George Grätzer
Friedrich Wehrung
Editors

Lattice Theory: Special Topics and Applications

Volume 2

Birkhäuser

Birkhäuser

George Grätzer • Friedrich Wehrung
Editors

Lattice Theory: Special Topics and Applications

Volume 2

Birkhäuser

Editors
George Grätzer
Department of Mathematics
University of Manitoba
Winnipeg, Manitoba, Canada

Friedrich Wehrung
Department of Mathematics
University of Caen
Caen, France

ISBN 978-3-319-44235-8 ISBN 978-3-319-44236-5 (eBook)
DOI 10.1007/978-3-319-44236-5

Library of Congress Control Number: 2016954128

Mathematics Subject Classification (2010): 06, 06B35, 06D10, 06D22, 06B35, 06D10, 06D22, 05B35, 05E45, 51D20, 06A15, 06B05, 06B10, 06B15, 06C10, 06C20, 16E50

Printed on acid-free paper

This book is published under the trade name Birkhäuser, www.birkhauser-science.com
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Introduction

George Grätzer started writing his *General Lattice Theory* in 1968. It was published in 1978. It set out "to discuss in depth the basics of general lattice theory." Almost 900 exercises, 193 research problems, and a detailed Further Topics and References for each chapter completed the picture.

As T. S. Blyth wrote in the Mathematical Reviews: "*General Lattice Theory* has become the lattice theorist's bible. Now, two decades on, we have the second edition, in which the old testament is augmented by a new testament that is epistolic. The new testament gospel is provided by leading and acknowledged experts in their fields."

Another decade later, Grätzer considered updating the second edition to reflect some exciting and deep developments. "When I started on this project, it did not take me very long to realize that what I attempted to accomplish in 1968–1978, I cannot even try in 2009. To lay the foundation, to survey the contemporary field, to pose research problems, would require more than one volume or more than one person. So I decided to cut back and concentrate in this volume on the *foundation*."

So *Lattice Theory: Foundation* (referenced in this volume as LTF) provides the foundation. Now we complete this project with *Lattice Theory: Special Topics and Applications*, written by a distinguished group of experts, to cover some of the vast areas not in LTF.

Volume 1 (cited as [209]) is divided into three parts and ten chapters:

This book, Volume 2, is divided into ten chapters:

Chapter 1. Varieties of Lattices
by P. Jipsen and H. Rose

Chapter 2. Free and Finitely Presented Lattices
by R. Freese and J. B. Nation

Chapter 3. Classes of Semidistributive Lattices
by K. Adaricheva and J. B. Nation

Chapter 4. Lattices of Algebraic Subsets and Implicational Classes
by K. Adaricheva and J. B. Nation

Chapter 5. Convex Geometries
by K. Adaricheva and J. B. Nation

Chapter 6. Bases of Closure Systems
by K. Adaricheva and J. B. Nation

Chapter 7. Permutohedra and Associahedra
by N. Caspard, L. Santocanale, and F. Wehrung

Chapter 8. Generalizations of the Permutohedron
by L. Santocanale and F. Wehrung

Chapter 9. Lattice Theory of the Poset of Regions
by N. Reading

Chapter 10. Finite Coxeter Groups and the Weak Order
by N. Reading

George Grätzer and Friedrich Wehrung, editors

Contents

Chapter 9

Lattice Theory of the Poset of Regions

N. Reading

Hyperplane arrangements (collections of codimension-1 subspaces) have long been an object of study in combinatorics, topology, and geometry. This chapter explores the lattice theory of the poset of regions of a (real) hyperplane arrangement. We discuss the open problem, first posed by Björner, Edelman, and Ziegler [70], of characterizing by local geometric conditions which posets of regions are lattices. We give a local geometric characterization ("tightness") of which posets of regions are semidistributive lattices. Along the way, we discuss a local condition for checking that a partially ordered set is a lattice, along with analogous local conditions for determining lattice-theoretic properties. In the case of simplicial arrangements (which are in particular tight), we characterize the regions of the arrangement in terms of two notions of combinatorial convexity.

We then turn our attention to lattice congruences on posets of regions, focusing in particular on the tight case. We begin with a discussion of lattice congruences from a combinatorial point of view. We then establish that tight posets of regions have the special property of being polygonal lattices. We discuss how to decompose the hyperplanes in a tight arrangement into shards and show how the polygonal property leads to a geometric characterization of

G. Grätzer, F. Wehrung (eds.), *Lattice Theory: Special Topics and Applications*,
DOI 10.1007/978-3-319-44236-5_9

lattice congruences of tight poset of regions in terms of shards. Finally, we discuss how the geometric characterization in terms of shards is inherited by lattice quotients of the poset of regions.

9-1. Basic notions

9-1.1 Hyperplane arrangements

Definition 9-1.1. A (*linear*) *hyperplane* in $\mathbb{R}^n$ is a linear subspace of dimension $n-1$ (i.e., codimension 1). An *arrangement* of hyperplanes is a finite collection of hyperplanes.

Other important notions of hyperplane arrangements exist in the literature.[1] In the standard terminology, our object of study is a *real, central hyperplane arrangement*, but we mostly omit the adjectives "real" and "central" for arrangements and omit the adjective "linear" for hyperplanes.

A hyperplane arrangement $\mathcal{A}$ is *essential* if the intersection $\bigcap_{H\in\mathcal{A}} H$ of all of the hyperplanes in $\mathcal{A}$ is the origin. There is no harm in requiring $\mathcal{A}$ to be essential; if it is not, then an essential hyperplane arrangement is obtained by taking the quotient of $\mathbb{R}^n$ by the subspace $\bigcap_{H\in\mathcal{A}} H$ and taking the quotient of each hyperplane by $\bigcap_{H\in\mathcal{A}} H$. In general, we do not make this requirement, because it is convenient in specific examples to have the freedom to construct non-essential arrangements. For those results which require an essential arrangement, the corresponding result for non-essential arrangements is easily obtained (but often much less convenient to state) as a corollary. The *rank* of $\mathcal{A}$ is the dimension of the quotient of $\mathbb{R}^n$ by the subspace $\bigcap_{H\in\mathcal{A}} H$, or equivalently the dimension of the linear span of normal vectors to the hyperplanes in $\mathcal{A}$.

Definition 9-1.2. The *complement* of $\mathcal{A}$ is the set $\mathbb{R}^n \setminus (\bigcup_{H\in\mathcal{A}} H)$ of all points not contained in any hyperplane of $\mathcal{A}$. The connected components of the complement are unbounded n-dimensional open sets. The closures of these connected components are called *regions* or $\mathcal{A}$-*regions* and we write $\mathcal{R}(\mathcal{A})$ for the set of regions of $\mathcal{A}$. (Some authors, including the authors of the foundational papers [70, 139, 144], use the term "region" for the connected components themselves rather than their closures.)

Example 9-1.3. Figure 9-1.1 represents a hyperplane arrangement $\mathcal{A}$ in $\mathbb{R}^3$. This picture is obtained as follows: Each hyperplane in $\mathcal{A}$ intersects the unit sphere in $\mathbb{R}^3$ in a great circle. The resulting collection of great circles is stereographically projected to the plane. Since a great circle maps to a circle under stereographic projection (or to a line if the great circle contains the poles), we obtain the diagram shown. The intersection of a region with the

[1] For example, collections of hyperplanes in finite-dimensional vector spaces over other fields, or collections of *affine* hyperplanes (translates of linear hyperplanes).

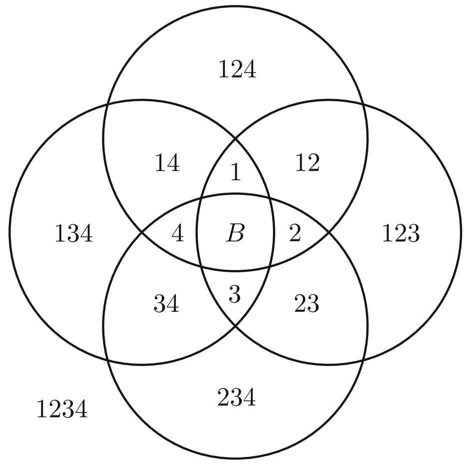

Figure 9-1.1: A hyperplane arrangement

unit sphere is a curvilinear polygon that projects to a possibly unbounded area defined by the projected circles. In Figure 9-1.1, there are 14 (projections of) regions, including the unbounded region outside all of the circles. The labels in the regions are explained later in Example 9-1.14.

9-1.2 Polyhedral geometry

To describe the regions, we need some terminology from polyhedral geometry. (For more information, see for example [464].)

We write bold symbols ($\mathbf{x}$, $\mathbf{y}$, etc.) for vectors in $\mathbb{R}^n$ and the corresponding non-bold symbols, with subscripts, for their entries, so that for example the symbol $\mathbf{x}$ stands for $(x_1, \ldots, x_n) \in \mathbb{R}^n$. We write $\langle \cdot, \cdot \rangle$ for the usual Euclidean inner product on $\mathbb{R}^n$. Given a hyperplane H in $\mathbb{R}^n$, there exists a vector $\mathbf{n}$ (unique up to nonzero scaling) such that $H = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{n}, \mathbf{x} \rangle = 0\}$. This is a *normal vector* to H, and $H = H_{\mathbf{n}}$ is the *hyperplane normal to* $\mathbf{n}$. The hyperplane H defines two *closed* (*linear*) *halfspaces* in $\mathbb{R}^n$, namely the sets $H_{\mathbf{n}}^- = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{n}, \mathbf{x} \rangle \le 0\}$ and $H_{\mathbf{n}}^+ = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{n}, \mathbf{x} \rangle \ge 0\}$, where $\mathbf{n}$ is a normal vector to H.

