2.2.

Remark 2.3.15. Suppose that F is a face of the canonical join complex of a finite semidistributive lattice L . Corollary 2.1.5 says that $\wedge \kappa(F)$ is a canonical meet representation. By taking the canonical join representation of $\wedge \kappa(F)$, we can view the map κ as an operation on the canonical join complex. Similarly, we can view κ_* as an operation on the canonical meet complex.

The main premise of [3] is that the action of Kreweras complementation on the set of noncrossing partitions and the action of Panyshchev complementation on the set of nonnesting

partitions (that is, the set of antichains in the root poset for a finite crystallographic root system) coincide. Indeed, both maps are an instance of the operation of κ (or κ_*) on the canonical join complex (or canonical meet complex).

On the one hand, the action of κ on the canonical join complex of the Tamari lattice coincides with Kreweras complementation (recall from Example 2.2.8 that canonical join representations in the Tamari lattice are essentially noncrossing partitions). On the other hand, Panyshchikov complementation is a special case of an operation on the set of antichains in a finite poset \mathcal{P} called *rowmotion*, as we now explain. When A is an antichain in \mathcal{P} , we write $\text{Row}(A)$ for the antichain $\{x \in \mathcal{P} : x \text{ is minimal among elements not in } I_A\}$. (Our notation is based on [86]. See also [3, 18, 19, 37, 62, 80].) So, we have $I_A = I^{\text{Row}(A)}$. It follows immediately from the definition of κ_* that $\kappa_*(I^y) \mapsto I_y$. We obtain the following result.

Proposition 2.3.16. *Suppose that \mathcal{P} is a finite poset, and A is an antichain in \mathcal{P} . Then the map κ_* , acting on faces of the canonical meet complex of $J(\mathcal{P})$, sends the order ideal I^A to the order ideal $I^{\text{Row}(A)}$.*

2.3.3 Crosscut-simplicial lattices

In this section, we prove Corollary 2.1.2. Recall that one direction of the proof was given as [57, Theorem 3.1]. Because it is easy, we give an alternative argument in the next paragraph.

Write A for the set of atoms in L . When L is a finite semidistributive lattice every join of two atoms is a canonical join representation. In particular, Theorem 2.1.1 implies that each distinct subset of atoms gives rise to a distinct element in L . Thus the crosscut complex for L is either the boundary of the simplex on A or equal to the simplex on A , depending on whether $\bigvee A = \hat{1}$ or $\bigvee A < \hat{1}$. Since each interval in L inherits semidistributivity, it follows that L is crosscut-simplicial.




Figure 2.9: A finite crosscut-simplicial lattice failing both SD_{\vee} and SD_{\wedge} .

Before we proceed with the proof of the converse, we point out that the join-semidistributivity hypothesis in Corollary 2.1.2 is crucial. (For example, consider the crosscut-simplicial lattice

shown in Figure 2.9. This lattice fails both SD_{\vee} and SD_{\wedge} .) Join-semidistributivity gives us a powerful restriction: A finite join-semidistributive lattice L fails SD_{\wedge} if and only if L contains the lattice shown in Figure 2.1 as a sublattice ([39, Theorem 5.56]).

We now begin our proof. The following lemmas will be useful. The first lemma is a local version of Theorem 2.3.1, and appears as [73, Lemma 9-2.5].

Lemma 2.3.17. *Suppose that L is a finite lattice satisfying the following property:*

If $x, y,$ and z are elements of L with $x \wedge y = x \wedge z$ and also, y and z cover a common element, then $x \wedge (y \vee z) = x \wedge y$.

Then L is meet-semidistributive.

Lemma 2.3.18. *Suppose that L is a finite join-semidistributive lattice that fails SD_{\wedge} . Then there exists $x, y,$ and z such that $y \vee z > x$ and $x, y,$ and z cover a common element.*

Proof. We prove the proposition by induction on the size of L . As mentioned above, L contains the lattice shown in Figure 2.1 as sublattice, and this proves the base case. By Lemma 2.3.17, we can assume that there exist $x, y,$ and z in L such that $x \wedge y = x \wedge z$, $x \wedge (y \vee z) \neq x \wedge y$, and $\text{cov}_{\downarrow}(y) \cap \text{cov}_{\downarrow}(z)$ is not empty. Among all such triples, we choose $\{x, y, z\}$ minimal in join-refinement. Write a for the element in $\text{cov}_{\downarrow}(y) \cap \text{cov}_{\downarrow}(z)$ (if there is more than one element in $\text{cov}_{\downarrow}(y) \cap \text{cov}_{\downarrow}(z)$, then $y \wedge z$ does not exist). If x also covers a , then we are done (because if $x \succ a$ and $y \vee z \not\prec x$, then $(y \vee z) \wedge x = a$, and that contradicts our assumption that $\{x, y, z\}$ fail SD_{\wedge}). So we assume that x does not cover a .




Figure 2.10: Dashed lines represent relations in L and solid lines represent cover relations.

We first prove that $x < y \vee z$. We write w for $x \wedge (y \vee z)$. See Figure 2.10. Observe that $x \wedge y < w$ (because x, y and z fail SD_{\wedge} , the inequality is strict). On the one hand $w \wedge (x \wedge y) = x \wedge y$. On the

other hand, $x \geq w$, so $(x \wedge w) \wedge y = w \wedge y$. By symmetry, $w \wedge z = x \wedge z$. Therefore, $w \wedge y = w \wedge z$. Note that $w \neq y \wedge w$ (otherwise $w \leq y \wedge x$, and that is absurd). Finally, we observe that $w \wedge (y \vee z) = w$. Thus, $\{w, y, z\}$ fails SD_\wedge . The minimality of $\{x, y, z\}$ in join-refinement implies that $w = x$. We have proved the claim that $y \vee z > x$. By induction, we may assume $y \vee z = \hat{1}$.




Figure 2.11: Dashed lines represent relations in L while solid lines represent cover relations.

Next, we claim that $x \vee y$ and $x \vee z$ are incomparable. By way of contradiction assume that $x \vee z \geq x \vee y$, so we have $x \vee z \geq x, y, z$. Therefore, $z \vee x = z \vee y$, as shown on the left in Figure 2.11. Observe that $z \vee (x \wedge y) = z$. Since L is join-semidistributive, we have $z = \hat{1}$. This contradicts the fact that $x \wedge z \neq x \wedge (y \vee z)$. We have proved the claim that $x \vee y$ and $x \vee z$ are incomparable.

Finally, we claim that there is some $w' \in \text{cov}^\uparrow(a) \setminus \{y, z\}$. Suppose that $\{y, z\} = \text{cov}^\uparrow(a)$, and consider the righthand of Figure 2.11. Either $y \leq a \vee x$ or $z \leq a \vee x$, but not both. Indeed, if $x \vee a \geq y, z$ then $x \vee a = \hat{1}$, so $x \vee a = x \vee y = x \vee z$. This contradicts the fact that $x \vee y$ and $x \vee z$ are incomparable. By symmetry, we assume that $y \leq x \vee a$. Then $y \leq x \vee a \leq x \vee z$. Thus we have $x \vee y \leq x \vee z$, also contradicting the fact that $x \vee y$ and $x \vee z$ are incomparable. We conclude that there exists some $w' \in \text{cov}^\uparrow(a) \setminus \{y, z\}$ (in particular, $w' \leq a \vee x$). The claim follows. The triple $\{w', y, z\}$ satisfies the statement of the proposition. \square

Proof of Theorem 2.1.2. We prove that if L is join-semidistributive and crosscut-simplicial then it is semidistributive. Assume that L fails SD_\wedge . Lemma 2.3.18 says that there exists x, y and z covering a common element $a \in L$ such that $y \vee z > x$. In particular, $[a, y \vee z]$ is not crosscut-simplicial because $\{y, z\}$ is not a face in the crosscut complex. That is a contradiction. Therefore, L is a finite semidistributive lattice, and the statement follows from Theorem 2.1.1. \square

2.4 Lattice-theoretic constructions

2.4.1 Sublattices and quotient lattices

A map $\phi : L \rightarrow L'$ between lattices L and L' is a **lattice homomorphism** if ϕ respects the meet and join operations. The image of ϕ is a **sublattice** of L' and a **lattice quotient** of L . It is immediate that each sublattice of a semidistributive lattice is also semidistributive. When L is finite, the image $\phi(L)$ also inherits semidistributivity (see [73, Proposition 1-5.24]). Outside of the finite case, it is not generally true that if L is semidistributive, then $\phi(L)$ is semidistributive. (Similarly statements hold for meet and join-semidistributivity.) We obtain the following result as an immediate corollary of Theorem 2.1.1.

Corollary 2.4.1. *Suppose that L is a finite join-semidistributive lattice whose canonical join complex is flag. Then, the canonical join complex of each sublattice and quotient lattice of L is also flag.*

An equivalence relation Θ on L is a **lattice congruence** if Θ satisfies the following: if $x \equiv_{\Theta} y$, then $x \vee t \equiv_{\Theta} y \vee t$ and $x \wedge t \equiv_{\Theta} y \wedge t$ for each x, y , and t in L (see [44, Lemma 8]). It is immediate that the fibers of the lattice homomorphism ϕ constitute a lattice congruence of L . Conversely, each lattice congruence also gives rise to a lattice quotient (see [44, Theorem 11]).

When L is finite, Θ is lattice congruence if and only if it satisfies the following: Each class is an interval; the map $\pi_{\downarrow}^{\Theta}$ sending $x \in L$ to the smallest element in its Θ -class is order preserving; the map π_{\uparrow}^{Θ} sending $x \in L$ to the largest element in its Θ -class is order preserving. Both $\pi_{\downarrow}^{\Theta}$ and π_{\uparrow}^{Θ} are lattice homomorphisms onto their images, and $\pi_{\downarrow}^{\Theta}(L)$ and $\pi_{\uparrow}^{\Theta}(L)$ are isomorphic lattice quotients of L . As a lattice quotients, both $\pi_{\downarrow}^{\Theta}(L)$ and $\pi_{\uparrow}^{\Theta}(L)$ are endowed with their own join and meet operations. So, for example, when we write $\vee A$ or $\wedge A$ for some subset $A \subset \pi_{\downarrow}^{\Theta}(L)$, we must indicate whether the join or meet is taken in L or in its lattice quotient. It turns out that $\pi_{\downarrow}^{\Theta}(L)$ is a sub-join-semilattice of L , meaning that the join operation in $\pi_{\downarrow}^{\Theta}(L)$ coincides with the join operation in L . However, in general, the expression $\wedge A$ may differ depending on whether the meet is taken in L or in $\pi_{\downarrow}^{\Theta}(L)$. (In other words, $\pi_{\downarrow}^{\Theta}(L)$ is generally *not* a sublattice of L .) Similar statements hold for $\pi_{\uparrow}^{\Theta}(L)$.

Below we quote [71, Proposition 6.3]. In the proposition, a join-irreducible element $j \in L$ is **contracted** by the congruence Θ if j is congruent to the unique element that it covers.

Proposition 2.4.2. *Suppose that L is a finite join-semidistributive lattice and Θ is a lattice congruence on L with associated projection map $\pi_{\downarrow}^{\Theta}$. Then, the element x belongs to $\pi_{\downarrow}^{\Theta}(L)$ if and only if no canonical joinand of x is contracted by Θ .*

Suppose that $x \in \pi_{\downarrow}^{\Theta}(L)$. Since $\pi_{\downarrow}^{\Theta}$ is a sub-join-semilattice of L , the canonical join representation of x taken in the lattice quotient $\pi_{\downarrow}^{\Theta}(L)$ is equal to the canonical join representation taken in L .

Corollary 2.4.3. *Suppose that L is a finite join-semidistributive lattice with canonical join complex Δ , and Θ is a lattice congruence on L . Then, the canonical join complex of $\pi_{\downarrow}^{\Theta}(L)$ is the induced subcomplex of Δ supported on the set of join-irreducible elements not contracted by Θ .*

Remark 2.4.4. The canonical join complex of a sublattice L' of L need not be an induced subcomplex of Δ . In fact, the sets $\text{Irr}(L')$ and $\text{Irr}(L)$ may be disjoint. For example, consider the canonical join complex of the sublattice $\{\hat{0}, \hat{1}\}$ in the boolean lattice B_n , where $n > 1$.




Figure 2.12: The Tamari lattice T_3 and its canonical join complex.

Remark 2.4.5. In general, not every induced subcomplex of Δ is the canonical join complex of a lattice quotient of L . Each lattice congruence is determined by the set of join-irreducible elements that it contracts. But, a given collection of join-irreducible elements may not correspond to a lattice congruence. For j and j' in $\text{Irr}(L)$, we say that j **forces** j' if every congruence that contracts j also contracts j' . In the Tamari lattice T_3 pictured in Figure 2.12 both a and b force c . So, for example, there is no quotient of T_3 whose canonical join complex is the subcomplex induced by $\{b, c\}$.

2.4.2 Products and sums

In the following easy propositions, we construct new semidistributive lattices from old ones, and give the corresponding construction for the canonical join complex. Recall that the **join** of the simplicial complexes Δ and Δ' is the complex $\Delta * \Delta' = \{F \cup F' : F \in \Delta \text{ and } F' \in \Delta'\}$.

Proposition 2.4.6. *Suppose that L_1 and L_2 are finite, join-semidistributive lattices with corresponding canonical join complex Δ_i for $i = 1, 2$. Then the canonical complex for $L_1 \times L_2$ is the join $\Delta_1 * \Delta_2$.*

The **ordinal sum** of lattices L_1 and L_2 written $L_1 \oplus L_2$ is the lattice whose set of elements is the disjoint union $L_1 \uplus L_2$, ordered as follows: $x \leq y$ if and only if $x \leq y$ in L_i , for $i = 1, 2$, or $x \in L_1$ and $y \in L_2$.

Proposition 2.4.7. *Suppose that L_1 and L_2 are finite, join-semidistributive lattices with corresponding canonical join complex Δ_i , for $i = 1, 2$. Then the canonical join complex of $L_1 \oplus L_2$ is equal to the disjoint union $\Delta_1 \uplus \Delta_2 \uplus \{v\}$, in which the vertex v corresponds to the minimal element of L_2 .*

We define the **wedge sum** $L_1 \vee L_2$ to be the lattice quotient of the ordinal sum $L_1 \oplus L_2$ in which the minimal element of L_2 is identified with the maximal element of L_1 . (Our nonstandard terminology is inspired by the wedge sum of topological spaces.)

Proposition 2.4.8. *Suppose that L_1 and L_2 are finite, join-semidistributive lattices with corresponding canonical join complex Δ_i , for $i = 1, 2$. Then the canonical join complex of $L_1 \vee L_2$ is equal to the disjoint union $\Delta_1 \uplus \Delta_2$.*

2.4.3 Day's doubling construction

A subset C of L is **order-convex** if for each $x, y \in C$ with $x \leq y$, we have that the interval (x, y) belongs to C . Suppose that $C \subseteq L$ is order convex, and let $\mathbf{2}$ be the two element chain $0 < 1$. We write X for the set of elements $x \in L$ such that $x \geq c$ for some $c \in C$. Define $L[C]$ to be the following induced subposet of $L \times \mathbf{2}$:

$$[((L \setminus X) \cup C) \times 0] \uplus (X \times 1)$$

We say that $L[C]$ is obtained by **doubling** L with respect to C . This procedure, due to Day [28], is defined more generally for all posets. If L is also a lattice, then $L[C]$ is a lattice and the map $\pi_C : L[C] \rightarrow L$ given by $(x, \epsilon) \mapsto x$ is a surjective lattice homomorphism (see [28] or [57, Lemma 6.1]). In the next proposition, we show that when C is an interval in L , doubling L with respect to C also preserves semidistributivity.

Proposition 2.4.9. *Suppose that L is a finite semidistributive lattice, $I = [a, b]$ is an interval in L , and write \mathcal{E} for the edge set of the canonical join graph for L . Then $L[I]$ is semidistributive, and the canonical join graph for $L[I]$ has edge set*

$$\mathcal{E}' \uplus \{(j, 0), (a, 1)\} : j \in \text{can}(w) \text{ for } w \in I \text{ and } j \not\leq a\},$$

where \mathcal{E}' is the set of pairs $\{(j, \epsilon), (j', \epsilon')\}$, such that $\{j, j'\} \in \mathcal{E}$, and (j, ϵ) and (j', ϵ') are the minimal elements of the fibers $\pi_I^{-1}(j)$ and $\pi_I^{-1}(j')$, respectively.

In the proof below we check that $L[I]$ satisfies (1) from Proposition 2.3.5. (The obvious dual argument gives meet-semidistributivity). One can also verify semidistributivity directly for $L[I]$ using [57, Lemma 6.1]. Our approach has the advantage of giving the canonical join

representation of each element of $L[I]$. In either case, the argument is tedious but, at least, elementary.

Proof. Suppose that (w, ϵ) is not in $I \times 1$, where $\epsilon = 0, 1$. Observe that the map $\pi_I : (y, \epsilon') \mapsto y$ is a bijection from $\text{cov}_\downarrow((w, \epsilon))$ to $\text{cov}_\downarrow(w)$. For each $y \in \text{cov}_\downarrow(w)$, write $\eta(y)$ for the unique minimal element of L satisfying $y \vee \eta(y) = w$. We write (y, ϵ') for the corresponding element in $\text{cov}_\downarrow((w, \epsilon))$. Let $(\eta(y), \epsilon'')$ be the minimal element of the fiber $\pi_I^{-1}(\eta(y))$ in $L[I]$. We claim that $(\eta(y), \epsilon'') \vee (y, \epsilon') = (w, \epsilon)$. If $\epsilon = 0$, the claim is immediate, and if $\epsilon = 1$ then the claim follows from the fact that $(w, 0) \notin L[I]$. It is straightforward, using the surjection π_I , to check that $(\eta(x), \epsilon'')$ is the unique minimal element of $L[C]$ whose join with (x, ϵ') is equal to (y, ϵ) .

Suppose that $(w, 1) \in I \times 1$. If $w = a$, then $(w, 1) = (a, 1)$ is join-irreducible. So, it is immediate that it satisfies condition (1) of Proposition 2.3.5. So we assume that $w > a$. Observe that the lower covers of $(w, 1)$ are $(y, 1)$ such that $y \in \text{cov}_\downarrow(w) \cap I$ and $(w, 0)$. We claim that the set $\{\eta(y) : y \in \text{cov}_\downarrow(w) \cap I\}$ is precisely the set of canonical joinands of w that are not weakly below a . If $y \in \text{cov}_\downarrow(w) \setminus I$, then $y \vee a = w$. By minimality of $\eta(y)$, we conclude that $\eta(y) \leq a$. If $y \in \text{cov}_\downarrow(w) \cap I$ and $\eta(y) \leq a$, then $\eta(y) \vee y = y$, which is a contradiction. The claim follows. As above, it is straightforward to check that $(\eta(y), 0)$ is the unique minimal element in $L[I]$ whose join with $(y, 1)$ is equal to $(w, 1)$, for each $y \in \text{cov}_\downarrow(w) \cap I$.

Suppose that $(w', \epsilon') \vee (w, 0) = (w, 1)$, where $\epsilon' \in \{0, 1\}$. Then $\epsilon' = 1$, and we have $w' \geq a$. Therefore, $(a, 1)$ is the unique minimal element whose join with $(w, 0)$ is equal to $(w, 1)$. Proposition 2.3.5 says that L is join-semidistributive. The second statement follows from Proposition 2.3.4. \square

Below we gather some useful facts that follow immediately from the proof of Proposition 2.4.9.

Proposition 2.4.10. *Suppose that L is a finite semidistributive lattice, $I = [a, b]$ is an interval in L , and $j \in \text{Irr}(L)$ such that $j \neq a$. For each $w \in L$ and $\epsilon, \epsilon' \in \{0, 1\}$ the following statements hold:*

1. *If (j, ϵ) is a canonical joinand of (w, ϵ') in $L[I]$, then j is a canonical joinand of w .*
2. *If (j, ϵ) is a canonical joinand of $(w, \epsilon') \in I \times \mathbf{2}$ then $\epsilon = 0$.*
3. *If (j, ϵ) is a canonical joinand of $(w, 0) \in I \times 0$ and $j \not\leq a$, then (j, ϵ) is also a canonical joinand of $(w, 1)$.*
4. *(w, ϵ') has $(a, 1)$ as canonical joinand if and only if $(w, \epsilon') \in I \times 1$.*

A lattice is **congruence uniform** if it is obtained from the one element lattice by a finite sequence of doublings of intervals. Suppose that L is a finite congruence uniform lattice.

Proposition 2.4.9 says that after each iteration of the doubling procedure, the resulting lattice has precisely one additional join-irreducible element, namely $(a, 1)$, where a is the smallest element of the interval that is doubled. Thus the canonical join graph of each congruence uniform lattice L has a natural labeling, in which the vertex labeled i is the join-irreducible element that is added in the i^{th} step of the doubling sequence for L .

Remark 2.4.11. Non-isomorphic congruence uniform lattices may have the same labeled canonical join graphs. For example, doubling the boolean lattice B_2 with respect to any singleton interval $I = \{x\}$, results in the labeled canonical join graph depicted in Figure 2.13 below. When x is equal to $\hat{0}$ or $\hat{1}$, we obtain the ordinal sums $B_0 \oplus B_2$ and $B_2 \oplus B_0$, respectively. When x is either join-irreducible element of B_2 , the resulting lattice is isomorphic to the Tamari lattice T_3 from Figure 2.12.




Figure 2.13: The canonical labeled join graph of three non-isomorphic congruence uniform lattices.

We conclude this subsection with various applications of Proposition 2.4.9. In each example, we discuss how to realize a labeled or unlabeled graph as the canonical join graph for some congruence uniform lattice.

Example 2.4.12 (Complete graphs). In our first example we consider the complete graph K_n on n vertices, which can be realized as the canonical join graph for the boolean lattice B_n . In fact, the boolean lattice is the only lattice whose canonical join graph is K_n .

Proposition 2.4.13. *Supposed that L is a finite semidistributive lattice with canonical join graph equal to the complete graph K_n . Then, L is isomorphic to B_n .*

Proof. Write x_S for the element with canonical join representation $\vee(\{j_i : i \in S\})$, where S is subset of $[n] = \{1, 2, \dots, n\}$. Suppose that $x_S \leq x_{S'}$ for some $S' \subseteq [n]$, and there exists $k \in S$ that is not in S' . Since $j_k \vee \vee(\{j_i : i \in S'\})$ is a canonical join representation, in particular this join is irredundant. So, $j_k \not\leq \vee(\{j_i : i \in S'\}) = x_{S'}$, and that is a contradiction. Therefore, the map $x_S \mapsto S$ is order preserving. It is immediate that the inverse map is order preserving. \square

Example 2.4.14 (Chordal graphs). Similar to the construction of the complete graph (as a labeled canonical join graph), one can construct certain chordal graphs as the canonical join graph for a congruence uniform lattice. In the construction, each doubling with respect to some interval I has $I \times \mathbf{2}$ isomorphic to a boolean lattice.

Suppose that G is a graph. The *closed neighborhood* $N[v]$ is the subgraph of G induced by the set of vertices v' adjacent to v , together with v . The *open neighborhood* $N(v)$ is the subgraph induced by the set $N[v] \setminus \{v\}$. A *perfect elimination ordering* for G is a linear ordering $v_1 < v_2 < \dots < v_n$ of the vertices of G such that for each $i = 1, 2, \dots, n$, the intersection of $N[v_i]$ with the set $\{v_i, v_{i+1}, \dots, v_n\}$ is a clique in G . Recall that a graph G is *chordal* if and only if it has a perfect elimination ordering.

Proposition 2.4.15. *Suppose that G is a labeled graph and $\mathcal{L} = v_n < v_{n-1} < \dots < v_1$ is a perfect elimination ordering. If $N(v_{i+1}) \subseteq N(v_i)$ for each $i \in [n-1]$, then there exists a congruence uniform lattice L such that G is its labeled canonical join graph.*

Proof. We prove the statement by induction on n . We write L' for a congruence uniform lattice whose labeled canonical join graph is the subgraph induced by the first $n-1$ vertices in G . In particular, L' is isomorphic to $L''[I]$ where L'' is congruence uniform, $I = [a, b]$ is an interval in L'' , and the vertex v_{n-1} corresponds to the join-irreducible element $(a, 1)$ in L' .

We give the argument for the case when that v_n and v_{n-1} are neighbors. The proof is similar when $v_n \notin N(v_{n-1})$. We write $\{v_{i_1}, \dots, v_{i_k}\}$ for the set of vertices $N(v_n) \setminus \{v_{n-1}\}$, and j_{i_1}, \dots, j_{i_k} for the corresponding join-irreducible elements in L' . Since \mathcal{L} is a perfect elimination order, the vertices $\{v_{i_1}, \dots, v_{i_k}, v_{n-1}\}$ form a clique in the subgraph induced by $V \setminus \{v_n\}$. By Theorem 2.1.1 the canonical join complex of L' is flag. Thus, the expression $(a, 1) \vee (\bigvee \{j_{i_1}, \dots, j_{i_k}\})$ is the canonical join representation for some element in (y, ϵ) in L' . (In particular, that $j_{i_l} \not\leq (a, 0)$ for each $l \in [1, k]$.) The fourth item of Proposition 2.4.10 implies that $y \in [a, b]$ and $\epsilon = 1$.

Consider the interval $I' = [(a, 0), (y, 1)]$. We claim that G is the labeled canonical join graph for $L'[I']$. It is straightforward (with Proposition 2.4.9) to verify that the new vertex, v_n , in the canonical join graph for $L'[I']$, is adjacent to v_{i_1}, \dots, v_{i_k} , and v_{n-1} . Conversely, assume that v is some vertex that is adjacent to v_n in the canonical join graph for $L'[I']$. Write (j, ϵ') for the corresponding join-irreducible element in L' . To prove the claim, we need to show that $(j, \epsilon') \in \{j_{i_1}, \dots, j_{i_k}, (a, 1)\}$. (That is, we need to check that v_n is adjacent to *only* the vertices $\{v_{i_1}, \dots, v_{i_k}, v_{n-1}\}$.) This is obvious if $(j, \epsilon') = (a, 1)$, so we assume that $j \neq a$. Proposition 2.4.9 implies that (j, ϵ') is the canonical joinand for some element $(w, \epsilon'') \in I'$ and also $(j, \epsilon') \not\leq (a, 0)$.

First we show that $\bigvee \{(j, \epsilon'), (a, 1)\}$ is a canonical join representation in L' . Observe that $(w, \epsilon'') \in [a, b] \times \mathbf{2}$ (because $I' \subseteq [a, b] \times \mathbf{2}$). The first item of Proposition 2.4.10 says that j is a canonical joinand of w in L'' (and in particular, $j \leq w$). The second item says that $\epsilon'' = 0$, so that $(j, \epsilon') = (j, 0)$. Therefore, $j \not\leq a$. Proposition 2.4.9 implies that (j, ϵ') and $(a, 1)$ join canonically

in L' , as desired. For the remainder of the proof, we write $(j, 0)$ instead of (j, ϵ') .

Because \mathcal{L} is a perfect elimination ordering (and the vertex v corresponding to $(j, 0)$ occurs strictly earlier than v_{n-1}) the set $\{j_{i_1}, \dots, j_{i_k}, (j, 0), (a, 1)\}$ is a face in the canonical join complex of L' . If we can show that

$$\bigvee(\{j_{i_1}, \dots, j_{i_k}, (a, 1), (j, 0)\}) = (y, 1),$$

then we will have proved the claim. Clearly, $(a, 0) < (j, 0) \vee (a, 1)$. Also, $(j, 0) \vee (a, 1) \leq (w, 1)$ (because $j \leq w$ and $a < w$). We conclude that $(j, 0) \vee (a, 1)$ belongs to I' . Since I' is a sublattice of L' , the following expression also belongs to I' :

$$((a, 1) \vee (j, 0)) \vee ((a, 1) \vee \bigvee\{j_{i_1}, \dots, j_{i_k}\}).$$

Therefore, $\bigvee(\{j_{i_1}, \dots, j_{i_k}, (a, 1), (j, 0)\}) = (y, 1)$ as desired.

A similar argument, replacing the interval $[(a, 0), (y, 1)]$ with $[(a, 1), (y, 1)]$, proves the case in which v_n and v_{n-1} are not adjacent. The argument is valid because $N(v_n) \subseteq N(v_{n-1})$. \square




Figure 2.14: The two leftmost graphs are isomorphic Hasse diagrams for the distributive lattice L . Rightmost is the lattice obtained by doubling the interval $[a, e]$ in L .

Example 2.4.16 (Cycle graphs). For each positive integer n , there is a finite congruence uniform lattice whose canonical join graph is isomorphic to the unlabeled cycle graph C_n on n vertices. We provide an illustration for $n = 5, 6$, and 7 . Leftmost in Figure 2.14 is the Hasse diagram for a distributive lattice L , and rightmost is the Hasse diagram obtained by doubling the interval $[a, e]$ in L . (The middle Hasse diagram is isomorphic to the leftmost Hasse diagram, and serves only to make the doubling as clear as possible.) Each distributive lattice is congruence uniform, so the rightmost lattice is congruence uniform, as desired. It is an easy exercise to verify that the canonical join graph for this right-most lattice is isomorphic to C_5 .

The analogous construction is given in Figure 2.15 for $n = 6$ and 7 . In these cases, the lattice




Figure 2.15: Doubling the interval $[a, e]$ in the leftmost congruence uniform lattice yields the left-middle lattice, whose canonical join graph is isomorphic to C_6 . Doubling the interval $[a, e]$ in the right-middle lattice yields the rightmost lattice, whose canonical join graph is isomorphic to C_7 .

L being doubled is not distributive. Because it is easy to check that L is congruence uniform, we leave the details to the reader. (Note that C_n , for $n \geq 5$ is among the minimal graphs excluded by Theorem 2.2.4, and so does not appear as the canonical join graph for a distributive lattice.)

2.5 Discussion and open problems

The discussion in Section 2.4 does not constitute a complete list of lattice theoretic operations which preserve (join)-semidistributivity. For example, the derived lattice $\mathbb{C}(L)$ discussed in [78], the box product defined in [45] (see also, [90, Corollary 8.2]), and the lattice of multichains from [58] all preserve (join)-semidistributivity.

Because it is relatively easy, we will discuss this last operation in a small example. Recall that an m -**multichain** in a poset \mathcal{P} is a collection of m elements satisfying $x_1 \leq x_2 \leq \dots \leq x_m$. We write an m -multichain as a tuple (x_1, \dots, x_m) or more compactly as a vector \vec{x} . We write the set of all m -multichains, partially ordered component-wise, as $\mathcal{P}^{[m]}$. When \mathcal{P} is a lattice, then $\mathcal{P}^{[m]}$ is a sublattice of the m -fold direct product of \mathcal{P}^m (see [58, Theorem 2.4]). It follows immediately that if L satisfies SD_\vee or SD_\wedge then $L^{[m]}$ also does, for each $m \in \mathbb{N}$ (see also, [58, Proposition 2.10]). In the proposition below, $(j)_k$ is the element $(\hat{0}, \dots, \hat{0}, j, \dots, j)$, where k is the left-most coordinate that is equal to j .

Proposition 2.5.1. *Suppose that L is a finite lattice. Then, $\text{Irr}(L^{[m]})$ is equal to the set $\{(j)_k : j \in \text{Irr}(L)\}$, where $k \in [m]$.*

Proof. We first show that $(j)_k$ is join-irreducible when $j \in \text{Irr}(L)$. Suppose that $\vec{w} \vee \vec{v} = (j)_k$. We have $w_i \vee v_i = j$, for each $i \geq k$. Since j is join-irreducible, we may assume that $w_k = j$. Since \vec{w} is a multichain, we have that $j \leq w_i$ for each $i \geq k$. Thus, $\vec{w} = (j)_k$, as desired.

Next, suppose that $\bar{w} \in \text{Irr}(L^{[m]})$. Let w_k be the first nonzero entry in \bar{w} , and assume that $w_k \notin \text{Irr}(L)$ so that there exist a and b in $L \setminus \{w_k\}$ with $w_k = a \vee b$. Then

$$\bar{w} = (\hat{0}, \dots, \hat{0}, a, w_{k+1}, \dots, w_m) \vee (\hat{0}, \dots, \hat{0}, b, w_{k+1}, \dots, w_m).$$

By this contradiction, we conclude that $w_k \in \text{Irr}(L)$. Next, suppose that $w_i \neq w_k$, for some $i > k$. Since $w_k < w_i$, there is an element $y \in \text{cov}_\downarrow(w_i)$ such that $w_k \leq y$. We have the following nontrivial join-representation of \bar{w} :

$$\bar{w} = (\hat{0}, \dots, \hat{0}, w_k, \dots, y, w_{i+1}, \dots, w_m) \vee (\hat{0}, \dots, \hat{0}, y'', w_{k+1}, \dots, w_i, \dots, w_m),$$

where $y'' \in \text{cov}_\downarrow(w_k)$. Therefore $w_i = w_k$, and the proposition follows. \square

Example 2.5.2. Let L be the weak order on the symmetric group S_3 , and consider $L^{[2]}$. The lattice L and $L^{[2]}$ are shown in Figure 2.16, and the corresponding canonical join complexes are shown in Figure 2.17. Observe that if $j \vee j'$ is a canonical join representation in L then both $(\hat{0}, j) \vee (\hat{0}, j')$ and $(j, j) \vee (j', j')$ are canonical join representations in $L^{[2]}$. This accounts for the edges $\{(\hat{0}, a), (\hat{0}, b)\}$ and $\{(a, a), (b, b)\}$ in the complex for $L^{[2]}$.

To see how we obtain the remaining edges in Figure 2.17, consider the canonical join representation of $(d, \hat{1})$. Observe that $\text{cov}_\downarrow((d, \hat{1})) = \{(d, d), (b, \hat{1})\}$. It is easily checked that (d, d) is the smallest element in $L^{[2]}$ whose join with $(b, \hat{1})$ is equal to $(d, \hat{1})$. Similarly, $(\hat{0}, a)$ is the smallest element whose join with (d, d) is equal to $(d, \hat{1})$. Therefore, the canonical join representation for $(d, \hat{1}) = (d, d) \vee (\hat{0}, a)$. The canonical join representations of the remaining elements in $L^{[2]}$ are computed similarly.

This example is emblematic of the general construction, as can be seen in the next proposition which describes the canonical join graph for $L^{[m]}$. We leave the details of proof to the reader.

Proposition 2.5.3. *Suppose that L is a finite semidistributive lattice with join-irreducible elements j and j' .*

1. *If $i < k$ then $\{(j)_i, (j')_k\}$ is a face in the canonical join complex of $L^{[m]}$ if and only if j' is a canonical joinand of $j \vee j'$ in L .*
2. *If $i = k$, then $\{(j)_i, (j')_k\}$ is a face in the canonical join complex of $L^{[m]}$ if and only if $\{j, j'\}$ is a face in the canonical join complex of L .*

Note that the operation on the canonical join complex corresponding to $L \mapsto L^{[m]}$ depends on the lattice L (not just the canonical join complex of L).




Figure 2.16: Left: The weak order for the symmetric group S_3 . Right: The lattice of 2-multichains.




Figure 2.17: Left: The canonical join complex of weak order for the symmetric group S_3 . Right: The canonical join complex of the lattice of 2-multichains.

Question 2.5.4. *What lattice theoretic operations (preserving join-semidistributivity) correspond to geometric operations on the canonical join complex that are independent of L ?*

Alternatively, it would be interesting to know which geometric operations (on the class of finite simplicial complexes) have a corresponding lattice theoretic analogue. We point out that conspicuously absent from the discussion in Section 2.4 is closure under taking induced subcomplexes (see Remark 2.4.5).

Question 2.5.5. *Let \mathcal{C} be the class of simplicial complexes that can be realized as the canonical join complex of some finite semidistributive lattice. Is \mathcal{C} closed under taking induced subcomplexes?*

Say that \mathcal{G}_n is the set of labeled graphs that can be realized the (labeled) canonical join graph for a congruence uniform lattice with n join-irreducible elements, and \mathcal{G} is the union $\bigcup_{n \in \mathbb{N}} \mathcal{G}_n$. Using Stembridge’s poset Maple package ([84]) and Proposition 2.4.9, we have counted the number of elements of \mathcal{G}_n for $n \leq 6$. While our computations indicate that not every labeled

graph appears, they also suggest that \mathcal{G} is closed under subgraphs (so that the corresponding class of simplicial complexes is closed under taking subcomplexes). We close this chapter by asking two related questions:

Question 2.5.6. *Which labeled graphs can be realized as the labeled canonical join graph for some congruence uniform lattice?*

Question 2.5.7. *Suppose that G is the canonical join graph for a congruence uniform lattice L . What data, in addition to G , is necessary in order to determine L up to isomorphism?*

Chapter 3

The Canonical Join Complex of the Tamari Lattice

3.1 Introduction

In this chapter, we study the canonical join complex of the Tamari lattice. Recall that, informally, the canonical join representation of an element w is the unique lowest irredundant expression $\vee A$ for w in terms of the join operation. (As in the previous chapter, we abuse the notation and say that A is a canonical join representation.) The *canonical join complex* is the abstract simplicial complex whose faces are the subsets A of L such that A is a canonical join representation. In general, the canonical join complex is not a pure complex. (In particular, the canonical join complex of the Tamari lattice is very different from the associahedron.) We recall the canonical join complex of the Tamari lattice T_3 in Figure 3.1.




Figure 3.1: The Tamari lattice T_3 and its canonical join complex.

The canonical join complex was first defined in [72] by Reading for the special case of the symmetric group S_n (ordered according to the weak order). Recall from Example 2.2.6 that canonical join representation of a permutation is encoded by a noncrossing arc diagram, a generalization of the bump diagram for a noncrossing partition. Each diagram consists of a

collection of curves, called arcs, that satisfy certain compatibility relations. For example, no two arcs may intersect in their interior. (See Section 3.2.2 for the complete definition.) Each arc corresponds to a vertex of the canonical join complex, and a collection of arcs corresponds to a face if and only if each pair of arcs is compatible. (This is [72, Corollary 3.5].) Figure 3.2 shows the noncrossing arc diagrams that correspond to the faces in the canonical join complex of the weak order on S_3 .




Figure 3.2: The faces in the canonical join complex of the weak order on S_3 .

Like the h -complex of the Coxeter complex defined in [31], the entries of the f -vector of the canonical join complex of the weak order on the symmetric group are equal to the Eulerian numbers. (However, in general, the canonical join complex of the symmetric group is not isomorphic, or even homotopy-equivalent, to the h -complex of the Coxeter complex.) Similar statements hold for the canonical join complex of the Tamari lattice and certain Tamari-like lattices called c -Cambrian lattices: The entries of the f -vector of the canonical join complex of the Tamari lattice (and each c -Cambrian lattice in type A) are equal to the Narayana numbers. (Recall, this is Proposition 2.2.9.)

For each finite Coxeter group W and each orientation c of its associated Coxeter diagram, recall that there is a lattice quotient of the weak order on W called the c -**Cambrian lattice**. The canonical join representation of its elements is closely related to the associated cluster algebra and to the noncrossing partition lattice $NC(W, c)$ [75]. In type A, each c -Cambrian lattice is a lattice quotient of the weak order on S_n , consisting of certain pattern avoiding permutations. In particular, when c is a linear orientation—an orientation in which all of the arrows point in the same direction—the corresponding c -Cambrian lattice is a Tamari lattice. For one choice of linear orientation, the elements of this quotient are the 312-avoiding permutations. Recall that, throughout, we write T_n for this realization of the Tamari lattice. (For the opposite orientation, the elements of the corresponding c -Cambrian lattice avoid the pattern 231.)

As with the classical Tamari lattice, the type-B Tamari lattice can be realized as a partial order on certain triangulations of a fixed convex polygon or certain bracket vectors. We realize the type-B Tamari lattice T_n^s as a c -Cambrian lattice for the type-B Coxeter group B_n (where c is a linear orientation for the type-B Coxeter diagram). See [68, Section 7] and [87].

Below we give our main results. In the following theorems and throughout the chapter, we

do not distinguish between an abstract simplicial complex and its geometric realization. In the statements below, $\text{Cat}(A_{r-1})$ is the classical Catalan number $\frac{1}{r+1}\binom{2r}{r}$, $\text{Cat}(B_r) = \binom{2r}{r}$ is the type-B Catalan number, and $\text{Cat}^+(B_r) = \binom{2r-1}{r-1}$ is the (type-B) positive Catalan number.

Theorem 3.1.1. *The canonical join complex of the Tamari lattice T_n is shellable. It is contractible when n is even and homotopy equivalent to a wedge of $\text{Cat}(A_{r-1})$ many spheres, all of dimension $r - 1$, when $n = 2r + 1$.*

Theorem 3.1.2. *The canonical join complex of the Tamari lattice T_n^s in type-B is shellable.*

1. *When $n = 2r$, the canonical join complex is homotopy equivalent to a wedge of $\text{Cat}(B_r)$ many spheres all of dimension $r - 1$.*
2. *When $n = 2r - 1$ for $r > 1$, the canonical join complex is homotopy equivalent to a wedge of $\text{Cat}^+(B_r) - \text{Cat}(A_{r-2}) = 2\binom{2r-2}{r-2}$ many spheres, equally distributed in dimensions $r - 1$ and $r - 2$.*

As an immediate consequence, we have that the alternating sum of the Narayana numbers is either zero or a signed Catalan number. For n even, the alternating sum of type-B Narayana numbers is a type-B Catalan number. These identities are well-known and also appear as specializations of Coker's identities. See [25], or [23, Equation 1.1] for the type-A case and [23, Equation 2.1] for type-B case.