Definition 9-1.4. A *cone* in $\mathbb{R}^n$ is a set that is closed under addition and positive scaling. A cone is in particular a *convex set*, meaning that if $\mathbf{x}$ and $\mathbf{y}$ are points in the cone, then the entire line segment $\overline{\mathbf{xy}}$ connecting $\mathbf{x}$ and $\mathbf{y}$ is contained in the cone. The *dimension* of a cone is the dimension of the smallest linear subspace containing the cone. The *relative interior* of a cone is its interior relative to the smallest linear subspace containing the cone. A *closed polyhedral cone* in $\mathbb{R}^n$ is a set that can be written as the intersection of finitely many closed linear halfspaces. Equivalently, a closed polyhedral cone is the *nonnegative linear span* (the set of all linear combinations with nonnegative coefficients) of a finite set of vectors. The equivalence of these two definitions of a closed polyhedral cone is non-trivial. (See, for example, [464, Lecture 1].) A closed polyhedral cone is in particular a cone, and thus a convex set. A closed polyhedral cone in $\mathbb{R}^n$ may have dimension less than n. For example, for any nonzero vector $\mathbf{n}$, the cone $H_{\mathbf{n}}^- \cap H_{\mathbf{n}}^+$ is the hyperplane $H_{\mathbf{n}}$.

We will now give a quick definition, sufficient for our needs, of the faces of a closed polyhedral cone. For a more standard definition, see for example [464, Lecture 2.1]. Given a full-dimensional closed polyhedral cone R in $\mathbb{R}^n$, an expression for R as $\bigcap_{\mathbf{n}\in N} H_{\mathbf{n}}^-$ is *non-redundant* if for any proper subset N' of N, the cone $\bigcap_{\mathbf{n}\in N'} H_{\mathbf{n}}^-$ is strictly larger than R. Given a non-redundant expression $\bigcap_{\mathbf{n}\in N} H_{\mathbf{n}}^-$ for R, the *facets* of R are the intersections $H_{\mathbf{n}} \cap R$ for $\mathbf{n} \in N$. Exercise 9.1 is to verify that for each $\mathbf{n} \in N$, the facet $H_{\mathbf{n}} \cap R$ is $(n-1)$-dimensional, and thus full-dimensional in $H_{\mathbf{n}}$. Exercise 9.2 verifies that for any $\mathbf{n}$ such that $H_{\mathbf{n}}^- \supseteq R$ and $H_{\mathbf{n}} \cap R$ is $(n-1)$-dimensional, $H_{\mathbf{n}} \cap R$ is a facet of R. For $\mathbf{n} \in N$, the hyperplane $H_{\mathbf{n}}$ is called a *facet-defining hyperplane* or *boundary hyperplane* for R, and the set of these hyperplanes is written $\mathcal{B}(R)$. Any intersection of facets is called a *face* of R. Any face is itself a closed polyhedral cone. By convention, the intersection of the empty set of facets is interpreted to be R, so that in particular R is a face of itself. It is possible that a given face may be written as an intersection of facets in several different ways. Faces of a closed polyhedral cone C not of full dimension can be defined by considering C as a full-dimensional cone in the smallest linear subspace containing C.

Definition 9-1.5. A full-dimensional closed polyhedral cone in $\mathbb{R}^n$ is *simplicial* if it has exactly n facets, or equivalently if it can be written as the intersection $\bigcap_{\mathbf{n}\in N} H_{\mathbf{n}}^-$ where N is a basis for $\mathbb{R}^n$. Equivalently again, a full-dimensional cone is simplicial if it can be written as the nonnegative linear span of some basis for $\mathbb{R}^n$. In a simplicial cone, each face has a unique expression as an intersection of facets and so there are $\binom{n}{k}$ faces of dimension k for each k from 0 to n. Alternately, each face is the nonnegative linear span of a subset of the basis whose nonnegative linear span is the cone.

9-1.3 Regions

Each region of a hyperplane arrangement is a closed polyhedral cone and is the closure of its interior, which is a connected component of $\mathbb{R}^n \setminus (\bigcup_{H\in\mathcal{A}} H)$. To see why, fix a connected component U of $\mathbb{R}^n \setminus (\bigcup_{H\in\mathcal{A}} H)$ and observe that, for every $H \in \mathcal{A}$, the set U is contained in one of the two open halfspaces defined by H (the two connected components of $\mathbb{R}^n \setminus H$). Indeed, this connected component is the intersection of open halfspaces, one for each $H \in \mathcal{A}$. It follows that the corresponding region is the intersection of closed halfspaces, one for each $H \in \mathcal{A}$.

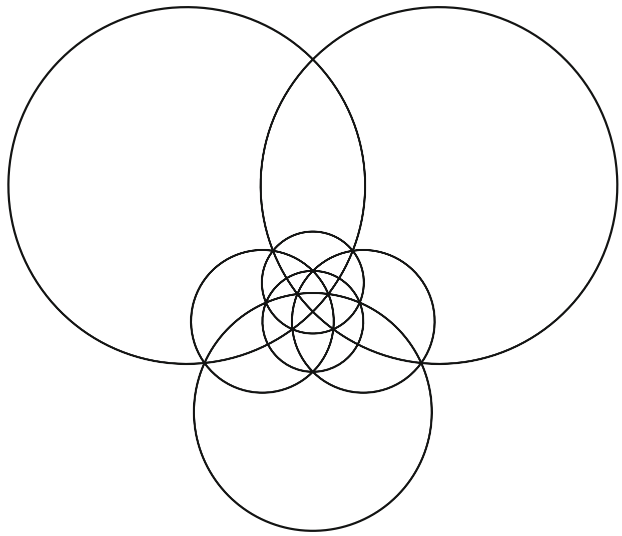

Figure 9-1.2: A simplicial hyperplane arrangement

Definition 9-1.6. A hyperplane arrangement is *simplicial* if every one of its regions is a simplicial cone in the sense of Definition 9-1.5.

Example 9-1.7. Figure 9-1.2 shows a simplicial hyperplane arrangement of rank 3 with 7 hyperplanes. The picture is a stereographic projection as explained in Example 9-1.3.

Proposition 9-1.8. *If Q and R are distinct regions of $\mathcal{A}$ and $Q \cap R$ is $(n-1)$-dimensional, then $Q \cap R$ is a facet of Q and a facet of R. If R is a region, then every facet of R is shared by a unique other region Q.*

Proof. If $Q \cap R$ is $(n-1)$-dimensional, then let $\mathbf{x}$ be a point in the relative interior of $Q \cap R$. Thus for small enough ε, the open ball of radius ε about $\mathbf{x}$ intersects $Q \cap R$ in an $(n-1)$-dimensional ball. By definition, the interiors of regions are disjoint, so this $(n-1)$-dimensional ball forms part of the boundary of Q and of R, and furthermore, the interiors of Q and of R are on opposite sides of the hyperplane H containing the $(n-1)$-dimensional ball. Exercise 9.2 implies that H is a facet-defining hyperplane of Q and of R. If $H' \neq H$ is a hyperplane intersecting the relative interior of the facet $F = R \cap H$ of R, then H' intersects the relative interior of R as well. We see in particular that no facet-defining hyperplane of Q (besides H) intersects the relative interior of F, so $F \subseteq Q \cap H$. By symmetry $Q \cap H \subseteq R \cap H = F$, and we have proved the first assertion.

If F is a facet of R, then let $\mathbf{x}$ be a point in the relative interior of F. Then for small enough ε, the open ball of radius ε about $\mathbf{x}$ intersects no hyperplane of $\mathcal{A}$ besides the facet-defining hyperplane H of F. The hyperplane H cuts the open ball into two halves, one of which is contained in R. The other half is contained in some region Q whose intersection with R is thus $(n-1)$-dimensional. By the first assertion, R and Q share a common facet. Since interiors of regions are disjoint, the region Q is the unique region sharing the facet F with R. □

Definition 9-1.9. A *fan* in $\mathbb{R}^n$ is a collection $\mathcal{F}$ of closed polyhedral cones in $\mathbb{R}^n$ such that (i) if C is a cone in $\mathcal{F}$ and D is a face of C, then D is in $\mathcal{F}$, and (ii) if C and D are cones in $\mathcal{F}$ then $C \cap D$ is a face of C and a face of D. The fan $\mathcal{F}$ is *complete* if the union of the cones in $\mathcal{F}$ is all of $\mathbb{R}^n$. See [464, Lecture 7] for more details about fans. The following theorem is a special case[2] of [373, Theorem 1.2], but can probably be attributed to folklore.

$\diamondsuit$ **Theorem 9-1.10.** *Suppose $\mathcal{M}$ is a collection of full-dimensional cones in $\mathbb{R}^n$ with disjoint interiors, such that $\bigcup \mathcal{M} = \mathbb{R}^n$. Suppose also that whenever two cones in $\mathcal{M}$ have an $(n-1)$-dimensional intersection, their intersection is a facet of each. Then the set $\mathcal{F}$ consisting of cones in $\mathcal{M}$ and faces of cones in $\mathcal{M}$ is a complete fan.*

Proposition 9-1.8 and Theorem 9-1.10 imply the following corollary.

Corollary 9-1.11. *The regions of $\mathcal{A}$ are the maximal cones of a complete fan.*

Two regions are *adjacent* if they share a facet in common. The *adjacency graph* of $\mathcal{A}$ is the graph $G(\mathcal{A})$ whose vertices are regions and whose edges are pairs of adjacent regions. The adjacency graph is connected. In fact, much more is true: Taking $\mathcal{A}$ to be essential, $G(\mathcal{A})$ is the graph consisting of the vertices and edges of an n-dimensional zonotope, so Balinski's Theorem says that the graph is n-connected (meaning connected even after removing any $n-1$ vertices). See, for example, [464, Lecture 7.3] for details on zonotopes and [464, Lecture 3.5] for details on Balinski's Theorem. For our purposes, the following lemma, which extends the assertion of connectivity in another direction, is more relevant.

Lemma 9-1.12. *Given regions Q and R of a hyperplane arrangement $\mathcal{A}$, there exists a sequence of regions $Q = R_0, \ldots, R_k = R$ with R_{i-1} adjacent to R_i for $i = 1, \ldots, k$. The sequence can be chosen so as to have an additional property: Moving from Q to R in the sequence, no hyperplane of $\mathcal{A}$ is crossed more than once.*

Proof. Choose a point $\mathbf{x}$ in the interior of Q and a point $\mathbf{y}$ in the interior of R such that the line segment $\overline{\mathbf{xy}}$ does not intersect any $(n-2)$-dimensional subspace of the form $H_1 \cap H_2$ for $H_1, H_2 \in \mathcal{A}$. To see why we can choose

[2]Take $k = n-2$ in [373, Theorem 1.2]. Condition (i) of [373, Theorem 1.2] holds because an open ball in $\mathbb{R}^n$ remains path connected, even after removing a finite number of sets of dimension at most $n-2$.

such $\mathbf{x}$ and $\mathbf{y}$, consider, for any $\mathbf{x}$, the set Y of vectors $\mathbf{y}$ such that the line containing $\mathbf{x}$ and $\mathbf{y}$ intersects some subspace $H_1 \cap H_2$. The set Y is a union of finitely many hyperplanes, and thus cannot cover the interior of R, because the latter is full-dimensional. Following the line segment $\overline{\mathbf{xy}}$, we cross only one hyperplane at a time, and thus visit a sequence of regions $Q = R_0, \ldots, R_k = R$ with R_{i-1} adjacent to R_i for $i = 1, \ldots, k$. The line segment is not contained in any hyperplane of $\mathcal{A}$, so it intersects each hyperplane of $\mathcal{A}$ in at most one point. We have constructed the desired sequence. □

9-1.4 The poset of regions

A hyperplane $H \in \mathcal{A}$ *separates* a region R from the base region B if some line segment (or equivalently, every line segment) from the interior of R to the interior of B intersects H. The *separating set* $S(R)$ of R (with respect to B) is the set of hyperplanes in $\mathcal{A}$ that separate R from B.

Definition 9-1.13. The *poset of regions* $\mathrm{Pos}(\mathcal{A}, B)$ is the set $\mathcal{R}(\mathcal{A})$ of regions, partially ordered with $Q \leq R$ if and only if $S(Q) \subseteq S(R)$. This is a valid partial order: reflexivity and transitivity are immediate, and antisymmetry is an easy exercise (Exercise 9.3). Typically, different choices of B, for the same arrangement $\mathcal{A}$, give non-isomorphic posets $\mathrm{Pos}(\mathcal{A}, B)$. (For example, consider the arrangement of Example 9-1.3. Two different posets of regions for this arrangement are shown in Figures 9-1.3 and 9-3.2. See also Examples 9-1.14 and 9-3.6.)

Example 9-1.14. For the arrangement of Figure 9-1.1, choose B to be the region that projects to the center (labeled in the figure). The other regions are labeled with their separating sets. The hyperplanes are numbered $1, 2, 3, 4$ and separating sets are shown without set braces and commas. The region $-B$, with separating set $\mathcal{A}$, projects to the unbounded area outside of all circles. The hyperplanes themselves are not labeled with their numbers, but the numbering is clear from the separating sets. The resulting poset of regions is shown in Figure 9-1.3.

The poset of regions $\mathrm{Pos}(\mathcal{A}, B)$ is self-dual. The anti-automorphism is $R \mapsto -R$. (See Exercise 9.4.) It is immediate that B is the unique minimal element of $\mathrm{Pos}(\mathcal{A}, B)$, and the antipodal region $-B$ is the unique maximal element. The proof of the following proposition is left as Exercise 9.7.

Proposition 9-1.15. *The cover relations in* $\mathrm{Pos}(\mathcal{A}, B)$ *are* $Q \prec R$ *if and only if* Q *and* R *are adjacent and* $|S(Q)| < |S(R)|$. *In this case,* $S(R) = S(Q) \cup \{H\}$, *where* H *is the hyperplane defining the common facet of* Q *and* R.

Using Proposition 9-1.15, one can show that $\mathrm{Pos}(\mathcal{A}, B)$ is graded, with rank function $|S(R)|$ (Exercise 9.8).

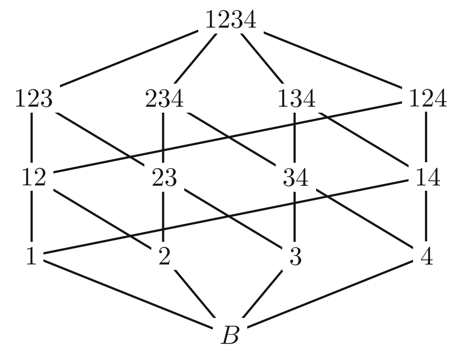

Figure 9-1.3: A poset of regions

Definition 9-1.16. Given a facet F of a region R, Proposition 9-1.8 says that there is a unique region $Q \neq R$ also containing F as a facet. Proposition 9-1.15 asserts that either $Q \prec R$ or $Q \succ R$. Accordingly, we define a *lower facet of* R *with respect to* B to be a facet that R shares with a region $Q \prec R$. Similarly, an *upper facet of* R *with respect to* B is a facet that R shares with a region $Q \succ R$. A *lower hyperplane* (with respect to B) of a region R is the facet-defining hyperplane of a lower facet of R, and an *upper hyperplane* is the facet-defining hyperplane of an upper facet of R. Write $\mathcal{L}(R)$ for the set of lower hyperplanes of R and $\mathcal{U}(R)$ for the set of upper hyperplanes of R.

The following lemmas are convenient restatements of Proposition 9-1.15.

Lemma 9-1.17. *Suppose* R *is a region.*

(i) *The regions covered by* R *in* $\mathrm{Pos}(\mathcal{A}, B)$ *are exactly the regions* Q *such that* $S(Q) = S(R) \setminus \{H\}$, *for some lower hyperplane* H *of* R.

(ii) *The regions covering* R *in* $\mathrm{Pos}(\mathcal{A}, B)$ *are exactly the regions* Q *such that* $S(Q) = S(R) \cup \{H\}$, *for some upper hyperplane* H *of* R.

Lemma 9-1.18. *Let* $\mathbf{b}$ *be a vector in the interior of* B. *Suppose* Q *and* R *are adjacent regions and let* $\mathbf{n}$ *be a normal vector to their shared facet with* $\langle \mathbf{x}, \mathbf{n} \rangle > 0$ *for all* $\mathbf{x}$ *in the interior of* Q. *Then* $Q \prec R$ *if and only if* $\langle \mathbf{b}, \mathbf{n} \rangle > 0$.

The following proposition is proved as Exercise 9.9.

Proposition 9-1.19. *Let* $\mathbf{b}$ *be a vector in the interior of* B, *and for each* $H \in \mathcal{A}$, *let* $\mathbf{n}_H$ *be a nonzero normal vector to* H *such that* $\langle \mathbf{b}, \mathbf{n}_H \rangle > 0$. *Suppose* R *is a region of* $\mathcal{A}$ *and choose* $\mathbf{r}$ *in the interior of* R. *Then* $S(R) = \{H \in \mathcal{A} \mid \langle \mathbf{r}, \mathbf{n}_H \rangle < 0\}$.

The choice of vectors $\mathbf{n}_H$ is unique up to positive scaling of each $\mathbf{n}_H$ and is independent of the choice of $\mathbf{b}$, as long as $\mathbf{b}$ is in the interior of B.

The following simple observation first appeared as part of the proof of [144, Corollary 2.4].

Lemma 9-1.20. *For any region R, the interval $[R, -B]$ in $\mathrm{Pos}(\mathcal{A}, B)$ is isomorphic to the interval $[R, -B]$ in $\mathrm{Pos}(\mathcal{A}, R)$. (The isomorphism is the identity map.)*

Proof. In this proof, we write $S_B(Q)$ for the separating set of a region Q with respect to the base region B. Then for $Q \in [R, -B]_{\mathrm{Pos}(\mathcal{A},B)}$ we have $S_R(Q) = S_B(Q) \setminus S_B(R)$. Therefore, the identity map is an isomorphism from $[R, -B]_{\mathrm{Pos}(\mathcal{A},B)}$ to $[R, -B]_{\mathrm{Pos}(\mathcal{A},R)}$. $\square$

We are interested in the case where $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. The following result of [70] establishes a necessary condition.

$\diamondsuit$ **Theorem 9-1.21.** *If $\mathcal{A}$ is essential and $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice, then B is a simplicial cone.*

The converse to Theorem 9-1.21 holds in a special case, described in Theorem 9-1.22 below, but does not hold in general. The following is [70, Theorem 3.2]. See also [70, Example 3.3].

$\diamondsuit$ **Theorem 9-1.22.** *If $\mathcal{A}$ has rank at most 3 and B is a simplicial cone, then $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice.*

As pointed out in [70, Section 3], there ought to be a necessary and sufficient *local* condition for $\mathrm{Pos}(\mathcal{A}, B)$ to be a lattice. More specifically, the condition should be based on local configurations of hyperplanes/regions, so for example, the property of being simplicial is local. The problem of finding such a local condition is open. (See Problem 9.1 at the end of this chapter.) In Section 9-3, we establish the lattice property for pairs $(\mathcal{A}, B)$ satisfying a certain local condition, more general than simpliciality, that we call tightness. The proof of the lattice property for that class relies on a shortcut that we call the BEZ Lemma, which we present later as Lemma 9-2.2. To use the BEZ lemma, we need to know that joins exist "locally." The next section is devoted to establishing key local properties of the poset of regions, including the needed local result about joins.

9-1.5 Faces, rank-two subarrangements, and intervals

Definition 9-1.23. Let $\mathcal{A}$ be an arrangement and fix a base region B. Let U be an $(n-2)$-dimensional subspace of $\mathbb{R}^n$ and write $\mathcal{A}'$ for $\{H \in \mathcal{A} \mid U \subset H\}$. If $|\mathcal{A}'| \geq 2$, then $\mathcal{A}'$ is called a *rank-two subarrangement* of $\mathcal{A}$. We emphasize that $\mathcal{A}'$ consists of **all** of the hyperplanes in $\mathcal{A}$ containing U. Write B' for the $\mathcal{A}'$-region containing B. The facet-defining hyperplanes of B' are called the *basic hyperplanes* of $\mathcal{A}'$. Given any two distinct hyperplanes H_1 and H_2 of $\mathcal{A}$, there is a unique rank-two subarrangement containing H_1 and H_2, namely the set of all hyperplanes in $\mathcal{A}$ that contain $H_1 \cap H_2$.

The straightforward proof of the following lemma is left as Exercise 9.10.