The topology of each Tamari-like c -Cambrian lattice in type A is similarly nice.

Theorem 3.1.3. *For each orientation c of the type-A Coxeter diagram, the canonical join complex of the corresponding c -Cambrian lattice is vertex decomposable.*

Since vertex decomposability implies shellability, and the Tamari lattice is an example of a c -Cambrian lattice, Theorem 3.1.3 implies the shellability assertion in Theorem 3.1.1. We pull out the special case of the Tamari lattice for two reasons: First, constructing a shelling of its facets is easy; and second, our work in type A will motivate the proof of Theorem 3.1.2.

We conclude this introduction by noting that the particularly nice topological results for the Tamari lattice (and the Tamari-like c -Cambrian lattices in type A) do not extend to other finite Coxeter groups. For each orientation c , the c -Cambrian lattice in the type- D_5 Coxeter group is not shellable.

3.2 Background

3.2.1 Lattice-theoretic background

In this section, we briefly review the necessary lattice-theoretic terminology. Much of this material is repeated from Section 2.3.1. We recall it here for the convenience of the reader. Readers

who are familiar with the definitions are welcome to skim or skip this section.

Throughout, we assume that L is a finite lattice. A **join-representation** for an element $w \in L$ is an expression $\vee A$ that evaluates to w , where A is a subset of L . A join-representation $\vee A$ is a **irredundant** if, for each proper subset $A' \subset A$, we have that $\vee A' < \vee A$. Observe that if $\vee A$ is irredundant, then A is an antichain. We write $\text{ijr}(w)$ for the collection of irredundant join-representations of w . We partially order $\text{ijr}(w)$ as follows: $A \leq B$ whenever the order ideal generated by A is contained in the order ideal generated by B . (This relation is also sometimes called **join-refinement** [39, Section I.3].) The **canonical join representation** of w is the unique minimal element of $\text{ijr}(w)$, when such an element exists. The elements $j \in A$ are called the **canonical joinands** of w . At times we say that the set A is a canonical join representation, although more precisely we mean that the expression $\vee A$ is a canonical join representation.

Recall that when L is finite and each element admits a canonical join representation, we say that L is **join-semidistributive**. If the dual lattice is also join-semidistributive, then we say that L is **semidistributive**. Suppose that L is a finite join-semidistributive lattice. We define the **canonical join complex** of L to be the collection of subsets A such that A is a canonical join representation. The next proposition is [72, Proposition 2.2], and it implies that the canonical join complex is an abstract simplicial complex.

Proposition 3.2.1. *Suppose L is a finite lattice and A is a canonical join representation in L . Then each subset of A is a canonical join representation.*

Recall that j is join-irreducible if, whenever $j = \vee A$, we have $j \in A$. (Equivalently, j is join-irreducible if and only if it covers precisely one element in L .) Thus, if j is join-irreducible then $\{j\}$ is its canonical join representation. On the other hand, if A is a canonical join representation, then each element $a \in A$ is join-irreducible. So, the vertex set for the canonical join complex of L is its set of join-irreducible elements.

Recall from Section 2.4.1 that **lattice congruence** Θ is an equivalence relation on the elements of L that respects the meet and join operations. That is, if $w \equiv_{\Theta} u$, then $w \vee t \equiv_{\Theta} u \vee t$ and $w \wedge t \equiv_{\Theta} u \wedge t$ for each w, u , and t in L . Equivalently, when L is finite, a lattice congruence is an equivalence relation satisfying the following three properties: First, each Θ -class is an interval. Second, the map $\pi_{\downarrow}^{\Theta} : L \rightarrow L$ that sends w to the smallest element in its Θ -class is order preserving. Third, the map $\pi_{\uparrow}^{\Theta} : L \rightarrow L$ that sends w to the largest element in its Θ -class is order preserving. When Θ is a lattice congruence, the image $\pi_{\downarrow}^{\Theta}(L)$ is a lattice in its own right, and the map $\pi_{\downarrow}^{\Theta} : L \rightarrow \pi_{\downarrow}^{\Theta}(L)$ respects the meet and join operations in both L and $\pi_{\downarrow}^{\Theta}(L)$. In general, a map $\phi : L' \rightarrow L''$ between lattices L' and L'' that respects the meet and join operations is called a **lattice homomorphism**, and the image $\phi(L')$ is a **lattice quotient** of L' . In particular, $\pi_{\downarrow}^{\Theta}(L)$ is a lattice quotient of L . Similarly the image of π_{\uparrow}^{Θ} is a lattice quotient of L , and it is isomorphic to $\pi_{\downarrow}^{\Theta}(L)$.

3.2.2 The noncrossing arc complex

In this section, we review the definition of a noncrossing arc diagram, establish some useful notation, and review the connection to canonical join representations. The definitions here are based on [72], where the reader will find additional examples. For the remainder of the chapter, we write $[n]$ for the set $\{1, 2, \dots, n\}$, and $[i, k]$ for the set $\{i, i + 1, \dots, k\}$ when $i < k$.

A **noncrossing arc diagram** consists of n nodes arranged vertically and labeled in increasing order from bottom to top, together with a (possibly empty) collection of curves called **arcs**. Each arc connects two distinct nodes and travels monotonically upward from its lower endpoint to its higher endpoint, passing either to the left or to the right of each node in between. In addition, each pair of arcs α and α' must satisfy:

(C1) α and α' do not share the same top endpoint or the same bottom endpoint;

(C2) α and α' do not intersect in their interiors.

The **support of an arc** α , written $\text{supp}(\alpha)$, with endpoints $i < l$ is the set of numbers $\{i, i + 1, \dots, l\}$. We write $\text{supp}(\alpha)^\circ$ for the set $\{i + 1, \dots, l - 1\}$. When $\text{supp}(\alpha)^\circ$ is empty, we say that α is a **simple arc**. We say that the arcs α and α' are **combinatorially equivalent** if α and α' have the same endpoints and for each $k \in \text{supp}(\alpha)^\circ$, α and α' pass on the same side (either left or right) of k . Each arc is considered only up to combinatorial equivalence. Two arcs are **compatible** if there is a noncrossing arc diagram that contains them. The next proposition is [72, Proposition 3.2].

Proposition 3.2.2. *Given any collection of pairwise compatible arcs, there is a noncrossing arc diagram whose arcs are combinatorially equivalent to the given arcs.*

The **noncrossing arc complex** on n nodes is the simplicial complex whose faces are the collections of pairwise compatible arcs. We view each collection of compatible arcs as a noncrossing arc diagram. For example, Figure 2.2 depicts some of the nonempty faces in the noncrossing arc complex on seven nodes. To avoid confusion, we will only use the word *vertex* to refer to a vertex of the noncrossing arc complex; that is, a diagram with that contains precisely one arc. The endpoint of an arc will always be referred to as a *node*.

Next, we describe a bijection δ from the symmetric group S_n to the faces of the noncrossing arc complex on n nodes. We write each a permutation in one-line notation as $a_1 \dots a_n$. For each **descent** $a_i > a_{i+1}$ there is a corresponding arc with endpoints $a_{i+1} < a_i$. This arc passes to the right (respectively left) of each number a_l with $l < i$ (respectively, $l > i$). For example, the noncrossing arc diagram for the permutation 4123 has a single arc connecting 1 and 4. Because the numbers 2 and 3 are on the right side of the descent (in the one-line notation), the second and third nodes also lie on the right side of this arc. See Figure 3.3. We can visualize this map




Figure 3.3: From left to right: $\delta(4123)$, $\delta(2413)$, $\delta(2341)$, and $\delta(3412)$.

as follows: Plot the points $\{(i, a_i) : i \in [n]\}$, and connect (i, a_i) with $(i + 1, a_{i+1})$ by a straight line segment whenever $a_i > a_{i+1}$. Then, deform the picture until the points lie in a vertical line (so that the y -coordinates label the nodes in the diagram). The lines connecting descents curve to become arcs in the diagram. (See [72, Figure 4] for an example.) The next theorem is [72, Theorem 3.1].

Theorem 3.2.3. *The map δ is a bijection from the set of permutations in S_n to the set of noncrossing arcs diagrams on n nodes.*

Recall that, in the weak order on the symmetric group, permutations are ordered by containment of inversion sets. (An *inversion* for $w_1 \dots w_n$ is a pair (w_i, w_j) with $w_i > w_j$, where $1 \leq i < j \leq n$. The *inversion set* of $w_1 \dots w_n$ is the set of all such pairs.) Each descent $w_i > w_{i+1}$ in the permutation $w_1 \dots w_n$ corresponds to a cover $w_1 \dots w_i w_{i+1} \dots w_n \succ w_1 \dots w_{i+1} w_i \dots w_n$ in which the positions of w_i and w_{i+1} are swapped. (These entries are highlighted in red.) In particular, a permutation is join-irreducible if and only if it has precisely one descent. Thus, the map δ restricts to a bijection from the set of join-irreducible permutations to the set of noncrossing arc diagrams with exactly one arc. The weak order on S_n is semidistributive [26], so each permutation has a canonical join representation.

Proposition 2.2.2, in particular, says that for each $w \in S_n$, there is a bijection from the set $\{y : w \succ y\}$ to the canonical join representation of w . The canonical joinand j associated to the element $y \in \{y : w \succ y\}$ satisfies: j is the unique minimal element in S_n whose join with y is equal to w . (In the context of Coxeter groups, Proposition 2.2.2 is a generalization of [72, Proposition 3.2] and [75, Theorem 8.1].) We can think of this bijection $y \mapsto j$ as a map from the set of descents of w to its canonical join representation. (It follows immediately that the number of faces with k vertices in the canonical join complex of the weak order on S_n is the Eulerian number $\left\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \right\rangle$.) Interpreted as a map from descents to canonical joinands, this bijection is exactly the same correspondence between descents and arcs that is induced by δ . The arcs appearing in $\delta(w)$ correspond to the join-irreducible permutations in its canonical join representation.

Example 3.2.4. Consider $w = 2431$ and the cover relation $2431 \succ 2413$ (associated to the $(3, 1)$ descent highlighted in red). Consider the set $\{u \in S_n : u \vee 2413 = 2431\}$. Observe that each such permutation u must have $(3, 1)$ as an inversion. The minimal elements that contain this inversion are the join-irreducible permutations 2314 and 3124 , in which $(3, 1)$ is a descent.



Figure 3.4: Left: $\delta(2431)$. Right: The arc corresponding to the canonical joinand 2314.

Only the former is below w . We conclude that 2314 is a canonical joinand of w , and indeed its arc appears in the noncrossing arc diagram $\delta(w)$. (More generally, see [72, Proposition 2.3].)

Example 3.2.4 motivates the proof of [72, Theorem 3.4], which we now quote.

Theorem 3.2.5. *The restriction of δ to the set of join-irreducible permutations in S_n induces an isomorphism from the canonical join complex of the weak order on S_n to the noncrossing arc complex on n nodes.*

Recall that a complex is **flag** if each of its minimal non-faces has size 2. As an immediate consequence of Proposition 3.2.2 and Theorem 3.2.5, the canonical join complex of the weak order on S_n is flag. (See [72, Corollary 3.6]. This is a special case of our main result, Theorem 2.1.1, from the previous chapter.)

3.2.3 The c -Cambrian congruence and the Tamari lattice

Recall that we realize the Tamari lattice T_n as the subposet of the weak order on S_n consisting of the permutations that avoid the pattern 312. In this section, we review the connection to the c -Cambrian congruence. Then, we characterize the canonical join complex of the Tamari lattice in terms of noncrossing arc diagrams.

Recall that a permutation $w = w_1, \dots, w_n$ avoids the pattern 312 if it has no subsequence of entries $w_i > w_l > w_k$, with $1 \leq i < k < l \leq n$. (In other words, w_i plays the role of 3, w_k plays the role of 1, and w_l plays the role of 2.) It is well-known that w contains an instance of the 312-pattern if and only if it contains such a triple in which the “3” and the “1” are adjacent. (Equivalently, w contains a subsequence $w_i > w_l > w_k$ with $k = i + 1$.)

Let c denote an orientation of the type-A Coxeter diagram. Each choice of orientation gives rise to a partition of $[2, n - 1]$ into a set R of right nodes and its complementary set L of left nodes as follows: Label the nodes of the Coxeter diagram in decreasing order from left to right. If the edge between $i - 1$ and i has a left (respectively right) arrow then $i \in L$ (respectively $i \in R$), where $i \in [2, n - 1]$. We decorate each element of $[2, n - 1]$ with either an under-bar or an over-bar as follows: We write \bar{i} for each $i \in L$ and $\overline{\bar{i}}$ for each $i \in R$.

Suppose that $w \succ v$ in the weak order on S_n . Recall that we obtain v by swapping the entries of a descent $w_i > w_{i+1}$ in the one-line notation for w . Thus, there is a corresponding arc $\alpha \in \delta(w)$

with endpoints $w_{i+1} < w_i$. We say that v is obtained from w by a $3\underline{12} \rightarrow 13\underline{2}$ *move* if α passes to the left of some node in L . (We think of $\underline{2}$ as representing an element of L .) Similarly, v is obtained from w by a $\bar{2}31 \rightarrow \bar{2}13$ *move* if α passes to the right of some node in R .

Example 3.2.6. Consider $w = 2357164$ and $v = 2351764$ in S_7 , where $R = \{2, 4, 6\}$ and $L = \{3, 5\}$. We have $2357164 \succ 2351764$, where we swap the fourth and fifth entries. Since the corresponding arc (with endpoints at $1 < 7$) passes to the right of 2, we conclude that v is obtained from w by a $\bar{2}31 \rightarrow \bar{2}13$ move.

We define a relation on the covering pairs $w \succ v$ in S_n as follows: $v \equiv w$ if and only if v is obtained from w by either a $\bar{2}31 \rightarrow \bar{2}13$ or a $3\underline{12} \rightarrow 13\underline{2}$ move. The transitive closure of this relation is a lattice congruence Θ_c (on the weak order) on S_n called the *c-Cambrian* congruence. (The fact that Θ_c is indeed a lattice congruence is [68, Theorem 5.1]. See also Proposition 4.3.1.) We have the following proposition.

Proposition 3.2.7. *A join-irreducible permutation is contracted by Θ_c precisely when its corresponding arc passes to the left of some element in L or to the right of some element in R .*

Suppose that c is the linear orientation, in which every arrow points left. Observe that a join-irreducible permutation is *not* contracted by Θ_c precisely when its corresponding arc does not pass to the left of any node. Equivalently, a join-irreducible permutation is not contracted by Θ_c if and only if it avoids 312. Proposition 2.4.3 implies that $w \in \pi_{\downarrow}^{\Theta_c}(S_n)$ if and only if w also avoids 312. Thus, the Tamari lattice T_n is equal to $\pi_{\downarrow}^{\Theta_c}(S_n)$. The next proposition follows immediately from Proposition 2.4.3 and Theorem 3.2.5. In the statement of the proposition, a *right arc* is an arc that does not pass to the left of any node.

Proposition 3.2.8. *The canonical join complex $\Delta(n)$ of the Tamari lattice T_n is isomorphic to the subcomplex of the noncrossing arc complex on n nodes induced by the set of right arcs.*

We write $\alpha_{i,k}$ for the right arc with endpoints $i < k$. (Observe that there is precisely one right arc for each pair of nodes $i, k \in [n]$.) Throughout the remainder of the chapter, we write $\Delta(n)$ for the complex of compatible right arcs on n nodes. At times it is convenient to restrict the node set to a contiguous subset of $[n]$. We write $\Delta([i, k])$ for subcomplex of $\Delta(n)$ induced by restricting to the nodes $[i, k]$.

3.2.4 The type-B Tamari lattice

In this section, we define the type-B Tamari lattice and characterize its canonical join complex. Throughout, we write $[\pm n]$ for the set $\{-n, \dots, -1, 1, \dots, n\}$ and $S_{\pm n}$ for the symmetric group on $[\pm n]$. A *signed permutation* (in full one-line notation) is a permutation $w_{-n} \dots w_{-1} w_1 \dots w_n$ satisfying $w_{-i} = -w_i$. As usual, we write B_n for the type-B Coxeter group of rank n . Recall that

the weak order on B_n can be realized as the sublattice of $S_{\pm n}$ induced by the set of signed permutations. The symmetry condition that defines a signed permutation implies that descents may come in pairs: When $i > -1$, each descent $w_i > w_{i+1}$ has a symmetric partner $w_{-i-1} > w_{-i}$. (One moves down by a cover relation in the weak order on B_n either by swapping the positions of *both* of these symmetric descents, or by swapping the -1 and 1 positions when $w_{-1} > w_1$.) It follows immediately that the noncrossing arc diagram (with nodes labeled $-n, \dots, -1, 1, \dots, n$ from bottom to top) corresponding to a signed permutation is fixed by a half-turn rotation through the middle of its diagram. We call such diagrams ***symmetric noncrossing arc diagrams***, and a ***symmetric arc*** is either a pair of arcs that are related by the half-turn rotation or a single arc this is fixed by the rotation. See Figure 3.6 for some examples.

The next proposition is a special case of Proposition 4.3.20. In the statement, a ***symmetric orientation*** for the Coxeter diagram of $S_{\pm n}$ is an orientation that is fixed by a half-turn rotation through the middle node.

Proposition 3.2.9. *Suppose that c is a symmetric orientation for the Coxeter diagram for $S_{\pm n}$. Then the c -Cambrian congruence Θ_c on $S_{\pm n}$ restricts to a lattice congruence Θ'_c on B_n . Moreover, a signed permutation w is in $\pi_{\downarrow}^{\Theta'_c}(B_n)$ if and only if it belongs to $\pi_{\downarrow}^{\Theta_c}(S_{\pm n})$.*

In particular, let c be the symmetric orientation for $S_{\pm n}$ with $L = [n-1]$ and $R = [-n+1, -1]$. Then, the congruence Θ'_c (from Proposition 3.2.9) on the weak order on B_n is the c -Cambrian congruence for a linear orientation of the type-B Coxeter diagram. (See Proposition 4.3.15.) We define the type-B Tamari lattice T_n^s to be the lattice quotient $\pi_{\downarrow}^{\Theta'_c}(B_n)$. Our definition agrees with the definition given in [68, Section 7]. (See also [87].)

The next result follows from Proposition 3.2.7 and Proposition 3.2.9.

Proposition 3.2.10. *The canonical join complex of the type-B Tamari lattice T_n^s is isomorphic to the subcomplex of symmetric noncrossing arc diagrams on $[\pm n]$ induced by set of symmetric arcs which do not pass to the left of any positive node or to the right of any negative node.*

We write $\Delta^s(n)$ for the canonical join complex of the type-B Tamari lattice T_n^s . There is precisely one symmetric arc for each pair of nodes in $[\pm n]$. Given a pair of arcs $\alpha_{i,k}$ and $\alpha_{-k,-i}$ that together comprise a symmetric arc in $\Delta^s(n)$, we write $\alpha_{i,k}^s$ for the corresponding symmetric arc, where $k > i$ and $k > -i$. When the endpoints of a symmetric arc are not specified, we simply write α^s . To distinguish the arc $\alpha_{i,k}$ from the symmetric arc $\alpha_{i,k}^s$, we sometimes refer to the former as an ***ordinary arc***. A ***simple symmetric arc*** is either a pair of simple arcs fixed by a half-turn rotation through the center of the diagram, or the ordinary simple arc with endpoints -1 and 1 .

3.2.5 Noncrossing perfect matchings

In this section, we introduce and count certain matchings on graphs. Recall that the *Coxeter-Catalan number* corresponding to the Coxeter group W is the number

$$\text{Cat}(W) = \prod_{i=1}^n \frac{e_i + h + 1}{e_i + 1},$$

where $\{e_1, \dots, e_n\}$ are the exponents for W , and h is its Coxeter number (see [51, Section 3.20]). When W is rank 0, the formula for $\text{Cat}(W)$ is the empty product, which we interpret to be 1. In the identities given below, it is convenient for us to write B_0 and B_1 . We interpret B_0 as the rank 0 Coxeter group, and we interpret B_1 as the rank 1 Coxeter group. So, we have $\text{Cat}(A_0) = \text{Cat}(B_0) = 1$, and also $\text{Cat}(B_1) = \text{Cat}(A_1) = 2$. The *positive Catalan number* is the number

$$\text{Cat}^+(W) = \prod_{i=1}^n \frac{e_i + h - 1}{e_i + 1}.$$

The next lemma follows immediately from the formulas for $\text{Cat}(B_n)$ and $\text{Cat}^+(B_n)$. (Recall that, the Coxeter number for B_n is equal to $2n$, and its exponents are $1, 3, \dots, 2n - 1$.)

Lemma 3.2.11. $\text{Cat}(B_n) = 2 \text{Cat}^+(B_n)$, for $n > 0$.

The next lemma is essentially [35, Proposition 3.7], specialized to $W = B_{n-1}$ and $k = n - 1$.

Lemma 3.2.12. $\text{Cat}^+(B_n) = \sum_{i=0}^{n-2} \text{Cat}(B_i) \text{Cat}(A_{n-i-2}) + \text{Cat}(B_{n-1})$

A *perfect matching* on the set $[2n]$ is a partition of $[2n]$ into blocks of size two. We will represent a perfect matching in the following way: Draw $2n$ nodes vertically, and label them in increasing order $1, 2, \dots, 2n$ from bottom to top. We draw the edges in a matching so that they pass to the right of the nodes between their endpoints. A perfect matching is *noncrossing* if no two edges cross in their interiors. It is well-known that there are $\text{Cat}(A_{n-1})$ many noncrossing perfect matchings on $\{1, 2, \dots, 2n\}$. (See for example [82, Exercise 6.19].)

Consider the set of noncrossing perfect matchings on $[\pm n]$. We label the nodes in increasing order $-n, \dots, -1, 1, \dots, n$. Such a matching M is *symmetric* if it is fixed by negation. That is, the set $\{a, b\}$ is an edge in M if and only if $\{-a, -b\}$ is also an edge. We say that the edge $\{-a, a\}$ is *fixed by negation*. Write O_r for the number of symmetric noncrossing perfect matchings on the set $[\pm(2r - 1)]$, and E_r for the number of symmetric noncrossing perfect matchings on the set $[\pm(2r)]$.

Proposition 3.2.13. *Let O_r and E_r be defined as above.*

1. $O_r = \text{Cat}^+(B_r)$, when $r \geq 1$

2. $E_r = \text{Cat}(B_r)$, when $r \geq 0$

Proof. We prove the proposition by induction. The base cases are trivial. We assume that for each $i < r$, both statements hold. We count symmetric noncrossing perfect matchings on the set $[\pm(2r-1)]$ in which j is the largest positive number such that $\{-j, j\}$ is an edge. Because the restriction to the (possibly empty) set $\{j+1, \dots, 2r-1\}$ is a noncrossing perfect matching, j is odd.

When $j = 2r-1$, we remove this edge to obtain a symmetric noncrossing perfect matching on $[\pm(2r-2)]$. By induction, there are $\text{Cat}(B_{r-1})$ many such matchings. In general, for $j = 2i+1$, we map our matching to a pair (A, B) where A is a noncrossing perfect matching on the nodes $\{j+1, \dots, 2r-1\}$, and B is a symmetric noncrossing perfect matching on the set $[\pm(j-1)]$. By induction there are $\text{Cat}(B_i) \text{Cat}(A_{r-i-2})$ many matchings. Summing over $i \geq 0$ (and adding the term where $j = 2r-1$), we obtain $O_r = \sum_{i=0}^{r-2} \text{Cat}(B_i) \text{Cat}(A_{r-i-2}) + \text{Cat}(B_{r-1})$. By Lemma 3.2.12, we conclude that $O_r = \text{Cat}^+(B_r)$.

For each symmetric noncrossing perfect matching on $[\pm 2r]$, write j for the number that is paired with $2r$. (So, j ranges over the set $\{1, 3, \dots, 2r-1\} \cup \{-2r\}$.) Our argument above implies that there are $\text{Cat}^+(B_n)$ many matchings with $j = -2r$. Take $j = 2i+1$, with $i \in [0, r-2]$. Each matching maps bijectively to a pair (A, B) , where A is a noncrossing perfect matching on the set $\{j+1, \dots, 2r-1\}$, and B is a symmetric noncrossing perfect matching on the set $[\pm(j-1)]$. By induction, we have $\text{Cat}(B_i) \text{Cat}(A_{r-i-2})$ many diagrams. Summing over all i , we obtain $\sum_{i=0}^{r-2} \text{Cat}(B_i) \text{Cat}(A_{r-i-2}) + \text{Cat}(B_{r-1})$. Finally, we add the term $\text{Cat}^+(B_n)$, where $j = -2r$. Together Lemma 3.2.12 and Lemma 3.2.11 imply $E_r = \text{Cat}(B_r)$. \square

We write \mathcal{M}_n for the set of symmetric noncrossing perfect matchings on $[\pm n]$ that satisfy either of the two conditions below:

- M contains an edge $\{k, -k\}$ where $k > 1$.
- There exists no edge in M that is fixed by negation.

Equivalently, $M \notin \mathcal{M}_n$ if and only if $\{-1, 1\} \in M$, and this is the *only* edge in M that is fixed by negation. When n is even, every noncrossing perfect matching belongs to \mathcal{M}_n . In the next proposition we count \mathcal{M}_n for n odd.

Proposition 3.2.14. *Suppose that $n = 2r-1$. The number of elements in \mathcal{M}_n is equal to $\text{Cat}^+(B_r) - \text{Cat}(A_{r-2}) = 2 \binom{2r-2}{r-2}$.*

Proof. Let j be the largest positive number satisfying $\{-j, j\} \in M$. Observe that, when n is odd, $M \notin \mathcal{M}_n$ precisely when the number $j = 1$. Each matching with $j = 1$ maps bijectively to a noncrossing perfect matching on $[2, n]$. The statement follows from Proposition 3.2.13. \square

3.3 Shellability of the Tamari lattices

3.3.1 The Tamari lattice in type A

Before we proceed with the proof of Theorem 3.1.1 we recall some terminology. A *d-complex* is a simplicial complex in which the maximal dimension of the faces is equal to d . A d -complex is *pure* if each of its facets has dimension d . It is not difficult to verify that $\Delta(n)$ is not pure.

A (not necessarily pure) complex is *shellable* if its facets can be arranged in a linear order F_1, \dots, F_m so that the subcomplex $(\bigcup_{i=1}^{k-1} \overline{F_i}) \cap \overline{F_k}$ is a pure simplicial complex of dimension $\dim(F_k) - 1$ for all $k \in [2, m]$. (We write $\overline{F_k}$ for the collection of faces in F_k .) Such a linear order is called a *shelling*. A facet F is a *homology facet* if $(\bigcup_{i=1}^{k-1} \overline{F_i}) \cap \overline{F_k}$ is equal to the entire boundary of F_k . The following theorem is a combination of [13, Theorem 3.4 and Theorem 4.1].

Theorem 3.3.1. *Suppose that Δ is a shellable complex. Then Δ is homotopy equivalent to a wedge of spheres where each r -dimensional sphere corresponds to an r -dimensional homology facet.*

Suppose that $\mathcal{L} = F_1, F_2, \dots, F_m$ is a shelling of the facets for a non-pure simplicial complex. The *rearrangement lemma* [13, Lemma 2.6], says that \mathcal{L} can be rearranged so that it satisfies the following condition. (We write *(DD)* for “decreasing dimension”.)

$$\text{For facets } F \text{ and } F', \text{ if } |F| > |F'| \text{ then } F \text{ precedes } F' \text{ in } \mathcal{L}. \quad (DD)$$

We will see that this condition is sufficient for shelling the facets of $\Delta(n)$.

Fix some non-simple right arc $\alpha_{i,k} \in \Delta(n)$. Suppose that α' is a right arc that is compatible with $\alpha_{i,k}$. Note that α' does not have i as its bottom endpoint, nor $i+1$ as its top endpoint (otherwise the two arcs share bottom endpoints or they cross). Also, since α' is a right arc, it does not pass between i and $i+1$. Thus, $\{\alpha', \alpha_{i,i+1}\}$ is a face in $\Delta(n)$. Similarly, $\{\alpha', \alpha_{k-1,k}\} \in \Delta(n)$. Since $\Delta(n)$ is a flag complex, we immediately obtain the following lemma.

Lemma 3.3.2. *Suppose that $\alpha_{i,k}$ is a right arc in $\Delta(n)$ with $1 \leq i < k - 1 \leq n - 1$. Then, for each face $F \cup \{\alpha_{i,k}\}$ in $\Delta(n)$, the set $F \cup \{\alpha_{i,i+1}, \alpha_{k-1,k}\}$ is in $\Delta(n)$.*

For each arc α in $\Delta(n)$ write $S(\alpha)$ for the set of simple arcs that are compatible with it. In the next lemma we show that the degree of a face J is determined by the set $\bigcap_{\alpha \in J} S(\alpha)$. Recall that the *degree* of F , denoted $\deg(F)$, is the maximal size of the faces containing F .

Lemma 3.3.3. *Suppose that J is a face in $\Delta(n)$, and write $S' = \bigcap_{\alpha \in J} S(\alpha)$. Then, $S' \cup J$ is a facet of $\Delta(n)$, and every other face F that contains J has size strictly smaller than $|J \cup S'|$. In particular, $\deg(J) = |J \cup S'|$.*

Proof. Observe that S' is the unique maximal set of simple arcs that are compatible with each arc in J . Since any two simple arcs are compatible, $S' \cup J$ is in $\Delta(n)$. Suppose that $\alpha_{i,k}$ is a non-simple right arc satisfying: the set $J \cup S' \cup \{\alpha_{i,k}\}$ is in $\Delta(n)$. (In particular, $\alpha_{i,k}$ is compatible with each arc in S' .) Then Lemma 3.3.2 implies that $J \cup S' \cup \{\alpha_{i,i+1}, \alpha_{k-1,k}\}$ is also in $\Delta(n)$. The maximality of S' implies that $\{\alpha_{i,i+1}, \alpha_{k-1,k}\} \in S'$. Since $\alpha_{i,k}$ is *not* compatible with either $\alpha_{i,i+1}$ or $\alpha_{k-1,k}$, we have reached a contradiction.

Suppose that F is a face in $\Delta(n)$ containing J , and $F \not\subseteq J \cup S'$. Thus, F contains some non-simple arc that does not belong to J . Applying Lemma 3.3.2, we replace each such non-simple arc (not in J) with a pair of simple arcs and obtain a chain of faces that is strictly increasing in size. This chain terminates in a face of the form $J \cup S''$, where S'' is a collection of simple arcs. Thus $S'' \subseteq S'$, and we conclude that $|F| < |J \cup S'|$. \square

Finally, we prove a more detailed version of Theorem 3.1.1.

Theorem 3.3.4. *Let $\mathcal{L} = F_1, \dots, F_m$ be a linear ordering of the facets of $\Delta(n)$ satisfying (DD). Then \mathcal{L} is a shelling for $\Delta(n)$, and F_k is a homology facet if and only if it contains no simple arcs. Moreover,*

- *when $n = 2r$, each facet contains a simple arc;*
- *and when $n = 2r + 1$, each homology facet has precisely r arcs and maps bijectively to a noncrossing perfect matching on $[2r]$.*

Proof of Theorem 3.3.4 and Theorem 3.1.1. Let F_1, \dots, F_m be a linear ordering of the facets of $\Delta(n)$ satisfying (DD), and consider the complex $\overline{F_k} \cap \left(\bigcup_{i=1}^{k-1} \overline{F_i} \right)$, where k ranges over the set $[2, m]$. We write J for the set of non-simple arcs in F_k and S' for the set of simple arcs in F_k . Lemma 3.3.3 implies that every other facet containing J occurs after F_k in this linear ordering. So, each face of $\overline{F_k} \cap \left(\bigcup_{i=1}^{k-1} \overline{F_i} \right)$ is contained in $(J \cup S') \setminus \{\alpha\}$, for some α belonging to J . Lemma 3.3.2 says that we can swap out α for a pair of simple arcs, and obtain a face with strictly larger size. We conclude that $(J \cup S') \setminus \{\alpha\}$ is a facet of $\overline{F_k} \cap \left(\bigcup_{i=1}^{k-1} \overline{F_i} \right)$ for each $\alpha \in J$. We have proved that F_1, \dots, F_m is a shelling of $\Delta(n)$, and F_k is a homology facet if and only if it contains no simple arcs. We write $\mathcal{H}(n)$ for the set of noncrossing arc diagrams that are facets in $\Delta(n)$ and that do not contain any simple arcs. In general, we write $\mathcal{H}([i, k])$ for the set of noncrossing arc diagrams that are facets in $\Delta([i, k])$ and that do not contain any simple arcs.

Suppose that $n = 2r$, and F is a facet of $\Delta(n)$. We prove by induction on r that F contains a simple arc. Since F is a facet, there is some arc that has 1 as its bottom endpoint and $l \leq n$ as its top endpoint. If l is equal to 2, then we are done; assume that l is greater than 2. We remove this arc and both of its endpoints. If some other arc α' in F had l as its bottom endpoint,

then we shift α' down so that it now has a bottom endpoint at the node $l-1$. (No other arc in F has $l-1$ as a bottom endpoint. Otherwise it would either cross the arc $\alpha_{1,l}$ or share a top endpoint with it.) We obtain a facet of $\Delta(n-2)$. Since this procedure preserves the size of the support of each arc in $F \setminus \{\alpha_{1,l}\}$, we are done by induction.




Figure 3.5: A demonstration of the map μ .

When $n = 2r + 1$, we define a map μ from $\mathcal{H}(n)$ to the set of noncrossing perfect matchings on the set $[n-1]$ as follows: Suppose that $F \in \mathcal{H}(n)$. Each pair of arcs in F that share an endpoint are pulled apart, and isolated nodes are deleted. See Figure 3.5. It is not difficult to check that μ is a well-defined bijection by induction. We describe the argument that μ is well-defined: As above, each facet contains an arc $\alpha_{1,l}$. This arc encloses a noncrossing arc diagram $F' \in \mathcal{H}([2, l-1])$. The remaining arcs form a noncrossing arc diagram $F'' \in \mathcal{H}([l, n])$. By induction $\mu(F')$ is a noncrossing perfect matching M' on $[2, l-2]$, and $\mu(F'')$ is a noncrossing perfect matching M'' on $[l, n-1]$. Thus, we have $\mu(F) = \{1, l-1\} \cup M' \cup M''$ is a noncrossing perfect matching on $[n-1]$.

Observe that each noncrossing perfect matching on $[2r]$ has r edges. Under the map μ , arcs become edges. The last item of Theorem 3.3.4 follows. \square

3.3.2 The Tamari lattice in type B

We now turn to the symmetric noncrossing arc complex $\Delta^s(n)$. We will break the proof of Theorem 3.1.2 into several steps: We begin by showing that $\Delta^s(n)$ has a shelling \mathcal{L} (Theorem 3.3.8). Then, we count the homology facets for $\Delta^s(n)$ (Proposition 3.3.10), and finally we determine their dimensions (Proposition 3.3.11 and Proposition 3.3.12). We will reuse much of our work from the previous section by choosing \mathcal{L} so that the homology facets are exactly those which do not contain any simple symmetric arcs. (Recall that a *simple symmetric arc* is either a pair of simple arcs fixed by a half-turn rotation through the center of the diagram, or

the ordinary simple arc with endpoints -1 and 1 .)

In the previous section, we had a key observation: Any non-simple arc in a face of $\Delta(n)$ could be swapped out for two simple arcs (this is Lemma 3.3.2). As a consequence, for any face F in $\Delta(n)$, the unique largest face containing F is constructed by adding as many simple arcs as possible. Both of these statements fail for $\Delta^s(n)$.




Figure 3.6: Each diagram contains two symmetric arcs.

Example 3.3.5. Consider the face in $\Delta^s(3)$ with arcs $\{\alpha_{-1,3}^s, \alpha_{-2,2}^s\}$, shown leftmost in Figure 3.6. We can swap out $\alpha_{-1,3}^s$ for one simple symmetric arc, namely $\alpha_{2,3}^s$ to obtain the face $\{\alpha_{2,3}^s, \alpha_{-2,2}^s\}$. But $\alpha_{-1,1}^s$ is not compatible with $\alpha_{-2,2}^s$. The reader can verify that the two faces $\{\alpha_{-1,3}^s, \alpha_{-2,2}^s\}$ and $\{\alpha_{2,3}^s, \alpha_{-2,2}^s\}$ are both maximal faces (containing the arc $\alpha_{-2,2}^s$). Similarly, we can swap out $\alpha_{-2,2}^s$ for only one simple symmetric arc. (The faces shown in the leftmost and rightmost diagrams in Figure 3.6 are the maximal faces containing $\alpha_{-1,3}^s$.)

For $\alpha_{i,k}^s$ in $\Delta^s(n)$, (the centrally symmetric analogue to) Lemma 3.3.2 may fail whenever $i = -1$ or $i = -k$; and in this case, there may be several maximal faces that contain $\alpha_{i,k}^s$. So, we require that our linear ordering \mathcal{L} satisfies one other condition, in addition to (DD) . First, we write down the centrally-symmetric analogues to Lemma 3.3.2 and Lemma 3.3.3.

Lemma 3.3.6. *Suppose that $\alpha_{i,k}^s$ is in $\Delta^s(n)$, and $\alpha_{i,k}^s$ is not a simple symmetric arc. Let $F \cup \{\alpha_{i,k}^s\}$ be face of $\Delta^s(n)$.*

1. *If $0 < i < k$, then the set $F \cup \{\alpha_{i,i+1}^s, \alpha_{k-1,k}^s\}$ is in $\Delta^s(n)$.*
2. *If $-k < i < -1$, then the set $F \cup \{\alpha_{-i-1,-i}^s, \alpha_{k-1,k}^s\}$ is in $\Delta^s(n)$.*
3. *If $i = -1$ or $i = -k$, then the set $F \cup \{\alpha_{k-1,k}^s\}$ is in $\Delta^s(n)$.*

Proof. The first item follows immediately from Lemma 3.3.2. For the second and third items, the same argument as given in the paragraph preceding Lemma 3.3.2 works here, with one main difference: While the arcs in $\Delta(n)$ do not pass between any nodes (because they are right arcs), the arcs in $\Delta^s(n)$ may have an “inflection”. If $i < 0$, then the ordinary arcs that comprise $\alpha_{i,k}^s$ pass between -1 and 1 . These are the only nodes that the arcs in $\Delta^s(n)$ pass between. In particular, for both of the cases above, there is no arc in F that passes between $k-1$ and k . \square

We write $S(\alpha^s)$ for the set of simple symmetric arcs that are compatible with a symmetric arc α^s in $\Delta^s(n)$.

Lemma 3.3.7. *Suppose that J is a face in $\Delta^s(n)$, and write $S' = \bigcap_{\alpha^s \in J} S(\alpha^s)$. Then, $S' \cup J$ is a facet of $\Delta^s(n)$ and $\deg(J) = |J \cup S'|$.*

Proof. The proof here is essentially the same as the proof of Lemma 3.3.3. The main difference is that $J \cup S'$ may not be the *unique* maximal face containing J (as we saw above in Figure 3.6). Suppose that F is a face in $\Delta^s(n)$ that contains J , and $F \not\subseteq J \cup S'$. Observe that S' is the unique maximal collection of simple symmetric arcs that are compatible with each arc in J . Thus, F contains some non-simple symmetric arc that is not in J . We use Lemma 3.3.6 to swap out these non-simple symmetric arcs for simple symmetric arcs. We obtain a chain of faces that is *weakly* increasing in size. This chain still terminates in a face of the form $J \cup S''$, where S'' is a collection of simple symmetric arcs, and we have $S'' \subseteq S$ as before. The statement follows. \square

Theorem 3.3.8. *Let $\mathcal{L} = F_1, \dots, F_m$ be a linear ordering of the facets of $\Delta^s(n)$ satisfying (DD) and the following condition: If F_i and F_k are facets with the same size and if the number of simple symmetric arcs in F_i is greater than the number of simple symmetric arcs in F_k , then $i < k$. Then \mathcal{L} is a shelling of $\Delta^s(n)$, and F_i is a homology facet if and only if it does not contain any simple symmetric arcs.*

Proof. For our proof here, we use the same argument that appears in the first paragraph of the proof for Theorem 3.3.4. Consider the complex $\overline{F}_k \cap \left(\bigcup_{i=1}^{k-1} \overline{F}_i \right)$, where $k \in [2, m]$. We write J for the set of non-simple symmetric arcs in F_k and S' for the set of simple symmetric arcs in F_k . Lemma 3.3.7 implies that every other facet in $\Delta^s(n)$ containing J occurs after F_k . So, each face of $\overline{F}_k \cap \left(\bigcup_{i=1}^{k-1} \overline{F}_i \right)$ is contained in $(J \cup S') \setminus \{\alpha^s\}$, for some element α^s belonging to J . Lemma 3.3.6 says that we can swap out α^s in $J \cup S'$ for at least one simple symmetric arc. The resulting face either has strictly larger size or it has strictly more simple symmetric arcs. We conclude that $(J \cup S') \setminus \{\alpha^s\}$ is a facet of $\overline{F}_k \cap \left(\bigcup_{i=1}^{k-1} \overline{F}_i \right)$ for each $\alpha^s \in J$. \square

Let $\mathcal{H}^s(n)$ denote the set of noncrossing arc diagrams that are facets in $\Delta^s(n)$ and that contain no simple symmetric arcs. Next, we define a map μ_s from $\mathcal{H}^s(n)$ to the set of symmetric noncrossing perfect matchings on $[\pm n]$.