Lemma 9-1.24. *Let $\mathcal{A}$ be an arrangement, fix a base region B, and let $\mathcal{A}'$ be a rank-two subarrangement of $\mathcal{A}$. The hyperplanes in $\mathcal{A}'$ can be totally ordered as $H_1, \ldots, H_k$ with the following property: For any region R, the set $S(R) \cap \mathcal{A}'$ is either $\{H_1, \ldots, H_i\}$ for some $i = 0, \ldots, k-1$ or $\{H_i, \ldots, H_k\}$ for some $i = 1, \ldots, k$. This total order is unique up to reversing the order. The hyperplanes H_1 and H_k are the basic hyperplanes of $\mathcal{A}'$.*

Define the set of *faces* of $\mathcal{A}$ to be the union of the sets of faces of all the regions of $\mathcal{A}$. Rank-two subarrangements arise naturally when one considers the regions containing a given codimension-2 face of $\mathcal{A}$.

Lemma 9-1.25. *Let $\mathcal{A}$ be an arrangement, fix a base region B, and let F be an $(n-2)$-dimensional face of $\mathcal{A}$.*

(i) *The set $\mathcal{A}' = \{H \in \mathcal{A} \mid F \subset H\}$ is a rank-two subarrangement.*

(ii) *The set of regions containing F is an interval $[Q, R]$ in $\mathrm{Pos}(\mathcal{A}, B)$.*

(iii) *$[Q, R]$ is isomorphic to $\mathrm{Pos}(\mathcal{A}', B')$, where B' is the $\mathcal{A}'$-region containing B.*

(iv) *$[Q, R]$ is the union of two chains, disjoint except at Q and R, each having $|\mathcal{A}'| + 1$ elements.*

(v) *$S(Q) \cap \mathcal{A}' = \varnothing$ and $S(R) = S(Q) \cup \mathcal{A}'$.*

Proof. Let F be an $(n-2)$-dimensional face of $\mathcal{A}$, and specifically, let R be some region having F as a face. Write U for the linear span of F, which is a linear subspace of dimension $n-2$ because F is a cone of dimension $n-2$. A hyperplane $H \in \mathcal{A}$ contains F if and only if it contains U. There are at least two hyperplanes in $\mathcal{A}'$, namely the facet-defining hyperplanes for the two facets of R whose intersection is F. This proves (i).

Let $\mathbf{x}$ be a point in the relative interior of F. We claim that $\mathcal{A}'$ is the set of hyperplanes containing $\mathbf{x}$. Since $\mathbf{x} \in F$, each $H \in \mathcal{A}'$ contains $\mathbf{x}$. On the other hand, suppose some hyperplane $H \in \mathcal{A} \setminus \mathcal{A}'$ contains $\mathbf{x}$. Since H is not in $\mathcal{A}'$, it does not contain U. Thus there is a 2-dimensional plane $P \subseteq H$ such that $\mathbb{R}^n$ is the direct sum of P with U.

Given a point $\mathbf{y}$ in the interior of R that is very close to $\mathbf{x}$, we can subtract a vector in U to obtain a point $\mathbf{y}' = \mathbf{x} + \mathbf{p}$ for $\mathbf{p} \in P$. Since $\mathbf{x}$ is in the relative interior of F and since U is the intersection of the facet-defining hyperplanes for the two facets of R whose intersection is F, if the initial $\mathbf{y}$ is close enough to $\mathbf{x}$, the point $\mathbf{y}'$ is in the interior of R. But $\mathbf{y}'$ is in H, so H intersects the interior of R. This contradicts the fact that R is a region, thus proving the claim.

We also claim that the set of regions containing F equals the set of regions containing $\mathbf{x}$. Every region containing F contains $\mathbf{x}$. If some region R' contains $\mathbf{x}$ but not F, then there is some other point $\mathbf{x}'$ in the relative interior of F such that $\mathbf{x}$ and $\mathbf{x}'$ are on the opposite sides of some facet-defining hyperplane H of R'. The line segment $\overline{\mathbf{x}\mathbf{x}'}$ intersects H in a point $\mathbf{x}''$ also in the relative interior of F. But then H contains $\mathbf{x}''$ but not F, contradicting the previous claim (with $\mathbf{x}''$ replacing $\mathbf{x}$ in the claim). This proves the second claim.

Now consider a ball of radius ε about $\mathbf{x}$. The second claim implies that for small enough ε, the regions intersecting the ball are exactly the regions containing F. The first claim implies that the separating sets of these regions differ only on the set $\mathcal{A}'$. Furthermore, these regions are in bijection with the regions of $\mathcal{A}'$. The bijection takes each region Q of $\mathcal{A}$ intersecting the ball to the unique region of $\mathcal{A}'$ containing Q. The induced subposet of $\mathrm{Pos}(\mathcal{A}, B)$ consisting of regions intersecting the ball is thus isomorphic to $\mathrm{Pos}(\mathcal{A}', B')$. Assertions (ii), (iii), and (v) follow, and then (iv) follows by Lemma 9-1.24. $\square$

Recall that Proposition 9-1.15 says that two regions that form a covering pair in $\mathrm{Pos}(\mathcal{A}, B)$ share a common facet.

Lemma 9-1.26. *Suppose R_1 and R_2 are distinct regions of $\mathcal{A}$, both covering a region Q in* $\mathrm{Pos}(\mathcal{A}, B)$*. Let F_1 be the facet shared by Q and R_1, let F_2 be the facet shared by Q and R_2, and suppose $F_1 \cap F_2$ is $(n-2)$-dimensional. Let H_1 be the hyperplane containing F_1, let H_2 be the hyperplane containing F_2, and let $\mathcal{A}'$ be the rank-two subarrangement containing H_1 and H_2.*

(i) *The basic hyperplanes of $\mathcal{A}'$ are H_1 and H_2.*

(ii) *$R_1 \vee R_2$ exists and has separating set $S(Q) \cup \mathcal{A}'$. This is a disjoint union.*

(iii) *The interval $[Q, R_1 \vee R_2]$ is the set of regions containing $F_1 \cap F_2$.*

(iv) *There exists a region R with $S(R) = S(R_1 \vee R_2) \setminus \{H_1\}$ (and thus $R \prec R_1 \vee R_2$).*

(v) *The interval $[Q, R_1 \vee R_2]$ is the union of two chains, disjoint except at Q and $R_1 \vee R_2$, each having $|\mathcal{A}'| + 1$ elements.*

Proof. The separating sets of R_1 and R_2 are $S(Q) \cup \{H_1\}$ and $S(Q) \cup \{H_2\}$ respectively, so Lemma 9-1.24 says that H_1 and H_2 are the basic hyperplanes of $\mathcal{A}'$. This is (i).

Lemma 9-1.24 also implies that any region with H_1 and H_2 in its separating set has $\mathcal{A}'$ contained in its separating set, so any upper bound for R_1 and R_2 has separating set containing $S(Q) \cup \mathcal{A}'$. Lemma 9-1.25 says that there exists an element with separating set exactly $S(Q) \cup \mathcal{A}'$. Therefore, this element is $R_1 \vee R_2$, and we have established (ii).

Lemma 9-1.25 also says that the set of regions containing $F_1 \cap F_2$ is an interval composed of two chains, disjoint except at the top and bottom of the interval. Since Q is covered by two distinct elements, it must be the bottom element of the interval. Also, $R_1 \vee R_2$ covers two distinct elements in the interval. (Otherwise, the unique element it covers is also an upper bound for

R_1 and R_2.) Thus $R_1 \vee R_2$ is the top element of the interval, and we have established (iii) and (v).

Lemma 9-1.25 says furthermore that $S(R_1 \vee R_2) = S(Q) \cup \mathcal{A}'$ and that the separating sets of regions in $[Q, R_1 \vee R_2]$ differ only by hyperplanes in $\mathcal{A}'$. Therefore, Lemma 9-1.24 allows only two possibilities for an element covered by $R_1 \vee R_2$ in the interval $[Q, R_1 \vee R_2]$. Such an element has separating set either $S(R) = S(Q) \cup (\mathcal{A}' \setminus \{H_1\})$ or $S(R) = S(Q) \cup (\mathcal{A}' \setminus \{H_2\})$. Since the region $R_1 \vee R_2$ covers two elements of $[Q, R_1 \vee R_2]$, both possible separating sets occur, and in particular, we have established (iv). □

Exercise 9.11 describes how much of Lemma 9-1.26 holds without the assumption that $F_1 \cap F_2$ is $(n-2)$-dimensional.

The dual statement to Lemma 9-1.26 holds by the dual argument.

Lemma 9-1.27. *Suppose Q_1 and Q_2 are distinct regions of $\mathcal{A}$, both covered by a region R in* $\mathrm{Pos}(\mathcal{A}, B)$. *Let F_1 be the facet shared by Q_1 and R, let F_2 be the facet shared by Q_2 and R, and suppose $F_1 \cap F_2$ is $(n-2)$-dimensional. Let H_1 be the hyperplane containing F_1, let H_2 be the hyperplane containing F_2, and let $\mathcal{A}'$ be the rank-two subarrangement containing H_1 and H_2.*

(i) *The basic hyperplanes of $\mathcal{A}'$ are H_1 and H_2.*

(ii) *$Q_1 \wedge Q_2$ exists and has separating set $S(R) \setminus \mathcal{A}'$. Also, $\mathcal{A}' \subseteq S(R)$.*

(iii) *The interval $[Q_1 \wedge Q_2, R]$ is the set of regions containing $F_1 \cap F_2$.*

(iv) *There exists a region Q with $S(Q) = S(Q_1 \wedge Q_2) \cup \{H_1\}$ (and thus $Q_1 \wedge Q_2 \prec Q$).*

(v) *The interval $[Q_1 \wedge Q_2, R]$ is the union of two chains, disjoint except at $Q_1 \wedge Q_2$ and R, each having $|\mathcal{A}'| + 1$ elements.*

9-2. Lattice-theoretic shortcuts

Combinatorialists often encounter lattices "in nature" as partial orders and must prove that the partial orders are indeed lattices. Here we discuss some shortcuts to proving the lattice property. We then broaden the discussion to consider various shortcuts in a similar spirit, including shortcuts to establish semidistributivity or to detect homomorphisms.

To begin, we give the simplest and best-known shortcut for proving the lattice property. Recall that 0 denotes the unique minimal element of a poset, if such exists, and that $\downarrow x$ denotes the set of elements weakly below x.

Lemma 9-2.1. *Suppose P is a finite join-semilattice with* 0. *Then P is a lattice.*

Proof. We verify that meets exist. Let x and y be elements of P. The set $U = (\downarrow x) \cap (\downarrow y)$ of lower bounds for $\{x, y\}$ is nonempty because it contains 0. Because P is finite, we can form[3] the join $\bigvee U$. Since x and y are both upper bounds for U, we have $\bigvee U \leq x$ and $\bigvee U \leq y$, so $(\bigvee U) \in U$. Thus $\bigvee U$ is the unique maximal lower bound for $\{x, y\}$, or in other words, it is $x \wedge y$. □

9-2.1 The BEZ Lemma and some extensions

There is a much more powerful shortcut that we call the BEZ Lemma after the authors (Björner, Edelman, and Ziegler) of the paper where it first appeared. We have weakened the hypotheses slightly by not requiring *a priori* that P has a 1.

Lemma 9-2.2 (BEZ Lemma)**.** *Suppose P is a finite poset with* 0*. Suppose also that, for all x and y in P such that x and y cover a common element z, the join $x \vee y$ exists. Then P is a lattice.*

Proof. We prove that P is a join-semilattice and apply Lemma 9-2.1. We argue by induction on the number of elements in P, with the base case $|P| = 1$ being trivial. Let x and y be elements of P. If x and y are comparable, then $x \vee y$ exists, so we assume that x and y are incomparable. In particular, they are both strictly above 0. Let a_x and a_y be elements of P such that $0 \prec a_x \leq x$ and $0 \prec a_y \leq y$. If $a_x = a_y$ then both x and y lie in the induced subposet $\uparrow a_x$ (the set of elements weakly above a_x), which is strictly smaller than P. Thus by induction, x and y have a join in $\uparrow a_x$. But any upper bound of x and y in P is in $\uparrow a_x$, so x and y have a unique minimal upper bound in P.

If $a_x \neq a_y$ then $a_x \vee a_y$ exists, because a_x and a_y both cover 0. Both x and $a_x \vee a_y$ are in $\uparrow a_x$, which is strictly smaller than P, so as before we conclude by induction that $x \vee (a_x \vee a_y)$ exists. Both y and $x \vee (a_x \vee a_y)$ are in $\uparrow a_y$, so we similarly conclude that $x \vee (a_x \vee a_y) \vee y$ exists. But since $a_x \leq x$ and $a_y \leq y$, the element $x \vee (a_x \vee a_y) \vee y$ is the desired element $x \vee y$. □

We next discuss extensions of Lemmas 9-2.1 and 9-2.2 beyond finite posets. We extend Lemmas 9-2.1 and 9-2.2 to establish the (meet-semi)lattice property for infinite posets that are "finite under going down." A poset P is *lower finite* if the downset $\downarrow x$ of every element $x \in P$ is finite. In some references, including [421, Section 3.4] and [383, Section 2], lower finite posets are referred to as *finitary* posets. A poset is *well-founded* if every nonempty subset U of P has at least one minimal element (an element x of U such that there exists no $y \in U$ with $y < x$). For example, every lower finite poset is well-founded. Exercise 9.13 shows that a poset is well-founded if and only if it satisfies the *Descending Chain Condition*, meaning that there is no infinite sequence $a_1 > a_2 > \cdots$ of elements of P.

[3] A standard argument (see for example LTF Lemma 2) shows that if the join exists for all pairs of elements of P, then every finite subset of P has a join. One shows that the join operation is associative, and that $x_1 \vee \cdots \vee x_k$ is the least upper bound for $\{x_1, \ldots, x_k\}$.

Lemma 9-2.3. *If P is a well-founded poset, then conditions* (i) *and* (ii) *below are equivalent. If P is a lower finite poset, then conditions* (i)–(iii) *below are equivalent.*

(i) *P is a meet-semilattice.*

(ii) *P has a unique minimal element 0 and every nonempty subset of P either has no upper bound or has a join.*

(iii) *P has a unique minimal element 0 and every pair $x, y \in P$ either has no upper bound or has a join.*

Proof. Suppose (i) holds. Since P is well-founded, it has at least one minimal element. The meet-semilattice property then implies that there is exactly one minimal element. Suppose that some nonempty subset $A \subseteq P$ has an upper bound b. Then since P is well-founded, there is an element $\leq b$ which is minimal among upper bounds for A. Suppose b_1 and b_2 are both minimal upper bounds for A. Then any $a \in A$ has $a \leq b_1$ and $a \leq b_2$ so that $a \leq b_1 \wedge b_2$. Thus $b_1 \wedge b_2$ is an upper bound for A. We must have $b_1 = b_2$, since b_1 and b_2 were both assumed to be minimal upper bounds for A. We have shown that A has a unique minimal upper bound, and this is $\bigvee A$. Thus (ii) holds.

Suppose (ii) holds. For any pair $x, y \in P$, let $\mathcal{B}$ be the set $(\downarrow x) \cap (\downarrow y)$. Since P has a unique minimal element, $\mathcal{B}$ is nonempty. Since x is an upper bound for $\mathcal{B}$, condition (ii) implies that $\bigvee \mathcal{B}$ exists. Since x is an upper bound for $\mathcal{B}$, we have $\bigvee \mathcal{B} \leq x$. Similarly, $\bigvee \mathcal{B} \leq y$. We conclude that $x \wedge y$ exists and equals $\bigvee \mathcal{B}$. This establishes (i).

With no additional hypotheses on P, (ii) implies (iii). Assuming lower finiteness, we show that (iii) implies (i). Suppose (iii) holds. Taking $x, y \in P$ and defining $\mathcal{B}$ as in the previous paragraph, we see that $\mathcal{B}$ is finite by the hypothesis of lower finiteness. Since $\mathcal{B}$ has an upper bound, every pair of elements in $\mathcal{B}$ has an upper bound, and thus has a join in P by (iii). Therefore $\mathcal{B}$ has a join[4] in P and $\bigvee \mathcal{B}$ equals $x \wedge y$ as in the previous paragraph. □

Exercise 9.14 asks for an example of a well-founded partial order satisfying condition (iii) of Lemma 9-2.3 but not conditions (i) and (ii).

As an immediate consequence of the implication (iii) $\Longrightarrow$ (i) in Lemma 9-2.3, we obtain an extension of Lemma 9-2.1 to lower finite posets:

Lemma 9-2.4. *Suppose P is a lower finite join-semilattice with 0. Then P is a lattice.*

We have extended Lemma 9-2.1 to lower finite posets, but unfortunately the BEZ Lemma becomes false when we replace "finite" with "lower finite." As an example, consider the lower finite poset of Figure 9-2.1, in which, for

[4]See the footnote on page 411.

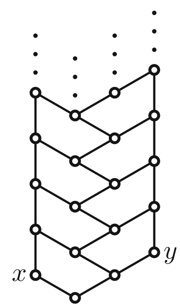

Figure 9-2.1: A counterexample to the BEZ Lemma for lower finite posets

example, the elements marked x and y do not have a join. If we apply the strategy of the proof of Lemma 9-2.2, the induction never terminates. However, we prove the following version of the BEZ Lemma.

Lemma 9-2.5 (BEZ Lemma for lower finite meet-semilattices)**.** *Suppose P is a lower finite poset with 0. Suppose also that, for any x and y in P such that x and y cover a common element z, either $\{x, y\}$ has no upper bound or the join $x \vee y$ exists. Then P is a meet-semilattice.*

Proof. We verify condition (iii) of Lemma 9-2.3. Suppose $x, y \in P$ have an upper bound. The proof that $x \vee y$ exists is identical to the proof of Lemma 9-2.2, except that we argue by induction on the minimum size of the set $\downarrow_P m$, where m ranges over all upper bounds of x and y. When $a_x = a_y$, then we pass to the induced subposet $P' = \uparrow_P a_x$ and note that a smallest $\downarrow_{P'} m$ is smaller than a smallest $\downarrow_P m$. The induction in the other cases works similarly. $\square$

A meet-semilattice L is *complete* if every subset of L (not just every finite subset of L) has a greatest lower bound. Completeness of join-semilattices is defined similarly, and a complete lattice is a lattice which is both a complete meet-semilattice and a complete join-semilattice. (See LTF Section I.3.14). The results of this section can be extended to assert completeness, as explored in Exercises 9.15, 9.16, and 9.17. One must be careful, however, because it is possible, for example, for a lattice to be a complete meet-semilattice but not a complete join-semilattice (Exercise 9.16).

9-2.2 More BEZ-type lemmas

The argument for Lemma 9-2.2 is very versatile, and we spend the rest of this section discussing other "BEZ-type" lemmas. Our first BEZ-type lemma simplifies the process of checking whether a lattice is semidistributive. We recall Definition 3-1.1: A lattice L is *meet-semidistributive* if any elements $x, y, z \in L$ with $x \wedge y = x \wedge z$ also satisfy $x \wedge (y \vee z) = x \wedge y$. The lattice is *join-semidistributive* if the dual condition holds: If $x \vee y = x \vee z$, then $x \vee (y \wedge z) = x \vee y$. If both conditions hold, then the lattice is called *semidistributive*.

Lemma 9-2.6 (BEZ Lemma for meet-semidistributivity)**.** *Suppose L is a finite lattice with the following property: If x, y, and z are elements of L with $x \wedge y = x \wedge z$ and if y and z cover a common element, then $x \wedge (y \vee z) = x \wedge y$. Then L is meet-semidistributive.*

Proof. Let x, y, and z be elements of L with $x \wedge y = x \wedge z$. We now show that $x \wedge (y \vee z) = x \wedge y$. We argue by induction on the size of $\uparrow (y \wedge z)$. If y and z are comparable, then the assertion is trivial, so assume not. Let a_y and a_z be elements covering $y \wedge z$ with $a_y \leq y$ and $a_z \leq z$. In particular, $a_y \neq a_z$, because if $a_y = a_z$, then we reach the contradiction that $a_y = y \wedge z$.