Suppose that $F \in \mathcal{H}^s(n)$. We would like to use the map μ (defined in the third paragraph of the proof for Theorem 3.3.4) whenever possible. To that end, we write $P(F)$ for the set of arcs $\alpha_{i,k}^s \in F$ with $0 < i < k$, and $N(F)$ for the set of arcs $\alpha_{i,k}^s \in F$ with $i < 0 < k$. We will see that the set $P(F)$ decomposes into a collection of smaller noncrossing arc diagrams, each of which is either a maximal collection of non-simple ordinary right arcs or, symmetrically, a maximal collection of non-simple ordinary “left arcs”. We will apply the map μ to each collection of right

arcs and, by symmetry, to each collection of left arcs. That leaves us with one main challenge: how to pair off the endpoints of the arcs in $N(F)$. Before we can describe μ_s more precisely, we need the following easy lemma.




Figure 3.7: An illustration for the proof of Lemma 3.3.9.

Lemma 3.3.9. *Suppose that α_{i_1, k_1}^s and α_{i_2, k_2}^s are symmetric arcs in $\Delta^s(n)$ with $i_j < 0 < k_j$, for $j \in \{1, 2\}$, and $k_1 < k_2$. Then, α_{i_1, k_1}^s and α_{i_2, k_2}^s are compatible if and only if $i_1 < i_2$.*

Proof. Without loss of generality, assume that $k_2 = n$. Observe that if $i_2 = -n$, then it is not compatible with any arc $\alpha_{i, k}$ where $i < 0 < k$. So, we assume that $-n < i_2$. In particular, α_{i_2, k_2}^s consists of two arcs, related by a half-turn rotation of the diagram. We write α for the arc with endpoints $i_2 < k_2$ and $-\alpha$ for its symmetric partner. Observe that the arc $-\alpha$ encloses the nodes $1, \dots, -i_2 - 1$, so that any arc that has one endpoint in this set also has its second endpoint in this set. (See Figure 3.7.) Thus $-i_1 > -i_2$.

For the converse, recall that 1 and -1 are the only nodes that α and $-\alpha$ pass between. Thus, any arc that begins at a node between $-i_2$ and k_2 and that ends at a negative node between i_2 and $-k_2$ (passing to the right of positive nodes and to the left of negative nodes) is compatible with α and with $-\alpha$. \square

We write $N(F) = \{\alpha_{i_1, k_1}^s, \alpha_{i_2, k_2}^s, \dots, \alpha_{i_l, k_l}^s\}$ where $k_1 < \dots < k_l$. The upshot of Lemma 3.3.9 is that $-i_l < \dots < -i_1 < k_1$. We make three similar observations. First, consider the restriction of F to the nodes $[k_l, n]$. This is a maximal collection of ordinary right arcs, none of which are simple. Thus, this restriction is an element of $\mathcal{H}([k_l, n])$. Second, when $-i_l > 1$, the restriction to $[(-i_l - 1)]$ is also maximal collection of ordinary right arcs. This diagram is an element of $\mathcal{H}(-i_l - 1)$. Third, the nodes between any consecutive pair in the set $\{-i_l < \dots < -i_1 < k_1 < \dots < k_l\}$ are filled in with a maximal collection of right arcs. (See Figure 3.8.) So, for example, the restriction of F to the set $[k_j, k_{j+1} - 1]$ belongs to $\mathcal{H}([k_j, k_{j+1} - 1])$, for each $j \in [1, l - 1]$. The analogous statement holds for the restriction to $[-i_j, -i_{j-1} - 1]$, where $j \in [2, l]$, and for the restriction to $[-i_1, k_1 - 1]$. We write $F_1, F_{-i_1}, \dots, F_{k_l}$ for the collection of these smaller noncrossing arc diagrams, where we index each diagram by its bottom node.




Figure 3.8: The nodes $[k_1, k_2 - 1]$ are filled in with a maximal collection of ordinary right arcs.

We are now prepared to describe the map μ_s . We apply the map μ to each of the diagrams $F_1, F_{-i_l}, \dots, F_{k_l}$. Recall that μ sends a maximal collection of non-simple right arcs on the nodes $[a, b]$ to a matching on the set $[a, b - 1]$. It is convenient for us to shift the indices of the image of μ up by one instead. For example, when $-i_l > 1$, we map the noncrossing arc diagram F_1 on $[1, -i_l - 1]$ to a noncrossing perfect matching on $[2, -i_l - 1]$, and we send F_{-i_j} (a noncrossing arc diagram on $[-i_j, -i_{j-1} - 1]$) to a matching on $[-i_j + 1, -i_{j-1} - 1]$. We carry out the symmetric process on the “left arcs” in the negative portion of the diagram.

Thus, we have paired off every number in $[\pm n]$, except $1 \leq -i_l < \dots < k_l$ (and the corresponding negative numbers). We complete the matching as follows: We add the edges $\{1, -i_l\}$ and $\{-1, i_l\}$, unless $-i_l = 1$. If $-i_l = 1$, then we add the edge $\{-1, 1\}$. Finally, for each number $a \in \{-i_{l-1} < \dots < k_{l-1}\}$ we add the edge $\{-a, a\}$. If $k_l \neq -i_l$, then we also add $\{-k_l, k_l\}$. (Note that if $k_l = -i_l$, then $l = 1$.) We can visualize the entire construction in three steps as follows:

First step. Cut every arc in $N(F)$ where it passes between -1 and 1 . We call the resulting curves, each of which have precisely one endpoint, **arc segments**. We write α_a for the arc segment whose endpoint is a . Reflect the negative half of the diagram about the vertical column of the nodes, so that each arc and arc segment passes to the right of each node.

Second step. Apply the map μ to each of the diagrams $F_1, F_{-i_l}, \dots, F_{k_l}$, as described above. The effect of shifting our indices up is to “pull apart” two adjacent arcs (as in Figure 3.5). For each $a \in \{-i_l, \dots, k_l\}$, the bottom arc in F_a is pulled apart from the arc segment α_a that is below it. The one exception occurs when we shift up the indices of $\mu(F_1)$: In this case, the shift makes 1 an isolated node. We carry out the symmetric process on the negative half of the diagram.

Third step. Anchor the arc segment α_{-i_l} to 1 and symmetrically anchor α_{i_l} to -1 , unless $i_l = -1$. If $i_l = -1$, then we glue the segments α_{-i_l} and α_{i_l} together between -1 and 1 . We glue each

remaining arc segment α_a to the corresponding negative segment α_{-a} . See Figure 3.9 and Figure 3.10. (In the pictures, we focus on the cutting and gluing. We leave out all but the bottom node from each of the diagrams $F_1, F_{-i_l}, \dots, F_{k_l}$.)




Figure 3.9: An illustration of the map μ_s when $i_l \neq -1$.




Figure 3.10: An illustration of the map μ_s when $i_l = -1$.

In the next proposition, we prove that μ_s is a bijection onto \mathcal{M}_n . Recall from Section 3.2.5 that $M \notin \mathcal{M}_n$ if and only if $\{-1, 1\}$ is the *unique* edge in M that is fixed by negation. (We say an edge $\{a, b\}$ is fixed by negation if $b = -a$.)

Proposition 3.3.10. *Let μ_s be the map from $\mathcal{H}^s(n)$ to the set of symmetric noncrossing perfect matchings on $[\pm n]$ defined in the preceding paragraphs. Then, μ_s is a bijection onto \mathcal{M}_n .*

Proof. For each diagram $F \in \mathcal{H}^s(n)$, the set $N(F)$ is nonempty. In particular, there is an arc $\alpha_{i,k}^s \in F$ with $i < 0 < k$, where $1 < k$. (Otherwise, $F \cup \{\alpha_{-1,1}^s\}$ is a face.) We have two possibilities: First, it is possible that $i = -k$, and $\alpha_{i,k}^s$ is the unique element in $N(F)$. In that case, μ_s sends F

to a matching M that does not contain any edges that are fixed by negation. Otherwise, μ_s sends F to a matching that contains the edge $\{k, -k\}$. Thus the image of μ_s is contained in \mathcal{M}_n .

We define an inverse map $\nu : \mathcal{M}_n \rightarrow \mathcal{H}_n$. Suppose that $M \in \mathcal{M}_n$, and collect the set of positive numbers $\{m_1 < \dots < m_t\}$ that satisfy $\{-m_i, m_i\} \in M$ and $1 < m_i$. We write m_0 for the number that is paired with 1. Either $m_0 = -1$ or $1 < m_0$. Also, $m_0 < m_1$, otherwise M has a crossing.

We make three similar observations. First, the restriction of M to $[m_t + 1, n]$ is a noncrossing perfect matching. Second, the restriction to the set $[m_i + 1, m_{i+1} - 1]$ is a noncrossing perfect matching, for each $i \in [1, t - 1]$. Third, if $m_0 \neq -1$, then the restriction to the set $[2, m_0 - 1]$ is also noncrossing perfect matching. We begin the construction of $\nu(M)$ by applying μ^{-1} to each of these matchings. Each time we apply μ^{-1} , we shift down the indices of the resulting diagram by one. So, for example, we map the matching on $[m_t + 1, n]$ to a maximal collection of right arcs on the nodes $[m_t, n]$. Then, we insert the symmetric left arcs. We complete the construction by adding the arcs $\alpha_{-m_i, m_{t-i}}^s$ where $0 \leq i \leq \lfloor \frac{t}{2} \rfloor$. If the set $\{m_1, \dots, m_t\}$ is empty, then we add the arc α_{-m_0, m_0}^s . (Recall that, for $M \in \mathcal{M}_n$, if $\{m_1, \dots, m_t\}$ is empty then $m_0 \neq -1$.) We visualize the construction in three steps.

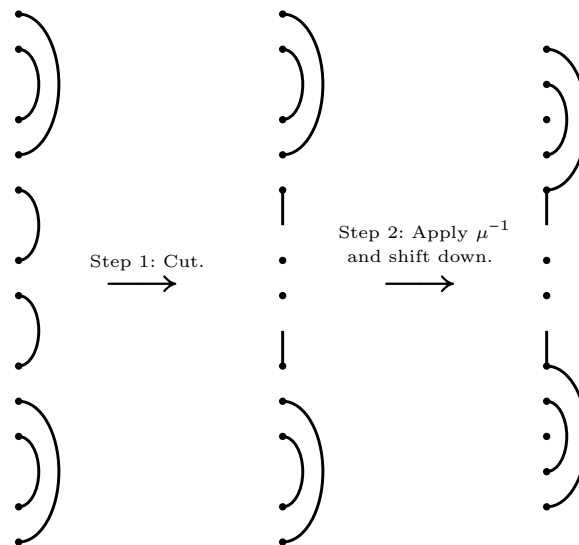


Figure 3.11: An illustration of the first two steps for the map ν . We curve some of the edges in the matching M to make them more suggestive of the arcs they will become in $\nu(M)$.

First step. For each $i \in [t]$, we cut the edge connecting $-m_i < m_i$ where it passes to the right of -1 and 1 . If $m_0 = -1$, then we also cut this edge. If $m_0 \neq -1$, then we detach the node 1 from the edge $\{1, m_0\}$. Symmetrically, we detach -1 from the edge $\{-m_0, -1\}$. (So, in this case, 1 and -1 are isolated nodes.) We call these edges, which now have precisely one endpoint, **segments**.

Second step. We apply μ^{-1} as described in the paragraph above. With each application of μ^{-1} we obtain a noncrossing arc diagram F' . The effect of shifting the indices in F' down is to glue the bottom node in F' to the top node of the segment below it. See Figure 3.11. The one exception occurs when we apply μ^{-1} to the matching on $[2, m_0 - 1]$: After shifting the indices down, the bottom node of the corresponding arc diagram is 1 . This process creates no crossings nor any incompatible shared endpoints. We carry out the symmetric process on the negative half of the diagram.

Third step. We reflect the negative half of the diagram about the vertical column of the nodes, so that each edge and segment now passes on the left side of the negative nodes. We call the segments in the negative half of the diagram **left segments**. The segments in the positive half of the diagram are called **right segments**. Then, we pair each right segment with a left segment, and we glue the pair together so that they pass between -1 and 1 . We pair the highest remaining right segment with highest remaining left segment. (We measure the height of a segment by the value of its endpoint.) See Figure 3.12.




Figure 3.12: An illustration of the map ν .

We write F for $\nu(M)$, and we verify that $F \in \mathcal{H}^s(n)$. We begin by checking that each pair of arcs in F is compatible. In our pictorial description above, we already noted that F does not have any endpoint incompatibilities. As we transform M into the noncrossing arc diagram F ,

the only possible place where we might possibly create a crossing is in the third step, where we glue together left and right segments. Lemma 3.3.9 guarantees that there are no crossings. Since, by our construction, F is symmetric, we conclude that it belongs to $\Delta^s(n)$. We claim that F is a facet. Observe that, after the second step in the construction, we have a maximal collection of ordinary right arcs between each consecutive pair of segments. The claim follows. Finally we check that F has no simple symmetric arcs. Since μ^{-1} is well-defined, the only possible simple symmetric arc in F is $\alpha_{-1,1}^s$. Suppose that there is an ordinary (though possibly not right) arc $\alpha \in F$ that has 1 as its top endpoint. According to our construction, its bottom endpoint is $-m_t$. Since $M \in \mathcal{M}_n$, we conclude that $-m_t < -1$. It is clear from the pictorial construction that ν is the inverse to μ_s . The proposition follows. \square

By Proposition 3.2.13 we conclude that when $n = 2r$, the complex $\Delta^s(n)$ is homotopy equivalent to a wedge of $\text{Cat}(B_r)$ many spheres. By Proposition 3.2.14, when $n = 2r - 1$, the complex $\Delta^s(n)$ is homotopy equivalent to a wedge of $\text{Cat}^+(B_r) - \text{Cat}(A_{r-2})$ many spheres. In the next two propositions we count the number of symmetric arcs in F , for each $F \in \mathcal{H}^s(n)$. The analogous computation for (ordinary) noncrossing arc diagrams in $\mathcal{H}(n)$ was easy: Under the map μ , arcs became edges, and there are r edges in each noncrossing perfect matching on $[2r]$. In the next proposition, we make essentially the same argument, replacing “arcs” and “edges” with “symmetric arcs” and “symmetric edges” in the sentence above.

Our first task is to make precise what is meant by “symmetric edges”. Write M for $\mu_s(F)$. As in the proof of Proposition 3.3.10, let $m_1 < \dots < m_t$ be the set of positive numbers satisfying $\{-m_i, m_i\} \in M$ and $m_i > 1$. Let m_0 be the number that is paired with 1. Informally, we say that a collection of edges is symmetric if the collection contributes precisely one symmetric arc to F .

More precisely, we say that the three edges $\{1, m_0\}$, $\{-m_0, 1\}$, and $\{-m_t, m_t\}$ are a ***symmetric triple***, when $m_0 \neq -1$ and $\{m_1, \dots, m_t\}$ is nonempty. We make two observations: First, under the map ν (after cutting/detaching, reflecting, and gluing), these three edges become the symmetric arc α_{-m_0, m_t}^s . Thus, the symmetric triple contributes precisely one symmetric arc to $F = \nu(M)$. Second, M contains at most one symmetric triple.

A ***symmetric pair of edges*** in M is any of the following:

- $\{a, b\}$ and $\{-a, -b\}$, where $1 < a < b$;
- $\{-m_i, m_i\}$ and $\{-m_{t-i}, m_{t-i}\}$, where $i > 0$;
- $\{-1, 1\}$ and $\{-m_t, m_t\}$, when $m_0 = -1$.

Observe that each symmetric pair of edges in M contributes precisely one symmetric arc to F . In particular, the pair $\{a, b\}$ and $\{-a, -b\}$ contributes the symmetric arc $\alpha_{a,b}^s$; the pair $\{-m_i, m_i\}$ and $\{-m_{t-i}, m_{t-i}\}$ contributes $\alpha_{-m_i, m_{t-i}}^s$; and $\{-1, 1\}$ and $\{-m_t, m_t\}$ contributes α_{-1, m_t}^s .

If $m_i = m_{t-i}$ then we say that $\{-m_i, m_i\}$ is a *symmetric edge*. We make two observations: First, after cutting, reflecting and re-gluing, $\{-m_i, m_i\}$ becomes the symmetric arc α_{-m_i, m_i}^s . Thus, the symmetric edge contributes precisely one symmetric arc. Second, the arc α_{-m_i, m_i}^s is fixed by the half-turn rotation of F . Since F has at most one such centrally symmetric arc, M also has at most one symmetric edge. We make one final observation: In our pictorial description of the map ν , no edge is deleted. Thus, every edge in M either belongs a symmetric triple or a symmetric pair; or it is a symmetric edge.

Proposition 3.3.11. *Suppose that $F \in \mathcal{H}^s(n)$.*

1. *If $n = 2r$, then F has precisely r symmetric arcs.*
2. *If $n = 2r - 1$ with $r > 1$, then F has precisely r symmetric arcs if and only if it contains a symmetric arc of the form $\alpha_{-1, k}$, where $1 < k \leq n$. Otherwise, F contains $r - 1$ symmetric arcs.*

Proof. Write M for $\mu_s(F)$. Let $m_1 < \dots < m_t$ be the set of positive numbers satisfying $\{-m_i, m_i\} \in M$ and $m_i > 1$. Let m_0 be the number that is paired with 1.

First, assume that M contains a symmetric triple: $\{1, m_0\}$, $\{-m_0, -1\}$, and $\{-m_t, m_t\}$. If $n = 2r - 1$, we claim that each of the remaining $2r - 4$ edges belongs to a symmetric pair. Since M has at most one symmetric edge and at most one symmetric triple, the claim follows. The symmetric triple contributes one symmetric arc to F , and the remaining edges contribute $r - 2$ symmetric arcs. We conclude that F has $r - 1$ symmetric arcs. If $n = 2r$, then, after accounting for the symmetric triple, there are $2r - 3$ remaining edges. Precisely one of these is a symmetric edge, and every other edge belongs to a symmetric pair. Together, the symmetric edge and the symmetric triple contribute two symmetric arcs to F . The remaining $2r - 4$ edges contribute $r - 2$ symmetric arcs. We conclude that F has r edges.

Now, assume that M does not contain a symmetric triple. So, each edge in M is either a symmetric edge or belongs to a symmetric pair. We conclude that the number of symmetric arcs in F is either $n/2$ if n is even or $(n + 1)/2$ if n is odd.

Observe that F contains a symmetric arc $\alpha_{-1, k}^s$ with $k > 1$ if and only if $m_0 = -1$. When n is odd, the set $\{m_1, \dots, m_t\}$ is nonempty (because $M \in \mathcal{M}_n$). Thus, when n is odd, F contains a symmetric arc $\alpha_{-1, k}^s$ if and only if M does *not* contain a symmetric triple. The statement follows. □

Finally, we complete the proof of Theorem 3.1.2 by showing that the diagrams in $\mathcal{H}^s(2r - 1)$ are equally distributed between sizes $r - 1$ and r .

Proposition 3.3.12. *There is an equal number of symmetric noncrossing arc diagrams in $\mathcal{H}^s(2r - 1)$ with size $r - 1$ and with size r .*

Proof. Write $n = 2r - 1$. As in the proof of Proposition 3.3.10, let $\{m_1, \dots, m_t\}$ be the set of positive numbers satisfying $\{-m_i, m_i\} \in M$ and $m_i > 1$. Let m_0 be the number that is paired with 1. Since n is odd and $M \in \mathcal{M}_n$, the set $\{m_1, \dots, m_t\}$ is nonempty. We define an involution on \mathcal{M}_n that toggles between the matchings in which $m_0 \neq -1$ and those in which $m_0 = -1$. Informally, this map interchanges the role of m_0 and m_1 .

First suppose that $m_0 \neq -1$. We map M to a matching \widetilde{M} in which the edges $\{1, m_0\}$ and $\{-m_0, -1\}$ are replaced with the edges $\{-1, 1\}$ and $\{-m_0, m_0\}$. Observe that m_0 is the smallest in $[2, n]$ that is paired with its negative in \widetilde{M} . (Thus, m_0 plays the role of \tilde{m}_1 .) If $m_0 = -1$, then we replace $\{-1, 1\} = \{-m_0, m_0\}$ and $\{-m_1, m_1\}$ with the edges $\{1, m_1\}$ and $\{-m_1, -1\}$. (Thus, m_1 plays the role of \tilde{m}_0 .) The statement follows. \square

3.4 Vertex Decomposability of the c -Cambrian lattices

In the final section of this chapter, we turn to the proof of Theorem 3.1.3. The proof involves three inductive arguments, one of which is complicated by breaking into cases. Before we dive in, we review the definition of the c -Cambrian lattices in type A, and for each c , we explain how to model the canonical join complex of the corresponding c -Cambrian lattice as a complex of noncrossing arc diagrams.

3.4.1 The c -Cambrian lattices

Each c -**Cambrian lattice** of type A_{n-1} is defined as the lattice quotient $\pi_{\downarrow}^{\Theta_c}(S_n)$ of the weak order on S_n where Θ_c is the c -Cambrian congruence (defined in Section 3.2.3). Recall that each c -Cambrian congruence Θ_c is determined by an orientation c that partitions the set $[2, n-1]$ into a set R of right nodes and a set L of left nodes. Proposition 3.2.7 says that a join-irreducible permutation is contracted by Θ_c precisely when its corresponding arc passes to the left of some element in L or to the right of some element in R . The next proposition generalizes Proposition 3.2.8. It follows immediately from Proposition 2.4.3 and Theorem 3.2.5. (See also [72, Example 4.9].)

Proposition 3.4.1. *Let c be any choice of orientation for the Coxeter diagram of A_{n-1} . Then the canonical join complex of the c -Cambrian lattice (of type A_{n-1}) is isomorphic to the subcomplex of the noncrossing arc complex on n nodes induced by the set of arcs that do not pass to the left of any element in L or to the right of any element in R .*

We write $\Delta(n, R, L)$ for the canonical join complex of the c -Cambrian lattice (of type A_{n-1}). Observe that there is a unique arc α in $\Delta(n, R, L)$ for each pair of endpoints $1 \leq i < k \leq n$. This arc passes to left of each node in $R \cap \text{supp}^\circ(\alpha)$ and to the right of each node in $L \cap \text{supp}^\circ(\alpha)$. We abuse notation and, for the remainder of the chapter, we write $\alpha_{i,k}$ for this arc.

Let J be any subset of $[n]$. We write $\Delta_J(n, R, L)$ for the complex of noncrossing arc diagrams obtained from $\Delta(n, R, L)$ by deleting the nodes in $[n] \setminus J$ and any arc that is adjacent to some node in $[n] \setminus J$.

Remark 3.4.2. Consider the complex $\Delta(n, R, L)$, where L is any subset of $[2, n-1]$ and $R = [2, n-1] \setminus L$. Upon restricting to the nodes in L , each arc in $\Delta_L(n, R, L)$ may be interpreted as a right arc (because each arc passes to right of every node between its endpoints). Thus, $\Delta_L(n, R, L)$ is isomorphic to $\Delta(l)$ where $l = |L|$.

Remark 3.4.2 motivates the proof of the next lemma. In the statement, the symbol $(*)$ denotes the join operation for simplicial complexes. Recall that the **join** of the simplicial complexes Δ and Δ' is the complex:

$$\Delta * \Delta' = \{F \cup F' : F \in \Delta \text{ and } F' \in \Delta'\}.$$

The **link** of a vertex v in Δ is the subcomplex:

$$\text{lk}(v) = \{F \in \Delta : v \notin F \text{ and } F \cup \{v\} \in \Delta\}.$$

Lemma 3.4.3. *Let L be any subset of $[2, n-1]$ and let R be the complementary subset. Then the link of $\alpha_{1,n}$ in $\Delta(n, R, L)$ is isomorphic to $\Delta_R(n, R, L) * \Delta_L(n, R, L)$.*

Proof. Observe that any arc $\alpha' \in \text{lk}(\alpha_{1,n})$ lies either strictly to the left or strictly to the right of $\alpha_{1,n}$. So, α' must either have both of its endpoints in L or both of its endpoints in R . Thus α' corresponds to an arc either in $\Delta_L(n, R, L)$ or $\Delta_R(n, R, L)$. Whenever the support of an arc in $\Delta_L(n, R, L)$ intersects the support of an arc in $\Delta_R(n, R, L)$, they have $\alpha_{1,n}$ between them. Thus, each arc in $\Delta_L(n, R, L)$ is compatible with each arc in $\Delta_R(n, R, L)$. We conclude that the link $\text{lk}(\alpha_{1,n})$ is isomorphic to $\Delta_R(n, R, L) * \Delta_L(n, R, L)$. \square

3.4.2 Vertex decomposability

Before we proceed with the proof of Theorem 3.1.3 we recall some terminology. A (not necessarily pure) simplicial complex Δ is **vertex decomposable** if either Δ is a simplex or there is a vertex v satisfying:

(VD1) The complex $\Delta \setminus v$ is vertex decomposable.

(VD2) The link $\text{lk}(v)$ is vertex decomposable.

(VD3) No facet of $\text{lk}(v)$ is a facet of $\Delta \setminus v$.

In this case, the vertex v is called the *shedding vertex*. (A nonpure vertex decomposable simplicial complex is sometimes called *semipure vertex decomposable*, see [54, Definition 3.29].) The following is [14, Theorem 11.3].

Theorem 3.4.4. *If Δ is a vertex decomposable complex then it is shellable.*

When Δ is vertex decomposable, with shedding vertex v , one constructs a shelling inductively: Write F_1, \dots, F_a for a shelling of $\Delta \setminus \{v\}$ and E_1, \dots, E_b for a shelling of $\text{lk}(v)$. Then, $F_1, \dots, F_a, E_1 \cup \{v\}, \dots, E_b \cup \{v\}$ is a shelling of Δ .

We will need the following well-known fact (see, for example, [54, Theorem 3.30]).

Lemma 3.4.5. *Suppose that Δ and Δ' are vertex decomposable complexes. Then, the join $\Delta * \Delta'$ is also vertex decomposable.*

Our proof of Theorem 3.1.3 is by induction on n . The base case, when $n = 3$ is obvious. (Figure 3.1 shows the canonical join complex when $n = 3$.) Throughout the remainder of the section, we assume that $\Delta(m', S', T')$ is vertex decomposable for all $m' < n$ and for each subset $T' \subset [2, m' - 1]$. (The set $S' = [2, m' - 1] \setminus T'$.) The next theorem is the main result in this section, and it implies Theorem 3.1.3.

Theorem 3.4.6. *Assume that $\Delta(m', S', T')$ is vertex decomposable for all $m' < n$ and all $T' \subseteq [2, m' - 1]$. Let L be any subset of $[2, n - 1]$ and let R be the complementary subset. The complex $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k}\}$ is vertex decomposable, where k ranges over the set $[3, n]$.*

Proposition 3.4.7. *If Theorem 3.4.6 holds, then $\Delta(n, R, L)$ is vertex decomposable with shedding vertex $\alpha_{1,n}$. In particular, Theorem 3.1.3 holds.*

Proof. By Theorem 3.4.6 and induction on n , we have $\Delta(n, R, L) \setminus \{\alpha_{1,n}\}$ is vertex decomposable. This is (VD1). By Lemma 3.4.3, our inductive hypothesis on n , and Lemma 3.4.5 we obtain (VD2). Suppose that $F \in \text{lk}(\alpha_{1,n})$. Each arc $\alpha_{i,k} \in F$ has $i > 1$. We conclude that the set $F \cup \{\alpha_{1,2}\}$ is a noncrossing arc diagram in $\Delta(n, R, L) \setminus \{\alpha_{1,n}\}$. In particular F is not a facet in $\Delta(n, R, L) \setminus \{\alpha_{1,n}\}$. Thus, we obtain (VD3). \square

To prove Theorem 3.4.6, we argue that $\alpha_{1,k}$ is a shedding vertex for the complex $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k+1}\}$, where $k \in [3, n - 1]$. The argument is by induction on k . The main difficulty is in verifying that (VD2) holds. Unfortunately, this requires a third induction that is further complicated by cases.

Our general strategy is to pull out any straightforward facts that we will need along the way. We organize the proof(s) of vertex decomposability by checking conditions (VD1), (VD2), and (VD3) in separate lemmas. We begin by checking that the base case holds for our induction on k .

Lemma 3.4.8. *Assume that $\Delta(m', S', T')$ is vertex decomposable for all $m' < n$ and for each $T' \subseteq [2, m-1]$. Then $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,3}\}$ is vertex decomposable.*

Proof. The only arc in $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,3}\}$ that has 1 as a bottom endpoint is $\alpha_{1,2}$. Also, $\alpha_{1,2}$ is the only arc that has 2 as a top endpoint. Thus, $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,3}\}$ is isomorphic to the join $\Delta_{[2,n]}(n, R, L) * \{\alpha_{1,2}\}$. The complex $\Delta_{[2,n-1]}(n, R, L)$ is vertex decomposable by our hypothesis. The statement follows from Lemma 3.4.5. \square

Now, we argue that $\alpha_{1,k}$ is a shedding vertex for $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k+1}\}$, where $k \in [3, n-1]$. Our inductive hypothesis implies that (VD1) holds. The next lemma will simplify some of our discussion of the link of $\alpha_{1,k}$. The statement follows from the fact that each pair of arcs in $\{\alpha_{1,n}, \dots, \alpha_{1,3}\}$ share a bottom endpoint at 1.

Lemma 3.4.9. *For each $k \in [3, n-1]$, the link of $\alpha_{1,k}$ taken in $\Delta(n, R, L)$ coincides with the link taken in $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k+1}\}$.*

In light of Lemma 3.4.9, we compute $\text{lk}(\alpha_{1,k})$ in $\Delta(n, R, L)$ for the remainder of the section. The next lemma implies that $\alpha_{1,k}$ satisfies (VD3).

Lemma 3.4.10. *Let L be any subset of $[2, n-1]$ and let R be the complementary subset. Suppose that $F \in \text{lk}(\alpha_{1,k})$, where $k \in [3, n-1]$. Then $F \cup \{\alpha_{1,2}\}$ is a noncrossing arc diagram in $\Delta \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k}\}$. In particular, F is not a facet in $\Delta(n, R, L) \setminus \{\alpha_{1,n}, \dots, \alpha_{1,k}\}$.*

Proof. It is enough to check that no arc in F has 1 as a bottom endpoint. This follows immediately from the fact that $F \in \text{lk}(\alpha_{1,k})$. \square

We have reduced the proof of Theorem 3.4.6 to checking that the link $\text{lk}(\alpha_{1,k})$ in $\Delta(n, R, L)$ is vertex decomposable. (By Lemma 3.4.9, this implies that condition (VD2) holds.) Recall from Lemma 3.4.3 that the link $\text{lk}(\alpha_{1,n})$ is isomorphic to $\Delta_R(n, R, L) * \Delta_L(n, R, L)$. The arcs in $\text{lk}(\alpha_{1,k})$ also split into two disjoint sets, depending on whether they pass on the left or the right side of $\alpha_{1,k}$, as we now make precise in Lemma 3.4.11. In the statement, and for the remainder of the section, we write $\Delta_J^{\setminus k}(n, R, L)$ for subcomplex induced by the set of arcs in $\Delta_J(n, R, L)$ that do *not* have k as a top endpoint, where $k \in J$.

Lemma 3.4.11. *Let L is any subset of $[2, n-1]$, $R = [2, n-1] \setminus L$, and consider the link of $\alpha_{1,k}$ in $\Delta(n, R, L)$, where $k \in [3, n-1]$.*

1. *If $k \in L$, then $\text{lk}(\alpha_{1,k})$ is isomorphic to $\Delta_{I_1}(n, R, L) * \Delta_{J_1}^{\setminus k}(n, R, L)$, where $I_1 = [2, k-1] \cap L$ and $J_1 = [k, n] \cup R$.*
2. *If $k \in R$, then $\text{lk}(\alpha_{1,k})$ is isomorphic to $\Delta_{I_2}(n, R, L) * \Delta_{J_2}^{\setminus k}(n, R, L)$, where $I_2 = [2, k-1] \cap R$ and $J_2 = [k, n] \cup L$.*

Proof. We prove the first item, where $k \in L$. The proof of the second item is symmetric. Suppose that $\alpha_{i,l} \in \text{lk}(\alpha_{1,k})$. If $i < k < l$, then $\alpha_{i,l}$ lies on the right side of $\alpha_{1,k}$ (because $k \in L$). We conclude that $i \in R$. Similarly, $i \in [2, k-1] \cap L$ if and only if $l \in [2, k-1] \cap L$. (If a compatible arc begins at a point on the left side of $\alpha_{1,k}$, it must end at a point on the left side. Otherwise it will intersect $\alpha_{1,k}$. See Figure 3.13 and Example 3.4.12.) Thus, each arc in the link of $\alpha_{1,k}$ corresponds to an arc in either $\Delta_{I_1}(n, R, L)$ or $\Delta_{J_1}^k(n, R, L)$. Note that whenever the support of an arc in $\Delta_{I_1}(n, R, L)$ intersects the support of an arc in $\Delta_{J_1}^k(n, R, L)$, they have $\alpha_{1,k}$ between them. Thus, each arc in $\Delta_{I_1}(n, R, L)$ is compatible with each arc in $\Delta_{J_1}^k(n, R, L)$. The statement follows. \square



Figure 3.13: The arcs $\alpha_{1,7}$ and $\alpha_{2,8}$ in $\Delta(9, \{4, 5, 8\}, \{2, 3, 6, 7\})$.

Example 3.4.12. In Lemma 3.4.11 we consider the link of $\alpha_{1,k}$, where $k \in L$. As an example, consider $\alpha_{1,7}$ in $\Delta(9, \{4, 5, 8\}, \{2, 3, 6, 7\})$. Note that any arc whose bottom endpoint lies in $[2, 6] \cap L$ also has its top endpoint in this set—otherwise it intersects $\alpha_{1,7}$ as shown in Figure 3.13.

With Lemma 3.4.11 in hand, we have further reduced the proof of Theorem 3.4.6 to showing that the complex $\Delta_J^k(n, R, L)$ is vertex decomposable, where $J = [k, n] \cup R$ or $J = [k, n] \cup L$. By symmetry, we assume that $J = [k, n] \cup R$. This is the case where $k \in L$ (the first item in Lemma 3.4.11). In particular, k is the smallest node in $J \cap L$. Write m for the cardinality of the set $[k, n] \cup R$, and $n_1 < n_2 < \dots < n_{r-1}$ for the numbers in $R \cap [2, k-1]$. After reindexing, we have that $\Delta_J^k(n, R, L)$ is isomorphic to a complex of the form $\Delta^r(m, S, T)$, where T is a subset of $[2, m-1]$, and r is its smallest element. (Throughout, $S = [2, m-1] \setminus T$.) If we can prove the next proposition, then we will have completed the proof of Theorem 3.4.6.

Proposition 3.4.13. *Assume that $\Delta(m', S', T')$ is vertex decomposable for each $m' \leq n-1$ and each subset $T' \subseteq [2, m']$. Take $m \leq n-1$, and let $r \in [2, m-1]$. If the subset $T \subseteq [2, m-1]$ is chosen so that r is its smallest element, then $\Delta^r(m, S, T)$ is vertex decomposable.*

We prove Proposition 3.4.13 by induction on m . (The base case where $m = 2$ is trivial.) The argument is complicated, and will occupy the remainder of the section. It requires three

separate cases, and we use the inductive hypothesis on m in each case.

We begin with the case when $r > 3$, where we can reuse almost all of our earlier work. Our goal is to prove the following proposition. The hypothesis in the statement is our inductive hypothesis on m for Proposition 3.4.13. (We assume that $\Delta(m'', S'', T'')$ is vertex decomposable as part of the inductive hypothesis on n for Theorem 3.1.3.)

Proposition 3.4.14. *Assume that Proposition 3.4.13 holds whenever m is replaced by $m'' < m$. In particular, we assume that $\Delta(m'', S'', T'')$ is vertex decomposable, for each $T'' \subseteq [2, m'' - 1]$. Take $m \leq n - 1$, and let $r \in [3, m - 1]$. Suppose that the subset $T \subseteq [2, m - 1]$ is chosen so that r is its smallest element. Then the complex $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k}\}$ is vertex decomposable, where $k \in [3, m] \setminus \{r\}$.*

We prove Proposition 3.4.14 by induction on k , just as we had been doing for Theorem 3.4.6. The analogues to Lemma 3.4.3 and Lemma 3.4.9 hold. Importantly, because $r > 3$, the analogue to Lemma 3.4.10 also holds. (The reader can verify that Lemma 3.4.10 fails in $\Delta^{\setminus 2}(m, S, T)$.) First we gather some useful facts that will help us to check (VD2).

Lemma 3.4.15. *Let $r \in [3, m - 1]$, and suppose that the subset $T \subseteq [2, m - 1]$ is chosen so that r is its smallest element.*

1. *Suppose that $k < r$. Then the link of $\alpha_{1,k}$ taken in $\Delta^{\setminus r}(m, S, T)$ is isomorphic to $\Delta_{[2,k-1]}(m, S, T) * \Delta_{[k,m]}^{\setminus r}(m, S, T)$.*
2. *Suppose that $k > r$. Then the link of $\alpha_{1,k}$ taken in $\Delta^{\setminus r}(m, S, T)$ coincides with the link taken in $\Delta(m, S, T)$.*

Proof. First, assume that $k < r$, and, for the moment, consider $\text{lk}(\alpha_{1,k})$ taken in $\Delta(m, S, T)$. Since $[2, r - 1] \subseteq S$, we have $k \in S$. Lemma 3.4.11 says that $\text{lk}(\alpha_{1,k})$ is isomorphic to the join $\Delta_{I_2}(m, S, T) * \Delta_{J_2}^{\setminus k}(m, S, T)$, where $I_2 = [2, k - 1] \cap S$ and $J_2 = [k, m] \cup T$. We observe that $[2, k - 1] \cap S = [2, k - 1]$ and $[k, m] \cup T = [k, m]$. Thus, the condition that no arc has k as a top endpoint in $\Delta_{J_2}^{\setminus k}(m, S, T)$ is vacuous. Therefore, $\Delta_{J_2}^{\setminus k}(m, S, T) = \Delta_{J_2}(m, S, T)$. We conclude that the link of $\alpha_{1,k}$ taken in $\Delta^{\setminus r}(m, S, T)$ is isomorphic to $\Delta_{[2,k-1]}(m, S, T) * \Delta_{[k,m]}^{\setminus r}(m, S, T)$.

Now, assume that $k > r$, and consider the link of $\alpha_{1,k}$ taken in $\Delta(m, S, T)$. No arc in $\text{lk}(\alpha_{1,k})$ has r as its top endpoint because r is the smallest element of T . (See Figure 3.14 and Example 3.4.16). Thus $\text{lk}(\alpha_{1,k})$ in $\Delta^{\setminus r}(m, S, T)$ coincides with the link taken in $\Delta(m, S, T)$. The second item follows. \square



Figure 3.14: The arc $\alpha_{1,7}$ in $\Delta(9, \{1, 2, 3, 4, 5, 7\}, \{6, 8\})$.

Example 3.4.16. Consider $\alpha_{1,7}$ in the complex $\Delta(9, \{1, 2, 3, 4, 5, 7\}, \{6, 8\})$ shown in Figure 3.14. There is no arc with 6 as a top endpoint that is compatible with $\alpha_{1,7}$. So, the link of $\alpha_{1,7}$ taken in $\Delta(9, \{1, 2, 3, 4, 5, 7\}, \{6, 8\})$ coincides with the link taken in $\Delta^{\setminus 6}(9, \{1, 2, 3, 4, 5, 7\}, \{6, 8\})$.

Proof of Proposition 3.4.14. We prove the proposition by induction on k . The proof for the base case is the same as the proof given in Lemma 3.4.8. The complex $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,3}\}$ is isomorphic to $\Delta_{[2,m]}^{\setminus r}(m, S, T) * \{\alpha_{1,2}\}$. The complex $\Delta_{[2,m]}^{\setminus r}(m, S, T)$ is vertex decomposable by our hypothesis.

We argue that $\alpha_{1,k}$ is a shedding vertex for $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k+1}\}$, where $k \in [3, m-1] \setminus \{r\}$. We have (VD1) by induction. Observe that the analogue to Lemma 3.4.9 holds: The link of $\alpha_{1,k}$ taken in $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k+1}\}$ coincides with the link taken in $\Delta^{\setminus r}(m, S, T)$. By Lemma 3.4.15, our hypothesis, and Lemma 3.4.5 we obtain (VD2). Since $r > 2$, the analogue to Lemma 3.4.10 also holds: For each $F \in \text{lk}(\alpha_{1,k})$, the set $F \cup \{\alpha_{1,2}\}$ belongs to $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k}\}$. Thus, we obtain (VD3). \square

The next proposition is analogous to Proposition 3.4.7 and completes the proof of the case where $r > 3$ for Proposition 3.4.13.