Since $a_y \leq y$, we have $x \wedge a_y \leq x \wedge y$. But $x \wedge y$ is a lower bound for y, and since $x \wedge y = x \wedge z$, a lower bound for z as well. Thus $x \wedge y \leq y \wedge z < a_y$, so $x \wedge y \leq x \wedge a_y$. We have shown that $x \wedge a_y = x \wedge y$. The same argument shows that $x \wedge a_z = x \wedge z$, so $x \wedge a_y = x \wedge a_z$. Since a_y and a_z both cover $y \wedge z$, by hypothesis, we have $x \wedge (a_y \vee a_z) = x \wedge a_y$. Thus $x \wedge (a_y \vee a_z) = x \wedge y$. Since y and $a_y \vee a_z$ are both above a_y, which is strictly greater than $y \wedge z$, by induction we have $x \wedge (y \vee a_y \vee a_z) = x \wedge y$. But this also equals $x \wedge z$. Now $(y \vee a_y \vee a_z)$ and z are both above a_z, which is strictly greater than $y \wedge z$. Again by induction, we have $x \wedge (y \vee a_y \vee a_z \vee z) = x \wedge z$. We rewrite this as $x \wedge (y \vee z) = x \wedge z$. □

The dual proof establishes a BEZ Lemma for join-semidistributivity. A similar argument proves a criterion for distributivity. We leave the proof of the following lemma to Exercise 9.18. One should be careful in dualizing Lemma 9-2.7, as illustrated in Exercise 9.19.

Lemma 9-2.7 (BEZ Lemma for distributivity)**.** *Suppose L is a finite lattice such that the distributive law $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ holds whenever y and z cover a common element. Then L is distributive.*

The now-familiar argument also establishes a criterion for detecting homomorphisms.

Lemma 9-2.8 (BEZ Lemma for join-homomorphisms)**.** *Suppose L is a lower finite lattice and suppose L' is a join-semilattice. Suppose $\eta \colon L \to L'$ is an order-preserving map such that $\eta(x \vee y) = \eta(x) \vee \eta(y)$ holds whenever there exists $z \in L$ with $z \prec x$ and $z \prec y$. Then η is a join-homomorphism.*

The hypothesis that η is order-preserving is also local: It is equivalent to requiring that $\eta(x) \leq \eta(y)$ whenever $x \prec y$. (See Exercise 9.21.)

Proof. We argue by induction on the size of $\downarrow (x \vee y)$ that $\eta(x \vee y) = \eta(x) \vee \eta(y)$. The set $\downarrow (x \vee y)$ is finite because L is lower finite. The base case, where $\downarrow (x \vee y)$ has one element, is trivial. If x and y are comparable, then we are done by the hypothesis that η is order-preserving. Assume x and y are incomparable, so that both are strictly above 0. Let a_x and a_y have $0 \prec a_x \leq x$ and $0 \prec a_y \leq y$.

If $a_x = a_y$ then x and y are in $\uparrow a_x$, and the number of elements of $\uparrow a_x$ below $x \vee y$ is strictly less than the number of elements of L below $x \vee y$. Thus by induction the restriction of η to $\uparrow a_x$ has $\eta(x \vee y) = \eta(x) \vee \eta(y)$, but then this equality holds for the unrestricted map η.

If $a_x \neq a_y$ then by hypothesis $\eta(a_x \vee a_y) = \eta(a_x) \vee \eta(a_y)$. Since $a_x \vee a_y$ and x are both in $\uparrow a_x$, and since $x \vee (a_x \vee a_y) \leq x \vee y$, we apply induction to conclude that $\eta(x \vee a_x \vee a_y) = \eta(x) \vee \eta(a_x) \vee \eta(a_y)$. Since y and $x \vee a_x \vee a_y$ are both in $\uparrow a_y$, we similarly conclude that $\eta(x \vee a_x \vee a_y \vee y) = \eta(x) \vee \eta(a_x) \vee \eta(a_y) \vee \eta(y)$. But $x \vee a_x \vee a_y \vee y = x \vee y$, and since η is order-preserving, we also have $\eta(a_x) \leq \eta(x)$ and $\eta(a_y) \leq \eta(y)$, so that $\eta(x) \vee \eta(a_x) \vee \eta(a_y) \vee \eta(y) = \eta(x) \vee \eta(y)$. □

Lemma 9-2.8 has other BEZ-type lemmas as corollaries. The first is immediate.

Lemma 9-2.9 (BEZ Lemma for homomorphisms)**.** *Suppose L is a finite lattice and suppose L' is a lattice. Suppose $\eta\colon L \to L'$ is an order-preserving map such that $\eta(x \vee y) = \eta(x) \vee \eta(y)$ whenever x and y cover a common element and such that $\eta(x \wedge y) = \eta(x) \wedge \eta(y)$ whenever x and y are covered by a common element. Then η is a lattice homomorphism.*

In the following lemmas, we use subscripts S and L to denote joins, meets, and cover relations in a lattice L and in an induced subposet S of L.

Lemma 9-2.10 (BEZ lemma for join-subsemilattices)**.** *Suppose S is a finite induced subposet of a join-semilattice L and suppose S has a unique minimal element. Suppose also that, whenever $x, y, z \in S$ have $z \prec_S x$ and $z \prec_S y$, the join $x \vee_L y$ is in S. Then S is a lattice and is a join-subsemilattice (but not necessarily a sublattice) of L.*

Proof. Lemma 9-2.2 implies that S is a lattice. Let $\eta : S \to L$ be the inclusion of S as a subset of L. By hypothesis, whenever $x, y, z \in S$ have $z \prec_S x$ and $z \prec_S y$, the join $x \vee_L y$ is in S, so $x \vee_L y = x \vee_S y$. In other words, $\eta(x \vee y) = \eta(x) \vee \eta(y)$. Lemma 9-2.8 implies that $\eta(x \vee y) = \eta(x) \vee \eta(y)$ for any $x, y \in S$. In other words, the join operation in S agrees with the join operation in L. □

In any partially ordered set, the transitive closure of the comparability relation is an equivalence relation. The partially ordered set is *connected* if this equivalence relation has only one equivalence class.

Lemma 9-2.11 (BEZ lemma for sublattices)**.** *Suppose S is a nonempty connected finite induced subposet of a lattice L. Suppose that, whenever $x, y, z \in S$ have $z \prec_S x$ and $z \prec_S y$, the join $x \vee_L y$ is in S. Suppose also that, whenever $x, y, z \in S$ have $x \prec_S z$ and $y \prec_S z$, the meet $x \wedge_L y$ is in S. Then S is a sublattice of L.*

Proof. We first show that S has a unique minimal element. Since S is finite and nonempty, it has a minimal element. Suppose it has two distinct minimal elements m_1 and m_2. Since S is connected, there exist elements $m_1 = x_0, \ldots, x_k = m_2$ of S such that x_{i-1} and x_i are comparable for each $i = 1, \ldots, k$. Suppose there is some i with $1 \leq i < k$ such that $x_{i-1} \leq x_i \geq x_{i+1}$. Then $S' = S \cap {\downarrow} x_i$ and $L' = L \cap {\downarrow} x_i$ satisfy the hypotheses to the dual of Lemma 9-2.10. The dual of the lemma implies that the meet of x_{i-1} and x_{i+1} is in S' and therefore is in S. In the sequence $x_0, \ldots, x_k$, we replace x_i by $x_{i-1} \wedge x_{i+1}$. Continuing in this manner, we eventually have a sequence $m_1 = x_0, \ldots, x_k = m_2$ of elements of S and some j such that $x_0 \geq x_1 \geq \cdots \geq x_j \leq \cdots \leq x_{k-1} \leq x_k$. This contradicts the supposition that m_1 and m_2 are distinct minimal elements of S, thus proving that S has a unique minimal element.

Now Lemma 9-2.10 implies that S is a join-subsemilattice of L and is a lattice. In particular S has a unique maximal element. Thus the dual to Lemma 9-2.10 implies that S is also a meet-subsemilattice of L. □

We conclude this section on lattice-theoretic shortcuts by mentioning a criterion [463, Criterion 2] that is similar in spirit to the BEZ-lemma but establishes the lattice property without requiring the existence of any meets or joins.

♢ **Lemma 9-2.12.** *Suppose P is a poset having 0 and 1 and having a finite upper bound on the length of chains. Then P is a lattice if and only if the following condition holds: Whenever x_1 and x_2 cover a common element, y_1 and y_2 are covered by a common element, and $x_i \leq y_j$ for all $i, j \in \{1, 2\}$, there exists an element z with $x_i \leq z \leq y_j$ for all $i, j \in \{1, 2\}$.*

9-3. Tight posets of regions

The BEZ Lemma was formulated in [70] for the purpose of proving that the poset of regions of a simplicial arrangement (see Definition 9-1.6) is a lattice. Here, we prove a generalization of that result based on the proof given in [70]. The generalization replaces simplicial arrangements with tight arrangements, as defined below. As an indication that the generalization is natural, we prove that tightness characterizes posets of regions that are semidistributive lattices.

Let $\mathcal{A}$ be a hyperplane arrangement and choose a base region B. Recall from Definition 9-1.16 the notion of lower facets and upper facets.

Definition 9-3.1. A region is *tight with respect to B* if every pair of its upper facets intersects in a face of dimension $n-2$ and every pair of its lower facets intersects in a face of dimension $n-2$. An arrangement is *tight with respect to B* if all of its regions are tight with respect to B. The phrase "with respect to B" is essential here, as we see, for example, in Figure 9-3.5: The arrangement shown there is tight with respect to the region labeled B but not tight with respect to the region labeled 1. Since the antipodal map is an anti-automorphism of $\mathrm{Pos}(\mathcal{A}, B)$, to check tightness, it is enough to check either all pairs of lower facets of all regions or all pairs of upper facets of all regions. For convenience, we sometimes say that $(\mathcal{A}, B)$ is tight or that $\mathrm{Pos}(\mathcal{A}, B)$ is tight, to mean that $\mathcal{A}$ is tight with respect to B.

Theorem 9-3.2. *If $\mathcal{A}$ is tight with respect to B, then $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice.*

Proof. We verify the hypotheses of the BEZ Lemma (9-2.2). Suppose R_1 and R_2 cover a common region Q in $\mathrm{Pos}(\mathcal{A}, B)$. Let F_1 be the facet shared by R_1 and Q and let F_2 be the facet shared by R_2 and Q. Since $\mathcal{A}$ is tight with respect to B, the face $F_1 \cap F_2$ is $(n-2)$-dimensional. Lemma 9-1.26 says that $R_1 \vee R_2$ exists, so Lemma 9-2.2 implies that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. □

The motivating example for Definition 9-3.1 is the example of a simplicial arrangement. It is well known (and easily verified in Exercise 9.23) that for every simplicial region R, every pair of facets of R intersects in a face of dimension $n-2$. Thus we have the following proposition.