Proposition 3.4.17. *Assume that Proposition 3.4.13 holds whenever m is replaced by $m'' < m$. In particular, we assume that $\Delta(m'', S'', T'')$ is vertex decomposable, for each $T'' \subseteq [2, m'' - 1]$. Take $m \leq n - 1$, and let $r \in [3, m - 1]$. Suppose that the subset $T \subset [2, m - 1]$ is chosen so that r is its smallest element. Then, $\Delta^{\setminus r}(m, S, T)$ is vertex decomposable with shedding vertex $\alpha_{1,m}$.*

Proof. The proof is almost word-for-word the same as the proof of Proposition 3.4.7. Proposition 3.4.14 implies that (VD1) holds. Observe that the analogue of Lemma 3.4.3 holds: The link of $\alpha_{1,m}$ in $\Delta^{\setminus r}(m, S, T)$ is isomorphic to $\Delta_S(m, S, T) * \Delta_T^{\setminus r}(m, S, T)$. (The same argument establishes its proof here.) Consider $\Delta_T^{\setminus r}(m, S, T)$. Since r is smallest element in T , the condition that no arc has r as its top endpoint is vacuous. We conclude that $\Delta_T^{\setminus r}(m, S, T)$ is isomorphic to $\Delta_T(m, S, T)$. Our hypothesis implies that both $\Delta_T(m, S, T)$ and $\Delta_S(m, S, T)$ are vertex decomposable. By Lemma 3.4.5, we have (VD2). For each $F \in \text{lk}(\alpha_{1,k})$, the set $F \cup \{\alpha_{1,2}\}$ is in $\Delta^{\setminus r}(m, S, T) \setminus \{\alpha_{1,m}\}$ (because $r > 2$). Thus, we have (VD3). \square

We have further reduced the proof of Theorem 3.4.6 to the following statement:

Proposition 3.4.18. *Assume that Proposition 3.4.13 holds whenever m is replaced by $m'' < m$. In particular, we assume that $\Delta(m'', S'', T'')$ is vertex decomposable, for each $T'' \subseteq [2, m'' - 1]$. Suppose that T is any subset of $[2, m - 1]$ containing 2. Then $\Delta^{\setminus 2}(m, S, T)$ is vertex decomposable.*

We break the proof of Proposition 3.4.18 into two cases depending whether or not $3 \in T$. (Two cases are necessary. See Remark 3.4.23.) When $3 \notin T$, our argument is very similar to the one used for Proposition 3.4.14. So, we present this case first. Our goal is to prove the following:

Proposition 3.4.19. *Assume that Proposition 3.4.13 holds whenever m is replaced by $m'' < m$. In particular, we assume that $\Delta(m'', S'', T'')$ is vertex decomposable, for each $T'' \subseteq [2, m'' - 1]$. Suppose $T \subset [2, m - 1]$, and T contains 2 but not 3.*

1. *The complex $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k}\}$ is vertex decomposable, where $k \in [4, m]$.*
2. *In particular, $\Delta^{\setminus 2}(m, S, T)$ is vertex decomposable, with $\alpha_{1,m}$ its shedding vertex.*

As above, we will prove the first item of Proposition 3.4.19 by induction on k . Before we dive into the proof, we gather some useful lemmas. In the base case of the induction on k , we argue that $\alpha_{2,3}$ is a shedding vertex for $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}\}$. The next lemma will be useful.

Lemma 3.4.20. *Suppose $T \subset [2, m - 1]$, and T contains 2 but not 3. Consider the arc $\alpha_{1,k}$ in $\Delta^{\setminus 2}(m, S, T)$, where $k \in [4, m]$. Then, $\alpha_{1,k}$ intersects $\alpha_{2,3}$. In particular, the link of $\alpha_{2,3}$ in $\Delta^{\setminus 2}(m, S, T)$ is isomorphic to $\Delta_{[3,m]}(m, S, T)$.*

Proof. Since $2 \in T$ and $3 \in S$, every arc $\alpha_{1,k}$ with $3 < k$ passes between the nodes 2 and 3. Thus, $\alpha_{1,k}$ crosses the simple arc $\alpha_{2,3}$. In particular, $\alpha \in \Delta^{\setminus 2}(m, S, T)$ is compatible with $\alpha_{2,3}$ if and only if its bottom endpoint is greater than or equal to 3. Thus, the link of $\alpha_{2,3}$ in $\Delta^{\setminus 2}(m, S, T)$ is isomorphic to $\Delta_{[3,m]}(m, S, T)$. \square

The next lemma is analogous to Lemma 3.4.10, and will help us check (VD3).

Lemma 3.4.21. *Suppose that $T \subset [2, m - 1]$, and T contains 2 but not 3. Consider the link of $\alpha_{1,k}$ taken in $\Delta^{\setminus 2}(m, S, T)$.*

1. *For each $F \in \text{lk}(\alpha_{1,k})$, the set $F \cup \{\alpha_{1,3}\} \in \Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k}\}$, where $k \in [4, m]$.*
2. *For each $F \in \text{lk}(\alpha_{2,3})$, the set $F \cup \{\alpha_{1,3}\} \in \Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}, \alpha_{2,3}\}$.*

Proof. Consider the first item. Since $\alpha_{1,k}$ is not compatible with $\alpha_{1,l}$ (as long as $l \neq k$), it is enough to show that $F \cup \{\alpha_{1,3}\}$ belongs to $\Delta^{\setminus 2}(m, S, T)$. Suppose that $\alpha_{i,l} \in \text{lk}(\alpha_{1,k})$. The only arcs that have 3 as a top endpoint are $\alpha_{2,3}$ and $\alpha_{1,3}$. Lemma 3.4.20 implies that $\alpha_{2,3} \notin \text{lk}(\alpha_{1,k})$. Also, $\alpha_{1,3} \notin \text{lk}(\alpha_{1,k})$ because $\alpha_{1,3}$ and $\alpha_{1,k}$ share a bottom endpoint at 1. Thus, $l \geq 4$. When $i \geq 3$, it is obvious that $\alpha_{i,l}$ and $\alpha_{1,3}$ are compatible. When $i = 2$, then it is enough to check that $\alpha_{1,3}$ and $\alpha_{2,4}$ are compatible. Since $\Delta^{\setminus 2}(m, S, T)$ is flag, we obtain the first item.

The second item follows immediately from Lemma 3.4.20. \square

The following lemma will help us check that (VD2) holds.

Lemma 3.4.22. *Suppose that $T \subseteq [2, m - 1]$. Then the link of $\alpha_{1,k}$ taken in $\Delta^{\setminus 2}(m, S, T)$ coincides with the link taken in $\Delta(m, S, T)$ for each $k \in [3, m]$.*

Proof. We note that $\alpha_{1,2}$ is the only arc that has 2 as its top endpoint. So, it is the only arc that we delete when we pass from $\Delta(m, S, T)$ to $\Delta^{\setminus 2}(m, S, T)$. Consider the link $\text{lk}(\alpha_{1,k})$ taken in $\Delta(m, S, T)$. The arcs $\alpha_{1,2}$ and $\alpha_{1,k}$ are not compatible because they share a bottom endpoint. So, we lose nothing by passing to $\Delta^{\setminus 2}(m, S, T)$. The statement follows. \square

Proof of Proposition 3.4.19. We prove the first item by induction on k . In the base case we need to show that $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}\}$ is vertex decomposable. We claim that $\alpha_{2,3}$ is a shedding vertex. Checking (VD1) is somewhat involved, so we consider this condition last. By Lemma 3.4.20, the link $\text{lk}(\alpha_{2,3})$ taken in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}\}$ coincides with the link taken in $\Delta^{\setminus 2}(m, S, T)$. In particular, link of $\alpha_{2,3}$ is isomorphic $\Delta_{[3,m]}(m, S, T)$. The complex $\Delta_{[3,m]}(m, S, T)$ is vertex decomposable by our hypothesis. Thus, we obtain (VD2). By the second item in Lemma 3.4.21, we obtain (VD3). We observe that $\alpha_{1,3}$ is the only arc in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}, \alpha_{2,3}\}$ with 3 as its top endpoint. There is no arc in this complex that has 1 as its bottom endpoint. Thus, the complex $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,4}, \alpha_{2,3}\}$ is isomorphic to $\Delta_{[2,m]}^3(m, S, T) * \{\alpha_{1,3}\}$. By our hypothesis and Lemma 3.4.5, we have (VD1). We have proved that the base case holds.

Next, we argue that $\alpha_{1,k}$ is a shedding vertex in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k+1}\}$, for each k in the set $[4, m - 1]$. We have (VD1) by induction. Note that the analogue to Lemma 3.4.9 holds: The link of $\alpha_{1,k}$ taken in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{1,m}, \dots, \alpha_{1,k+1}\}$ coincides with the link taken in $\Delta^{\setminus 2}(m, S, T)$. By Lemma 3.4.22, our hypothesis, and Lemma 3.4.5, we obtain (VD2). Condition (VD3) follows immediately from the first item of Lemma 3.4.21. We have completed the proof of the first item.

Finally, we argue that $\alpha_{1,m}$ is a shedding vertex for $\Delta^{\setminus 2}(m, S, T)$. The first item in the proposition gives us (VD1). By Lemma 3.4.22, the link of $\alpha_{1,m}$ taken in $\Delta^{\setminus 2}(m, S, T)$ coincides with its link in $\Delta(m, S, T)$. Lemma 3.4.3 implies that the link $\text{lk}(\alpha_{1,m})$ is isomorphic to

$\Delta_S(m, S, T) * \Delta_T(m, S, T)$. By our hypothesis and Lemma 3.4.5, we have (VD2). By the first item in Lemma 3.4.21 we obtain (VD3). \square



Figure 3.15: The set $\{\alpha_{3,5}, \alpha_{2,3}\} \in \text{lk}(\alpha_{1,4})$, and it is a facet in $\Delta^{\setminus 2}(5, \{4\}, \{2, 3\}) \setminus \{\alpha_{1,5}, \alpha_{1,4}\}$.

Remark 3.4.23. Observe that our proof for Lemma 3.4.19 fundamentally fails if both 2 and 3 are in T . As a specific example, consider $\Delta^{\setminus 2}(5, \{4\}, \{2, 3\})$. We claim that $\alpha_{1,4}$ is *not* a shedding vertex for $\Delta^{\setminus 2}(5, \{4\}, \{2, 3\}) \setminus \{\alpha_{1,5}\}$. The set $\{\alpha_{3,5}, \alpha_{2,3}\}$ belongs to $\text{lk}(\alpha_{1,4})$, and it is a facet in $\Delta^{\setminus 2}(5, \{4\}, \{2, 3\}) \setminus \{\alpha_{1,5}, \alpha_{1,4}\}$. See Figure 3.15. Thus, the condition (VD3) fails.

We have reduced the proof of Theorem 3.4.6 to one final proposition. This is the last case in the proof of Proposition 3.4.13.

Proposition 3.4.24. *Assume that Proposition 3.4.13 holds whenever m is replaced by $m'' < m$. In particular, we assume that $\Delta(m'', S'', T'')$ is vertex decomposable, for each $T'' \subseteq [2, m'' - 1]$. Suppose that $T \subseteq [2, m - 1]$ and $2, 3 \in T$. Then $\Delta^{\setminus 2}(m, S, T)$ is vertex decomposable. In particular,*

1. $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,k}\}$ is vertex decomposable, where $k \in [4, m]$; and
2. $\alpha_{2,m}$ is a shedding vertex for $\Delta^{\setminus 2}(m, S, T)$.

As before, we prove the first item of Proposition 3.4.24 by induction on k . In the base case, we will argue that $\alpha_{1,3}$ is a shedding vertex for $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}\}$. The next lemma will be useful.

Lemma 3.4.25. *Suppose that $T \subseteq [2, m - 1]$ and $2, 3 \in T$, and consider $\alpha_{2,k}$ in $\Delta^{\setminus 2}(m, S, T)$, where $k \in [4, m]$. Then:*

1. $\alpha_{2,k}$ intersects the arc $\alpha_{1,3}$.
2. In particular, the link of $\alpha_{1,3}$ taken in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}\}$ coincides with the link taken in $\Delta(m, S, T)$.

Proof. The first item is easily verified by drawing the arcs $\alpha_{1,3}$ and $\alpha_{2,k}$. For example, see Figure 3.16 below. It follows immediately that the link of $\alpha_{1,3}$ taken in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}\}$ coincides with the link taken in $\Delta^{\setminus 2}(m, S, T)$. Recall that in passing from $\Delta(m, S, T)$ to



Figure 3.16: In $\Delta(9, \{4, 5, 8\}, \{2, 3, 6, 7\})$, arcs $\alpha_{2,8}$ and $\alpha_{1,3}$ intersect.

$\Delta^{\setminus 2}(m, S, T)$ we simply delete $\alpha_{1,2}$. Since $\alpha_{1,2} \notin \text{lk}(\alpha_{1,3})$, we obtain the second statement of the lemma. \square

The next lemma is the analogue to Lemma 3.4.9. (The statement holds because each arc in $\{\alpha_{2,m}, \dots, \alpha_{2,4}\}$ has a bottom endpoint at 2.)

Lemma 3.4.26. *Suppose that $T \subseteq [2, m-1]$ and $k \in [4, m-1]$. The link of $\alpha_{2,k}$ taken in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,k+1}\}$ coincides with the link taken in $\Delta^{\setminus 2}(m, S, T)$.*

In light of Lemma 3.4.26, we will take the link of $\alpha_{2,k}$ in $\Delta^{\setminus 2}(m, S, T)$. The next lemma is analogous to Lemma 3.4.10 and will help us check (VD3).

Lemma 3.4.27. *Suppose that $T \subseteq [2, m-1]$, $2, 3 \in T$, and $k \in [4, m]$. Consider the link of $\alpha_{2,k}$ taken in $\Delta^{\setminus 2}(m, S, T)$ and the link of $\alpha_{1,3}$ taken in $\Delta(m, S, T)$.*

1. For each $F \in \text{lk}(\alpha_{2,k})$ the set $F \cup \{\alpha_{2,3}\} \in \Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,k}\}$.
2. For each $F \in \text{lk}(\alpha_{1,3})$ the set $F \cup \{\alpha_{2,3}\} \in \Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}\}$.

Proof. If $\alpha \in \text{lk}(\alpha_{2,k})$, then α 's bottom endpoint is not equal to 2. By the first item in Lemma 3.4.25, no arc in $\text{lk}(\alpha_{2,k})$ has 3 as a top endpoint. In addition, there is no arc in $\Delta^{\setminus 2}(m, S, T)$ that passes between the nodes 2 and 3 (because both 2 and 3 belong to T). The first statement follows.

Suppose that $\alpha_{i,l} \in \text{lk}(\alpha_{1,3})$. Lemma 3.4.25 implies that $i \geq 3$. The second statement follows. \square

In the following lemma we compute the link of $\alpha_{2,k}$ in $\Delta^{\setminus 2}(m, S, T)$. The first two items in the statement are analogous to Lemma 3.4.11, and their proofs are similar. The third item is analogous to Lemma 3.4.3. Since its proof is identical to the argument given in Lemma 3.4.3, we do not repeat it here.

Lemma 3.4.28. *Suppose that $T \subseteq [2, m-1]$, $2, 3 \in T$, and $k \in [4, m]$. Consider the link of $\alpha_{2,k}$ in $\Delta^{\setminus 2}(m, S, T)$.*

1. If $k \in T$, then $\text{lk}(\alpha_{2,k})$ is isomorphic to $\Delta_{I_1}(m, S, T) * \Delta_{J_1}^{\setminus k}(m, S, T)$, where $I_1 = [3, k-1] \cap T$ and $J_1 = [k, m] \cup S \cup \{1\}$.
2. If $k \in S$, then $\text{lk}(\alpha_{2,k})$ is isomorphic to $\Delta_{I_2}(m, S, T) * \Delta_{J_2}^{\setminus k}(m, S, T)$, where the set I_2 is equal to $([4, k-1] \cap S) \cup \{1\}$ and $J_2 = [k, m] \cup (T \setminus \{2\})$.
3. If $k = m$, then $\text{lk}(\alpha_{2,k})$ is isomorphic to $\Delta_{K_1}(m, S, T) * \Delta_{K_2}(m, S, T)$, where $K_1 = S \cup \{1\}$ and $K_2 = T \setminus \{2\}$.

Proof. First assume that $k \in T$. Each arc $\alpha_{i,l}$ in $\text{lk}(\alpha_{2,k})$ with $i < k < l$ lies on the right side of $\alpha_{2,k}$. So, i belongs to S or is equal to 1. Conversely, every arc in $\text{lk}(\alpha_{2,k})$ that has 1 as a bottom endpoint must pass to the right of $\alpha_{2,k}$ (because $2 \in T$) and end in S or above k . Every arc in $\text{lk}(\alpha_{2,k})$ that has one endpoint in $[3, k-1] \cap T$ must have the other endpoint in $[3, k-1] \cap T$ (because every arc that begins on the left side of $\alpha_{2,k}$ also ends on the left side). The first statement follows.

When $k \in S$, each arc $\alpha_{i,l}$ in $\text{lk}(\alpha_{2,k})$ with $i < k < l$ lies on the left of $\alpha_{2,k}$. Thus, $i \in T$. Since both 2 and 3 are in T , the smallest possible bottom endpoint for such an arc is 3. On the other hand, suppose that $\alpha_{i,l}$ satisfies $l \in [4, k-1] \cap S$. Because $2, 3 \in T$ (and $\alpha_{i,l}$ lies on the right side of $\alpha_{2,k}$), we have $i = 1$ or $i \in [4, k-1] \cap S$. The second item follows. \square

In the next lemma, we argue that each of the complexes, $\Delta_{J_1}^{\setminus k}(m, S, T)$ and $\Delta_{J_2}^{\setminus k}(m, S, T)$, from Lemma 3.4.28 satisfies the hypotheses of Proposition 3.4.13, with an eye toward establishing (VD2).

Lemma 3.4.29. *Let $i = 1, 2$. Suppose that $k_i \in [4, m]$, and $T_i \subseteq [2, m-1]$ satisfies:*

- $k_1 \in T_1$;
- $k_2 \notin T_2$;
- $2, 3 \in T_i$ for both $i = 1, 2$.

Let $J_1 = [k_1, m] \cup S_1 \cup \{1\}$, and let $J_2 = [k_2, m] \cup (T_2 \setminus \{2\})$. Then, $\Delta_{J_i}^{\setminus k_i}(m, S_i, T_i)$ is isomorphic to a complex $\Delta^{\setminus r''_i}(m''_i, S''_i, T''_i)$ satisfying the hypotheses of Proposition 3.4.13 with $m''_i < m$, for both $i = 1, 2$.

Proof. First, we consider $\Delta_{J_1}^{\setminus k_1}(m, S_1, T_1)$. Observe that k_1 is the smallest node in $T_1 \cap J_1$. Since S_1 does not contain 2 or 3, the cardinality of J_1 is strictly less than m . Thus, after reindexing the node set J_1 , we obtain the desired isomorphism.

Next, consider $\Delta_{J_2}^{\setminus k_2}(m, S_2, T_2)$. Observe that k_2 is the smallest node in $S_2 \cap J_2$. Since $k_2 \geq 4$, the cardinality of J_2 is strictly less than m . Also, $\Delta_{J_2}^{\setminus k_2}(m, S_2, T_2)$ is isomorphic to $\Delta_{J_2}^{\setminus k_2}(m, S', T')$, where $S' = T_2$ and $T' = S_2$. After reindexing the node set for $\Delta_{J_2}^{\setminus k_2}(m, S', T')$, we obtain the desired isomorphism. \square

Proof of Proposition 3.4.24. In the base case for our induction on k , we consider the complex $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}\}$. We claim that $\alpha_{1,3}$ is a shedding vertex, and we begin by checking (VD2) and (VD3). By the second item in Lemma 3.4.25, we may take the link of $\alpha_{1,3}$ in $\Delta(m, S, T)$. The first item in Lemma 3.4.25 implies that each arc in the link of $\text{lk}(\alpha_{1,3})$ has a bottom endpoint that is greater than or equal to 3. Thus, the link $\text{lk}(\alpha_{1,3})$ is isomorphic to $\Delta_{[3,m]}(m, S, T)$. By our hypothesis this complex is vertex decomposable, and we obtain (VD2). By second item in Lemma 3.4.27, the link $\text{lk}(\alpha_{1,3})$ satisfies (VD3).

Consider $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}\}$. The only arc in this complex that has 2 as a bottom endpoint is $\alpha_{2,3}$, and similarly, $\alpha_{2,3}$ is the only arc that has 3 as a top endpoint. Since no arc passes between 2 and 3, we conclude that each arc in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}\}$ is compatible with $\alpha_{2,3}$. Thus, the complex $\Delta^{\setminus 2}(n, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}\}$ is isomorphic to the join

$$\left(\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}, \alpha_{2,3}\} \right) * \{\alpha_{2,3}\}.$$

The complex $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}, \alpha_{2,3}\}$ is isomorphic to $\Delta_J^{\setminus 3}(m, S, T)$, where $J = [m] \setminus \{2\}$. (The isomorphism is obtained by deleting the node 2 in each diagram F in $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,4}, \alpha_{1,3}, \alpha_{2,3}\}$.) By our hypothesis, $\Delta_J^{\setminus 3}(m, S, T)$ is vertex decomposable. By Lemma 3.4.5, we obtain (VD1). We have proved that the base case holds.

We argue that $\alpha_{2,k}$ is a shedding vertex for $\Delta^{\setminus 2}(n, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,k+1}\}$, for $k \in [4, m-1]$. We have (VD1) by induction. We consider the link of $\alpha_{2,k}$ as computed in Lemma 3.4.28. By Lemma 3.4.29, our hypothesis, and Lemma 3.4.5, we have (VD2). By the first item in Lemma 3.4.27, the link $\text{lk}(\alpha_{2,k})$ satisfies (VD3). We conclude that $\alpha_{2,k}$ is a shedding vertex for $\Delta^{\setminus 2}(m, S, T) \setminus \{\alpha_{2,m}, \dots, \alpha_{2,k+1}\}$, as desired. This completes the proof of the first item.

For the second item, we claim that $\alpha_{2,m}$ is shedding vertex for $\Delta^{\setminus 2}(m, S, T)$. The first item of this proposition gives us (VD1). The third item in Lemma 3.4.28 says that the link of $\alpha_{2,m}$ is isomorphic to $\Delta_{K_1}(m, S, T) * \Delta_{K_2}(m, S, T)$, where $K_1 = S \cup \{1\}$ and $K_2 = T \setminus \{2\}$. For each $i = 1, 2$, the complex $\Delta_{K_i}(m, S, T)$ is vertex decomposable by our hypothesis. By Lemma 3.4.5, we obtain (VD2). (VD3) follows from the first item in Lemma 3.4.27. \square

This completes the inductive argument (on m) that Proposition 3.4.13 holds. Thus, we obtain (VD2) in our inductive argument (on k) for Theorem 3.4.6. In turn, this completes the inductive argument (on n) for Theorem 3.1.3.

Chapter 4

Coxeter BiCatalan Combinatorics

4.1 Introduction

This chapter¹ considers enumeration problems closely related to Coxeter-Catalan combinatorics. (For background on Coxeter-Catalan combinatorics, see for example [2, 32]). Each enumeration problem can be thought of as counting pairs of “twin” Coxeter-Catalan objects—twin sortable elements or twin nonnesting partitions, etc. Many of the terms used in this introductory section are new to this chapter and will be explained in Section 4.2.

In the setting of sortable elements and Cambrian lattices/fans, the enumeration problem is to count the following families of objects:

- maximal cones in the bipartite biCambrian fan (the common refinement of two bipartite Cambrian fans);
- pairs of twin c -sortable elements for bipartite c ;
- classes in the bipartite biCambrian congruence (the meet of two bipartite Cambrian congruences);
- elements of the bipartite biCambrian lattice;
- c -bisortable elements for bipartite c .

In type A, c -bisortable elements for bipartite c are in bijection with permutations avoiding a set of four bivincular patterns in the sense of [15, Section 2] and with alternating arc diagrams,

¹ The content of this chapter, Coxeter-biCatalan Combinatorics, will appear in the Journal of Algebraic Combinatorics under the same title and with authorship Emily Barnard and Nathan Reading. Barnard was the initial writer for the following: Sections 4.3.3–4.3.5 and Sections 4.4.3–4.4.5. In particular, the counting arguments for the (alternating) noncrossing arc diagrams and for the c -bisortable elements were originally constructed by Barnard. Portions of Section 4.3.6 were initially written by both authors. In particular, the folding arguments were written by Barnard. Collaborative revisions were made throughout.

as will be explained in Sections 4.3.1–4.3.3. In type B, similar bijections exist with certain signed permutations and with centrally symmetric alternating arc diagrams, as described in Section 4.3.5.

In the setting of nonnesting partitions (antichains in the root poset), the enumeration problem is to count two families of objects:

- antichains in the doubled root poset;
- pairs of twin nonnesting partitions.

In the setting of clusters of almost positive roots (in the sense of [35]), the problem is to count two families of objects:

- maximal cones in the bicluster fan (the common refinement of the cluster fan, in the original bipartite sense of Fomin and Zelevinsky, and its antipodal opposite);
- pairs of twin clusters, again in the bipartite sense.

In the setting of noncrossing partitions, the problem is to count the following families of objects:

- pairs of twin bipartite c -noncrossing partitions;
- pairs of twin bipartite (c, c^{-1}) -noncrossing partitions.

The main result of this chapter is the following.

Theorem 4.1.1. *For each finite Coxeter group/root system, all of the enumeration problems posed above have the same answer.*

In all of the settings above except the nonnesting setting, the objects described above can be defined for arbitrary choices of a Coxeter element. However, the enumerations depend on the choice of Coxeter element, and we emphasize that Theorem 4.1.1 is an assertion about the enumeration in the case where the Coxeter element is chosen to be bipartite. See Section 4.2.2 for the definition of Coxeter elements and bipartite Coxeter elements.

The enumeration problems in the nonnesting setting require a crystallographic root system, but Theorem 4.1.1 still holds in the other settings for noncrystallographic types. See also Remark 4.2.1.

We will see in Section 4.2 that within each group of bullet points above, the various enumeration problems have the same answer essentially by definition. Using known uniform correspondences from the usual Coxeter-Catalan combinatorics, it is straightforward to give (in Theorems 4.2.18 and 4.2.21) uniform bijections connecting the Cambrian/sortable setting to the noncrossing and cluster settings. The difficult part of the main result is the following theorem which connects the nonnesting setting to the other settings.

Theorem 4.1.2. *For crystallographic W , c -bisortable elements for bipartite c are in bijection with antichains in the doubled root poset.*

More specifically, we have the following refined version of Theorem 4.1.2.

Theorem 4.1.3. *For crystallographic W and for any k , the number of bipartite c -bisortable elements with k descents equals the number of k -element antichains in the doubled root poset.*

Our proof of Theorems 4.1.2 and 4.1.3 in Section 4.4 would be uniform if a uniform proof were known connecting the nonnesting setting to the other settings of the usual Coxeter-Catalan combinatorics. Indeed, the opposite is true: A well-behaved uniform bijection proving Theorem 4.1.2 or Theorem 4.1.3 would imply a uniform proof of the analogous Coxeter-Catalan statement. (See Remark 4.4.30 for details.) However, the proofs of these theorems are far from a trivial recasting of Coxeter-biCatalan combinatorics in terms of Coxeter-Catalan combinatorics. Instead, it requires a count of antichains in the doubled root poset indirectly in terms of the Coxeter-Catalan numbers and a nontrivial proof that the same formula holds for bipartite c -bisortable elements. The formula uses a notion of “double-positive” Catalan and Narayana numbers, which already appeared in [7] as the local h -polynomials of the positive cluster complex. (See Remark 4.4.7.)

We propose the terms *W -biCatalan number* and *W -biNarayana number* and the symbols $\text{biCat}(W)$ and $\text{biNar}_k(W)$ for the numbers appearing in Theorems 4.1.1 and 4.1.3.

Theorem 4.1.4. *The W -biCatalan numbers for irreducible finite Coxeter groups are listed in Table 4.1.*

Table 4.1: The W -biCatalan numbers

W	A_n	B_n	D_n	E_6	E_7	E_8	F_4	H_3	H_4	$I_2(m)$
$\text{biCat}(W)$	$\binom{2n}{n}$	2^{2n-1}	$6 \cdot 4^{n-2} - 2\binom{2n-4}{n-2}$	1700	8872	54066	196	56	550	$2m$

The type-A and type-B cases of Theorem 4.1.4 are proved, in the nonnesting setting, in Section 4.2.1 by recasting the antichain count as a count of lattice paths. The same cases can also be established in the setting of c -bisortable elements by recasting the problem in terms of alternating arc diagrams. Although the latter approach is more difficult, we carry out the type-A and type-B enumeration by the latter approach in Section 4.3, because the combinatorial models for bipartite c -bisortable elements in types A and B are of independent interest, and because the enumeration of alternating arc diagrams provides the crucial insight

which leads to the recursive proof of Theorem 4.1.1. (See Remark 4.3.13.) The type-D case of Theorem 4.1.4 is much more difficult, and involves solving the type-D case of the recursion used in the proof of Theorem 4.1.2. The formula in type D was first guessed using the package `GFUN` [77]. The enumerations in the exceptional types were obtained using Stembridge’s `posets` and `coxeter/weyl` packages [84].

We also obtain formulas for the W -biNarayana numbers outside of type D. Generating functions for biNarayana numbers for some type-D Coxeter groups are shown in Table 4.4. At present we have no conjectured formula for the D_n -biNarayana numbers. See Section 4.4.9 for a modest conjecture.

Theorem 4.1.5. *The biNarayana numbers for each of the irreducible finite Coxeter groups, except in type D, are given by the generating functions shown in Table 4.2.*

Table 4.2: The biNarayana numbers

W	$\sum_{k=0}^n \text{biNar}_k(W) q^k$
A_n	$\sum_{k=0}^n \binom{n}{k}^2 q^k$
B_n	$\sum_{k=0}^n \binom{2n}{2k} q^k$
E_6	$1 + 66q + 415q^2 + 736q^3 + 415q^4 + 66q^5 + q^6$
E_7	$1 + 119q + 1139q^2 + 3177q^3 + 3177q^4 + 1139q^5 + 119q^6 + q^7$
E_8	$1 + 232q + 3226q^2 + 13210q^3 + 20728q^4 + 13210q^5 + 3226q^6 + 232q^7 + q^8$
F_4	$1 + 44q + 106q^2 + 44q^3 + q^4$
G_2	$1 + 10q + q^2$
H_3	$1 + 27q + 27q^2 + q^3$
H_4	$1 + 116q + 316q^2 + 116q^3 + q^4$
$I_2(m)$	$1 + (2m - 2)q + q^2$

Naturally, one would like a uniform formula for the W -biCatalan number, but we have not found one. A tantalizing near-miss is the *non-formula* $\prod_{i=1}^n \frac{h+e_i-1}{e_i}$, where h is the Coxeter number and the e_i are the exponents. This expression captures the W -biCatalan numbers for W of types A_n , B_n , H_3 , and $I_2(m)$ —the “coincidental types” of [92]—but fails to even be an integer in some other types. In every case, the expression is a surprisingly good estimate of the W -biCatalan number.

Section 4.2 is devoted to filling in definitions and details for the discussion above and proving the easy parts of Theorem 4.1.1. In Section 4.3, we explain why, in type A, the bipartite

Table 4.4: The type-D biNarayana numbers

D_4	$1 + 20q + 42q^2 + 20q^3 + q^4$
D_5	$1 + 35q + 136q^2 + 136q^3 + 35q^4 + q^5$
D_6	$1 + 54q + 343q^2 + 600q^3 + 343q^4 + 54q^5 + q^6$
D_7	$1 + 77q + 731q^2 + 2011q^3 + 2011q^4 + 731q^5 + 77q^6 + q^7$
D_8	$1 + 104q + 1384q^2 + 5556q^3 + 8638q^4 + 5556q^5 + 1384q^6 + 104q^7 + q^8$
D_9	$1 + 135q + 2402q^2 + 13314q^3 + 29868q^4$ $+ 29868q^5 + 13314q^6 + 2402q^7 + 135q^8 + q^9$
D_{10}	$1 + 170q + 3901q^2 + 28624q^3 + 87874q^4 + 126336q^5$ $+ 87874q^6 + 28624q^7 + 3901q^8 + 170q^9 + q^{10}$

bisortable elements are in bijection with alternating arc diagrams and carry out the enumeration of alternating arc diagrams. We carry out a similar enumeration in type B, in terms of centrally symmetric alternating arc diagrams. We conjecture that the bipartite biCambrian fan is simplicial (and thus that its dual polytope is simple), and prove the conjecture in types A and B. In Section 4.4, we discuss double-positive Coxeter-Catalan numbers and establish a formula counting antichains in the doubled root poset in terms of double-positive Coxeter-Catalan numbers. We then show that bipartite c -bisortable elements satisfy the same recursion, thus proving Theorem 4.1.3 and completing the proof of Theorem 4.1.1. Finally, we establish some additional formulas involving double-positive Coxeter-Catalan numbers, Coxeter-Catalan numbers, and Coxeter-biCatalan numbers and use them to prove the formula for $\text{biCat}(D_n)$ and thus complete the proof of Theorem 4.1.4.

4.2 BiCatalan objects

In this section, we fill in the definitions and details behind the enumeration problems discussed in the introduction. An exposition in full detail would require reviewing Coxeter-Catalan combinatorics in full detail, so we leave some details to the references.

4.2.1 Antichains in the doubled root poset and twin nonnesting partitions

The *root poset* of a finite crystallographic root system Φ is the set of positive roots in Φ , partially ordered by setting $\alpha \leq \beta$ if and only if $\beta - \alpha$ is in the nonnegative span of the simple roots. Recall that the dual of a poset (X, \leq) is the poset (X, \geq) . That is, the dual has the same ground set, with $x \leq y$ in the dual poset if and only if $x \geq y$ in the original poset. The *doubled*




Figure 4.1: Some doubled root posets

root poset consists of the root poset, together with a disjoint copy of the dual poset, identified on the simple roots. Figure 4.1 shows some doubled root posets.

The antichain counts in types A and B are easy and known, in the guise of lattice path enumeration. Antichains in the doubled root poset of type A_n are in an easy bijection with lattice paths from $(0,0)$ to (n,n) with steps $(1,0)$ and $(0,1)$. The bijection can be made so that the number of elements in the antichain corresponds to the number of right turns in the path (the number of times a $(1,0)$ -step immediately follows a $(0,1)$ -step). To specify a path with k right turns, we need only specify where the right turns are. This means choosing $0 \leq x_1 < \dots < x_k \leq n-1$ and $1 \leq y_1 < \dots < y_k \leq n$ and placing right turns at $(x_1, y_1), \dots, (x_k, y_k)$. Thus, as is well-known, there are $\binom{n}{k}^2$ paths with k right turns.

Antichains in the doubled root poset of type B_n are similarly in bijection with lattice paths from $(-2n+1, -2n+1)$ to $(2n-1, 2n-1)$ with steps $(2,0)$ and $(0,2)$ that are symmetric with respect to the reflection through the line $y = -x$. The k -element antichains correspond to paths with either $2k$ right turns, (k of which are to the left of the line $y = -x$) or $2k-1$ right turns ($k-1$ of which are left of the line $y = -x$ and one of which is on the line $y = -x$). Each path

is uniquely determined by its first $2n - 1$ steps, whereupon the path intersects the line $y = -x$. Thus, the paths map bijectively to words of length $2n - 1$ in the letters N and E (for North steps $(0, 2)$ and East steps $(2, 0)$). Appending the letter E to the end of each word, the k -element antichains correspond to the words having exactly k positions where an E appears immediately after an N . (The number of right turns in the path is odd if and only if one of these is position $2n$.) The $2n$ -letter words ending in E and having exactly k instances of an E following an N are in bijection with $2k$ -element subsets of $\{1, \dots, 2n\}$. (Given such a word, take the set of positions where the letter changes, with the convention that an N in the first position is a change but an E in the first position is not. So, for example, $ENNEEE$ gives the subset $\{2, 4\}$ and $NEEENE$ gives $\{1, 2, 5, 6\}$.) We see that there are $\binom{2n}{2k}$ k -element antichains, and 2^{2n-1} total antichains, in the doubled root poset of type B_n .

Remark 4.2.1. It is not clear in general how one should define a “root poset” for a noncrystallographic root system. See [2, Section 5.4.1] for a discussion. In type $I_2(m)$, there is an obvious way to define an unlabeled poset generalizing the root posets of types A_2 , B_2 , and G_2 . We say “unlabeled” here because it is obvious how the poset should look but not obvious how the poset elements should correspond to roots. There is also an unlabeled type- H_3 root poset suggested in [2, Section 5.4.1]. For these choices of root posets, one can verify that Theorem 4.1.1 holds in these types as well.

Remark 4.2.2. The doubled root poset, and similar posets, were probably first considered by Proctor (see [85, Remark 4.8(a)]) and then by Stembridge, as a tool for counting reduced expressions for certain elements of finite Coxeter groups. In the simply-laced types (A, D, and E), the doubled root poset corresponds to the *smashed Cayley order* defined by Stembridge in [85, Section 4]. In the non-simply laced types, the smashed Cayley order is disconnected and is a strictly weaker partial order than the doubled root poset. Stembridge [85, Theorem 4.6] shows that the component whose elements are short roots is a distributive lattice. Thus in particular the doubled root posets of types A, D, and E are distributive lattices. One can easily check distributivity in the remaining crystallographic types B, F, and G (and in fact in types H_3 and $I_2(m)$). By the Fundamental Theorem of Distributive Lattices [81, Theorem 3.4.1], the doubled root poset is isomorphic to the poset of order ideals in its subposet of join-irreducible elements. These posets of join-irreducible elements are shown in Figure 4.2 for several types. An explicit root-theoretic description of the poset of join-irreducible elements in the simply-laced types also appears in [85, Theorem 4.6].

The *support* of a root β is the set of simple roots appearing with nonzero coefficient in the expansion of β in the basis of simple roots. The support of a set of roots is the union of the supports of the roots in the set. We write Δ for the simple roots and, given a set A of roots, we write A° for the set of non-simple roots in A . If A_1 and A_2 are nonnesting partitions (i.e.




Figure 4.2: Some posets of join-irreducibles of doubled root posets

antichains in the root poset), then (A_1, A_2) is a pair of *twin nonnesting partitions* if and only if $A_1 \cap \Delta = A_2 \cap \Delta$, and $\text{supp}(A_1^\circ) \cap \text{supp}(A_2^\circ) = \emptyset$.

Given an antichain A in the doubled root poset, define $\text{top}(A)$ to be the intersection of A with the root poset that forms the top of the doubled root poset. Define $\text{bottom}(A)$ to be the intersection of A with the dual root poset that forms the bottom of the doubled root poset. Both $\text{top}(A)$ and $\text{bottom}(A)$ are sets of positive roots. The following proposition is an immediate consequence of the observation that a root β in the top part of the doubled root poset is related to a root γ in the bottom part of the doubled root poset if and only if the supports of β and γ overlap.

Proposition 4.2.3. *The map $A \mapsto (\text{top}(A), \text{bottom}(A))$ is a bijection from antichains in the doubled root poset to pairs of twin nonnesting partitions.*

We pause to observe that the first biNarayana number (the number of elements of the doubled root poset) is the number of roots minus the rank of W .

Proposition 4.2.4. *If W is an irreducible finite Coxeter group with Coxeter number h and rank n , then $\text{biNar}_1(W) = n(h - 1)$.*

4.2.2 BiCambrian fans

The Cambrian fan is a complete simplicial fan whose maximal faces are naturally in bijection [69, 74] with seeds in an associated cluster algebra of finite type and with noncrossing partitions. Furthermore, the Cambrian fan is the normal fan [48, 49] to a simple polytope called the *generalized associahedron* [20, 35], which encodes much of the combinatorics of the associated cluster algebra. More directly, the Cambrian fan is the \mathbf{g} -vector fan of the cluster algebra. (This was conjectured, and proved in a special case, in [74, Section 10] and then proved in general in [93].)

The defining data of a Cambrian fan is a finite Coxeter group W and a Coxeter element c of W . We emphasize that the results discussed in Section 4.1 concern a special “bipartite” choice of c , as explained below, but for now we proceed with a discussion for general c . A *Coxeter element* is the product of a permutation of the simple generators of W , or equivalently it is an orientation of the Coxeter diagram of W . Given a choice of W , we will assume the usual representation of W as a reflection group acting with trivial fixed subspace. The collection of reflecting hyperplanes in this representation is the *Coxeter arrangement* of W . The hyperplanes in the Coxeter arrangement cut space into cones, which constitute a fan called the *Coxeter fan* $\mathcal{F}(W)$. The maximal cones of the Coxeter fan are in bijection with the elements of W . The Cambrian fan $\mathcal{C}(W, c)$ is the coarsening of the Coxeter fan obtained by gluing




Figure 4.3: Cambrian fans and the biCambrian fan in type B_2

together maximal cones according to an equivalence relation on W called the c -Cambrian congruence. Further details on the c -Cambrian congruence appear in Section 4.2.3. For fixed W , all choices of c give distinct but combinatorially isomorphic Cambrian fans.

For each Coxeter element c , the inverse element c^{-1} is also a Coxeter element, corresponding to the opposite orientation of the diagram. We define the **biCambrian fan** $\mathcal{BC}(W, c)$ to be the coarsest common refinement of the Cambrian fans $\mathcal{C}(W, c)$ and $\mathcal{C}(W, c^{-1})$. Since both $\mathcal{C}(W, c)$ and $\mathcal{C}(W, c^{-1})$ are coarsenings of $\mathcal{F}(W)$, so is $\mathcal{BC}(W, c)$. Naturally, $\mathcal{BC}(W, c^{-1}) = \mathcal{BC}(W, c)$.

Example 4.2.5. To illustrate the definition, take W of type B_2 with simple generators s_1 and s_2 . Figure 4.3 shows, from left to right, the s_1s_2 -Cambrian fan, the s_2s_1 -Cambrian fan, and the s_1s_2 -biCambrian fan. Observe that the s_1s_2 -biCambrian fan coincides with the B_2 Coxeter fan. In general, when W is rank 2, the c -biCambrian fan for any choice of Coxeter element c is equal to the Coxeter fan $\mathcal{F}(W)$.

Example 4.2.6. For W of type A_3 , there are two non-isomorphic c -biCambrian fans, shown in Figures 4.4 and 4.5 respectively. Each figure can be understood as follows: Intersecting the c -biCambrian fan with a unit sphere about the origin, we obtain a decomposition of the sphere into spherical convex polygons. The picture shows a stereographic projection of this polygonal decomposition to the plane. In each case, the walls of one Cambrian fan are shown in red and the walls of the opposite Cambrian fan are shown in blue. Walls that are in both Cambrian fans are shown dashed red and blue.