Proposition 9-3.3. *A simplicial arrangement is tight with respect to any choice of base region B.*

In light of Proposition 9-3.3, Theorem 9-3.2 has the following corollary.

Corollary 9-3.4. *If $\mathcal{A}$ is simplicial and B is any region of $\mathcal{A}$, then $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice.*

Example 9-3.5. Figure 9-3.1 shows a hyperplane arrangement that is not simplicial but is tight with respect to the region marked B.

Example 9-3.6. The poset of regions of Examples 9-1.3 and 9-1.14 (Figures 9-1.1 and 9-1.3) is not a lattice. If we take R_1 and R_2 to be the regions with separating sets 1 and 3 and try to push through the argument in the proof of Theorem 9-3.2, we see how tightness is crucial.

Example 9-3.7. On the other hand, the lattice property may hold even in the non-tight case. Figure 9-3.2 shows the same hyperplane arrangement as Figure 9-1.1, but with a different choice of base region. The arrangement is not tight with respect to this choice of base region. For example, the region whose separating set is 1 is not tight with respect to B. However, this poset of regions is still a lattice, as one can check directly. As an alternative to checking directly, one can appeal to Theorem 9-1.22.

9-3.1 Tightness and semidistributivity

Given that there are posets of regions $\mathrm{Pos}(\mathcal{A}, B)$ that are lattices, even though $\mathcal{A}$ is not tight with respect to B, one might consider the tightness hypothesis

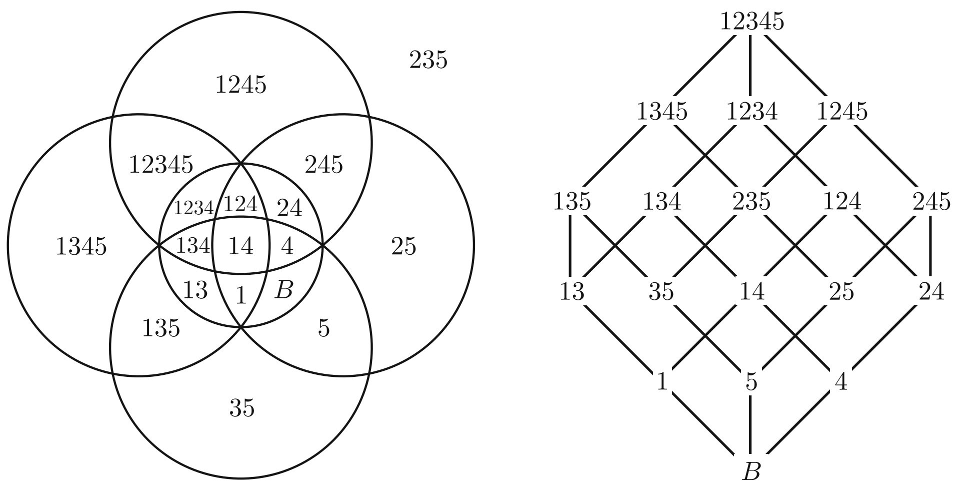

Figure 9-3.1: A hyperplane arrangement $\mathcal{A}$, tight with respect to B, and its lattice of regions

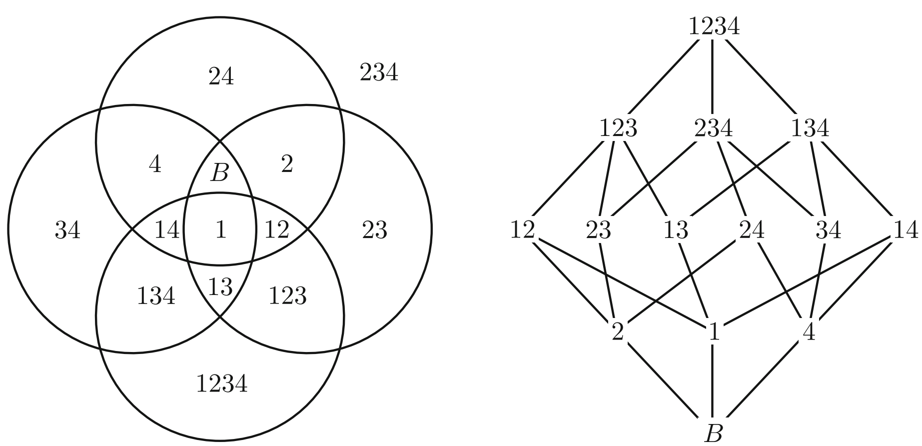

Figure 9-3.2: A hyperplane arrangement, not tight with respect to B, and its lattice of regions

in Theorem 9-3.2 to constitute an artificial generalization of the more natural-seeming simplicial hypothesis of Corollary 9-3.4. It turns out, however, that the tightness condition is lattice-theoretically quite natural: It characterizes semidistributivity.

Theorem 9-3.8. *The poset of regions* $\mathrm{Pos}(\mathcal{A}, B)$ *is a semidistributive lattice if and only if* $\mathcal{A}$ *is tight with respect to* B*.*

By Proposition 9-3.3, Theorem 9-3.8 has the following corollary.

Corollary 9-3.9. *If* $\mathcal{A}$ *is simplicial and* B *is any region of* $\mathcal{A}$*, then* $\mathrm{Pos}(\mathcal{A}, B)$ *is a semidistributive lattice.*

If $\mathcal{A}$ has the property that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice for every choice of B, then Theorem 9-1.21 implies that $\mathcal{A}$ is simplicial. Thus we have the following additional corollary to Theorem 9-3.8.

Corollary 9-3.10. *For any hyperplane arrangement* $\mathcal{A}$*, the following are equivalent:*

(i) $\mathrm{Pos}(\mathcal{A}, B)$ *is a lattice for every region* B *of* $\mathcal{A}$*.*

(ii) $\mathrm{Pos}(\mathcal{A}, B)$ *is a semidistributive lattice for every region* B *of* $\mathcal{A}$*.*

(iii) $\mathcal{A}$ *is simplicial.*

Using Theorem 9-3.8, we will also prove the following property of tight arrangements. Recall from Definition 9-1.16 the definitions of lower hyperplanes and upper hyperplanes of a region and the sets $\mathcal{L}(R)$ and $\mathcal{U}(R)$.

Proposition 9-3.11. *Let* $\mathbf{b}$ *be a vector in the interior of* B*, and for each* $H \in \mathcal{A}$*, let* $\mathbf{n}_H$ *be a nonzero normal vector to* H *such that* $\langle \mathbf{b}, \mathbf{n}_H \rangle > 0$*. If* $(\mathcal{A}, B)$ *is tight, then for each region* R *of* $\mathcal{A}$*, the vectors* $\{\mathbf{n}_H \mid H \in \mathcal{L}(R)\}$ *are linearly independent and the vectors* $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$ *are linearly independent.*

Remark 9-3.12. Proposition 9-3.11 does not give an alternate characterization of the property of tightness, because there exist non-tight pairs $(\mathcal{A}, B)$ satisfying the conclusions of Proposition 9-3.11. One such pair is found in Example 9-3.7.

We now prepare to prove Theorem 9-3.8 by proving two lemmas.

Lemma 9-3.13. *If a region* R *has exactly two lower facets with respect to* B*, then the intersection of the two facets is* $(n-2)$*-dimensional.*

Proof. Let F and G be the two lower facets, let $\mathbf{x}$ be a point in the relative interior of F, let $\mathbf{y}$ be a point in the relative interior of G, and let $\mathbf{b}$ be a point in the interior of B. Lemma 9-1.18 implies that, for small enough ε, the points $\mathbf{x}' = \mathbf{x} - \varepsilon\mathbf{b}$ and $\mathbf{y}' = \mathbf{y} - \varepsilon\mathbf{b}$ are in the interior of R. The convexity of R implies that the line segment connecting $\mathbf{x}'$ and $\mathbf{y}'$ is contained in the interior of R.

Now, for any point $\mathbf{z}$ in the interior of R, define $\mathbf{p}(\mathbf{z})$ to be the point forming the intersection of the boundary of R with the set $\{\mathbf{z} + \lambda\mathbf{b} \mid \lambda \geq 0\}$. Lemma 9-1.18 implies that the point $\mathbf{p}(\mathbf{z})$ is contained in one or more lower facets of R but is not contained in any upper facets of R. We have $\mathbf{p}(\mathbf{x}') = \mathbf{x}$ and $\mathbf{p}(\mathbf{y}') = \mathbf{y}$.

Let H_F and H_G be the hyperplanes containing F and G respectively. Since $\mathbf{x}$ is in the relative interior of F, it is in H_F but not in H_G. Similarly, $\mathbf{y}$ is in H_G but not in H_F. Since each $\mathbf{p}(\mathbf{z})$ is in a lower facet of R and since H_F and H_G are the facet-defining hyperplanes of the only lower facets of R, there is a point $\mathbf{z}'$ on the line segment from $\mathbf{x}'$ to $\mathbf{y}'$ such that $\mathbf{p}(\mathbf{z}')$ is in $H_F \cap H_G$. Since $\mathbf{z}'$ is in the interior of R, there is an open ball about $\mathbf{z}'$ contained in R. The map $\mathbf{p}$ takes this ball to a relatively open neighborhood U of $\mathbf{p}(\mathbf{z}')$ in the boundary of R. But U is also contained in the union of the lower facets of R. Thus, the intersection of U with $H_F \cap H_G$ is an open neighborhood of $\mathbf{p}(\mathbf{z}')$ in $H_F \cap H_G \cap R$. In particular, the set $F \cap G = H_F \cap H_G \cap R$ is $(n-2)$-dimensional. $\square$

Lemma 9-3.14. *Let C be a collection of regions of $\mathcal{A}$, suppose $\bigvee C$ exists and let R be an upper bound for C in $\mathrm{Pos}(\mathcal{A}, B)$. Then R is a minimal upper bound for R if and only if, for every lower hyperplane H of R, there exists $Q \in C$ such that $H \in S(Q)$.*

Proof. Suppose there exists $H \in \mathcal{L}(R)$ such that $H \notin \bigcup_{Q \in C} S(Q)$. By Lemma 9-1.17, there is a region whose separating set is $S(R) \setminus \{H\}$, and this region is an upper bound for C. Therefore R is not a minimal upper bound.