Remark 4.2.7. We observe that in Examples 4.2.5 and 4.2.6 that the common walls of $\mathcal{C}(W, c)$ and $\mathcal{C}(W, c^{-1})$ are exactly the reflecting hyperplanes for the simple generators of W . This fact true in general, and the simplest proof involves **shards**. We will not define shards here, but definitions and results can be found, for example, in [71]. Assuming for a moment that background, we sketch a proof. First, recast [75, Theorem 8.3] as the statement that the c -Cambrian congruence removes all but one shard from each reflecting hyperplane of W . As explained in the argument for [70, Proposition 1.3] (located in [70, Section 3] just after the proof of [70, Theorem 1.1]), the antipodal map sends the shard that is not removed by the c -Cambrian congruence to the shard that is not removed by the c^{-1} -Cambrian congruence. The




Figure 4.4: The linear biCambrian fan in type A_3




Figure 4.5: The bipartite biCambrian fan in type A_3

only shards that are fixed by the antipodal map are shards that consist of an entire reflecting hyperplane, and [71, Lemma 3.11] says that these are exactly the reflecting hyperplanes for the simple generators.

The c^{-1} -Cambrian fan $\mathcal{C}(W, c^{-1})$ coincides with $-\mathcal{C}(W, c)$, the image of the c -Cambrian fan under the antipodal map. This is an immediate corollary of [70, Proposition 1.3], which is a statement about the c -Cambrian congruence. See also [76, Remark 3.26]. Thus we have the following proposition which amounts to an alternate definition of the biCambrian fan.

Proposition 4.2.8. *The biCambrian fan $\mathcal{BC}(W, c)$ is the coarsest common refinement of $\mathcal{C}(W, c)$ and $-\mathcal{C}(W, c)$.*

Since $\mathcal{C}(W, c)$ and $\mathcal{C}(W, c^{-1})$ are the normal fans of two generalized associahedra, a standard fact (see [94, Proposition 7.12]) yields the following result.

Proposition 4.2.9. *For any W and c , the fan $\mathcal{BC}(W, c)$ is the normal fan of a polytope, specifically, the Minkowski sum of the generalized associahedra dual to $\mathcal{C}(W, c)$ and $\mathcal{C}(W, c^{-1})$.*

The definition of $\mathcal{BC}(W, c)$ seems strange *a priori*, but it is well-motivated *a posteriori* by enumerative results. The first such result is Theorem 4.2.10 below. When W is the symmetric group S_n (i.e. when W is of type A_{n-1}), the Coxeter diagram of W is a path. A **linear Coxeter element** of S_n is the product of the generators in order along the path.

Theorem 4.2.10. *When W is the symmetric group S_n and c is the linear Coxeter element, the number of maximal cones in $\mathcal{BC}(W, c)$ is the Baxter number.*

For more on the Baxter number, see [10, 24]. Theorem 4.2.10 was observed empirically (in the language of lattice congruences) in [67, Section 10] and then proven by J. West [91]. See also [43, 55]. The theorem is also related to the observation by Dulucq and Guibert [30] that pairs of twin binary trees are counted by the Baxter number.

Once one sees that the Baxter number counts maximal cones of $\mathcal{BC}(W, c)$ for W of type A and for a particular c , it is natural to look at other types of finite Coxeter group W , with the idea of defining a “ W -Baxter number” for each finite Coxeter group W . Indeed, there is a good notion of a “type-B Baxter number” discovered by Dilks [29]. The Coxeter diagram of type B is also a path, and taking c to be a linear Coxeter element, the maximal cones of $\mathcal{BC}(W, c)$ are counted by the type-B Baxter number. Despite the nice type-B result, there seems to be little hope for a reasonable definition of the W -Baxter number, because some types of Coxeter diagrams are not paths and thus it is not clear how to generalize the notion of a *linear* Coxeter element.

There is, however, a choice of Coxeter element that can be made uniformly for all finite Coxeter groups. Since the Coxeter diagram of any finite Coxeter group is acyclic, the diagram

is in particular bipartite. Thus we can fix a bipartition $S_+ \cup S_-$ of the diagram and orient each edge of the diagram from its vertex in S_- to its vertex in S_+ . The resulting Coxeter element is called a *bipartite Coxeter element*, and if c is a bipartite Coxeter element of W , we call $\mathcal{BC}(W, c)$ a *bipartite biCambrian fan*.

Proposition 4.2.9 says that $\mathcal{BC}(W, c)$ is the normal fan of a polytope, but does not guarantee that this polytope is simple (equivalently, that this fan is simplicial). In fact, simpleness fails for the linear Coxeter element of S_n , and this failure can be seen already in S_4 . (See Figure 4.4, and also [55, Figure 13]. The latter shows the 1-skeleton of this polytope disguised as the Hasse diagram of a certain lattice.) We conjecture that the situation is better in the bipartite case.

Conjecture 4.2.11. *If W is a bipartite Coxeter element, then $\mathcal{BC}(W, c)$ is a simplicial fan. (Equivalently, its dual polytope is simple.)*

We have verified Conjecture 4.2.11, with the aid of Stembridge’s packages [84], up to rank 6. Also, in Section 4.3.6, we prove the following theorem using alternating arc diagrams, by appealing to some results of [52] linking the lattice theory of the weak order to the representation theory of finite-dimensional algebras, and then applying a folding argument.

Theorem 4.2.12. *Conjecture 4.2.11 holds in types A and B.*

In Section 4.2.3, we will prove the following theorem.

Theorem 4.2.13. *If Conjecture 4.2.11 holds for a Coxeter group W , then the h -vector of the simplicial sphere underlying $\mathcal{BC}(W, c)$, for c bipartite, has entries $\text{biNar}_k(W)$.*

In light of the evidence for Conjecture 4.2.11 and in light of Theorem 4.2.13, we propose the term *simplicial W -biassociahedron* for the polytope whose face fan is $\mathcal{BC}(W, c)$ for c bipartite, and *simple W -biassociahedron* for the polytope whose normal fan is $\mathcal{BC}(W, c)$ for c bipartite.

Remark 4.2.14. Theorems 4.1.5, 4.2.12, and 4.2.13 imply that the A_n -biassociahedron has the same h -vector as the B_n -associahedron (also known as the cyclohedron). One is naturally led to ask whether these two polytopes are combinatorially isomorphic. The answer is no already for $n = 3$. The normal fan to the A_3 -biassociahedron is shown in Figure 4.5. The dual graph to this fan has a vertex that is incident to two hexagons and a quadrilateral. The graph of the B_3 -associahedron (shown for example in [32, Figure 3.9]) has no such vertex.

4.2.3 The biCambrian congruence, twin sortable elements, and bisortable elements

A *congruence* Θ on a lattice L is an equivalence relation respecting the meet and join operations. As in previous chapters, we consider only finite lattices, and some results quoted in this

section can fail for infinite lattices. On a finite lattice, congruences are characterized by three properties: congruence classes are intervals; the projection $\pi_{\downarrow}^{\Theta}$, mapping each element to the bottom element of its congruence class, is order preserving; and the projection π_{Θ}^{\uparrow} , mapping each element to the top element of its congruence class, is order preserving. The Θ -classes are exactly the fibers of $\pi_{\downarrow}^{\Theta}$. The quotient L/Θ of a finite lattice L modulo a congruence Θ is a lattice isomorphic to the subposet induced by the set $\pi_{\downarrow}^{\Theta}(L)$ of elements that are the bottoms of their congruence classes. The congruence Θ is determined by the set $\pi_{\downarrow}^{\Theta}(L)$: Specifically $x \equiv y$ modulo Θ if and only if the unique maximal element of $\pi_{\downarrow}^{\Theta}(L)$ below x equals the unique maximal element of $\pi_{\downarrow}^{\Theta}(L)$ below y .

The map $\pi_{\downarrow}^{\Theta}$ is a lattice homomorphism from L onto the subposet $\pi_{\downarrow}^{\Theta}(L)$, but care must be taken to avoid misinterpreting this fact. Literally, the fact that $\pi_{\downarrow}^{\Theta}$ is a lattice homomorphism means that for any $U \subseteq L$, we have $\pi_{\downarrow}^{\Theta}(\vee U) = \vee_{x \in U} \pi_{\downarrow}^{\Theta}(x)$ and $\pi_{\downarrow}^{\Theta}(\wedge U) = \wedge_{x \in U} \pi_{\downarrow}^{\Theta}(x)$, but in each identity, the join on the left side occurs in L while the join on the right side occurs in $\pi_{\downarrow}^{\Theta}(L)$. It is easy to check that $\pi_{\downarrow}^{\Theta}(L)$ is also a join-sublattice of L , so the distinction between the join in L and the join in $\pi_{\downarrow}^{\Theta}(L)$ is unnecessary. However, in general, $\pi_{\downarrow}^{\Theta}(L)$ need not be a meet-sublattice of L , so in interpreting the identity $\pi_{\downarrow}^{\Theta}(\wedge U) = \wedge_{x \in U} \pi_{\downarrow}^{\Theta}(x)$, it is crucial to be clear on where the meets occur.

The maximal cones of the Coxeter fan $\mathcal{F}(W)$, partially ordered according to a suitable linear functional, form a lattice isomorphic to the weak order on W . (This fact is true either for the right or left weak order. We will work with the right weak order.) Any lattice congruence Θ on the weak order on W defines a fan $\mathcal{F}_{\Theta}(W)$ coarsening $\mathcal{F}(W)$. (See [67, Theorem 1.1] and [67, Section 5].) Specifically, for each Θ -class, the union of the corresponding maximal cones in $\mathcal{F}(W)$ is itself a convex cone, and the collection of all these convex cones and their faces is the fan $\mathcal{F}_{\Theta}(W)$. Each Coxeter element c specifies a congruence Θ_c on the weak order called the *c -Cambrian congruence*. (See [68] for the definition.) The fan $\mathcal{F}_{\Theta_c}(W)$ is the *c -Cambrian fan* $\mathcal{C}(W, c)$ described earlier.

The set $\text{Con}(L)$ of all congruences on a given lattice L is itself a sublattice of the lattice of set partitions of L . In particular, the meet of two congruences is the coarsest set partition of L refining both congruences. We define the *c -biCambrian congruence* to be the meet, in $\text{Con}(W)$, of the Cambrian congruences Θ_c and $\Theta_{c^{-1}}$. The fan $\mathcal{F}_{\Theta}(W)$ for $\Theta = \Theta_c \wedge \Theta_{c^{-1}}$ is the coarsest common refinement of $\mathcal{F}(\Theta_c(W))$ and $\mathcal{F}(\Theta_{c^{-1}}(W))$. Thus the *c -biCambrian fan* $\mathcal{BC}(W, c)$ is the fan $\mathcal{F}_{\Theta}(W)$ for $\Theta = \Theta_c \wedge \Theta_{c^{-1}}$. In particular, the *c -biCambrian congruence classes* are in bijection with the maximal cones of $\mathcal{BC}(W, c)$. We define the *c -biCambrian lattice* to be the quotient of the weak order modulo the *c -biCambrian congruence*. The elements of the *c -biCambrian lattice* are thus in bijection with the maximal cones of $\mathcal{BC}(W, c)$.

We write π_{\downarrow}^c for the projection taking each element of W to the bottom element of its *c -Cambrian congruence class*, and similarly $\pi_{\downarrow}^{c^{-1}}$. (That is, π_{\downarrow}^c stands for $\pi_{\downarrow}^{\Theta}$ where $\Theta = \Theta_c$.)

Consider the map that sends each c -biCambrian congruence class to the pair $(\pi_{\downarrow}^c(w), \pi_{\downarrow}^{c^{-1}}(w))$, where w is any representative of the class. Because the c -biCambrian congruence Θ is the meet $\Theta_c \wedge \Theta_{c^{-1}}$, two elements u and v are congruent in the c -biCambrian congruence if and only if $\pi_{\downarrow}^c(u) = \pi_{\downarrow}^c(v)$ and $\pi_{\downarrow}^{c^{-1}}(u) = \pi_{\downarrow}^{c^{-1}}(v)$. Thus, the map from classes to pairs is a well-defined bijection from c -biCambrian congruence classes to its image.

The bottom elements of the c -Cambrian congruence are called *c -sortable elements*. (In fact c -sortable elements have an independent combinatorial definition [69, Section 2], but were shown to be the bottom elements of c -Cambrian congruences in [70, Theorems 1.1 and 1.4].) Given elements u and v of W , we define the pair (u, v) to be a pair of *twin (c, c^{-1}) -sortable elements* of W if there exists $w \in W$ such that $u = \pi_{\downarrow}^c(w)$ and $v = \pi_{\downarrow}^{c^{-1}}(w)$. The map considered in the previous paragraph is a bijection between c -biCambrian congruence classes and pairs of twin (c, c^{-1}) -sortable elements of W . The twin sortable elements are similar in spirit to the twin binary trees of [30], which were already mentioned in connection with Theorem 4.2.10. Indeed, for W of type A and c linear, the connection is implicit in the construction in [55] of a diagonal rectangulation from a pair of binary trees. (See also [55, Remark 6.6].) Also in type A, but for general c , the twin binary trees are generalized in [22] to *twin Cambrian trees*, which correspond explicitly to pairs of twin (c, c^{-1}) -sortable elements. Indeed, [22, Proposition 57] amounts to another computation of the type-A biCatalan number, quite different from the two given here (in Sections 4.2.1 and 4.3.4).

Another set of objects naturally in bijection with c -biCambrian congruence classes are the bottom elements of c -biCambrian congruence classes. We coin the term *c -bisortable elements* for these bottom elements. Although the c -sortable elements have a direct combinatorial characterization [69, Section 2], we currently have no direct combinatorial characterization of c -bisortable elements. We do offer the following indirect characterization of c -bisortable elements in terms of c -sortable elements and c^{-1} -sortable elements.

Proposition 4.2.15. *For any c , an element $w \in W$ is c -bisortable if and only if there exists a c -sortable element u and a c^{-1} -sortable element v such that $w = u \vee v$ in the weak order. When w is c -bisortable, we can take $u = \pi_{\downarrow}^c(w)$ and $v = \pi_{\downarrow}^{c^{-1}}(w)$.*

Proof. Given c -bisortable w , take $u = \pi_{\downarrow}^c(w)$ and $v = \pi_{\downarrow}^{c^{-1}}(w)$. Then $u \leq w$ and $v \leq w$. Since Cambrian congruence classes are intervals, any upper bound w' for u and v with $w' \leq w$ is congruent to u modulo Θ_c and congruent to v modulo $\Theta_{c^{-1}}$. Thus w' is congruent to w in the c -biCambrian congruence. Since w is the bottom element of its c -biCambrian congruence class, we conclude that $w' = w$. We have shown that $w = u \vee v$.

Suppose $w = u \vee v$ for some c -sortable element u and some c^{-1} -sortable element v . Since $\pi_{\downarrow}^c(w)$ is the unique maximal c -sortable element below w , we have $\pi_{\downarrow}^c(w) \geq u$. Similarly, $\pi_{\downarrow}^{c^{-1}}(w) \geq v$. If there exists $w' < w$ in the same c -biCambrian congruence class as w , then $w' \geq \pi_{\downarrow}^c(w')$,

$\pi_{\downarrow}^c(w') = \pi_{\downarrow}^c(w) \geq u$, and $w' \geq \pi_{\downarrow}^{c^{-1}}(w') = \pi_{\downarrow}^{c^{-1}}(w) \geq v$. This contradicts the fact that $w = u \vee v$, and we conclude that w is c -bisortable. \square

Recall that for any congruence Θ on a finite lattice L , the set $\pi_{\downarrow}^{\Theta}(L)$ is a join-sublattice of L . The Cambrian congruences have a stronger property: For any Coxeter element c , the c -sortable elements constitute a sublattice [70, Theorem 1.2] of the weak order on W . It is natural to ask whether the same is true for c -bisortable elements, but the answer is no. We give an example for $W = S_5$ and bipartite c : The permutations 45312 and 53142 are both c -bisortable but their meet 31452 is not. (To check this example, Proposition 4.3.6 will be very helpful.)

Each c -bisortable element v covers some number of elements in the c -biCambrian lattice. By a general fact on lattice quotients (see for example [71, Proposition 6.4]), v covers the same number of elements in the weak order on W . This number is $\text{des}(v)$, the number of descents of v . (We will define descents in Section 4.4.5.) The **descent generating function of c -bisortable elements** is the sum $\sum x^{\text{des}(v)}$ over all c -bisortable elements v . We will show that its coefficients are the W -biNarayana numbers. Its coefficients are the W -biNarayana numbers. A general fact about lattice quotients of the weak order [67, Proposition 3.5] implies that, when $\mathcal{BC}(W, c)$ is simplicial, the descent generating function of c -bisortable elements equals the h -polynomial of $\mathcal{BC}(W, c)$. In the bipartite case, Theorem 4.2.13 follows immediately.

4.2.4 Twin clusters and bicluster fans

Clusters of almost positive roots were introduced in [35], where they were used to define the generalized associahedra. In [36], clusters of almost positive roots were used to model cluster algebras of finite type. Here, we will not need the cluster-algebraic background, which can be found in [36]. Instead, we define almost positive roots and c -compatibility and quote some results about c -clusters and their relationship to c -sortable elements. We will also not need the more refined notion of “compatibility degree.”

In a finite root system, the **almost positive roots** are those roots which either are positive, or are the negatives of simple roots. The definition of compatibility in [35] is a special case (namely the bipartite case) of what we here call c -compatibility. The general definition was given in [56], but here we give a rephrasing found in [69, Section 7], translated into the language of almost positive roots.

We write $\{\alpha_1, \dots, \alpha_n\}$ for the simple roots and $\{s_1, \dots, s_n\}$ for the simple reflections. For each i in $\{1, \dots, n\}$, we define an involution σ_i on the set of almost positive roots by

$$\sigma_i(\beta) := \begin{cases} \beta & \text{if } \beta = -\alpha_j \text{ with } j \neq i, \text{ or} \\ s_i\beta & \text{otherwise.} \end{cases} \quad (4.2.1)$$

We write $[\beta : \alpha_i]$ for the coefficient of α_i in the expansion of β in the basis of simple roots. A simple reflection s_i is **initial** in a Coxeter element c if c has a reduced word starting with s_i . If s_i is initial in c , then $s_i c s_i$ is another Coxeter element.

The **c -compatibility** relations are a family of symmetric binary relations \parallel_c on the almost positive roots. They are the unique family of relations with

- (i) For any i in $\{1, \dots, n\}$, and Coxeter element c ,

$$-\alpha_i \parallel_c \beta \text{ if and only if } [\beta : \alpha_i] = 0.$$

- (ii) For each pair of almost positive roots β_1 and β_2 , each Coxeter element c , and each s_i initial in c ,

$$\beta_1 \parallel_c \beta_2 \text{ if and only if } \sigma_i(\beta_1) \parallel_{s_i c s_i} \sigma_i(\beta_2).$$

The **c -clusters** are the maximal sets of pairwise c -compatible almost positive roots. By [35, Theorem 1.8] and [56, Proposition 3.5], for fixed W , all c -clusters are of the same size, and furthermore, each is a basis for the root space (the span of the roots). Write $\mathbb{R}_{\geq 0}C$ for the nonnegative linear span of a c -cluster C . Then [35, Theorem 1.10] and [56, Theorem 3.7] state that the cones $\mathbb{R}_{\geq 0}C$, for all c -clusters C , are the maximal cones of a complete simplicial fan. We call this fan the **c -cluster fan**.

We define the **c -bicluster fan** to be the coarsest common refinement of the c -cluster fan and its antipodal opposite. A pair (C_1, C_2) of c -clusters is called a pair of **twin c -clusters** if the cones $\mathbb{R}_{\geq 0}C_1$ and $-\mathbb{R}_{\geq 0}C_2$ (the nonpositive linear span of C_2) intersect in a full-dimensional cone. It is immediate that maximal cones in the c -bicluster fan are in bijection with pairs of twin c -clusters.

Example 4.2.16. For W of type A_3 , up to symmetry there are two different c -bicluster fans: one for linear c and one for bipartite c , shown in Figures 4.6 and 4.7 respectively. These are again stereographic projections as explained in Example 4.2.6.

The two fans in Example 4.2.16 are combinatorially isomorphic. Despite this tantalizing fact, in this chapter, we only explore bicluster fans in the special case of bipartite Coxeter elements (the original setting of [35, 36]), where they are easily related to biCambrian fans. For the bipartite choice of c , [74, Theorem 9.1] says that the c -Cambrian fan is *linearly* isomorphic to the cluster fan. Combining this fact with Proposition 4.2.8, we have the following theorem.

Theorem 4.2.17. For all finite Coxeter groups W and for bipartite of c , the c -bicluster fan is linearly isomorphic to the c -biCambrian fan.




Figure 4.6: The linear bicluster fan in type A_3




Figure 4.7: The bipartite bicluster fan in type A_3

Because of the bijection between c -bisortable elements and maximal cones in $\mathcal{BC}(W, c)$ and the bijection between maximal cones in the c -bicluster fan and pairs of twin c -clusters, we have the following immediate consequence of Theorem 4.2.17.

Theorem 4.2.18. *For all finite Coxeter groups W , c -bisortable elements for bipartite c are in bijection with pairs of twin c -clusters.*

Combining Theorems 4.2.13 and 4.2.17, we obtain the following theorem.

Theorem 4.2.19. *If Conjecture 4.2.11 holds for a Coxeter group W , then the bipartite c -bicluster fan is simplicial and the h -vector of the underlying simplicial sphere has entries $\text{biNar}_k(W)$.*

4.2.5 Twin noncrossing partitions

The absolute order on a finite Coxeter group W is the prefix order (or equivalently the subword order) on W relative to the generating set T , the set of reflections in W . (By contrast, the prefix order relative to the simple reflections S is the weak order, while the subword order relative to S is the Bruhat order.) We will use the symbol \leq_T for the absolute order. The c -**noncrossing partitions** in a finite Coxeter group W are the elements of W contained in the interval $[1, c]_T$ in the absolute order on W . For details on the absolute order and noncrossing partitions, see for example [2, Chapter 2]. For our purposes, the key fact is a theorem of Brady and Watt.

Let W be a finite Coxeter group of rank n represented as a reflection group in \mathbb{R}^n and let T be the set of reflections of W . For each reflection $t \in T$, let β_t be the corresponding positive root. Given $w \in [1, c]_T$, define a cone

$$F_c(w) = \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} \cdot \beta_t \leq 0 \ \forall t \leq_T w, \ \mathbf{x} \cdot \beta_t \geq 0 \ \forall t \leq_T cw^{-1} \right\}.$$

The following theorem combines [17, Theorem 1.1] with [17, Theorem 5.5].

Theorem 4.2.20. *For c bipartite, the map F_c is a bijection from $[1, c]_T$ to the set of maximal cones in the c -Cambrian fan.*

The astute reader will notice a difference between our definition of F_c and the definition appearing in [17, Section 1]. The set of reflections t such that $t \leq_T w$ is the intersection of T with some (non necessarily standard) parabolic subgroup of W . The definition in [17] imposes inequalities $\mathbf{x} \cdot \beta_t \leq 0$ only for those β_t that are simple roots for that parabolic subgroup. Our definition imposes additional inequalities, all of which are implied by the inequalities for the simple roots. We similarly add additional redundant inequalities of the form $x \cdot \beta_t \geq 0$.

Theorem 4.2.20 suggests a definition of twin noncrossing partitions. In fact, given Proposition 4.2.8, two natural definitions suggest themselves. Given $u, v \in [1, c]_T$, we call (u, v) a pair

of *twin c -noncrossing partitions* if $F_c(u) \cap (-\mathcal{F}_c(v))$ is full-dimensional. Similarly, given $u \in [1, c]_T$ and $v \in [1, c^{-1}]_T$, we call (u, v) a pair of *twin (c, c^{-1}) -noncrossing partitions* if $F_c(u) \cap \mathcal{F}_{c^{-1}}(v)$ is full-dimensional. Theorem 4.2.20 now immediately implies the following theorem.

Theorem 4.2.21. *For all W and bipartite c , the c -bisortable elements are in bijection with pairs of twin c -noncrossing partitions and with pairs of twin (c, c^{-1}) -noncrossing partitions.*

4.3 Bipartite c -bisortable elements and alternating arc diagrams

In this section, we show how bipartite c -bisortable elements of type A are in bijection with certain objects called alternating arc diagrams. We then prove the type-A enumeration of bipartite c -bisortable elements in Theorem 4.1.1 by counting alternating arc diagrams and prove the type-B enumeration by counting centrally symmetric alternating arc diagrams.

4.3.1 Pattern avoidance

The Coxeter group of type A_n is the symmetric group S_{n+1} . We will write permutations x in S_{n+1} in their one-line notations $x_1 \cdots x_{n+1}$. In the weak order on permutations in S_{n+1} , there is a cover $x_1 \cdots x_{n+1} \lessdot y_1 \cdots y_{n+1}$ if and only if there exists i such that $y_i = x_{i+1} > x_i = y_{i+1}$ and $y_j = x_j$ for $j \notin \{i, i+1\}$. We say that x is covered by y via a swap in positions i and $i+1$.

The Cambrian congruences on S_{n+1} are described in detail in [68]. We quote part of the description here. The simple generator s_i for A_n is the transposition $(i \ i+1)$, for $i = 1, 2, \dots, n$. Each Coxeter element c can be encoded by a coloring of the elements $2, \dots, n$ that we call a *barring*. Each element i is either *overbarred* and marked \bar{i} if s_i occurs before s_{i-1} in every reduced word for c , or *underbarred* and marked \underline{i} if s_i occurs after s_{i-1} in every reduced word for c . Passing from c to c^{-1} means swapping overbarring with underbarring.

We say x is obtained from y by a $\bar{2}31 \rightarrow \bar{2}13$ *move* if x is covered by y via a swap in positions i and $i+1$, for some i , and if there exists an *overbarred* element x_j with $j < i$ and $x_i < x_j < x_{i+1}$. Similarly, x is obtained from y by a $3\underline{1}2 \rightarrow 13\underline{2}$ *move* if x is covered by y via a swap in positions i and $i+1$, for some i , and if there exists an *underbarred* element x_j with $i+1 < j$ and $x_i < x_j < x_{i+1}$. Combining [68, Proposition 5.3] and [68, Theorem 6.2], we obtain the following proposition:

Proposition 4.3.1. *Suppose x and y are permutations in S_{n+1} with $x \lessdot y$ in the weak order, and assume that the numbers $2, \dots, n$ have been barred according to c . Then x and y are in the same c -Cambrian congruence class if and only if x is obtained from y by a $\bar{2}31 \rightarrow \bar{2}13$ move or a $3\underline{1}2 \rightarrow 13\underline{2}$ move.*

As an immediate corollary, we see that a permutation y is the bottom element of its c -Cambrian congruence class (i.e. is c -sortable) if and only if none of the permutations covered by y are obtained from y by a $\overline{231} \rightarrow \overline{213}$ move or a $31\underline{2} \rightarrow 13\underline{2}$ move. In other words, there is no subsequence \overline{bca} of y with $a < b < c$, with c immediately preceding a , and with b overbarred and no subsequence cab of y with $a < b < c$, with c immediately preceding a , and with b underbarred. In this case, we say that y *avoids* $\overline{231}$ and $31\underline{2}$.

We can similarly describe bottom elements of c -biCambrian congruence classes (the c -bisortable elements), keeping in mind that passing from c to c^{-1} means swapping overbaring with underbaring: An element y is the bottom element of its c -biCambrian congruence class if and only if none of the permutations covered by y are obtained from y by a $\overline{231} \rightarrow \overline{213}$ or $31\underline{2} \rightarrow 13\underline{2}$ move that is *also* a $\underline{231} \rightarrow \underline{213}$ or $31\overline{2} \rightarrow 13\overline{2}$ move. Thus c -bisortable permutations are described by a complicated pattern-avoidance condition that we will only describe, later, for the case of bipartite c , where it becomes much simpler.

4.3.2 Noncrossing arc diagrams

We now review the notion of *noncrossing arc diagrams* from [72]. Beginning with $n + 1$ distinct points on a vertical line, numbered $1, \dots, n + 1$ from bottom to top, we draw some (or no) curves called *arcs* connecting the points. Each arc moves monotone upwards from one of the points to another, passing either to the left or to the right of each point in between. Furthermore no two arcs may intersect in their interiors, no two arcs share the same upper endpoint, and no two arcs may share the same lower endpoint. We consider arc diagrams only up to their combinatorics, i.e. which pairs of points are joined by an arc and which points are left and right of each arc.

Given a permutation $x_1 \cdots x_{n+1}$ in S_{n+1} , we define a noncrossing arc diagram $\delta(x_1 \cdots x_{n+1})$. Each descent $x_i > x_{i+1}$ becomes an arc α in $\delta(x_1 \cdots x_{n+1})$ with lower endpoint x_{i+1} and upper endpoint x_i . For each integer j with $x_{i+1} < j < x_i$ that occurs to the *left* of x_i in $x_1 \cdots x_{n+1}$, the point j is *left* of the arc α . For each integer j with $x_{i+1} < j < x_i$ that occurs to the *right* of x_{i+1} in $x_1 \cdots x_{n+1}$, the point j is *right* of the arc α . It was shown in [72, Theorem 3.1] that δ is a bijection from permutations to noncrossing arc diagrams. More specifically, for each k , the map δ restricts to a bijection from permutations with k descents to noncrossing arc diagrams with k arcs.

A *c -sortable arc* is an arc that belongs to $\delta(v)$ for some c -sortable permutation v . The following characterization of c -sortable arcs in terms of the barring associated to c is immediate from the pattern-avoidance description above. (Compare [72, Example 4.9].)

Proposition 4.3.2. *For $W = A_n$ and any c , the c -sortable arcs are the arcs that do not pass to the left of any underbarred element of $\{2, \dots, n\}$ and do not pass to the right of any overbarred*

element of $\{2, \dots, n\}$.

In particular, since c and c^{-1} correspond to opposite barrings, the only arcs that are both c and c^{-1} -sortable are the arcs that connect adjacent endpoints i and $i + 1$.

Combining the above descriptions of c -sortable and c -bisortable elements in terms of overbarred and underbarred elements, we obtain the following proposition.

Proposition 4.3.3. *For $W = A_n$ and any c , the map δ restricts to a bijection from c -bisortable permutations with k descents to noncrossing arc diagrams on $n + 1$ vertices with k arcs, each of which is either c or c^{-1} -sortable.*

Proof. Suppose $x = x_1 \cdots x_n$ is a permutation such that $\delta(x)$ has an arc that is neither c -sortable nor c^{-1} -sortable. This arc has upper endpoint x_i and lower endpoint x_{i+1} for some i and it fails the conclusion of Proposition 4.3.2 for c and for c^{-1} . That is, it either passes left of an underbarred element or right of an overbarred element *and* it either passes left of an overbarred element or right of an underbarred element. Thus, switching x_i with x_{i+1} is both a $\bar{2}31 \rightarrow \bar{2}13$ or $31\bar{2} \rightarrow 13\bar{2}$ move *and* a $\underline{2}31 \rightarrow \underline{2}13$ or $31\bar{2} \rightarrow 13\bar{2}$ move. Therefore, x is not c -bisortable. The argument is easily reversed to prove the converse. \square

Alternately, Proposition 4.3.3 follows from the description of the c -biCambrian congruence as the meet of the c -Cambrian and c^{-1} -Cambrian congruences.

4.3.3 Alternating arc diagrams

We now consider the case where c is bipartite. Let c_+ be the product of the simple generators s_i where i is even, and c_- be the product of the simple generators s_i where i is odd. The bipartite Coxeter elements in A_n are c_+c_- and its inverse c_-c_+ . The barring associated to c_+c_- has all even numbers overbarred and all odd numbers underbarred. A ***right-even alternating arc*** is an arc that passes to the right of even vertices and to the left of odd vertices. A ***left-even alternating arc*** is an arc that passes to the left of even vertices and to the right of odd vertices. A ***right-even alternating arc diagram*** is a noncrossing arc diagram all of whose arcs are right-even alternating, and ***left-even alternating arc diagrams*** are defined analogously. The following proposition is an immediate consequence of Proposition 4.3.2.

Proposition 4.3.4. *Suppose $W = A_n$ and c is the bipartite Coxeter element c_+c_- .*

1. *The map δ restricts to a bijection from c -sortable permutations to right-even alternating arc diagrams.*
2. *The map δ restricts to a bijection from c^{-1} -sortable permutations to left-even alternating arc diagrams.*




Figure 4.8: Some alternating noncrossing arc diagrams

In each case, δ restricts further to send permutations with k descents bijectively to arc diagrams with k arcs.

An **alternating arc** is an arc that is either right-even alternating or left-even alternating or both. We call a noncrossing arc diagram consisting of alternating arcs an **alternating arc diagram**. Figure 4.8 shows several alternating noncrossing arc diagrams. From left to right, they are $\delta(5371624)$, $\delta(6473125)$, and $\delta(4275136)$. The following proposition is the bipartite case of Proposition 4.3.3.

Proposition 4.3.5. *For $W = A_n$ and c bipartite, the map δ restricts to a bijection from c -bisortable permutations with k descents to alternating arc diagrams on $n+1$ points with k arcs.*

Observe that an arc fails to be alternating if and only if it passes on the same side of two consecutive numbers. Thus, we obtain the following simpler description of the pattern avoidance condition defining bipartite c -bisortable elements.

Proposition 4.3.6. *If c is the bipartite Coxeter element c_+c_- of A_n , a permutation $x = x_1 \cdots x_{n+1}$ is c -bisortable if and only if, for every descent $x_i > x_{i+1}$, there exists no k with $x_{i+1} < k < k+1 < x_i$ such that k and $k+1$ are on the same side of the descent (i.e. k and $k+1$ both left of x_i or both right of x_{i+1}).*

The condition in Proposition 4.3.6 is that x avoids subsequences $dabc$, $dacb$, $bcda$, and $cbda$ with $a < b < c < d$, with d and a adjacent in *position*, and with b and c being adjacent in *value*. This is an instance of **bivincular** pattern avoidance in the sense of [15, Section 2]. We will not review the notation for bivincular patterns from [15], but we restate Proposition 4.3.6 in that notation as follows:

Proposition 4.3.7. *For c bipartite, a permutation is c -bisortable if and only if it avoids the bivincular patterns $(2341, \{3\}, \{2\})$, $(3241, \{3\}, \{2\})$, $(4123, \{1\}, \{2\})$, and $(4132, \{1\}, \{2\})$.*

4.3.4 Counting alternating arc diagrams

Let $[n]$ denote the set $\{1, 2, \dots, n\}$. To prove the type-A enumeration of bipartite c -bisortable elements in Theorem 4.1.1, we give a bijection π from noncrossing alternating arc diagrams on $n + 1$ vertices with k arcs to pairs (S, T) of subsets of $[n]$ with $|S| = |T| = k$.

To describe the bijection, we begin with the case $k = 1$. Recall that we number the endpoints in a diagram $1, \dots, n + 1$ from bottom to top. Suppose Σ is an alternating arc diagram whose only arc connects i to j with $i < j$. If the arc is right-even alternating, define $\pi(\Sigma)$ to be $(\{i\}, \{j - 1\})$. If the arc is left-even alternating, define $\pi(\Sigma)$ to be $(\{j - 1\}, \{i\})$. (Any arc that is both right-even alternating and left-even alternating has $j = i + 1$. The bijection sends this arc to $(\{i\}, \{i\})$.)

Now suppose that Σ is a noncrossing arc diagram with more than one arc. Whenever we encounter a right-even alternating arc in Σ with endpoints $i < j$, we put i into S and $j - 1$ into T ; whenever we encounter a left-even alternating arc with endpoints $i < j$ we put $j - 1$ into S and i into T . More precisely, suppose that Σ is an alternating arc diagram with k arcs. Let S' denote the set of numbers i such that i is bottom endpoint of a right-even alternating arc in Σ and let S'' denote the set of numbers $j - 1$ such that j is the top endpoint of a left-even alternating arc in Σ . Let T' denote the set of numbers $j' - 1$ such that j' is the top endpoint of a right-even alternating arc in Σ and let T'' denote the set of numbers i' such that i' is the bottom endpoint of a left-even alternating arc. The map π sends Σ to the pair $(S' \cup S'', T' \cup T'')$.

Theorem 4.3.8. *The map π is a bijection from the set of alternating arc diagrams on $n + 1$ points to the set of pairs of subsets of $[n]$ of the same size. For each k , the bijection restricts to a bijection from alternating arc diagrams with k arcs to pairs of subsets of size k .*

In preparation for the proof of Theorem 4.3.8, we will break each alternating diagram into smaller pieces. Two alternating arcs with endpoints $i < j$ and $i' < j'$ **overlap** if the intersection of the sets $\{i, \dots, j - 1\}$ and $\{i', \dots, j' - 1\}$ is nonempty. Informally, the arcs overlap if some part of one arc passes along side of the other arc. (If they only touch at their endpoints but don't pass along side one another, then they do *not* overlap). Given a collection \mathcal{E} of arcs, we can define an “overlap graph” with vertices \mathcal{E} and edges given by overlapping pairs in \mathcal{E} . We say that the collection \mathcal{E} is **overlapping** if this overlap graph is connected. Each noncrossing diagram can be broken into overlapping components, maximal overlapping collections of arcs. The definition of alternating arc diagrams and the definition of right-even and left-even alternating arcs let us immediately conclude that two distinct arcs appearing in the same alternating arc diagram, one right-even alternating and one left-even alternating, cannot overlap. We have proved the following fact.

Proposition 4.3.9. *Each overlapping component of an alternating arc diagram fits exactly one of the following descriptions: (1) It consists of right-even alternating arcs that are not left-even*

alternating; (2) It consists of left-even alternating arcs that are not right-even alternating; or (3) it consists of a single arc that is right-even and left-even alternating (and thus connects two adjacent points).

Proposition 4.3.9 implies that, on each overlapping component, the map π collects all of the top endpoints of the arcs into one set, and all of the bottom endpoints into the other set.

Now we describe how to break an alternating diagram Σ into its overlapping components. Let $P(\Sigma)$ be the set of numbers $p \in [n+1]$ such that no arc in Σ passes left or right of p . (A point $p \in P(\Sigma)$ may still be an endpoint of one or two arcs.) Write $P(\Sigma) = \{p_0, \dots, p_m\}$ with $p_0 < \dots < p_m$. In every case, $p_0 = 1$ and $p_m = n+1$. For each i , we claim that an arc in Σ has its lower endpoint in $\{p_{i-1}, p_{i-1}+1, \dots, p_i-1\}$ if and only if it has its upper endpoint in the set $\{p_{i-1}+1, p_{i-1}+2, \dots, p_i\}$. Indeed, if an arc has a lower endpoint in $\{p_{i-1}, p_{i-1}+1, \dots, p_i-1\}$, then since it cannot pass on either side of p_i , it must end at a number in $\{p_{i-1}+1, \dots, p_i\}$. A similar argument proves the converse, so we have established the claim. Let Σ_i denote the set of arcs with lower endpoints in $\{p_{i-1}, p_{i-1}+1, \dots, p_i-1\}$ (and thus with upper endpoints in the set $\{p_{i-1}+1, p_{i-1}+2, \dots, p_i\}$). By construction, Σ_i is an overlapping component, and all overlapping components are Σ_i for some i . Let (S_i, T_i) be the image of Σ_i under π , so that $\pi(\Sigma) = (\bigcup_{i=1}^m S_i, \bigcup_{i=1}^m T_i)$.

We say that two arcs are *compatible* if there is a noncrossing arc diagram containing both arcs. Our next task is to understand for which pairs (s, t) and (s', t') there exists an overlapping pair of compatible *alternating arcs*, one with endpoints s and $t+1$ and one with endpoints s' and $t'+1$. Since the arcs must overlap but may not share the same bottom endpoint and may not share the same top endpoint, and taking without loss of generality $s < s'$, there are only two cases. These cases are covered by the following two lemmas, which are easily verified.

Lemma 4.3.10. *Suppose $s < s' \leq t < t'$. Then there exist two compatible alternating arcs, one with endpoints s and $t+1$ and one with endpoints s' and $t'+1$ if and only if s' and t have the same parity. The pair of arcs can be chosen in exactly two ways, either both as right-even alternating arcs or both as left-even alternating arcs.*

Lemma 4.3.11. *Suppose $s < s' < t' < t$. Then there exist two compatible alternating arcs, one with endpoints s and $t+1$ and one with endpoints s' and $t'+1$ if and only if s' and t' have opposite parity. The pair of arcs can be chosen in exactly two ways, either both as right-even alternating arcs or both as left-even alternating arcs.*

Given a pair (S, T) of k -subsets of $[n]$, we will always write $S = \{s_1, \dots, s_k\}$ with $s_1 < \dots < s_k$ and $T = \{t_1, \dots, t_k\}$ with $t_1 < \dots < t_k$. Define $Q(S, T)$ to be the set of numbers $q \in [n+1]$ such that, for all j from 1 to k , neither $s_j < q \leq t_j$, nor $t_j < q \leq s_j$.

Lemma 4.3.12. *Let Σ be an alternating arc diagram. Then $Q(\pi(\Sigma)) = P(\Sigma)$.*

Proof. Write (S, T) for $\pi(\Sigma)$. If $p \in P(\Sigma)$, then no arc passes left or right of p . Thus there exists k such that s_j and t_j are less than p for all $j \leq k$ and s_j and t_j are greater than or equal to p for all $j > k$. We see that $p \in Q(S, T)$.

Suppose that $q \in (Q, S)$, and there exists some arc α that passes to the left or right of q . The arc α belongs to some overlapping component of Σ , and each pair s_i, t_i in the image of a different component satisfies $s_i, t_i < q$ or $s_i, t_i > q$. Thus, we may as well assume that Σ consists of a single overlapping component. Write $\pi(\Sigma) = (\{s_1, \dots, s_k\}, \{t_1, \dots, t_k\})$ with $s_1 < \dots < s_k$ and $t_1 < \dots < t_k$. Lemma 4.3.9 says that Σ consists of either right-even overlapping arcs or left-even overlapping arcs. Without loss of generality, we assume that Σ consists of only right-even overlapping arcs, so that $\{s_1, \dots, s_k\}$ is the set of bottom endpoints of those arcs. Thus, $s_i \leq t_i$ for each $i = 1, 2, \dots, k$. Let s_l be the bottom endpoint of α , and let l be the largest number such that $s_l < q$. We make two observations. First, α must connect s_i with $t_j + 1$, where j is strictly greater than i (otherwise $s_i < q < t_j + 1 \leq t_i$), and j is strictly greater than l (otherwise $s_j < q \leq t_j$). Second, $t_{l+1} \geq q > t_l$, because $t_{l+1} \geq s_{l+1} \geq q > t_l \geq s_l$. We conclude that each number in the set of bottom endpoints $\{s_{l+1}, s_{l+2}, \dots, s_k\}$ must connect with a number in the set $\{t_{l+1} + 1, \dots, t_k + 1\}$. Since $t_j + 1$ is already connected to s_i , there is some number in the set $\{t_{l+1} + 1, \dots, t_k + 1\}$ that is the top endpoint of two arcs, and that is a contradiction. \square

We are now prepared to prove the main theorem of this section.

Proof of Theorem 4.3.8. We first show that π is well-defined. Since each arc in Σ contributes exactly one of its endpoints to $S' \cup S''$ and the other to $T' \cup T''$, both $S' \cup S''$ and $T' \cup T''$ have size k as long as each contribution to $S' \cup S''$ is distinct and each contribution to $T' \cup T''$ is distinct. Each contribution to S' is distinct because no two arcs share the same lower endpoint, and each contribution to S'' is distinct because no two arcs share the same upper endpoint. Proposition 4.3.9 implies that a right-even alternating arc with bottom endpoint i and a *distinct* left-even alternating arc with top endpoint $i + 1$ are not compatible. Thus the only elements of $S' \cap S''$ come from arcs that are both right-even alternating and left-even alternating, and we see that each contribution to $S' \cup S''$ is distinct. The symmetric argument shows that each contribution to $T' \cup T''$ is distinct. We have shown that π is a well-defined map from alternating arc diagrams with k arcs to pairs of k -element subsets of $[n]$.

We complete the proof by exhibiting an inverse η to π . Let (S, T) be a pair of k -element subsets of $[n]$. Write $Q(S, T) = \{q_0, \dots, q_m\}$ with $q_0 < \dots < q_m$. For each i from 1 to m , define $S_i = S \cap \{q_{i-1}, q_{i-1} + 1, \dots, q_i - 1\}$ and $T_i = T \cap \{q_{i-1}, q_{i-1} + 1, \dots, q_i - 1\}$. We claim that $|S_i| = |T_i|$, and more specifically, that $s_j \in S_i$ if and only if $t_j \in T_i$. Indeed, suppose $s_j \in S_i$, so that $q_{i-1} \leq s_j < q_i$. If $t_j < q_{i-1}$, then $t_j < q_{i-1} \leq s_j$, contradicting the fact that $q_{i-1} \in Q(S, T)$. If $t_j \geq q_i$, then $s_j < q_i \leq t_j$, contradicting the fact that $q_i \in Q(S, T)$. We conclude that $t_j \in T_i$. The symmetric argument completes the proof of the claim.

Now, in light of Lemma 4.3.12 and the definition of π , by subtracting $q_{i-1} - 1$ from each element of S_i and T_i , we reduce to the case where $m = 1$ and thus $Q = \{1, n + 1\}$ and $(S_1, T_1) = (S, T)$. In particular, all of the arcs in the diagram $\eta(S, T)$ are right-even alternating, or all of the arcs are left-even alternating. If $n = 1$, then either $(S, T) = (\emptyset, \emptyset)$, in which case $\eta(S, T)$ has no arc, or $(S, T) = (\{1\}, \{1\})$, in which case $\eta(S, T)$ has an arc connecting 1 and 2.

If $n > 1$, then we observe that the element 1 must be in S or in T but must not be in both. Indeed, if 1 is in neither set or in both, we see that $2 \in Q(S, T)$, and this is a contradiction. In particular, we will need to construct an arc whose lower endpoint is 1 and whose upper endpoint is above 2. This arc will pass by 2, and so it is either right-even alternating or left-even alternating (but not both). If $1 \in S$, then the corresponding arc is right-even alternating, and if $1 \in T$ this arc is left-even alternating. Without loss of generality, we assume $1 \in S$, so that each i in S is a bottom endpoint and for each j in T , $j + 1$ is a top endpoint of a right-even alternating arc in $\eta(S, T)$. To complete the proof, we show that there is a unique way to pair off each bottom endpoint in S with a top endpoint in T so that the union of the resulting arcs is a noncrossing arc diagram. Since the arcs in the diagram are all right-even alternating, we must pair each element of S with a larger element of T .

We first decide which element of T we should pair with s_k . Because s_k is the maximum element of S , Lemma 4.3.10 implies that we must pair s_k with some t' such that the set $\{t \in T : s_k < t < t', t - s_k \text{ odd}\}$ is empty. Similarly, Lemma 4.3.11 implies that we must either pair s_k with t_k or pair s_k with some t' such that $t' - s_k$ is odd. Furthermore, if we choose t' according to those two rules, no matter how we pair the remaining elements of S and T , the arcs produced will be compatible with the arc whose bottom endpoint is s_k . We are forced to pair s_k with $\min\{t \in T : t \geq s_k, t - s_k \text{ odd}\}$, or with t_k if $\{t \in T : t \geq s_k, t - s_k \text{ odd}\} = \emptyset$. By induction on k , there is a unique way to pair the elements of $S \setminus \{s_k\}$ with the elements of $T \setminus \{t'\}$ to make a noncrossing alternating diagram. Putting in the pair (s_k, t') we obtain the unique pairing of elements of S with elements of T to make a noncrossing alternating diagram. The base of the induction is where $k = 1$. Here existence of a pairing is trivial and uniqueness comes from the requirement that the arc whose bottom endpoint is 1 must be right-even alternating. \square

Remark 4.3.13. The proof of Theorem 4.3.8 provides key insights that lead to our proof of Theorems 4.1.2 and 4.1.3. When we generalize beyond type A, the role of the arcs in an alternating arc diagram will be played by the *canonical joinands* of a bipartite c -bisortable element. (The latter are defined in Section 4.4.3. For the connection between arcs and canonical join representations, see [72, Section 3].) The fact that distinct right-even alternating and left-even alternating arcs do not overlap translates into the fact that, for c bipartite, c and c^{-1} -sortable join-irreducible permutations have disjoint support—a fact that we will prove uniformly in Section 4.4.5.

In light of the proof of Theorem 4.3.8, we might count alternating arc diagrams Σ with $n+1$ points in the following way: First, choose the set $P(\Sigma) = \{p_0, \dots, p_m\}$ with $p_0 < \dots < p_m$. When $p_{i+1} = p_i + 1$, choose either to connect p_i to $p_i + 1$ with an arc or not. When $p_{i+1} > p_i + 1$, choose either to use right-even alternating or left-even alternating arcs, and construct a diagram on the points p_i, \dots, p_{i+1} , such that for every j with $p_i + 1 \leq j \leq p_{i+1} - 1$, some arc passes left or right of j . Once we fix the type of arc (right-even or left-even alternating), the number of such diagrams on p_i, \dots, p_{i+1} is the number that in Section 4.4 will be called the double-positive Catalan number $\text{Cat}^+(A_m)$ for $m = p_{i+1} - p_i$. The proof of Theorems 4.1.2 and 4.1.3 generalizes this method of counting and shows that a corresponding method also counts antichains in the doubled root poset.

Remark 4.3.14. Looking ahead to Section 4.4, the previous remark implies an interpretation of the type-A double-positive Narayana number which—after some combinatorial manipulations that amount to changing from a bipartite Coxeter element to a linear Coxeter element—coincides with the interpretation given in [7, Theorem 1.1].

4.3.5 Enumerating bipartite c -bisortable elements in type B

In this section, we use certain alternating arc diagrams to prove the enumeration of bipartite c -bisortable elements of type B given in Theorem 4.1.1. We first analyze the c -biCambrian congruence on the weak order for B_n . In order to reuse much of our work from Section 4.3.4, we realize the weak order on B_n as a sublattice of the weak order on A_{2n-1} , through the usual signed-permutation model.

In any finite Coxeter group, the map $y \mapsto w_0 y w_0$ is a rank-preserving automorphism of weak order (where w_0 is the longest element). It is a well-known fact (and an easy exercise) that for any lattice automorphism, the set of fixed points of the automorphism is a sublattice. When W is A_{2n-1} it is easy to check that the set of fixed points of this map is a Coxeter group isomorphic to B_n . (For example, take as simple generators the fixed points $s_i s_{2n-1-i}$ for $i < n$, and s_n .)

Writing each x in A_{2n-1} as a permutation of the set $\{\pm 1, \pm 2, \dots, \pm n\}$ with full one-line notation $x_{-n} x_{-n+1} \dots x_{-1} x_1 \dots x_{n-1} x_n$, conjugation by w_0 acts by negating all of the entries of x and reversing its order. The fixed points of this automorphism are the signed permutations on $\{\pm 1, \pm 2, \dots, \pm n\}$, meaning the permutations which satisfy $x_i = -x_{-i}$. The subposet of the weak order on A_{2n-1} induced by the signed permutation is a sublattice, isomorphic to the weak order on B_n . Because $x_i = -x_{-i}$, it is convenient to write signed permutations in an abbreviated notation as $x_1 x_2 \dots x_n$.

It is easy to check that the signed permutation $y_1 \dots y_n$ (written in abbreviated notation) covers $x_1 \dots x_n$ in the weak order on B_n if and only if one of the two following conditions is satisfied: Either $y_i = x_{i+1} > x_i = y_{i+1}$ for $i, i+1 \in [n]$ and $y_j = x_j$ for each $j \notin \{i, i+1\}$, or

$0 < x_1 = -y_1$ and $x_j = y_j$ for all $j \in \{2, 3, \dots, n\}$. In the former case, the symmetry $y_i = -y_{-i}$ implies that $y_{-i-1} = x_{-i} > x_{-i-1} = y_{-i}$, so that the full one-line notation of y_{-n}, \dots, y_n has two descents: $y_i > y_{i+1}$ and $y_{-i-1} > y_{-i}$. (For more information on this realization of the weak order on the type- B Coxeter group see [12, Section 8.1]).

To define noncrossing diagrams of type B , we place $2n$ points on a vertical line, labeled from bottom to top by the integers $-n, -n+1, \dots, -1, 1, \dots, n-1, n$ such that there is a central symmetry that, for each i , maps the point labeled i to the point labeled $-i$. A **centrally symmetric noncrossing diagram** is a noncrossing arc diagram that is fixed by the central symmetry. The map δ restricts to a bijection from signed permutations to centrally symmetric noncrossing diagrams. We use the term **centrally symmetric arc** to describe either an arc that is fixed by the central symmetry or a pair of arcs that form an orbit under the symmetry. For each k , the map δ restricts further to a bijection between signed permutations with k descents and centrally symmetric noncrossing diagrams with k centrally symmetric arcs.

The simple generators of B_n , are $s_0 = (-1 \ 1)$ and $s_i = (-i-1 \ -i)(i \ i+1)$ for $i = 1, \dots, n-1$, written in cycle notation as permutations of $\{\pm 1, \dots, \pm n\}$. A **symmetric Coxeter element** of A_{2n-1} is a Coxeter element that is fixed by the automorphism $y \mapsto w_0 y w_0$. Equivalently, the Coxeter element can be written as a product of some permutation of the elements s_0, \dots, s_{n-1} defined above. This product in A_{2n-1} can be interpreted as a Coxeter element of B_n , which we denote by \tilde{c} . A Coxeter element is symmetric if and only if it corresponds to a barring of $\{\pm 1, \dots, \pm(n-1)\}$ with the property that i is overbarred if and only if $-i$ is underbarred. Thus, a signed permutation avoids the pattern $\bar{2}31$ if and only if it also avoids the pattern $31\bar{2}$ (in its full one-line notation). The signed permutations avoiding $\bar{2}31$ (and equivalently $31\bar{2}$) in their full notation are exactly the \tilde{c} -sortable elements by [68, Theorem 7.5]. Comparing with the description of c -sortable permutations following Proposition 4.3.1, we obtain the following proposition.

Proposition 4.3.15. *Suppose c is a symmetric Coxeter element of A_{2n-1} and suppose \tilde{c} is the corresponding Coxeter element of B_n . A signed permutation is \tilde{c} -sortable in B_n if and only if it is c -sortable as an element of A_{2n-1} .*

The analogous result holds for \tilde{c} -bisortable elements.

Proposition 4.3.16. *Suppose c is a symmetric Coxeter element of A_{2n-1} and suppose \tilde{c} is the corresponding Coxeter element of B_n . A signed permutation is \tilde{c} -bisortable in B_n if and only if it is c -bisortable as an element of A_{2n-1} .*

Proof. Suppose w is a signed permutation. If w is \tilde{c} -bisortable, then Proposition 4.2.15 says that $w = u \vee v$ for some \tilde{c} -sortable signed permutation u and some \tilde{c}^{-1} -sortable signed permutation v . Proposition 4.3.15 says that, as elements of A_{2n-1} , u is a c -sortable permutation and v is a c^{-1} -sortable permutation. Since the weak order on B_n is a sublattice of the weak order on A_{2n-1} ,

the join $u \vee v$ is the same in A_{2n-1} as in B_n , and thus Proposition 4.2.15 implies that w is c -bisortable.

On the other hand, if w is c -bisortable as an element of A_{2n-1} , then as in Proposition 4.2.15, we can write w as $u \vee v$, where u is the c -sortable permutation $\pi_{\downarrow}^c(w)$ and v is the c^{-1} -sortable permutation $\pi_{\downarrow}^{c^{-1}}(w)$. Since conjugation by w_0 is a lattice automorphism fixing w , we obtain $w = (w_0 u w_0) \vee (w_0 v w_0)$. But $w_0 u w_0$ is c -sortable and below w , so $w_0 u w_0 \leq u$. Since conjugation by w_0 is order preserving, we conclude that $w_0 u w_0 = u$. Similarly $w_0 v w_0 = v$. Thus, by Proposition 4.3.15, u is \tilde{c} -sortable and v is \tilde{c}^{-1} -sortable in B_n . Since the weak order on B_n is a sublattice of the weak order on A_{2n-1} , Proposition 4.2.15 says that w is \tilde{c} -bisortable. \square

A bipartite Coxeter element \tilde{c} of B_n is a symmetric, bipartite Coxeter element of A_{2n-1} , so combining Propositions 4.3.5 and 4.3.16, we immediately obtain the following proposition.

Proposition 4.3.17. *For $W = B_n$ and \tilde{c} a bipartite Coxeter element, the map δ restricts to a bijection from \tilde{c} -bisortable signed permutations with k descents to centrally symmetric alternating arc diagrams on $2n$ points with k centrally symmetric alternating arcs.*

Thus, to count the bipartite c -bisortable elements in B_n , it remains only to count centrally symmetric alternating arc diagrams. The points in the noncrossing arc diagram for a permutation in S_n are labeled $1, \dots, 2n$ from bottom to top. If we instead label the points $-n, \dots, -1, 1, \dots, n$ from bottom to top, we can interpret the map π as returning an ordered pair of subsets of $\{-n, \dots, -1, 1, \dots, n-1\}$. Define π_B to be the map on centrally symmetric alternating arc diagrams with $2n$ vertices that first does the map π to obtain (S, T) and then ignores T and outputs only S . The following theorem shows that the number of centrally symmetric alternating arc diagrams with k centrally symmetric arcs is $\binom{2n-1}{2k} + \binom{2n-1}{2k-1} = \binom{2n}{2k}$ as desired.

Theorem 4.3.18. *For each k , the map π_B restricts to a bijection from centrally symmetric alternating arc diagrams with k centrally symmetric arcs to subsets of $\{-n, \dots, -1, 1, \dots, n-1\}$ of size $2k$ or $2k-1$.*

Proof. We first show that π_B is a bijection from centrally symmetric alternating arc diagrams to subsets of $\{-n, \dots, -1, 1, \dots, n-1\}$. Given $S \subseteq \{\pm 1, \dots, \pm n\}$, we write $-S-1$ for the set $\{-i-1 : i \in S\}$, where we interpret $1-1$ to mean -1 in order to make $-S-1$ a subset of $\{\pm 1, \dots, \pm n\}$. Showing that π_B is a bijection is equivalent to showing that an alternating diagram Σ is centrally symmetric if and only if $\pi(\Sigma) = (S, -S-1)$ for some S .

The terms “right-even alternating” and “left-even alternating” should be understood in terms of the labeling of points as $1, \dots, 2n$. These terms become problematic when we label points as $-n, \dots, -1, 1, \dots, n$. (For example, whether a right-even alternating arc passes left or

right of the point labeled i depends on the sign of i , the parity of i , and the parity of n .) Without worrying about these details, we make two easy observations: First, an alternating arc is right-even alternating if and only if its image under the central symmetry is right-even alternating. Second, the central symmetry swaps top with bottom endpoints and positive with negative endpoints. These observations immediately imply that π maps centrally symmetric alternating arc diagrams to pairs of the form $(S, -S - 1)$.

These observations also immediately imply that if π maps an alternating arc diagram Σ to (S, T) and Σ' is the image of Σ under the central symmetry, then π maps Σ' to $(-T + 1, -S - 1)$, where $-T + 1$ is the set $\{-i + 1 : i \in T\}$, where we interpret $-1 + 1$ to mean 1. In particular, if π maps Σ to $(S, -S - 1)$, then π also maps Σ' to $(S, -S - 1)$. Since we already know that π is a bijection, we conclude that in this case Σ must be centrally symmetric. We have shown that Σ is centrally symmetric if and only if $\pi(\Sigma)$ is of the form $(S, -S - 1)$. Therefore π_B is a bijection.

It is now immediate that π_B maps a centrally symmetric alternating arc diagrams with k centrally symmetric arcs to a $(2k - 1)$ -element set if the diagram has an arc that is fixed by the central symmetry or to a $2k$ -element set if all of the arcs in the diagram come in symmetric pairs. \square

4.3.6 Simpliciality of the bipartite biCambrian fan in types A and B

We now prove Theorem 4.2.12, which states that the bipartite biCambrian fan is simplicial in types A and B. The proof of the type-A case of Theorem 4.2.12 proceeds by combining results of [52] and [72].

Some collections of noncrossing arc diagrams (including, we will see, the alternating arc diagrams), correspond to lattice quotients of the weak order. More specifically, a collection of noncrossing arc diagrams may be the image, under δ , of the bottom elements of congruence classes of some congruence. To describe when and how such a situation arises, we need the notion of a subarc. For $i < j$ and $i' < j'$, an arc α connecting i to j is a **subarc** of an arc α' connecting i' to j' if $i' \leq i$ and $j' \geq j$ and if α and α' pass to the same side of every point between i and j . It follows from [72, Theorem 4.1] and [72, Theorem 4.4] that a subset D of the noncrossing arc diagrams on $n + 1$ points is the image, under δ , of the set of bottom elements for some congruence Θ if and only if all of the following conditions hold.

- (i) There exists a set U of arcs such that a noncrossing diagram Σ is in D if and only if all arcs in Σ are in U .
- (ii) If an arc α is not in U and α is a subarc of some arc α' , then α' is also not in U .

We will call U the set of **unremoved arcs** of the congruence Θ . If C is any set of arcs and U is the maximal set such that $U \cap C = \emptyset$ and condition (4.3.6) above holds, then we say that the

congruence Θ is *generated by removing the arcs* C .

An element j of a finite lattice L is join-irreducible if it covers exactly one element j_* . A lattice congruence on L *contracts* a join-irreducible element j if the congruence has $j \equiv j_*$. A congruence is uniquely determined by the set of join-irreducible elements it contracts. The join-irreducible elements of the weak order on A_n are the permutations in S_{n+1} with exactly one descent. In particular, the map δ restricts to a bijection between join-irreducible elements in S_{n+1} and noncrossing arc diagrams with exactly one arc. (We will think of this restriction as mapping join-irreducible elements to arcs, rather than to singletons of arcs.) Under this bijection, the join-irreducible elements *not* contracted by a congruence Θ correspond to the arcs in U , where U is the set of unremoved arcs of Θ . The congruence is *generated by contracting* a set J of join-irreducible elements if and only if it is generated by removing the arcs $\delta(J)$.

We call j a *double join-irreducible* element if it is join-irreducible and if the unique element j_* covered by j is either the bottom element of the lattice or is itself join-irreducible. The following is part of the main result of [52].

Theorem 4.3.19. *Suppose Θ is a lattice congruence on the weak order on A_n . Then the following three conditions are equivalent.*

- (i) *The undirected Hasse diagram of the quotient lattice A_n/Θ is a regular graph.*
- (ii) *$\mathcal{F}_\Theta(A_n)$ is a simplicial fan.*
- (iii) *Θ is generated by contracting a set of double join-irreducible elements.*

We now apply these considerations to alternating arc diagrams. First, it is apparent that the set of alternating arc diagrams is the image of δ restricted to the set of bottom elements of a congruence. (Indeed, this is the bipartite c -biCambrian congruence.) It is also apparent that the congruence is generated by removing the arcs that connect i to $i+3$ and that *do not* alternate. (That is they pass to the same side of $i+1$ and $i+2$.) Applying the inverse of δ , we see that the congruence is generated by contracting the join-irreducible elements

$$1 \cdots (i-1)(i+1)(i+2)(i+3)i(i+4) \cdots (n+1)$$

and

$$1 \cdots (i-1)(i+3)i(i+1)(i+2)(i+4) \cdots (n+1)$$

for $i = 1, \dots, n-2$. These are both double join-irreducible elements, and thus Theorem 4.3.19 implies the type-A case of Theorem 4.2.12.

We now move to the type-B case of Theorem 4.2.12. Just as in type-A, there is a correspondence between congruences on the weak order and certain sets of (centrally symmetric)

noncrossing diagrams. However, there is currently no analogue to Theorem 4.3.19 in type B. Therefore, instead of arguing the type-B case as we argued the type-A case, we will use a folding argument to show that the type-A case implies the type-B case.

Say a lattice congruence of the weak order on A_{2n-1} is **symmetric under conjugation** by w_0 if for all $x, y \in A_{2n-1}$ we have $x \equiv y$ modulo Θ if and only if $w_0 x w_0 \equiv w_0 y w_0$ modulo Θ .

Proposition 4.3.20. *If Θ is a lattice congruence of the weak order on A_{2n-1} that is symmetric under conjugation by w_0 , then its restriction to the sublattice B_n is a congruence Θ' . An element of B_n is the bottom element of its Θ' -class if and only if it is the bottom element of its Θ -class.*

Proof. It is also a well-known and easy fact that the restriction of a lattice congruence to any sublattice is a congruence on the sublattice, and the first assertion of the proposition follows. One implication in the second assertion is immediate. For the other implication, suppose $x \in B_n$ is the bottom element of its Θ' -class and let $y = \pi_{\downarrow}^{\Theta}(x)$, so that in particular $x \equiv y$ modulo Θ . Then because Θ is symmetric under conjugation by w_0 , also $x = w_0 x w_0 \equiv w_0 y w_0$ modulo Θ . Since y is the bottom element of its Θ -class, $y \leq w_0 y w_0$. Since conjugation by w_0 is order preserving, also $w_0 y w_0 \leq y$, so $y = w_0 y w_0$. Thus y is in the Θ' -class of x , and we conclude that $y = x$, so that x is also the bottom element of its Θ -class. \square

Proposition 4.3.21. *Suppose that Θ is a lattice congruence of the weak order on A_{2n-1} and let Θ' denote its restriction to the weak order on B_n . If $\mathcal{F}_{\Theta}(A_{2n-1})$ is simplicial and Θ is symmetric under conjugation by w_0 , then $\mathcal{F}_{\Theta'}(B_n)$ is simplicial.*

Before we proceed with the proof of Proposition 4.3.21 we define some useful terminology. Recall that there is a linear functional λ that orients the adjacency graph on maximal cones in $\mathcal{F}(W)$ to yield a partial order isomorphic to the weak order on W . A facet of a maximal cone is a **lower wall** (with respect to λ) if passing through it to an adjacent maximal cone is the same as moving down by a cover in the weak order. **Upper walls** are defined dually. The maximal cones of $\mathcal{F}_{\Theta}(W)$ similarly have lower and upper walls with respect to λ ; passing from one cone to an adjacent cone through a lower wall corresponds to moving down by a cover in the lattice quotient induced by Θ . The lower walls of a maximal cone in $\mathcal{F}_{\Theta}(W)$ are the lower walls of the smallest element in the corresponding Θ -congruence class. (Recall that each maximal cone in $\mathcal{F}_{\Theta}(W)$ is the union of the set of maximal cones in $\mathcal{F}(W)$ in the same Θ -congruence class.) Dually, the upper walls of a maximal cone in $\mathcal{F}_{\Theta}(W)$ are the upper walls of the cone corresponding to the largest element in the Θ -congruence class.

Proof of Proposition 4.3.21. We begin by considering type A_{2n-1} in the usual geometric representation in \mathbb{R}^{2n} . However, to prepare for the type-B construction, we index the standard unit basis vectors of \mathbb{R}^{2n} as $-n, \dots, -1, 1, \dots, n$. In this representation, there is a reflecting hyperplane H_{ji} , with normal vector $e_j - e_i$, for each $i < j$ with $i, j \in \{\pm 1, \dots, \pm n\}$. The maximal

cone corresponding to the permutation $x_{-n}\cdots x_{-1}x_1\cdots x_n$ has a lower (respectively upper) wall contained in H_{ji} if and only if there exists $r \in \{-n, \dots, -1, 1, \dots, n-1\}$ such that $x_r = j$ and $x_{r+1} = i$ (respectively, $x_{r+1} = j$ and $x_r = i$). As the price for our choice of indices, when $r = -1$, we must interpret $r + 1$ here to mean 1.

Recall that the signed permutations of B_n are exactly the permutations in A_{2n-1} that are fixed under conjugation by w_0 and that the restriction of weak order to these w_0 -fixed permutations is weak order on B_n . As an abuse of terminology, the linear map on \mathbb{R}^{2n} that sends each vector $(v_{-n}, \dots, v_{-1}, v_1, \dots, v_n)$ to $-(v_n, \dots, v_1, v_{-1}, \dots, v_{-n})$ will be called the conjugation action of w_0 on \mathbb{R}^{2n} . Let L be the linear subspace of \mathbb{R}^{2n} consisting of vectors fixed by this action. These are the vectors with $v_i = -v_{-i}$ for all i . A permutation in A_{2n-1} is fixed under conjugation by w_0 if and only if its corresponding cone in $\mathcal{F}(A_{2n-1})$ intersects L in its relative interior, in which case the cone is also fixed under conjugation by w_0 . Thus, we obtain $\mathcal{F}(B_n)$ as the fan induced on L by $\mathcal{F}(A_{2n-1})$, and the weak order on B_n arises from that induced fan, ordered by the same linear functional λ as $\mathcal{F}(A_{2n-1})$. Moreover, $\mathcal{F}_{\Theta'}(B_n)$ is the fan induced on L by $\mathcal{F}_{\Theta}(A_{2n-1})$.

Almost all of the lower walls of a w_0 -fixed maximal cone C in $\mathcal{F}_{\Theta}(A_{2n-1})$ intersect L in pairs. Specifically, Proposition 4.3.20 implies that any such cone is associated to a signed permutation $x = x_{-n}\cdots x_{-1}x_1\cdots x_n$ that is the bottom element of its Θ -class. A descent $x_{-1}x_1$ of x contributes a single lower wall to C , and thus a single lower wall to $C \cap L$. We will say that such a lower wall is centrally symmetric. All other descents of x come in symmetric pairs $x_{-i-1}x_{-i}$ and $x_i x_{i+1}$, contributing two lower walls to C . However, these two walls have the same intersection with L and thus contribute only one lower wall to $C \cap L$. Similar dual statements hold for the upper walls. Most importantly, among all of the walls of $C \cap L$, there are at most two that are centrally symmetric: at most one among the set of lower walls, and at most one among set of upper walls.

Since $\mathcal{F}_{\Theta}(A_{2n-1})$ is simplicial, C has an odd number of walls. In particular, this implies that among all of the walls for C , there is exactly one that is centrally symmetric wall. Suppose that this wall is a lower wall. Then, C has an odd number of lower walls, say $2k - 1$, and their intersection with L yields k lower walls for the corresponding cone $C \cap L$ in $\mathcal{F}_{\Theta'}(B_n)$. Since $\mathcal{F}_{\Theta}(A_{2n-1})$ is simplicial, there are $2n - 2k$ upper walls, which intersect L in pairs, to form $n - k$ upper walls in $\mathcal{F}_{\Theta'}(B_n)$. Thus the cone associated to C in $\mathcal{F}_{\Theta'}(B_n)$ has a total of n walls. The same argument (switching lower walls with upper walls) shows that if the centrally symmetric wall is an upper wall, the cone associated to C in $\mathcal{F}_{\Theta'}(B_n)$ has n walls. We conclude that $\mathcal{F}_{\Theta'}(B_n)$ is simplicial. \square

Proof of the type-B case of Theorem 4.2.12. Let c be a bipartite Coxeter element in A_{2n-1} and let \tilde{c} be the same element thought of as a Coxeter element of B_n . Recall that \tilde{c} is also bipartite.

Using the bipartite case of Proposition 4.3.1 (with n replaced by $2n - 1$), it is easily checked

that $x \equiv y$ modulo Θ_c if and only if $w_0 x w_0 \equiv w_0 y w_0$ modulo Θ_c . It follows that the c -biCambrian congruence is symmetric under conjugation by w_0 . Since a congruence is uniquely determined by the set of bottom elements of its classes, Proposition 4.3.16 implies that the restriction of the c -biCambrian congruence to B_n is the \tilde{c} -biCambrian congruence. Thus the type-B case of the theorem follows from Proposition 4.3.21 and the type-A case of the theorem. \square

4.4 Double-positive Catalan numbers and biCatalan numbers

For each finite Coxeter group W , the positive W -Catalan number $\text{Cat}^+(W)$ is defined from the W -Catalan number $\text{Cat}(W)$ by inclusion-exclusion. In this section, we review the definition of the positive W -Catalan number and define the double-positive W -Catalan number $\text{Cat}^{++}(W)$ from the positive W -Catalan number by inclusion-exclusion. We then prove Theorems 4.1.2 and 4.1.3 by showing how to count both antichains in the doubled root poset and bipartite c -bisortable elements by the same formula involving double-positive Catalan numbers. Recall that these two theorems in particular establish that the terms “biCatalan number” and “biNarayana number” make sense. As we prove these theorems, we obtain as a by-product a formula for the W -biCatalan numbers in terms of the double-positive Catalan numbers of parabolic subgroups of W . This formula leads to a recursion for the W -biCatalan numbers. Using a similar recursion for the W -Catalan numbers and a few other enumerative facts, we solve that recursion for $\text{biCat}(D_n)$ to complete the proof of Theorem 4.1.4. The recursions discussed here all have Narayana q -analogues, but we are not at this time able to solve the recursion to find a formula for $\text{biCat}(D_n; q)$. See Section 4.4.9 for a brief discussion of the type-D biNarayana numbers.

The positive W -Catalan and positive W -Narayana numbers have interpretations in each setting of Coxeter-Catalan combinatorics. (See for example [4, 6, 8, 35, 46, 61, 69, 65].) In this chapter, we give the usual interpretations in the settings of nonnesting partitions and c -sortable elements, specifically in Sections 4.4.2 and 4.4.5. We give interpretations of the double-positive W -Catalan and W -Narayana numbers in the settings of nonnesting partitions and c -sortable elements.

The double-positive W -Narayana numbers appeared in [7] as the local h -vector of the positive part of the cluster complex. (See Remark 4.4.7.) As far as we know, [7] was the first appearance of the double-positive W -Catalan/Narayana numbers and the only appearance before the current chapter.

4.4.1 Double-positivity

We write S for the set of simple reflections generating W . Given $J \subseteq S$, the notation W_J stands for the subgroup of W generated by J . The subgroup W_J is called a *standard parabolic*

subgroup of W and is a Coxeter group in its own right with simple reflections J . In particular, each W_J has a Catalan number. As usual, we define the **positive W -Catalan number** to be

$$\text{Cat}^+(W) = \sum_{J \subseteq S} (-1)^{|S|-|J|} \text{Cat}(W_J). \quad (4.4.1)$$

We define the **double-positive W -Catalan number** to be

$$\text{Cat}^{++}(W) = \sum_{J \subseteq S} (-1)^{|S|-|J|} \text{Cat}^+(W_J). \quad (4.4.2)$$

We will prove the following formula for the biCatalan numbers.

Theorem 4.4.1. *For any finite Coxeter group W with simple generators S ,*

$$\text{biCat}(W) = \sum 2^{|S|-|I|-|J|} \text{Cat}^{++}(W_I) \text{Cat}^{++}(W_J), \quad (4.4.3)$$

where the sum is over all ordered pairs (I, J) of disjoint subsets of S .

We can prove a refinement of Theorem 4.4.1 using the usual notion of positive Narayana numbers and a notion of double-positive Narayana numbers. The **positive W -Narayana numbers** are

$$\text{Nar}_k^+(W) = \sum_{J \subseteq S} (-1)^{|S|-|J|} \text{Nar}_k(W_J). \quad (4.4.4)$$

We define the **double-positive W -Narayana number** to be

$$\text{Nar}_k^{++}(W) = \sum_{J \subseteq S} (-1)^{|S|-|J|} \text{Nar}_{k-|S|+|J|}^+(W_J). \quad (4.4.5)$$

In all of the settings where the Narayana numbers appear, it is apparent that $\text{Nar}_k(W) = 0$ whenever $k < 0$ or k is greater than the rank of W . These definitions establish that $\text{Nar}_k^+(W) = \text{Nar}_k^{++}(W) = 0$ as well for those values of k .

Defining $\text{Cat}^+(W; q) = \sum_k \text{Nar}_k^+(W) q^k$ and $\text{Cat}^{++}(W; q) = \sum_k \text{Nar}_k^{++}(W) q^k$, equations (4.4.4) and (4.4.5) correspond to

$$\text{Cat}^+(W; q) = \sum_{J \subseteq S} (-1)^{|S|-|J|} \text{Cat}(W_J; q). \quad (4.4.6)$$

and

$$\text{Cat}^{++}(W; q) = \sum_{J \subseteq S} (-q)^{|S|-|J|} \text{Cat}^+(W_J; q). \quad (4.4.7)$$

Taking $\text{biCat}(W; q) = \sum_k \text{biNar}_k(W) q^k$, we will prove the following q -analog of Theorem 4.4.1.

Theorem 4.4.2. *For any finite Coxeter group W with simple generators S ,*

$$\text{biCat}(W; q) = \sum q^{|M|} \text{Cat}^+(W_I; q) \text{Cat}^+(W_J; q), \quad (4.4.8)$$

where the sum is over all ordered triples (I, J, M) of pairwise disjoint subsets of S .

The following theorem is equivalent to Theorem 4.4.2.

Theorem 4.4.3. *For any finite Coxeter group W with simple generators S and any k ,*

$$\text{biNar}_k(W) = \sum \sum_{i=0}^{k-|M|} \text{Nar}_i^+(W_I) \text{Nar}_{k-|M|-i}^+(W_J), \quad (4.4.9)$$

where the outer sum is over all ordered triples (I, J, M) of pairwise disjoint subsets of S . (If $|M| > k$, then the inner sum is interpreted to be zero.)

To prove these theorems, as well as Theorems 4.1.2 and 4.1.3, we establish (in Propositions 4.4.8 and 4.4.29) that the right side of (4.4.8) counts antichains A in the doubled root poset with weight $q^{|A|}$ and also counts bipartite c -bisortable elements v with weight $q^{\text{des}(v)}$. Once these counts are established, Theorems 4.1.2 and 4.1.3 follow, and in particular the definitions of the biCatalan and biNarayana numbers are validated. Also, Theorem 4.4.2 holds, leading immediately to Theorems 4.4.1 and 4.4.3.

4.4.2 Counting twin nonnesting partitions

We now recall the interpretations of the positive Catalan and Narayana numbers and give the interpretations of double-positive Catalan and Narayana numbers in the nonnesting setting. (Results in [8, 65] give the same interpretations, but accomplish much more, by establishing bijections and counting formulas. By contrast, here we are only making simple assertions about inclusion-exclusion.) After giving these interpretations, we prove that the formula in Theorem 4.4.3 counts k -element antichains in the doubled root poset.

Since it is customary to talk about the “ W -Catalan number” rather than the “ Φ -Catalan number,” we will make statements about “the root poset of W ,” when W is a crystallographic Coxeter group. This is harmless because, although the map from crystallographic root systems to Coxeter groups is not one-to-one, for each crystallographic Coxeter group, all corresponding crystallographic root systems have isomorphic root posets. Correspondingly, when W_J is a standard parabolic subgroup of W , we will say that a root or set of roots is “contained in W_J ” if it is contained in the subset of Φ forming a root system for W_J . An antichain that is not contained in any proper parabolic W_J has full support, in the sense of Section 4.2.1.

For any $J \subseteq S$, the number of antichains in the root poset for W that are contained in W_J is $\text{Cat}(W_J)$. By inclusion-exclusion, we conclude that:

Proposition 4.4.4. *The number of antichains in the root poset for W with full support is $\text{Cat}^+(W)$. The number of k -element antichains in the root poset for W with full support is $\text{Nar}_k^+(W)$.*

For $J \subseteq S$, the map $A \mapsto A \setminus \{\alpha_i : i \in J\}$ is a bijection from the set of antichains containing the simple roots $\{\alpha_i : i \in J\}$ to the set of antichains in the root poset for $W_{S \setminus J}$.

Using this bijection, we prove the following proposition.

Proposition 4.4.5. *The number of antichains in the root poset for W containing no simple roots is $\text{Cat}^+(W)$. The number of k -element antichains in the root poset for W containing no simple roots is $\text{Nar}_{n-k}^+(W)$.*

Proof. The bijection mentioned above implies that the generating function for antichains containing the simple roots $\{\alpha_i : i \in J\}$ (and possibly additional simple roots) is $q^{|J|} \text{Cat}(W_{S \setminus J}; q)$. By inclusion-exclusion, the generating function for k -element antichains containing no simple roots is $\sum_{J \subseteq S} (-q)^{|S| - |J|} \text{Cat}(W_J; q)$. On the other hand, starting with (4.4.6), replacing q by q^{-1} , multiplying through by $q^{|S|}$ (i.e. q^n), and using the known symmetry the $q^{|J|} \text{Cat}(W_J; q^{-1})$ is equal to $\text{Cat}(W_J; q)$ of the coefficients of $\text{Cat}(W_J; q)$, we obtain

$$\sum_k \text{Nar}_{n-k}^+(W) q^k = \sum_{J \subseteq S} (-q)^{|S| - |J|} \text{Cat}(W_J; q). \quad \square$$

The bijection described above restricts to a bijection from the set of antichains *with full support* containing the simple roots $\{\alpha_i : i \in J\}$ to the set of antichains *with full support* in the root poset for $W_{S \setminus J}$. Thus, a similar inclusion-exclusion argument yields the following proposition.

Proposition 4.4.6. *The number of antichains in the root poset for W with full support containing no simple roots is $\text{Cat}^{++}(W)$. The number of k -element antichains in the root poset for W with full support containing no simple roots is $\text{Nar}_k^{++}(W)$.*

Remark 4.4.7. The polynomials $\text{Cat}^{++}(W; q)$ appeared in [7], where Athanasiadis and Savvi-dou showed that $\text{Cat}^{++}(W; q)$ is the local h -vector of the positive part of the cluster complex, as we now explain. We refer to [7] for the relevant definitions, which we will not need here. In light of [8, Theorem 1.5] and Proposition 4.4.5, the polynomial $h(\Delta_+(\Phi), x)$ appearing in [7] is $x^{|S|} \text{Cat}^+(W; x^{-1})$, where (W, S) is the Coxeter system associated to Φ . Thus the assertion of [7, Proposition 2.5] is that the local h -vector of the positive part of the cluster complex is $\sum_{J \subseteq S} (-1)^{|S| - |J|} x^{|J|} \text{Cat}^+(W; x^{-1})$. But since the local h -vector is symmetric by [79, Theorem 3.3], we can replace x by x^{-1} and multiply by $x^{|S|}$ to show that the local h -vector is $\sum_{J \subseteq S} (-x)^{|S| - |J|} \text{Cat}^+(W; x) = \text{Cat}^{++}(W; x)$.

We now prove the key result on antichains in the doubled root poset.

Proposition 4.4.8. *For any finite Coxeter group W with simple generators S , the generating function $\sum_A q^{|A|}$ for antichains A in the doubled root poset is*

$$\sum q^{|M|} \text{Cat}^+(W_I; q) \text{Cat}^+(W_J; q), \quad (4.4.10)$$

where the sum is over all ordered triples (I, J, M) of pairwise disjoint subsets of S .

Proof. In light of Proposition 4.4.6, the proposition amounts to the following assertions: First, there is a bijection from antichains A in the doubled root poset to triples (B, C, M) such that B and C are antichains in the root poset for W , each containing no simple roots, and the sets $I = \text{supp}(B)$, $J = \text{supp}(C)$ and M are pairwise disjoint. Second, under this bijection, $|B| + |C| + |M| = |A|$. Every antichain A in the doubled root poset consists of some set B of positive non-simple roots in the top root poset, some set C of positive non-simple roots in the bottom root poset, and some set M of simple roots. The sets I , J , and M are pairwise disjoint because A is an antichain. The map $A \mapsto (B, C, M)$ is the desired bijection. \square

It will be useful to have a similar formula for antichains in the (not doubled) root poset, which are known to be counted by $\text{Cat}(W)$.

Theorem 4.4.9. *For any finite Coxeter group W with simple generators S .*

$$\text{Cat}(W; q) = \sum q^{|J|} \text{Cat}^+(W_I; q), \quad (4.4.11)$$

where the sum is over all ordered pairs (I, J) of disjoint subsets of S .

Proof. Every antichain A in the root poset consists of some set B of positive non-simple roots and some set C of simple roots. Writing I and J for the supports of B and C , again I and J are disjoint. By Proposition 4.4.6, each pair (I, J) of disjoint subsets of S contributes $q^{|J|} \text{Cat}^+(W_I; q)$ to the count. \square

The following is an immediate consequence of Proposition 4.4.6 and will also be useful.

Proposition 4.4.10. *If W is reducible as $W_1 \times W_2$, then*

$$\text{Cat}^+(W; q) = \text{Cat}^+(W_1; q) \text{Cat}^+(W_2; q). \quad (4.4.12)$$

4.4.3 Canonical join representations and lattice congruences

To count bipartite c -bisortable elements, we will use a canonical factorization in the weak order called the canonical join representation. In this section, we focus exclusively on the lattice-theoretic tools that we will use in the following sections to complete the proof of Theorem 4.4.3.

The canonical join representation is a “minimal” expression for an element as a join of join-irreducible elements. The construction is somewhat analogous to prime factorizations of integers. Indeed, in the divisibility poset for positive integers, where $p \leq q$ if and only if $p|q$, the canonical join representation coincides with prime factorization. For our purposes, the canonical join representation is useful because of how it interacts with lattice congruences. Recall that a lattice congruence Θ contracts a join-irreducible element j if j is equivalent modulo Θ to the unique element that it covers. Each congruence Θ of a finite lattice is determined by the set of join-irreducible elements that it contracts. In particular, we can see which elements of W are c -sortable or c -bisortable by looking at their canonical join representations (much as we looked at the arcs in their arc diagrams in types A and B).

The canonical join representation of an element a is an expression $a = \vee A$ such that A is minimal in two senses, among sets joining to a . First, the join $\vee A$ is *irredundant*, meaning that there is no proper subset $A' \subset A$ with $\vee A' = \vee A$. Second, A has the smallest possible elements (in terms of the partial order on L). Specifically, a subset A of L *join-refines* a subset B of L if for each $a \in A$ there is an element $b \in B$ such that $a \leq b$. Join-refinement is a preorder on the subsets of L that restricts to a partial order on the set of antichains. The *canonical join representation* of a , if it exists, is the unique minimal antichain A , in the sense of join-refinement, that joins irredundantly to a . We sometimes write $\text{can}(a)$ for A . The elements of A are called the *canonical joinands* of a . It follows immediately that each canonical joinand is join-irreducible.

Not every finite lattice admits a canonical join representation for each of its elements. For example, in the diamond lattice M_3 , which has five elements, three of which are atoms, the largest element does not have a canonical join representation. Many interesting lattices do admit canonical join representations, including all finite distributive lattices and, as we will see, the weak order on finite Coxeter groups. The next proposition establishes the promised connection between canonical join representations and lattice congruences. (The last assertion in the proposition also follows from [71, Proposition 6.3].)

Proposition 4.4.11. *Suppose L is a finite lattice such that each element in L has a canonical join representation, and suppose that Θ is a lattice congruence on L . If j is a canonical joinand of $a \in L$ and j is not contracted by Θ , then j is a canonical joinand of $\pi_{\downarrow}^{\Theta}(a)$ in L . Moreover, if $\pi_{\downarrow}^{\Theta}(a) = a$ then none of the canonical joinands of a are contracted by Θ .*

The assertion that j is a canonical joinand of $\pi_{\downarrow}^{\Theta}(a)$ in L implies also that j is a canonical joinand of $\pi_{\downarrow}^{\Theta}(a)$ in $\pi_{\downarrow}^{\Theta}(L)$. (Since $\pi_{\downarrow}^{\Theta}(L)$ is a join-sublattice of L , every join-representation of $\pi_{\downarrow}^{\Theta}(a)$ in $\pi_{\downarrow}^{\Theta}(L)$ is also a join-representation of $\pi_{\downarrow}^{\Theta}(a)$ in L .)

Proof. Throughout the proof, we write $\{j_1, \dots, j_k\}$ for $\text{can}(a)$ with $j = j_1$. Recall that the lattice quotient L/Θ is isomorphic to the subposet of L induced by the set $\pi_{\downarrow}^{\Theta}(L)$. Suppose j is not

contracted by Θ , so that $\pi_{\downarrow}^{\Theta}(j) = j$. Recall that $\pi_{\downarrow}^{\Theta}$ is a lattice homomorphism, so $\pi_{\downarrow}^{\Theta}(a)$ is equal to $\bigvee_{i=1}^k \pi_{\downarrow}^{\Theta}(j_i) = j \vee \left(\bigvee_{i=2}^k \pi_{\downarrow}^{\Theta}(j_i) \right)$, (where the joins are all taken in the lattice quotient L/Θ). Since L/Θ is also a join-sublattice of L , the join in L/Θ coincides with the join in L . Thus $\pi_{\downarrow}^{\Theta}(a)$ is equal to $j \vee \left(\bigvee_{i=2}^k \pi_{\downarrow}^{\Theta}(j_i) \right)$ in L . Write B for the set $\text{can}(\pi_{\downarrow}^{\Theta}(a))$. Thus B join-refines $\{j\} \cup \{\pi_{\downarrow}^{\Theta}(j_2), \dots, \pi_{\downarrow}^{\Theta}(j_k)\}$. If no element of B is less or equal to j , then this join-refinement implies that each element of B is below some element of $\{\pi_{\downarrow}^{\Theta}(j_2), \dots, \pi_{\downarrow}^{\Theta}(j_k)\}$, so that $\pi_{\downarrow}^{\Theta}(a) \leq \bigvee_{i=2}^k \pi_{\downarrow}^{\Theta}(j_i)$. Since also $\pi_{\downarrow}^{\Theta}(a)$ is equal to $j \vee \left(\bigvee_{i=2}^k \pi_{\downarrow}^{\Theta}(j_i) \right)$, we see that $j \leq \bigvee_{i=2}^k \pi_{\downarrow}^{\Theta}(j_i)$. Recall that $\pi_{\downarrow}^{\Theta}(j_i) \leq j_i$ for each i , so we have $j \leq \bigvee_{i=2}^k j_i$. This contradicts the fact that $\bigvee_{i=1}^k j_i$ is irredundant. We conclude that there is some $j' \in B$ with $j' \leq j$. Observe that $(\bigvee B) \vee \left(\bigvee_{i=2}^k j_i \right) = a$ because $j_1 = j \leq \pi_{\downarrow}^{\Theta}(a) \leq a$. Thus, $\{j_1, \dots, j_k\}$ join-refines $B \cup \{j_2, \dots, j_k\}$. Since j is incomparable to each j_i , there is some $j'' \in B$ such that $j \leq j''$. But B is an antichain, so $j' = j'' = j$, and thus $j \in B$ as desired.

Suppose $\pi_{\downarrow}^{\Theta}(a) = a$. Then $a = \bigvee_{i=1}^n \pi_{\downarrow}^{\Theta}(j_i)$, so $\{j_1, \dots, j_k\}$ join-refines $\{\pi_{\downarrow}^{\Theta}(j_1), \dots, \pi_{\downarrow}^{\Theta}(j_k)\}$. Thus, for each j_i , there is some j_m with $j_i \leq \pi_{\downarrow}^{\Theta}(j_m)$. But $\pi_{\downarrow}^{\Theta}(j_m) \leq j_m$, and since $\{j_1, \dots, j_k\}$ is an antichain, we have $j_i = j_m$, and thus also $j_i = \pi_{\downarrow}^{\Theta}(j_i)$. \square

We will use the following easy proposition, which appears as [72, Proposition 2.2].

Proposition 4.4.12. *Suppose L is a finite lattice and $J \subset L$. If $\bigvee J$ is the canonical join representation of some element of L and if $J' \subseteq J$, then $\bigvee J'$ is the canonical join representation of some element of L .*

Next we consider canonical join representations in the weak order. Before we begin, we briefly review some relevant terminology. For each $w \in W$, the **length** of w , denoted $l(w)$, is the number of letters in a reduced (that is, a shortest possible) word for w in the alphabet S . The covers in the (right) weak order on W are $w \succ ws$ whenever $w \in W$ and $s \in S$ have $l(ws) < l(w)$. In this case, the simple generator s is a **descent** of w . Let T denote the set of reflections in W . An **inversion** of w is a reflection t such that $l(tw) < l(w)$. We denote the set of inversions of w by $\text{inv}(w)$. A **cover reflection** of w is an inversion t of w such that $tw = ws$ for some $s \in S$. Thus, the cover reflections of w are in bijection with the descents of w . We write $\text{cov}(w)$ for the set of cover reflections of w . The following proposition is quoted from [75, Theorem 8.1].

Proposition 4.4.13. *Fix a finite Coxeter group W , and an element $w \in W$. The canonical join representation of w exists and is equal to $\bigvee j_t$ where t ranges over the set of cover reflections of w , and j_t is the unique smallest element below w that has t as an inversion. In particular, w has $\text{des}(w)$ many canonical joinands.*

Recall that the **support** of w , written $\text{supp}(w)$, is the set of simple reflections appearing in a reduced word for w , and is independent of the choice of reduced word for w . The following

lemma is an immediate consequence of the fact that every standard parabolic subgroup W_J is a lower interval in the weak order on W .

Lemma 4.4.14. *For each $w \in W$, the support of w equals $\bigcup_{j \in \text{can}(w)} \text{supp}(j)$.*

For each element w and standard parabolic subgroup W_J , there is a unique largest element below w that belongs to W_J . We write w_J for this element and π_{\downarrow}^J for the map that sends w to w_J . In [66, Corollary 6.10], it was shown that the fibers of π_{\downarrow}^J constitute a lattice congruence of the weak order. We write Θ_J for this congruence. Since π_{\downarrow}^J sends each element to the bottom of its fiber, it is a lattice homomorphism from W to $\pi_{\downarrow}^J(W)$, which equals W_J .

Lemma 4.4.15. *Suppose that A_1 and A_2 are antichains with disjoint support such that $\vee A_1$ and $\vee A_2$ are both canonical join representations in the weak order on W . Then $\vee(A_1 \cup A_2)$ is a canonical join representation.*

Proof. We write A for $A_1 \cup A_2$. First we show that $\vee A$ is irredundant. By way of contradiction, assume that there is some $j \in A$ such that $\vee A = \vee(A \setminus \{j\})$. We may as well take $j \in A_1$. We write J for the support of A_1 . Since the support of each join-irreducible element j' in A_2 is disjoint from J , and since support decreases weakly in the weak order, we conclude that $\pi_{\downarrow}^J(j')$ is the identity element. Since π_{\downarrow}^J is a lattice homomorphism, we have $\pi_{\downarrow}^J(\vee A) = \vee A_1$ and $\pi_{\downarrow}^J(\vee(A \setminus \{j\})) = \vee(A_1 \setminus \{j\})$. We conclude that $\vee A_1 = \vee(A_1 \setminus \{j\})$, contradicting the fact that $\vee A_1$ is a canonical join representation.

Next we show that $\text{can}(\vee A)$ is contained in A . Assume that j'' is a canonical joinand of $\vee A$. There is some $j \in A$ such that $j'' \leq j$. Assume that $j \in A_1$, so that $\text{supp}(j'') \subset J$. Thus, $\pi_{\downarrow}^J(j'') = j''$. Proposition 4.4.11 says j'' is a canonical joinand of $\pi_{\downarrow}^J(\vee A) = \vee A_1$. Because A is an antichain, $j'' = j$. Since $\vee A$ is irredundant, and A contains $\text{can}(\vee A)$, we conclude that A is equal to $\text{can}(\vee A)$. \square

Observe that if $s \in S$ is a cover reflection of w then Proposition 4.4.13 implies that s is also a canonical joinand of w because simple reflections are atoms in the weak order. We immediately obtain the following useful fact.

Lemma 4.4.16. *Each $w \in W$ has $\text{can}(w) \cap S = \text{cov}(w) \cap S$.*

In much of what follows, for $s \in S$, we will use the abbreviation $\langle s \rangle$ to mean $S \setminus \{s\}$. It is known (see for example [70, Lemma 2.8]) that if $w \in W_{\langle s \rangle}$, then $\text{cov}(w \vee s) = \text{cov}(w) \cup \{s\}$. We close this section with a lemma that extends this statement to canonical join representations.

Lemma 4.4.17. *If $w \in W_{\langle s \rangle}$, then $\text{can}(w \vee s) = \text{can}(w) \cup \{s\}$.*

Proof. Since support is weakly decreasing in the weak order, each $j \in \text{can}(w)$ has support contained in $\langle s \rangle$. Lemma 4.4.15 says that $\vee(\text{can}(w) \cup \{s\})$ is a canonical join representation. \square

4.4.4 Canonical join representations of c -bisortable elements

In this section we focus on canonical join representations of c -sortable elements and c -bisortable elements. Our goal is to prove the following result:

Proposition 4.4.18. *Fix a bipartite c -bisortable element w and the corresponding twin (c, c^{-1}) -sortable elements $(u, v) = (\pi_{\downarrow}^c(w), \pi_{\downarrow}^{c^{-1}}(w))$. Then*

- (1) $\text{can}(w) \cap S = \text{can}(u) \cap \text{can}(v)$
- (2) $\text{can}(w)$ is the disjoint union $(\text{can}(u) \setminus S) \uplus (\text{can}(v) \setminus S) \uplus (\text{can}(w) \cap S)$
- (3) The sets $\text{supp}(\text{can}(u) \setminus S)$, $\text{supp}(\text{can}(v) \setminus S)$ and $\text{can}(w) \cap S$ are pairwise disjoint.

We begin with an easy application of Proposition 4.4.11 (the first item below can also be found as [75, Proposition 8.2]).

Proposition 4.4.19. *For any Coxeter element c and $w \in W$:*

1. w is c -sortable if and only if each of its canonical joinands is c -sortable.
2. w is c -bisortable if and only if each of its canonical joinands is either c - or c^{-1} -sortable.

Proof. The first assertion follows immediately from Proposition 4.4.11. Recall the notation Θ_c for the c -Cambrian congruence and write Θ for the c -biCambrian congruence. Since Θ is the meet $\Theta_c \wedge \Theta_{c^{-1}}$, a join-irreducible element in W is contracted by Θ if and only if it is contracted by Θ_c and by $\Theta_{c^{-1}}$. The second assertion follows. \square

Recall from Section 4.2.4 that a simple reflection s is initial in a Coxeter element c if there is a reduced word $a_1 \dots a_n$ for c with $a_1 = s$. Similarly s is **final** in c if there is a reduced word $a_1 \dots a_n$ for c with $a_n = s$. In much of what follows, the key property of a bipartite Coxeter element is that every $s \in S$ is either initial or final in c .

The following lemma is the combination of [75, Propositions 3.13, 5.3, and 5.4]. Recall that $v_{\langle s \rangle}$ is the largest element in W below v that belongs to $W_{\langle s \rangle}$.

Lemma 4.4.20. *Fix a c -sortable element v in W and a simple reflection $s \in S$.*

1. If s is final in c and $v \geq s$, then $v_{\langle s \rangle}$ is cs -sortable and $v = s \vee v_{\langle s \rangle}$.
2. If s be initial in c and $s \in \text{cov}(v)$, then $v_{\langle s \rangle}$ is sc -sortable and $v = s \vee v_{\langle s \rangle}$.

If v satisfies the conditions of either item in Lemma 4.4.20, then Lemma 4.4.17 implies that $\text{can}(v) = \{s\} \cup \text{can}(v_{\langle s \rangle})$. The following two lemmas are an easy application of Lemma 4.4.20.

Lemma 4.4.21. *If j is a c -sortable join-irreducible element and s is final in c with $j \geq s$, then $j = s$.*

Proof. The first assertion of Lemma 4.4.20 says that $j = s \vee j_{(s)}$. Since j is join-irreducible and not equal to $j_{(s)}$, we conclude that $j = s$. \square

Lemma 4.4.22. *If c is a bipartite Coxeter element and j is a join-irreducible element that is both c -sortable and c^{-1} -sortable, then j is a simple reflection.*

Proof. Because j is join-irreducible, it is not the identity, so there is some $s \in S$ such that $j \geq s$. Since c is bipartite, we can assume without loss of generality that s is final in c . (If not, then replace c with c^{-1} .) Thus $j = s$ by Lemma 4.4.21. \square

Putting together Lemma 4.4.21 and Lemma 4.4.22, we get an explicit description of $\pi_{\downarrow}^{c^{-1}}(j)$, for bipartite c -sortable join-irreducible elements.

Lemma 4.4.23. *Suppose that c is a bipartite Coxeter element and j is a c -sortable join-irreducible element. Let S' denote the set of simple reflections s such that $j \geq s$. Then $\pi_{\downarrow}^{c^{-1}}(j)$ is equal to $\vee S'$ which, in this case, is the product $\prod S'$ in W . Moreover, this join is a canonical join representation.*

Proof. The statement of the lemma is obvious if j is a simple reflection, so we assume that j is not simple. Thus, Lemma 4.4.22 implies that j is not c^{-1} -sortable, so $\pi_{\downarrow}^{c^{-1}}(j)$ is strictly less than j .

If any $s \in S'$ is final in c , then Lemma 4.4.21 says that $j = s$, contradicting our assumption. Thus, since c is bipartite, each $s \in S'$ is initial. In particular, the elements of S' pairwise commute, so that the notation $\prod S'$ makes sense and equals $\vee S'$. Moreover, since $\vee S'$ is an irredundant join of atoms, it is a canonical join representation. Since each simple reflection is both c - and c^{-1} -sortable, Proposition 4.4.19 says that this element is c^{-1} -sortable. We conclude that $\pi_{\downarrow}^{c^{-1}}(j) \geq \vee S'$.

Suppose that j' is a canonical joinand of $\pi_{\downarrow}^{c^{-1}}(j)$. There is some simple reflection s such that $j' \geq s$. Since also $j' \leq \pi_{\downarrow}^{c^{-1}}(j) \leq j$, we conclude that $s \in S'$. Every element of S' is initial in c and thus final in c^{-1} , so again by Lemma 4.4.21, $j' = s$. We conclude that $\text{can}(\pi_{\downarrow}^{c^{-1}}(j)) \subseteq S'$. Thus $\pi_{\downarrow}^{c^{-1}}(j) = \vee S'$. \square

Recall that Lemma 4.4.15 says that if j and j' are join-irreducible elements with disjoint support, then $j \vee j'$ is canonical. In Lemma 4.4.25 below, we prove that when j is bipartite c -sortable and j' is bipartite c^{-1} -sortable, the converse is also true. We begin with the case when j' is a simple reflection.

Lemma 4.4.24. *Given a bipartite Coxeter element c , a c -sortable join-irreducible element j and a simple reflection $s \in \text{supp}(j)$, there exists no element $w \in W$ with both s and j in $\text{can}(w)$.*

Proof. In light of Proposition 4.4.12, to prove this proposition, it is enough to show that no element can have $s \vee j$ as its canonical join representation. Suppose to the contrary that there is an element v with canonical join representation $s \vee j$. By Proposition 4.4.19, v is c -sortable. Also $s \vee j$ is irredundant, so j and s are incomparable. Since c is bipartite, s is either initial or final in c , so Lemma 4.4.20 says that $v = s \vee v_{\langle s \rangle}$. Since $v = s \vee j$ is a canonical join representation, we see that $j \leq v_{\langle s \rangle}$, contradicting the hypothesis that s is in the support of j . \square

Lemma 4.4.25. *Fix a bipartite Coxeter element c in W . Suppose that j is a c -sortable join-irreducible element and that j' is a c^{-1} -sortable join-irreducible element. Suppose that $j \vee j'$ is a canonical join representation for some element of W . Then j and j' have disjoint support.*

Proof. Suppose that $s \in \text{supp}(j) \cap \text{supp}(j')$, and assume without loss of generality that s is initial in c . It is immediate from the definition of c -sortable elements that $s \leq j$. (See for example [75, Proposition 2.29].) Since s is a c^{-1} -sortable element, also $s \leq \pi_{\downarrow}^{c^{-1}}(j \vee j')$. By Lemma 4.4.20(1) and Lemma 4.4.17, s is a canonical joinand of $\pi_{\downarrow}^{c^{-1}}(j \vee j')$. But also Proposition 4.4.11 says that j' is a canonical joinand of $\pi_{\downarrow}^{c^{-1}}(j \vee j')$. We have reached a contradiction to Lemma 4.4.24, and we conclude that $\text{supp}(j) \cap \text{supp}(j') = \emptyset$. \square

Finally, we prove Proposition 4.4.18.

Proof of Proposition 4.4.18. Lemma 4.4.22 implies that $\text{can}(w) \setminus S$ is the disjoint union

$$(\text{can}(w) \cap S) \uplus J_+ \uplus J_-$$

where J_+ is the set of c -sortable join-irreducible elements in $\text{can}(w) \setminus S$ and J_- is the set of c^{-1} -sortable join-irreducible elements in $\text{can}(w) \setminus S$. Moreover, by Lemma 4.4.25, these sets have pairwise disjoint support. For each $j \in J_-$, write S'_j for the set of simple reflections s such that $s \leq j$, and $S' = \bigcup S'_j$, where the union ranges over all $j \in J_-$. Lemma 4.4.23 says that $\pi_{\downarrow}^c(j) = \bigvee S'_j$. Since π_{\downarrow}^c is a join-homomorphism, $\pi_{\downarrow}^c(\bigvee J_-) = \bigvee S'$. Thus, applying the map π_{\downarrow}^c to the join $\bigvee [(\text{can}(w) \cap S) \uplus J_+ \uplus J_-]$, we see that $\bigvee [(\text{can}(w) \cap S) \uplus J_+ \uplus S']$ is a join representation of u . Since S' is contained in the support of J_- , the sets $\text{can}(w) \cap S$, J_+ , and S' also have pairwise disjoint support. Proposition 4.4.12 says that both $\bigvee \text{can}(w) \cap S$ and $\bigvee J_+$ are canonical join representations. Since $\bigvee S'$ is an irredundant join of atoms, it is also a canonical join representation. Thus, by Lemma 4.4.15, $\bigvee [(\text{can}(w) \cap S) \uplus J_+ \uplus S']$ is the canonical join representation of u . The symmetric argument gives the canonical join representation of v . We conclude that $\text{can}(w) \cap S = \text{can}(u) \cap \text{can}(v)$, $J_+ = \text{can}(u) \setminus S$, and $J_- = \text{can}(v) \setminus S$. The proposition follows. \square

4.4.5 Counting bipartite c -bisortable elements

In this section, we prove that the formulas in Theorem 4.4.3 counts bipartite c -bisortable elements, thus completing the proofs of Theorems 4.1.2, 4.1.3, 4.4.1, 4.4.2 and 4.4.3. We begin by interpreting the double-positive Catalan and Narayana numbers in the c -sortable setting. We define **positive c -sortable elements** to be the set of c -sortable elements not contained in any standard parabolic subgroup of W . Equivalently, these are the c -sortable elements whose support is not contained in any proper subset of S . As the name suggests, positive c -sortable elements are counted by the positive Catalan numbers. The following analogue of Proposition 4.4.4 is the combination of [69, Corollary 9.2] and [69, Corollary 9.3].

Proposition 4.4.26. *For any Coxeter element c of W , the number of positive c -sortable elements in W is $\text{Cat}^+(W)$. The number positive c -sortable elements with k descents is $\text{Nar}_k^+(W)$.*

We define **clever c -sortable elements** to be c -sortable elements which have no simple canonical joinands. We continue to let $\langle s \rangle$ stand for $S \setminus \{s\}$. To count clever c -sortable elements we will use Lemma 4.4.20 to define a map from c -sortable elements v with simple cover reflection s to c' -sortable elements in the standard parabolic subgroup $W_{\langle s \rangle}$, where c' is the restriction of c to $W_{\langle s \rangle}$. Our next task is to show that, for bipartite c , clever c -sortable elements are analogous, enumeratively, to antichains in the root poset having no simple roots:

Proposition 4.4.27. *Fix a bipartite Coxeter element c of W .*

1. *The number of clever c -sortable elements is $\text{Cat}^+(W)$.*
2. *The number of positive, clever c -sortable elements is $\text{Cat}^{++}(W)$.*
3. *The number of positive, clever c -sortable elements with exactly k descents is $\text{Nar}_k^{++}(W)$.*

We emphasize that while Proposition 4.4.26 holds for arbitrary c , Proposition 4.4.27 holds only for bipartite c . The proof of Proposition 4.4.27 will use inclusion-exclusion and the following technical lemma.

Lemma 4.4.28. *For bipartite c and $J \subseteq S$, let c' be the restriction of c to $W_{S \setminus J}$.*

1. *The map $\pi_{\downarrow}^{S \setminus J} : v \mapsto v_{S \setminus J}$ is a bijection from c -sortable elements of W with $J \subseteq \text{can}(v)$ to c' -sortable elements of $W_{S \setminus J}$. Also, $\text{can}(v_{S \setminus J}) = \text{can}(v) \setminus J$.*
2. *The map restricts to a bijection from positive c -sortable elements of W with $J \subseteq \text{can}(v)$ to positive c' -sortable elements of $W_{S \setminus J}$.*
3. *The map restricts to a bijection from positive c -sortable elements of W with $J \subseteq \text{can}(v)$ and with exactly k descents to positive c' -sortable elements of $W_{S \setminus J}$ with exactly $k - |J|$ descents.*

Proof. Suppose that v is c -sortable, and $J \subseteq \text{can}(v)$. Lemma 4.4.24 says that the support of each canonical joinand j in $\text{can}(v) \setminus J$ is contained in $S \setminus J$. (Lemma 4.4.24 applies to the non-simple elements of $\text{can}(v)$. Clearly, each simple reflection $s \in \text{can}(v) \setminus J$ is supported on the set $S \setminus J$.) On the one hand, $\pi_{\downarrow}^{S \setminus J}(j) = j$ for each $j \in \text{can}(v) \setminus J$. On the other hand, $\pi_{\downarrow}^{S \setminus J}(s)$ is the identity element for each s in J . Since $\pi_{\downarrow}^{S \setminus J}$ is a lattice homomorphism, we have $\pi_{\downarrow}^{S \setminus J}(\vee \text{can}(v)) = \vee[\text{can}(v) \setminus J]$. Proposition 4.4.11 implies that $\vee[\text{can}(v) \setminus J]$ is the canonical join representation of $\pi_{\downarrow}^{S \setminus J}(v) = v_{S \setminus J}$. Lemma 4.4.19 says that $v_{S \setminus J}$ is c' -sortable.

To complete the proof of the first assertion, we construct an inverse map. Suppose that v' is a c' -sortable element in $W_{S \setminus J}$. Lemma 4.4.14 says that the support of each canonical joinand $j \in \text{can}(v')$ is contained in $S \setminus J$. Lemma 4.4.15 says that the join $\vee[\text{can}(v') \cup J]$ is a canonical join representation for some element $v \in W$. Lemma 4.4.19 says that v is c -sortable. We conclude that the map sending v' to $\vee[\text{can}(v') \cup J]$ is a well-defined inverse.

Lemma 4.4.14, Lemma 4.4.24, and the fact that $\text{can}(v_{S \setminus J}) = \text{can}(v) \setminus J$ imply that v is positive in W if and only if $v_{S \setminus J}$ is positive in $W_{S \setminus J}$. The second assertion follows. The third assertion then follows from Proposition 4.4.13 and the fact that $\text{can}(v_{S \setminus J}) = \text{can}(v) \setminus J$. \square

Finally, we complete the proof of that bipartite c -bisortable elements are counted by the formula in Theorem 4.4.3.

Proposition 4.4.29. *For any finite Coxeter group W with simple generators S , the generating function $\sum_v q^{\text{des}(v)}$ for bipartite c -bisortable elements is*

$$\sum q^{|M|} \text{Cat}^+(W_I; q) \text{Cat}^+(W_J; q),$$

where the sum is over all ordered triples (I, J, M) of pairwise disjoint subsets of S .

Proof. Similarly to the proof of Proposition 4.4.8, the proposition amounts to establishing a bijection from bipartite c -bisortable elements w to triples (u', v', M) such that u' is a clever c -sortable element, v' is a clever c^{-1} -sortable element, and the sets $I = \text{supp}(u')$, $J = \text{supp}(v')$, and M are disjoint subsets of S , and then showing that $\text{des}(w) = \text{des}(u') + \text{des}(v') + |M|$.

Given a bipartite c -bisortable element w , write (u, v) for the pair $(\pi_{\downarrow}^c(w), \pi_{\downarrow}^{c^{-1}}(w))$ of twin (c, c^{-1}) -sortable elements. Proposition 4.4.18(2) says that $\text{can}(w)$ is the disjoint union $(\text{can}(u) \setminus S) \uplus (\text{can}(v) \setminus S) \uplus (\text{can}(w) \cap S)$. Proposition 4.4.18(4.4.18) says that the sets $I = \text{supp}(\text{can}(u) \setminus S)$, $J = \text{supp}(\text{can}(v) \setminus S)$, and $M = \text{can}(w) \cap S$ are pairwise disjoint subsets of S . By Proposition 4.4.12, $\vee \text{can}(u) \setminus S$ is the canonical join representation of a positive, clever c -sortable element u' in W_I . Similarly, $\vee \text{can}(v) \setminus S$ is the canonical join representation of a positive, clever c^{-1} -sortable element v' in W_J . Applying Proposition 4.4.13 several times, we see that $\text{des}(w) = \text{des}(u') + \text{des}(v') + |M|$.

We will show that this map $w \mapsto (u', v', M)$ is a bijection by showing that the map $(u', v', M) \mapsto u' \vee v' \vee (\vee M)$ is the inverse. On one hand, given w , construct (u', v', M) as above. Then w equals $\vee \text{can}(w)$, which equals

$$(\vee \text{can}(u) \setminus S) \vee (\vee \text{can}(v) \setminus S) \vee (\vee \text{can}(w) \cap S) = u' \vee v' \vee (\vee M).$$

On the other hand, given (u', v', M) satisfying the description above, set $w = u' \vee v' \vee (\vee M)$. Since u' , v' and M have pairwise disjoint support, we conclude that $\text{can}(u')$, $\text{can}(v')$, and M also have pairwise disjoint support. Lemma 4.4.15 says that $\vee \text{can}(u') \uplus \text{can}(v') \uplus M$ is the canonical join representation of w . By Lemma 4.4.19(1), each canonical joinand of u' is c -sortable and each canonical joinand of v' is c^{-1} -sortable. Since each simple generator is both c - and c^{-1} -sortable, we conclude that each canonical joinand of w either either c - or c^{-1} -sortable. By Lemma 4.4.19(2), w is c -bisortable. Thus, the map $(u', v', M) \mapsto u' \vee v' \vee (\vee M)$ is a well-defined.

Lemma 4.4.22 says that $\text{can}(u') \uplus M$ is equal to the set of c -sortable canonical joinands of w . Since u' is clever, $\text{can}(u')$ is equal to the set of c -sortable canonical joinands in $\text{can}(w) \setminus S$. Similarly, $\text{can}(v')$ is the set of c^{-1} -sortable canonical joinands in $\text{can}(w) \setminus S$, and $\text{can}(w) \cap S = M$. Define $u = \pi_{\downarrow}^c(w)$ and $v = \pi_{\downarrow}^{c^{-1}}(w)$. Proposition 4.4.18(2) says that $\text{can}(w) = (\text{can}(u) \setminus S) \uplus (\text{can}(v) \setminus S) \uplus (\text{can}(w) \cap S)$. Comparing this to the expression $\text{can}(w) = \text{can}(u') \uplus \text{can}(v') \uplus M$, we see that $\text{can}(u) \setminus S = \text{can}(u')$, that $\text{can}(v) \setminus S = \text{can}(v')$, and that $\text{can}(w) \cap S = M$. Thus the map described above takes w back to (u', v', M) . \square

Remark 4.4.30. The proof given here that twin nonnesting partitions are in bijection with bipartite c -bisortable elements would be uniform if there were a uniform proof connecting c -sortable elements and nonnesting partitions. The opposite is true as well: Suppose one proved uniformly that a given map ϕ is a bijection from antichains in the doubled root poset to bipartite c -bisortable elements and also that ϕ preserves the triples (I, J, M) appearing in Propositions 4.4.8 and 4.4.29. Then the restriction of ϕ to antichains in the root poset (i.e. those with $J = \emptyset$) is a bijection from antichains in the root poset to c -sortable elements.

Remark 4.4.31. The methods of this section don't apply well to the case where c is not bipartite, because the main structural results of the section, Propositions 4.4.18 and 4.4.27, can fail when c is not bipartite. This can already be seen in A_3 for the linear Coxeter element.

4.4.6 BiCatalan and Catalan formulas

In this section and the next, we prepare to prove the formula for $\text{biCat}(D_n)$ in Theorem 4.1.4, thus completing the proof of that theorem. Specifically, the proof requires combining a very large

number of identities relating q -analogs of biCatalan numbers, Catalan numbers, and double-positive Catalan numbers that we quote or prove here. In this section, we give recursions for the q -analogs of W -biCatalan and W -Catalan numbers for irreducible finite Coxeter groups, in which q -analogs of double-positive Catalan numbers appear as coefficients.

Proposition 4.4.32. *For an irreducible finite Coxeter group W and a simple generator $s \in S$, the q -analog of the W -biCatalan number satisfies*

$$\begin{aligned} \text{biCat}(W; q) &= (1 + q) \text{biCat}(W_{S \setminus \{s\}}; q) \\ &+ 2 \sum_{S_0} \text{Cat}^{++}(W_{S_0}; q) \prod_{i=1}^m \left[\frac{1}{2} \text{biCat}(W_{S_i}; q) + \frac{1+q}{2} \text{biCat}(W_{S_i \setminus \{s_i\}}; q) \right], \end{aligned} \quad (4.4.13)$$

where the sum is over all connected subgraphs S_0 of the diagram for W with $s \in S_0$, the connected components of the complement of S_0 in the diagram are S_1, \dots, S_m , and each s_i is the unique vertex in S_i that is connected by an edge to a vertex in S_0 .

Proof. For fixed s , we break the formula in Theorem 4.4.2 into four sums, according to whether s is in $S \setminus (I \cup J \cup M)$, in M , in I , or in J . The sum of terms with $s \in S \setminus (I \cup J \cup M)$ equals $\text{biCat}(W_{S \setminus \{s\}}; q)$. The sum of terms with $s \in M$ equals $q \cdot \text{biCat}(W_{S \setminus \{s\}}; q)$.

Consider next the sum of terms with $s \in I$, and in each term let S_0 be the connected component of the diagram containing s . Using (4.4.12), we can reorganize the sum according to S_0 to obtain

$$\sum_{S_0} \text{Cat}^{++}(W_{S_0}; q) \sum q^{|M|} \text{Cat}^{++}(W_{I'}; q) \text{Cat}^{++}(W_J; q),$$

where the S_0 -sum is as described in the statement of the proposition and the inner sum is over all ordered triples (I', J, M) of disjoint subsets of $S \setminus S_0$ such that no element of I' is connected by an edge of the diagram to an element of S_0 . Again using (4.4.12), we factor the inner sum further to obtain

$$\sum_{S_0} \text{Cat}^{++}(W_{S_0}; q) \prod_{i=1}^m \left[\sum q^{|M_i|} \text{Cat}^{++}(W_{I_i}; q) \text{Cat}^{++}(W_{J_i}; q) \right],$$

where the S_i and s_i are as in the statement of the proposition and the inner sum runs over all ordered triples (I_i, J_i, M_i) of pairwise disjoint subsets of S_i with $s_i \notin I_i$. The sum for each i can be broken up into a sum over terms with $s_i \in J_i$ and terms with $s_i \notin J_i$. Splitting the sum

over terms with $s_i \notin J_i$ in half, we obtain three sums:

$$\begin{aligned} & \sum_{s_i \in J_i} q^{|M_i|} \text{Cat}^{++}(W_{I_i}; q) \text{Cat}^{++}(W_{J_i}; q) \\ & + \frac{1}{2} \sum_{s_i \notin J_i} q^{|M_i|} \text{Cat}^{++}(W_{I_i}; q) \text{Cat}^{++}(W_{J_i}; q) \\ & + \frac{1}{2} \sum_{s_i \notin J_i} q^{|M_i|} \text{Cat}^{++}(W_{I_i}; q) \text{Cat}^{++}(W_{J_i}; q) \end{aligned}$$

The symmetry between I and J on the right side of Theorem 4.4.2 lets us recognize the sum of the first two terms as $\frac{1}{2} \text{biCat}(W_{S_i}; q)$, recalling that $s \notin I_i$ throughout. The third term is $\frac{1+q}{2} \text{biCat}(W_{S_i \setminus \{s_i\}}; q)$. We see that the sum of terms with $s \in I$ is the sum in the proposed formula, without the factor 2 in front. By symmetry, the sum of terms with $s \in J$ is the same sum, so we obtain the factor 2 in the sum and we have established the desired formula. \square

We obtain the following recursion for $\text{biCat}(D_n; q)$ from Proposition 4.4.32. The notation D_2 means $A_1 \times A_1$ and D_3 means A_3 .

Proposition 4.4.33. *For $n \geq 3$,*

$$\begin{aligned} \text{biCat}(D_n; q) &= (1+q) \text{biCat}(D_{n-1}; q) \\ &+ \sum_{i=1}^{n-3} \text{Cat}^{++}(A_i; q) \left(\text{biCat}(D_{n-i}; q) + (1+q) \text{biCat}(D_{n-i-1}; q) \right) \\ &+ 2(1+q)^2 \text{Cat}^{++}(A_{n-2}; q) + 4(1+q) \text{Cat}^{++}(A_{n-1}; q) + 2 \text{Cat}^{++}(D_n; q) \end{aligned} \quad (4.4.14)$$

Proof. In Proposition 4.4.32, take s to be a leaf of the D_n diagram whose removal leaves the diagram for D_{n-1} . The sum over S_0 splits into several pieces. First, the S_0 for which the diagram on $S \setminus \{S_0\}$ is of type D_k for $k \geq 3$ give rise to terms

$$\sum_{i=1}^{n-3} \text{Cat}^{++}(A_i; q) \left(\text{biCat}(D_{n-i}; q) + (1+q) \text{biCat}(D_{n-i-1}; q) \right).$$

Next, the term for which the diagram on $S \setminus \{S_0\}$ is of type D_2 is $2 \text{Cat}^{++}(A_{n-2}) \left(\frac{1}{2}(1+q) + \frac{1+q}{2} \cdot 1 \right)^2$, which simplifies to $2(1+q)^2 \text{Cat}^{++}(A_{n-2})$. The two terms for which the diagram on $S \setminus \{S_0\}$ is of type A_1 each contribute $2(1+q) \text{Cat}^{++}(A_{n-1})$. Finally, the term with $S_0 = S$ is $2 \text{Cat}^{++}(D_n; q)$. \square

We obtain the following recursion for $\text{biCat}(B_n; q)$ from Proposition 4.4.32 similarly. Here and throughout the chapter, we interpret B_0 and B_1 to be synonyms for A_0 and A_1 .

Proposition 4.4.34. For $n \geq 1$,

$$\begin{aligned} \text{biCat}(B_n; q) = & \\ & (1 + q) \text{biCat}(B_{n-1}; q) + 2 \text{Cat}^+(B_n; q) + 2(1 + q) \text{Cat}^+(A_{n-1}; q) \\ & + \sum_{i=1}^{n-2} \text{Cat}^+(A_i; q) [\text{biCat}(B_{n-i}; q) + (1 + q) \text{biCat}(B_{n-i-1}; q)]. \end{aligned} \quad (4.4.15)$$

Proof. In Proposition 4.4.32, take s to be a leaf of the B_n diagram whose removal leaves the diagram for B_{n-1} . The terms with $|S_0|$ from 1 to $n - 2$ are in the summation in (4.4.15), but we separate out the terms with $|S_0| = n - 1$ and $|S_0| = n$. For the term with $|S_0| = n - 1$, we use the facts that $\text{biCat}(B_1; q) = (1 + q)$ and that $\text{biCat}(B_0; q) = 1$. \square

Similarly, we obtain the following recursion for $\text{biCat}(A_n)$ by taking s to be either leaf of the diagram.

Proposition 4.4.35. For $n \geq 1$,

$$\begin{aligned} \text{biCat}(A_n; q) = & (1 + q) \text{biCat}(A_{n-1}; q) + 2 \text{Cat}^+(A_n; q) \\ & + \sum_{i=1}^{n-1} \text{Cat}^+(A_i; q) [\text{biCat}(A_{n-i}; q) + (1 + q) \text{biCat}(A_{n-i-1}; q)]. \end{aligned} \quad (4.4.16)$$

Next we gather some formulas involving the q -Catalan numbers. We begin with the usual recursion for the type-A Catalan numbers, although this q -version may be less widely familiar. It is easily obtained through the interpretation of $\text{Cat}(A_n; q)$ as the descent generating function for 231-avoiding permutations in S_{n+1} , by breaking up the count according to the first entry in the permutation. We omit the details.

Proposition 4.4.36. For $n \geq 1$,

$$\text{Cat}(A_n; q) = (1 + q) \text{Cat}(A_{n-1}; q) + q \sum_{i=1}^{n-1} \text{Cat}(A_{i-1}; q) \text{Cat}(A_{n-i-1}; q). \quad (4.4.17)$$

Furthermore, using known formulas for the Narayana numbers, we obtain a recursion that relates the q -Catalan number in types A and D.

Proposition 4.4.37. For $n \geq 2$,

$$\text{Cat}(D_n; q) = \frac{n+1}{2} (1 + q) \text{Cat}(A_{n-1}; q) - \left(\frac{n-1}{2} + q + \frac{n-1}{2} q^2 \right) \text{Cat}(A_{n-2}). \quad (4.4.18)$$

Proof. Taking the coefficient of q^k on both sides, we see that (4.4.18) is equivalent to

$$\begin{aligned} \text{Nar}_k(D_n) = & \frac{n+1}{2} (\text{Nar}_k(A_{n-1}) + \text{Nar}_{k-1}(A_{n-1})) \\ & - \frac{n-1}{2} \text{Nar}_k(A_{n-2}) - \text{Nar}_{k-1}(A_{n-2}) - \frac{n-1}{2} \text{Nar}_{k-2}(A_{n-2}). \end{aligned} \quad (4.4.19)$$

This can be verified using the known formulas for the type-A and type-D Narayana numbers. (See for example in [34, (9.1)] and [34, (9.3)], putting $m = 1$ in both formulas). \square

Next, we give a recursion for $\text{Cat}(W; q)$ analogous to (4.4.13). The proof follows the outline of the proof of Proposition 4.4.32, using Theorem 4.4.9 instead of Theorem 4.4.2. This proof is simpler than the proof of Proposition 4.4.32, so we omit the details.

Proposition 4.4.38. *For an irreducible finite Coxeter group W and a simple generator s , the q -analog of the W -Catalan number satisfies*

$$\begin{aligned} \text{Cat}(W; q) = & (1 + q) \text{Cat}(W_{S \setminus \{s\}}; q) \\ & + \sum_{S_0} \text{Cat}^+(W_{S_0}; q) \prod_{i=1}^m (1 + q) \text{Cat}(W_{S_i \setminus \{s_i\}}; q), \end{aligned} \quad (4.4.20)$$

where the sum is over all connected subgraphs S_0 of the diagram for W with $s \in S_0$, the connected components of the complement of S_0 in the diagram are S_1, \dots, S_m , and each s_i is the unique vertex in S_i that is connected by an edge to a vertex in S_0 .

The following three propositions give the type-A, type-B, and type-D cases of (4.4.20).

Proposition 4.4.39. *For $n \geq 0$,*

$$\text{Cat}(A_n; q) = \text{Cat}^+(A_n; q) + (1 + q) \sum_{i=0}^{n-1} \text{Cat}^+(A_i; q) \text{Cat}(A_{n-i-1}; q). \quad (4.4.21)$$

Proof. If $n = 0$, then the formula is $\text{Cat}(A_0; q) = \text{Cat}^+(A_0; q)$, which says $1 = 1$. Otherwise, taking s to be a leaf of the A_n diagram in (4.4.20), the sum over S_0 has the following terms: $\sum_{i=1}^{n-1} \text{Cat}^+(A_i; q)(1+q) \text{Cat}(A_{n-i-1}; q)$ and $\text{Cat}^+(A_n; q)$. Because $\text{Cat}^+(A_0; q) = 1$, we can merge the first term into the sum. \square

Proposition 4.4.40. *For $n \geq 0$,*

$$\text{Cat}(B_n; q) = \text{Cat}^+(B_n; q) + (1 + q) \sum_{i=0}^{n-1} \text{Cat}^+(A_i; q) \text{Cat}(B_{n-i-1}; q) \quad (4.4.22)$$

Proof. The formula holds for $n = 0$ and $n = 1$. For $n > 1$, take s to be the leaf whose deletion leaves a diagram of type B_{n-1} in (4.4.20), and rearrange the formula as in the proof of Proposition 4.4.21. \square

Proposition 4.4.41. *For $n \geq 3$,*

$$\begin{aligned}
\text{Cat}(D_n; q) = & \\
& (1+q) \text{Cat}(A_{n-1}; q) + (1+q) \text{Cat}^+(A_{n-1}; q) + \text{Cat}^+(D_n; q) \\
& + (1+q)^2 \sum_{i=1}^{n-2} \text{Cat}^+(A_i; q) \text{Cat}(A_{n-i-2}; q) \\
& + (1+q) \sum_{i=3}^{n-1} \text{Cat}^+(D_i; q) \text{Cat}(A_{n-i-1}; q).
\end{aligned} \tag{4.4.23}$$

Proof. Start with Proposition 4.4.38, taking s to be a leaf of the D_n diagram whose removal leaves the diagram for A_{n-1} . For S_0 not containing the leaf symmetric to s , we get terms $(1+q)^2 \sum_{i=1}^{n-2} \text{Cat}^+(A_i; q) \text{Cat}(A_{n-i-2}; q)$ and $(1+q) \text{Cat}^+(A_{n-1}; q)$. (The $i = 1$ term in the sum would be wrong, except that $\text{Cat}^+(A_1) = 0$.) For S_0 containing the leaf symmetric to s , we get $(1+q) \sum_{i=3}^{n-1} \text{Cat}^+(D_i; q) \text{Cat}(A_{n-i-1}; q)$ and $\text{Cat}^+(D_n; q)$. \square

4.4.7 The double-positive Catalan numbers

In this section, we consider the double-positive Catalan numbers for the classical reflection groups, and establish some identities for $\text{Cat}^+(A_n; q)$, $\text{Cat}^+(B_n; q)$, and $\text{Cat}^+(D_n; q)$ that will be useful for proving the type-D case of Theorem 4.1.4.

Remark 4.4.42. Athanasiadis and Savvidou, in [7, Theorem 1.2], gave formulas for the polynomials $\text{Cat}^+(W; q)$ for each W of finite type by explicitly determining coefficients ξ_i such that $\text{Cat}^+(W; q) = \sum_{i=0}^{\lfloor n/2 \rfloor} \xi_i q^i (1+q)^{n-2i}$. Similar formulas for the relevant polynomials $\text{Cat}(W; q)$ are known [64, Propositions 11.14–11.15], so the identities we need can in principle be obtained by manipulating the formulas from [7, 64]. Indeed, Proposition 4.4.44 is easily obtained in this way, but such proofs of Propositions 4.4.43 and 4.4.51 appear to be more complicated.

Table 4.6: Some double positive Catalan numbers

W	A_0	A_1	A_2	A_3	A_4	A_5	A_6	B_2	B_3	B_4	B_5	B_6	D_4	D_5	D_6	D_7
$\text{Cat}^+(W)$	1	0	1	2	6	18	57	2	6	22	80	296	10	42	168	660

In Table 4.6, we list some examples of the double-positive Catalan numbers for the classical reflection groups. From inspection of these numbers, several interesting relationships appear. First, the data suggests that $2 \text{Cat}^+(A_n) + \text{Cat}^+(A_{n-1}) = \text{Cat}(A_{n-1})$. Below, we establish a q -analog of this identity.

Proposition 4.4.43. *For $n \geq 1$,*

$$(1+q) \text{Cat}^+(A_n; q) + q \text{Cat}^+(A_{n-1}; q) = q \text{Cat}(A_{n-1}; q). \quad (4.4.24)$$

Proof. If $n = 1$, then the identity is $(1+q) \cdot 0 + q \cdot 1 = q \cdot 1$. If $n > 1$, then by induction, we can replace $(1+q) \text{Cat}^+(A_i; q)$ with $q(\text{Cat}(A_{i-1}; q) - \text{Cat}^+(A_{i-1}; q))$ in the terms $i > 1$ of (4.4.21) and observe that $\text{Cat}^+(A_0; q) = 1$ to obtain

$$\begin{aligned} \text{Cat}(A_n; q) &= \text{Cat}^+(A_n; q) + (1+q) \text{Cat}(A_{n-1}; q) \\ &\quad + q \sum_{i=1}^{n-1} \text{Cat}(A_{i-1}; q) \text{Cat}(A_{n-i-1}; q) \\ &\quad - q \sum_{i=1}^{n-1} \text{Cat}^+(A_{i-1}; q) \text{Cat}(A_{n-i-1}; q). \end{aligned}$$

The first sum, by Proposition 4.4.36, is $(\text{Cat}(A_n; q) - (1+q) \text{Cat}(A_{n-1}; q))$. The second sum can be reindexed to $q \sum_{i=0}^{n-2} \text{Cat}^+(A_i; q) \text{Cat}(A_{n-i-2}; q)$, which, by Proposition 4.4.39, equals $\frac{q}{1+q}(\text{Cat}(A_{n-1}; q) - \text{Cat}^+(A_{n-1}; q))$. We obtain

$$\begin{aligned} \text{Cat}(A_n; q) &= \text{Cat}^+(A_n; q) + (1+q) \text{Cat}(A_{n-1}; q) \\ &\quad + \text{Cat}(A_n; q) - (1+q) \text{Cat}(A_{n-1}; q) \\ &\quad - \frac{q}{1+q}(\text{Cat}(A_{n-1}; q) - \text{Cat}^+(A_{n-1}; q)), \end{aligned}$$

which simplifies to the desired identity. \square

The data also suggests that $\text{Cat}^+(D_n) = (n-2) \text{Cat}(A_{n-2})$. Indeed, the following is a q -analog.

Proposition 4.4.44. *For $n \geq 2$,*

$$\text{Cat}^+(D_n; q) = (n-2)q \text{Cat}(A_{n-2}; q). \quad (4.4.25)$$

Proof. For $n = 2$, the identity is $q + q^2 = (3-2)q(1+q)$. If $n \geq 3$, then we start with (4.4.23).

The first summation in the formula can be rewritten, using (4.4.21), as

$$(1+q)(\text{Cat}(A_{n-1};q) - (1+q)\text{Cat}(A_{n-2};q) - \text{Cat}^+(A_{n-1};q)). \quad (4.4.26)$$

By induction, the second summation can be rewritten as

$$(1+q)q \sum_{i=3}^{n-1} (i-2) \text{Cat}^+(A_{i-2};q) \text{Cat}(A_{n-i-1};q). \quad (4.4.27)$$

To further simplify (4.4.27), we use (4.4.17) to calculate

$$\begin{aligned} & (n-3)(\text{Cat}(A_{n-1};q) - (1+q)\text{Cat}(A_{n-2};q)) \\ &= (n-3)q \sum_{i=1}^{n-2} \text{Cat}(A_{i-1};q) \text{Cat}(A_{n-i-2};q) \\ &= q \sum_{i=1}^{n-2} ((i-1) \text{Cat}(A_{i-1};q) \text{Cat}(A_{n-i-2};q) \\ & \quad + (n-i-2) \text{Cat}(A_{i-1};q) \text{Cat}(A_{n-i-2};q)) \\ &= q \sum_{i=1}^{n-2} (i-1) \text{Cat}(A_{i-1};q) \text{Cat}(A_{n-i-2};q) \\ & \quad + q \sum_{i=1}^{n-2} (n-i-2) \text{Cat}(A_{i-1};q) \text{Cat}(A_{n-i-2};q) \end{aligned}$$

Both sums can be reindexed to agree with (4.4.27), except for the initial factor $(1+q)$. Thus (4.4.27) equals $\frac{n-3}{2}(1+q)(\text{Cat}(A_{n-1};q) - (1+q)\text{Cat}(A_{n-2};q))$. Finally, we use (4.4.18) to rewrite the $\text{Cat}(D_n; q)$. We obtain

$$\begin{aligned} & \frac{n+1}{2}(1+q)\text{Cat}(A_{n-1};q) - \left(\frac{n-1}{2} + q + \frac{n-1}{2}q^2\right)\text{Cat}(A_{n-2}) = \\ & \quad (1+q)\text{Cat}(A_{n-1};q) + (1+q)\text{Cat}^+(A_{n-1};q) + \text{Cat}^+(D_n; q) \\ & \quad + (1+q)(\text{Cat}(A_{n-1};q) - (1+q)\text{Cat}(A_{n-2};q) - \text{Cat}^+(A_{n-1};q)) \\ & \quad + \frac{n-3}{2}(1+q)(\text{Cat}(A_{n-1};q) - (1+q)\text{Cat}(A_{n-2};q)). \quad (4.4.28) \end{aligned}$$

This can be rearranged to say $\text{Cat}^+(D_n; q) = (n-2)q \text{Cat}(A_{n-2}; q)$. \square

In order to establish a needed identity for double-positive Catalan numbers of type B, we need a recursion for the q -Catalan number that comes from a completely different direction. The q -Catalan numbers $\text{Cat}(W; q)$ encode the h -vector of the generalized associahedron for W . (See, for example, [32, Section 5.2].) For each Coxeter group W of rank n and each i from 0 to n , define f_i to be the number of simplices in the simplicial generalized associahedron having

exactly i vertices (and thus dimension $i - 1$). Define a polynomial

$$f(W; x) = \sum_{k=0}^n f_k(W) x^k.$$

The following is [35, Proposition 3.7].

Proposition 4.4.45. *If W is reducible as $W_1 \times W_2$, then $f(W; x) = f(W_1; x)f(W_2; x)$. If W is irreducible with Coxeter number h , then*

$$\frac{df(W; x)}{dx} = \frac{h+2}{2} \sum_{s \in S} f(W_{S \setminus \{s\}}; x) \quad (4.4.29)$$

Since $f(W)$ encodes the f -vector of the generalized associahedron and $\text{Cat}(W; q)$ encodes the h -vector, (4.4.29) implies a formula for $\text{Cat}(W; q)$. Since $f(W)$ has coefficients reversed from the f -polynomial usually used to define h -vectors, the formula for $\text{Cat}(W; q)$ is somewhat more complicated than (4.4.29).

Proposition 4.4.46. *For an irreducible Coxeter group W with rank $n \geq 0$ and Coxeter number h , the q -analog of the Catalan number satisfies*

$$n \text{Cat}(W; q) + (1 - q) \frac{d}{dq} \text{Cat}(W; q) = \frac{h+2}{2} \sum_{s \in S} \text{Cat}(W_{S \setminus \{s\}}; q). \quad (4.4.30)$$

Proof. We begin with the right side of (4.4.30) and replace q by $x + 1$ throughout. The result is $\frac{h+2}{2} \sum_{s \in S} \text{rev}(f(W_{S \setminus \{s\}}; x))$, where rev is the operator that reverses the coefficients of a polynomial. In other symbols: $x^{n-1} \frac{h+2}{2} \sum_{s \in S} f(W_{S \setminus \{s\}}; x^{-1})$ Using (4.4.29), the quantity becomes $x^{n-1} \frac{df(W; x^{-1})}{d(x^{-1})}$.

Similarly, $\text{Cat}(W; x + 1) = \text{rev}(f(W; x)) = x^n f(W; x^{-1})$, so $f(W; x^{-1}) = x^{-n} \text{Cat}(W; x + 1)$. Thus the right side of (4.4.30) equals

$$\begin{aligned} & x^{n-1} \frac{d}{d(x^{-1})} [x^{-n} \text{Cat}(W; x + 1)] \\ &= x^{n-1} \frac{d}{dx} [x^{-n} \text{Cat}(W; x + 1)] (-x^2) \\ &= -x^{n+1} \left[-nx^{-n-1} \text{Cat}(W; x + 1) + x^{-n} \frac{d}{dx} \text{Cat}(W; x + 1) \right] \\ &= n \text{Cat}(W; x + 1) - x \frac{d}{dx} \text{Cat}(W; x + 1) \\ &= n \text{Cat}(W; x + 1) - x \frac{d}{d(x+1)} \text{Cat}(W; x + 1) \end{aligned}$$

Replacing x by $q - 1$ throughout, we obtain the left side of (4.4.30). □

The type-B version of (4.4.30) is the following recursion:

Proposition 4.4.47. *For $n \geq 0$,*

$$n \operatorname{Cat}(B_n; q) + (1 - q) \frac{d}{dq} \operatorname{Cat}(B_n; q) = (n + 1) \sum_{i=1}^n \operatorname{Cat}(A_{i-1}; q) \operatorname{Cat}(B_{n-i}; q). \quad (4.4.31)$$

The following formula is obtained using known formulas for Narayana numbers of types A and B.

Proposition 4.4.48. *For $n \geq 0$,*

$$\sum_{i=1}^n \operatorname{Cat}(A_{i-1}; q) \operatorname{Cat}(B_{n-i}; q) = n \operatorname{Cat}(A_{n-1}; q). \quad (4.4.32)$$

Proof. By (4.4.31), the assertion is equivalent to

$$n \operatorname{Cat}(B_n; q) + (1 - q) \frac{d}{dq} \operatorname{Cat}(B_n; q) = n(n + 1) \operatorname{Cat}(A_{n-1}; q). \quad (4.4.33)$$

Taking the coefficient of q^k on both sides, we see that (4.4.33) is equivalent to

$$(n - k) \operatorname{Nar}_k(B_n) + (k + 1) \operatorname{Nar}_{k+1}(B_n) = n(n + 1) \operatorname{Nar}_k(A_{n-1}). \quad (4.4.34)$$

This can be verified using the formulas for the type-A and type-B Narayana numbers, found for example in [34, (9.1)] and [34, (9.2)] (setting $m = 1$ in both formulas from [34]). \square

Using (4.4.16) and the observation that $\operatorname{biCat}(A_n; q) = \operatorname{Cat}(B_n; q)$, then applying (4.4.22) twice, (where, in the first instance n is replaced by $n + 1$ in (4.4.22)), we obtain the following formula.

Proposition 4.4.49. *For $n \geq 1$,*

$$\begin{aligned} \operatorname{Cat}(B_n; q) &= (1 + q) \operatorname{Cat}(B_{n-1}; q) - (1 + q) \operatorname{Cat}^{++}(A_{n-1}; q) \\ &\quad + \operatorname{Cat}^{++}(B_n; q) + (1 + q) \operatorname{Cat}^{++}(B_{n-1}; q). \end{aligned} \quad (4.4.35)$$

Next, we obtain the following formula.

Proposition 4.4.50. *For $n \geq 2$,*

$$\begin{aligned} (1 + q) \operatorname{Cat}(B_n; q) &= \\ &= (1 + q + q^2) \operatorname{Cat}(B_{n-1}; q) + (n - 1)q(1 + q) \operatorname{Cat}(A_{n-2}; q) \\ &\quad + q \operatorname{Cat}^{++}(B_{n-1}; q) + (1 + q) \operatorname{Cat}^{++}(B_n; q). \end{aligned} \quad (4.4.36)$$

Proof. Using (4.4.24) to replace each instance of $(1+q) \text{Cat}^+(A_i; q)$ in (4.4.22) with the difference $q(\text{Cat}(A_{i-1}; q) - \text{Cat}^+(A_{i-1}; q))$ for $i > 0$ and splitting into two sums, we obtain:

$$\begin{aligned} \text{Cat}(B_n; q) &= (1+q) \text{Cat}(B_{n-1}; q) + q \sum_{i=1}^{n-1} \text{Cat}(A_{i-1}; q) \text{Cat}(B_{n-i-1}; q) \\ &\quad - q \sum_{i=1}^{n-1} \text{Cat}^+(A_{i-1}; q) \text{Cat}(B_{n-i-1}; q) + \text{Cat}^+(B_n; q) \end{aligned}$$

We use (4.4.32) with n replaced by $n-1$ to evaluate the first sum. We reindex the second sum and evaluate it using (4.4.22) with n replaced by $n-1$.

$$\begin{aligned} \text{Cat}(B_n; q) &= (1+q) \text{Cat}(B_{n-1}; q) + q(n-1) \text{Cat}(A_{n-2}; q) \\ &\quad - \frac{q}{1+q} (\text{Cat}(B_{n-1}; q) - \text{Cat}^+(B_{n-1}; q)) + \text{Cat}^+(B_n; q). \end{aligned}$$

We multiply through by $(1+q)$ and simplify to obtain (4.4.36). □

Solving both (4.4.36) and (4.4.35) for $(1+q) \text{Cat}^+(B_n; q)$ and combining them, then solving for $(1+q+q^2) \text{Cat}^+(B_{n-1}; q)$, we obtain the key result for $\text{Cat}^+(B_{n-1})$.

Proposition 4.4.51.

$$\begin{aligned} (1+q+q^2) \text{Cat}^+(B_{n-1}; q) &= -q \text{Cat}(B_{n-1}; q) \\ &\quad + (n-1)q(1+q) \text{Cat}(A_{n-2}; q) + (1+q)^2 \text{Cat}^+(A_{n-1}) \end{aligned} \quad (4.4.37)$$

4.4.8 The Type D biCatalan number

We now complete the proof of Theorem 4.1.4 by proving the following theorem.

Theorem 4.4.52. *For $n \geq 2$, the D_n -biCatalan number is*

$$\text{biCat}(D_n) = 6 \cdot 4^{n-2} - 2 \binom{2n-4}{n-2}. \quad (4.4.38)$$

Since we have already established the type-A and type-B cases of Theorem 4.1.4, Theorem 4.4.52 is the assertion that $\text{biCat}(D_n) = 3 \text{biCat}(B_{n-1}) - 2 \text{biCat}(A_{n-2})$. In preparation for the proof, we let $X = X(q)$ and $Y = Y(q)$ be any rational functions of q and define, for each $n \geq 2$, a rational function $Z_n = Z_n(q)$ given by

$$Z_n = \text{biCat}(D_n; q) - X \text{biCat}(B_{n-1}; q) + Y \text{biCat}(A_{n-2}; q).$$

Combining (4.4.14), (4.4.15), and (4.4.16), we obtain the following recursion for Z_n for $n \geq 3$.

$$\begin{aligned} Z_n &= (1+q)Z_{n-1} + \sum_{i=1}^{n-3} \text{Cat}^+(A_i; q) (Z_{n-i} + (1+q)Z_{n-i-1}) \\ &\quad + 2((1+q)^2 - X(1+q) + Y) \text{Cat}^+(A_{n-2}) + 4(1+q) \text{Cat}^+(A_{n-1}) \\ &\quad + 2 \text{Cat}^+(D_n; q) - 2X \text{Cat}^+(B_{n-1}; q) \end{aligned} \quad (4.4.39)$$

One way to obtain a formula for q -biCatalan numbers $\text{biCat}(D_n; q)$ would be to find a choice of X and Y that makes this recursion for Z_n into something that can be solved. We have thus far been unable to find a choice of X and Y that works. Instead, we will prove Theorem 4.4.52 by showing that if $X(1) = 3$ and $Y(1) = 2$, then $Z_n(1) = 0$ for all $n \geq 2$. In the proof that follows, we take convenient choices of X and Y but delay specializing q to 1 until the end, because specializing earlier does not make the manipulations much easier, and because we hope that perhaps we are still getting closer to a formula for $\text{biCat}(D_n; q)$.

Proof of Theorem 4.4.52. Substituting (4.4.37) and (4.4.25) into (4.4.39), taking $X = 1 + q + q^2$, and taking $Y = 2q - q^2 + q^3$, we obtain

$$\begin{aligned} Z_n &= (1+q)Z_{n-1} + \sum_{i=1}^{n-3} \text{Cat}^+(A_i; q) (Z_{n-i} + (1+q)Z_{n-i-1}) \\ &\quad + 2q(1-q) \text{Cat}^+(A_{n-2}; q) + 2(1-q)(1+q) \text{Cat}^+(A_{n-1}; q) \\ &\quad - 2q(1+(n-1)q) \text{Cat}(A_{n-2}; q) + 2q \text{Cat}(B_{n-1}; q) \end{aligned} \quad (4.4.40)$$

We next apply (4.4.24) to rewrite the two double-positive q -Catalan numbers in (4.4.40) as a single q -Catalan number.

$$\begin{aligned} Z_n &= (1+q)Z_{n-1} + \sum_{i=1}^{n-3} \text{Cat}^+(A_i; q) (Z_{n-i} + (1+q)Z_{n-i-1}) \\ &\quad + 2q(1-q) \text{Cat}(A_{n-1}; q) - 2q(1+(n-1)q) \text{Cat}(A_{n-2}; q) + 2q \text{Cat}(B_{n-1}; q) \end{aligned} \quad (4.4.41)$$

Finally specializing q to 1 and using the fact that $\text{Cat}(B_{n-1}) = n \text{Cat}(A_{n-2})$ for $n \geq 3$ (which is immediate from the well-known formulas for the type-A and type-B Catalan numbers), we see that

$$Z_n(1) = 2Z_{n-1}(1) + \sum_{i=1}^{n-3} \text{Cat}^+(A_i) (Z_{n-i}(1) + 2Z_{n-i-1}(1)) \quad (4.4.42)$$

We easily verify that $Z_2(1) = 0$, and thus we have a simple inductive proof that $Z_n(1) = 0$ for all $n \geq 2$. Since we chose X and Y to have $X(1) = 3$ and $Y(1) = 2$, we obtain the desired identity $\text{biCat}(D_n) = 3 \text{biCat}(B_{n-1}) - 2 \text{biCat}(A_{n-2})$. \square

4.4.9 Type-D biNarayana numbers

Computational evidence suggests the following modest conjecture on the type-D biNarayana number $\text{biNar}_k(D_n)$.

Conjecture 4.4.53. *The type-D biNarayana number $\text{biNar}_k(D_n)$ is a polynomial in n (for $n \geq 2$) of degree $2k$ and leading coefficient $\frac{4^k}{(2k)!}$.*

If Conjecture 4.4.53 is true, then Table 4.7 shows $\frac{(2k)!}{2^k} \cdot \text{biNar}_k(D_n)$ for small k . The $k = 1$ case is verified by Proposition 4.2.4, and with some effort, the $k = 2$ case can be proved as well.

Table 4.7: $\frac{(2k)!}{2^k} \cdot \text{biNar}_k(D_n)$ for small k

k	$\frac{(2k)!}{2^k} \cdot \text{biNar}_k(D_n)$
0	1
1	$2n^2 - 3n$
2	$4n^4 - 20n^3 + 35n^2 - 7n - 24$
3	$8n^6 - 84n^5 + 365n^4 - 705n^3 + 212n^2 + 1104n - 1080$
4	$16n^8 - 288n^7 + 2268n^6 - 9576n^5 + 20349n^4 - 8022n^3 - 54133n^2 + 104826n - 60480$

REFERENCES

- [1] K. Adaricheva and J.B. Nation, *Classes of Semidistributive Lattices*. in Lattice Theory: Special Topics and Applications, ed. G. Grätzer and F. Wehrung, Birkhäuser, to appear.
- [2] D. Armstrong, *Generalized Noncrossing Partitions and Combinatorics of Coxeter Groups*. Mem. Amer. Math. Soc. **202** (2009), no. 949.
- [3] D. Armstrong, C. Stump and H. Thomas, *A uniform bijection between nonnesting and noncrossing partitions*. Trans. Amer. Math. Soc. **365**(8), 2013
- [4] C. A. Athanasiadis, *Generalized Catalan numbers, Weyl groups and arrangements of hyperplanes*. Bull. London Math. Soc. **36** (2004), no. 3, 294–302.
- [5] C. Athanasiadis. *Some combinatorial properties of flag simplicial pseudomanifolds and spheres*. Ark. Mat. **49** (2011), no. 1, 17–29.
- [6] C. A. Athanasiadis, T. Brady, J. McCammond and C. Watt, *h -vectors of generalized associahedra and noncrossing partitions*. Int. Math. Res. Not. **2006** Art. ID 69705, 28 pp.
- [7] C. A. Athanasiadis and C. Savvidou, *The local h -vector of the cluster subdivision of a simplex*. Sémin. Lothar. Combin. **66** (2011/12), Art. B66c, 21 pp.
- [8] C. A. Athanasiadis and E. Tzanaki, *On the enumeration of positive cells in generalized cluster complexes and Catalan hyperplane arrangements*. J. Algebraic Combin. **23** (2006), no. 4, 355–375.
- [9] E. Barnard, A. Carrol, and S. Zhu. *Minimal inclusions of torsion classes*. In preparation.
- [10] G. Baxter, *On fixed points of the composite of commuting functions*. Proc. Amer. Math. Soc. **15** (1964) 851–855.
- [11] A. Björner, *Topological methods in combinatorics, Handbook of Combinatorics*, Vol. 2, 1819–1872, Elsevier, Amsterdam, 1995.
- [12] A. Björner, Brenti, *The Combinatorics of Coxeter Groups*. Graduate Texts in Mathematics, 231, Springer, New York, 2005.
- [13] A. Björner and M. Wachs, *Shellable Nonpure Complexes and Posets I*. Trans. Amer. Math. Soc. **348** (1996), no. 4, 1299–1327.
- [14] A. Björner and M. Wachs, *Shellable nonpure complexes and posets. II*. Trans. Amer. Math. Soc. **349** (1997), no. 10, 3945–3975.
- [15] M. Bousquet-Mélou, A. Claesson, Anders, 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.
- [16] M. Bousquet-Mélou, E. Fusy and L.F Prévaille-Ratelle, *The number of intervals in the m -Tamari lattices*. Electron. J. Combin. **18** (2011), no. 2, Paper 31, 26 pp.

- [17] T. Brady and C. Watt, *From permutahedron to associahedron*. Proc. Edinb. Math. Soc. (2) **53** (2010), no. 2, 299–310.
- [18] A. Brouwer and A. Schrijver, *On the period of an operator, defined on antichains*. Math Centrum report ZW **24/74** (1974).
- [19] P. Cameron and D. Fon-Der-Flaass, *Orbits of antichains revisited*. European J. Combin. **16** (1995), no. 6, 545–554.
- [20] F. Chapoton, S. Fomin and A. Zelevinsky, *Polytopal realizations of generalized associahedra*. Canad. Math. Bull. **45** (2002) no. 4, 537–566.
- [21] R. Charney and M. Davis, *Euler characteristic of a nonpositively curved, piecewise Euclidean manifold*. Pac. J. Math., **171** (1995), 117–137.
- [22] G. Châtel and V. Pilaud, *The Cambrian and Baxter-Cambrian Hopf algebras*. Preprint, 2014. (arXiv:1411.3704)
- [23] W. Y. C. Chen, A. Y. Z. Wang and A. F. Y. Zhao, *Identities Derived from Noncrossing Partitions of Type B*. Electron. J. Combin. **18** (2011), no. 1, Paper 129, 17 pp.
- [24] F. R. K. Chung, R. L. Graham, V. E. Hoggatt Jr. and M. Kleiman, *The number of Baxter permutations*. J. Combin. Theory Ser. A **24** (1978), no. 3, 382–394.
- [25] C. Coker *Enumerating a class of lattice paths*. Discrete Math. **271** (2003), no. 1-3, 13–28.
- [26] C. Le Conte de Poly-Barbut, *Sur les Treillis de Coxeter Finis* (French), Math. Inf. Sci.hum. 32 no. **125** (1994), 41-57.
- [27] A. Day, *Splitting lattices generate all lattices*. Algebra Universalis **7** (1977), no. 2, 163–169.
- [28] A. Day, *A Simple Solution to the Word Problem for Lattices*. Canad. Math. Bull. **13** (1970), 253–254.
- [29] K. Dilks, *Involutions on Baxter Objects*. Preprint, 2014. (arXiv:1402.2961)
- [30] S. Dulucq and O. Guibert, *Stack words, standard tableaux and Baxter permutations*. Discrete Math. **157** (1996), no. 1–3, 91–106.
- [31] P. Edelman and V. Reiner, *H-Shellings and h Complexes*. Advances in Mathematics. **106**, (1994), 36–64.
- [32] S. Fomin and N. Reading, *Root systems and generalized associahedra*. IAS/Park City Math. Ser. **13**, 63–131.
- [33] S. Fomin, and N. Reading, *Root systems and generalized associahedra*. Geometric combinatorics, 63–131, IAS/Park City Math. Ser., 13, Amer. Math. Soc., Providence, RI, 2007.
- [34] S. Fomin and N. Reading, *Generalized cluster complexes and Coxeter combinatorics*. Int. Math. Res. Not. **2005**, no. 44, 2709–2757.

- [35] S. Fomin and A. Zelevinsky, *Y-systems and generalized associahedra*. Ann. of Math. (2) **158** (2003), no. 3, 977–1018.
- [36] S. Fomin and A. Zelevinsky, *Cluster Algebras II: Finite Type Classification*. Inventiones Mathematicae **154** (2003), 63–121.
- [37] D. Fon-Der-Flaass, *Orbits of antichains in ranked posets*. European J. Combin. **14** (1993), no. 1, 17–22.
- [38] R. Forman, *Topics in combinatorial differential topology and geometry*. Geometric combinatorics, 133–205, IAS/Park City Math. Ser., 13, Amer. Math. Soc., Providence, RI, 2007.
- [39] R. Freese, J. Jezek and J. Nation, *Free lattices*. Mathematical Surveys and Monographs, **42**, American Mathematical Society, 1995.
- [40] H. Friedman and D. Tamari, *Problèmes d’associativité: Une structure de treillis finis induite par une loi demi-associative*. J. Combinatorial Theory **2** 1967 215–242.
- [41] A. Garver, T. McConville, *Lattice properties of oriented exchange graphs and torsion classes*. Preprint available: arxiv:1507.04268v1.
- [42] W. Geyer, *On Tamari lattices*. Discrete Math. **133** (1994), no. 1-3, 99–122.
- [43] S. Giraudo, *Algebraic and combinatorial structures on pairs of twin binary trees*. J. of Algebra **360** (2012) 115–157.
- [44] G. Grätzer, *Lattice theory: foundation*. Birkhäuser/Springer Basel AG, Basel, 2011.
- [45] G. Grätzer and F. Wehrung, *A new lattice construction: the box product*. J. Algebra **221** (1999), no. 1, 315–344.
- [46] M. D. Haiman, *Conjectures on the quotient ring by diagonal invariants*. J. Algebraic Combin. **3** (1994), no. 1, 17–76.
- [47] P. Hersh and K. Mészáros *SB-labelings and posets with each interval homotopy equivalent to a sphere or a ball*. arXiv:1407.5311.
- [48] C. Hohlweg and C. E. M. C Lange, *Realizations of the associahedron and cyclohedron*. Discrete Comput. Geom. **37** (2007), no. 4, 517–543.
- [49] C. Hohlweg, C. Lange and H. Thomas, *Permutahedra and generalized associahedra*. Adv. Math. **226** (2011), no. 1, 608–640.
- [50] S. Huang and D. Tamari, *Problems of associativity: A simple proof for the lattice property of systems ordered by a semi-associative law*. J. Combinatorial Theory Ser. A **13** (1972), 7–13.
- [51] J. Humphreys, *Reflection groups and Coxeter groups*. Cambridge Studies in Advanced Mathematics, 29. Cambridge University Press, Cambridge, 1990.

- [52] O. Iyama, N. Reading, I. Reiten, and H. Thomas, *Algebraic lattice quotients of Weyl groups coming from preprojective algebras*. In preparation.
- [53] B. Jónsson and J.E. Kiefer, *Finite sublattices of a free lattice*. *Canad. J. Math.* **14** 1962 487–497.
- [54] J. Jonsson, *Simplicial complexes of graphs*. Lectures Notes in Mathematics, Springer-Verlag, Berlin, 2008.
- [55] S. E. Law and N. Reading, *The Hopf algebra of diagonal rectangulations*. *J. Combin. Theory Ser. A* **119** (2012), no. 3, 788–824.
- [56] R. Marsh, M. Reineke and A. Zelevinsky, *Generalized associahedra via quiver representations*. *Trans. Amer. Math. Soc.* **355** (2003) no. 10, 4171–4186.
- [57] T. McConville *Crosscut-Simplicial Lattices*. arXiv:1409.6269v1
- [58] H. Mühle, *On the Poset of Multichains*. Preprint available: arXiv:1506.04276
- [59] *Associahedra, Tamari lattices and related structures*. Tamari memorial Festschrift. Edited by F. Muller-Hoissen, J. M. Pallo, J. Stasheff, Progress in Mathematical Physics, 299. Birkhuser/Springer, Basel, 2012.
- [60] Pak, I. *A History of the Catalan numbers*. Appendix B in *Catalan numbers*. Cambridge University Press, New York, 2015. viii+215 pp.
- [61] D. I. Panyushev, *ad-Nilpotent ideals of a Borel subalgebra: generators and duality*. *J. Algebra* **274** (2004), no. 2, 822–846.
- [62] D. Panyushev, *On orbits of antichains of positive roots*. *European J. Combin.* **30** (2009), no. 2, 586–594.
- [63] K. Petersen, *Eulerian numbers*. Birkhäuser Advanced Texts: Basler Lehrbücher. Birkhäuser/Springer, New York, 2015.
- [64] A. Postnikov, V. Reiner, and L. Williams, *Faces of generalized permutohedra*. *Doc. Math.* **13** (2008), 207–273.
- [65] E. N. Sommers, *B-stable ideals in the nilradical of a Borel subalgebra*. *Canad. Math. Bull.* **48** (2005), no. 3, 460–472.
- [66] N. Reading, *Lattice congruences of the weak order*. *Order* **21** (2004) no. 4, 315–344.
- [67] N. Reading, *Lattice congruences, fans and Hopf algebras*. *J. Combin. Theory Ser. A* **110** (2005) no. 2, 237–273.
- [68] N. Reading, *Cambrian lattices*. *Adv. Math.* **205** (2006) no. 2, 313–353.
- [69] N. Reading, *Clusters, Coxeter-sortable elements and noncrossing partitions*. *Trans. Amer. Math. Soc.* **359** (2007), no. 12, 5931–5958.

- [70] N. Reading, *Sortable elements and Cambrian lattices*. Algebra Universalis **56** (2007), no. 3-4, 411–437.
- [71] N. Reading, *Noncrossing partitions and the shard intersection order*. J. Algebraic Combin. **33** (2011), no. 4, 483–530.
- [72] N. Reading, *Noncrossing arc diagrams and canonical join representations*. SIAM J. Discrete Math. **29** (2015), no. 2, 736–750.
- [73] N. Reading *Lattice Theory of the Poset of Regions*. in Lattice Theory: Special Topics and Applications, ed. G. Grätzer and F. Wehrung, Birkhäuser, to appear.
- [74] N. Reading and D. E. Speyer, *Cambrian Fans* J. Eur. Math. Soc. (JEMS) **11** no. 2, 407–447.
- [75] N. Reading and D. E. Speyer, *Sortable elements in infinite Coxeter groups* Trans. Amer. Math. Soc. **363** (2011) no. 2, 699–761.
- [76] N. Reading and D. E. Speyer, *Cambrian frameworks for cluster algebras of affine type*. Preprint, 2015. (arXiv:1504.00260)
- [77] B. Salvy and P. Zimmermann. *GFUN: a Maple package for the manipulation of generating and holonomic functions in one variable*. ACM Trans. Math. Softw. **20** (1994), no. 2, 163–177.
- [78] L. Santocanale, *Derived semidistributive lattices*. Algebra Universalis **63** (2010), no. 2-3, 101–130.
- [79] R. P. Stanley, *Subdivisions and local h-vectors*. J. Amer. Math. Soc. **5** (1992), 805–851.
- [80] R. Stanley, *Promotion and evacuation*, Electron. J. Combin. **16** (2009), no. 2, R9.
- [81] R. P. Stanley, Enumerative combinatorics. Vol. 1, second edition. Cambridge Studies in Advanced Mathematics **49**. Cambridge University Press, Cambridge, 2012.
- [82] R. P. Stanley, Enumerative combinatorics. Vol. 2, second edition. Cambridge Studies in Advanced Mathematics **49**. Cambridge University Press, Cambridge, 2012.
- [83] R.P. Stanley, *Catalan numbers*. Cambridge University Press, New York, 2015. viii+215 pp.
- [84] J. Stembridge, *Maple packages for symmetric functions, posets, root systems, and finite Coxeter groups*. Available at <http://www.math.lsa.umich.edu/~jrs/maple.html>.
- [85] J. Stembridge, *Quasi-minuscule quotients and reduced words for reflections*. J. Algebraic Combin. **13** (2001), no. 3, 275–293.
- [86] J. Striker and N. Williams, *Promotion and Rowmotion*. Eur. J. Of Combinatorics, **33**, (2012) no. 8, 1919–1942.
- [87] H. Thomas, *Tamari lattices and noncrossing partitions in type B*. Discrete Mathematics, **306** (2006), no. 21, 2711–2723.

- [88] W. Trotter, *Combinatorics and Partially Ordered Sets: Dimension Theory*. The Johns Hopkins University Press, 1992.
- [89] M. Wachs, *Poset Topology: Tools and Applications*, Geometric combinatorics, 497–615, IAS/Park City Math. Ser., 13, Amer. Math. Soc., Providence, RI, 2007.
- [90] F. Wehrung, *From join-irreducibles to dimension theory for lattices with chain conditions*. J. Algebra Appl. **1** (2002), no. 2, 215–242.
- [91] J. West, personal communication, 2006.
- [92] N. Williams, *Cataland*. Ph.D. Thesis, University of Minnesota, 2013.
- [93] S. Yang and A. Zelevinsky *Cluster algebras of finite type via Coxeter elements and principal minors*. Transformation Groups **13** (2008), no. 3–4, 855–895.
- [94] G. Ziegler, *Lectures on polytopes*. Graduate Texts in Mathematics, **152**. Springer-Verlag, New York, 1995.