Conversely, suppose that for every $H \in \mathcal{L}(R)$, there exists $Q \in C$ such that $H \in S(Q)$. Then Lemma 9-1.17 implies that no region covered by R is an upper bound for C. We conclude that R is a minimal upper bound. $\square$

Proof of Theorem 9-3.8. Suppose $\mathcal{A}$ is tight with respect to B. We know by Theorem 9-3.2 that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. We verify the hypotheses of Lemma 9-2.6 (the BEZ lemma for meet-semidistributivity). Suppose W, X, Y, and Z are regions such that $X \wedge Y = X \wedge Z$ and such that $W \prec Y$ and $W \prec Z$. Let M be the region $X \wedge Y = X \wedge Z$. In particular, M is a lower bound for Y and Z, so $M \leq W = Y \wedge Z$. We need to show that $X \wedge (Y \vee Z) = M$.

Let F_Y be the facet shared by Y and W and let F_Z be the facet shared by Z and W. Let H_Y and H_Z be the hyperplanes containing these facets. We have $S(Y) = S(W) \cup \{H_Y\}$ and $S(Z) = S(W) \cup \{H_Z\}$. Since $(\mathcal{A}, B)$ is tight, $F_Y \cap F_Z$ is $(n-2)$-dimensional. Let $\mathcal{A}'$ be the set of hyperplanes of $\mathcal{A}$ containing $F_Y \cap F_Z$. Lemma 9-1.26 says that $\mathcal{A}'$ is a rank-two subarrangement with basic hyperplanes H_Y and H_Z and that $S(Y \vee Z) = S(W) \cup \mathcal{A}'$. Also, $S(W) \cap \mathcal{A}' = \varnothing$, and since $M \leq W$, we have $S(M) \cap \mathcal{A}' = \varnothing$ also.

Now let $M' = X \wedge (Y \vee Z)$. In any case $M' \geq M$. Suppose for the sake of contradiction that $M' > M$ and consider a region N with $M \prec N \leq M'$. Since $N \leq M'$, we have $N \leq X$. Therefore, since $M = X \wedge Y$, we have $N \not\leq Y$.

Let H be the hyperplane separating M from N, so that $S(N) = S(M) \cup \{H\}$. But $N \le M' \le (Y \vee Z)$, so $H \in S(W) \cup \mathcal{A}'$. But $H \notin S(W)$ because otherwise $N \le W < Y$. Thus $H \in \mathcal{A}'$. We also rule out the possibility that $H = H_Y$, because if so, then $S(N) = S(M) \cup \{H_Y\} \subseteq S(W) \cup \{H_Y\} = S(Y)$. Arguing symmetrically (exchanging Y and Z), we rule out the possibility that $H = H_Z$. Since $S(M) \cap \mathcal{A}' = \varnothing$ and $H \in \mathcal{A}' \setminus \{H_Y, H_Z\}$, we consider the set $S(N)$ to obtain a contradiction to Lemma 9-1.24. This contradiction shows that $X \wedge (Y \vee Z) = M$. Lemma 9-2.6 now says that $\mathrm{Pos}(\mathcal{A}, B)$ is meet-semidistributive. Since all posets of regions are self-dual (Exercise 9.4), we conclude that $\mathrm{Pos}(\mathcal{A}, B)$ is semidistributive.

Now suppose $\mathcal{A}$ is not tight with respect to B. If $\mathrm{Pos}(\mathcal{A}, B)$ is not a lattice, then we are done, so we assume that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. Since $(\mathcal{A}, B)$ is not tight, there is a region W with two upper facets whose intersection is of dimension lower than $n - 2$. Let Y and Z be the regions covering W through these facets, and let H_Y and H_Z be the facet-defining hyperplanes. (It may be helpful to consider the poset of regions shown in Figure 9-3.2, taking W, Y, and Z to be the regions with separating sets 1, 12 and 14.) Let $R = Y \vee Z$.

For most of the rest of the proof, we will consider separating sets with respect to W and thus order relations in the poset of regions $\mathrm{Pos}(\mathcal{A}, W)$. In $\mathrm{Pos}(\mathcal{A}, W)$, the region R need not be the join of Y and Z, but in any case it is a minimal upper bound for Y and Z. The separating sets of Y and Z, with respect to W, are $S(Y) = \{H_Y\}$ and $S(Z) = \{H_Z\}$, so Lemma 9-3.14 implies that the set $\mathcal{L}(R)$ of lower hyperplanes of R, with respect to W, is contained in $\{H_Y, H_Z\}$. If $\mathcal{L}(R) = \varnothing$, then we obtain a contradiction: As an easy consequence (proved as Exercise 9.24) of Lemma 9-1.18 we see that $R = W$. If $|\mathcal{L}(R)| = 1$, then R covers a unique element Q, which is therefore an upper bound for Y and Z, contradicting the fact that R is a minimal upper bound. Thus $\mathcal{L}(R) = \{H_Y, H_Z\}$.

By Lemma 9-1.17, there are exactly two distinct regions, Q_Y and Q_Z, that are covered by R in $\mathrm{Pos}(\mathcal{A}, W)$, separated from R by the hyperplanes H_Y and H_Z respectively. Lemma 9-3.13 implies that the two facets of R associated to these covers intersect in an $(n-2)$-dimensional face F of R. Lemma 9-1.27 says that $Q_Y \wedge Q_Z$ exists in $\mathrm{Pos}(\mathcal{A}, W)$. Writing X for $Q_Y \wedge Q_Z$, Lemma 9-1.27 says further that $S(X) = S(R) \setminus \mathcal{A}'$ and $S(R) = S(X) \cup \mathcal{A}'$, where $\mathcal{A}'$ is the set of all hyperplanes in $\mathcal{A}$ containing $H_Y \cap H_Z$. In particular, $\{H_Y, H_Z\} \cap S(X) = \varnothing$ and therefore $X \wedge Y = X \wedge Z = W$. If $X = W$, then Lemma 9-1.27 says that all regions in the interval $[X, R]$, including Y and Z, contain F, and we obtain a contradiction to the assumption that the facets $Y \cap W$ and $Z \cap W$ of W intersect in dimension lower than $n - 2$. Thus $X \neq W$.

We now return to the poset of regions $\mathrm{Pos}(\mathcal{A}, B)$, rather than $\mathrm{Pos}(\mathcal{A}, W)$. Lemma 9-1.20 and the fact that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice imply that the assertions $X \wedge Y = X \wedge Z = W$ and $X \le R$ proved in $\mathrm{Pos}(\mathcal{A}, B)$ also hold in $\mathrm{Pos}(\mathcal{A}, B)$. But then $X \wedge (Y \vee Z) = X \wedge R = X \neq W = X \wedge Y$. This is a counterexample to the meet-semidistributive law in $\mathrm{Pos}(\mathcal{A}, B)$. $\square$

We now give the proof of Proposition 9-3.11, which is modeled after the proof of Theorem 9-1.21 given in [70], but uses semidistributivity. Indeed, Proposition 9-3.11 establishes Theorem 9-1.21 in the case where $(\mathcal{A}, B)$ is tight.

Proof of Proposition 9-3.11. We will show that $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$ is linearly independent. Linear independence of $\{\mathbf{n}_H \mid H \in \mathcal{L}(R)\}$ follows by Exercise 9.4.

If $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$ is not linearly independent, then there exists a linear relation $\sum_{\mathbf{y}\in U} c_\mathbf{y}\mathbf{y} = \sum_{\mathbf{z}\in V} c_\mathbf{z}\mathbf{z}$ for U and V disjoint subsets, not both empty, of $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$ and all $c_\mathbf{y}$ and $c_\mathbf{z}$ positive. For each vector $\mathbf{v}$ in $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$, let $R_\mathbf{v}$ be the region whose separating set is $S(R) \cup \{H\}$ for $\mathbf{v} = \mathbf{n}_H$. Let $J = \bigvee_{\mathbf{y}\in U} R_\mathbf{y}$. Since $R_\mathbf{y} \wedge R_\mathbf{z} = R$ for any $\mathbf{y} \in U$ and $\mathbf{z} \in V$, semidistributivity implies that $R_\mathbf{z} \wedge J = R$ and thus $R_\mathbf{z} \not\leq J$ for each $\mathbf{z} \in V$. Thus by Proposition 9-1.19, any vector $\mathbf{r}$ in the interior of J has $\langle \mathbf{r}, \mathbf{y}\rangle < 0$ for all $\mathbf{y} \in U$ and has $\langle \mathbf{r}, \mathbf{z}\rangle > 0$ for all $\mathbf{z} \in V$. Applying the linear functional $\langle \mathbf{r}, \cdot\rangle$ to both sides of $\sum_{\mathbf{y}\in U} c_\mathbf{y}\mathbf{y} = \sum_{\mathbf{z}\in V} c_\mathbf{z}\mathbf{z}$ yields a nonpositive number on the left side and a nonnegative number on the right side. Since U and V are not both empty, the two numbers are not both zero, so we obtain a contradiction. We conclude that $\{\mathbf{n}_H \mid H \in \mathcal{U}(R)\}$ is linearly independent. □

9-3.2 Simplicial arrangements

Recall that the *faces* of an arrangement $\mathcal{A}$ are the faces of the regions of $\mathcal{A}$. In particular, the *rays* of $\mathcal{A}$ are the 1-dimensional faces of $\mathcal{A}$ and the *facets* of $\mathcal{A}$ are the $(n-1)$-dimensional faces of $\mathcal{A}$. There is exactly one 0-dimensional face, the origin, and the n-dimensional faces are the regions. Every facet of $\mathcal{A}$ is a facet of exactly two regions of $\mathcal{A}$.

We now show that the rays of a simplicial arrangement admit a particularly nice coloring. The coloring result (Theorem 9-3.15 below) can be stated in standard terminology as follows. A simplicial arrangement defines an abstract simplicial complex whose vertices are the rays of the arrangement and whose maximal faces are the sets of rays contained in regions. Theorem 9-3.15 is precisely the statement that this complex is *balanced*. For more information on balanced complexes, see for example [418, Chapter III.4].

Theorem 9-3.15. *If $\mathcal{A}$ is a simplicial arrangement in $\mathbb{R}^n$, then the rays of $\mathcal{A}$ can be colored with n colors such that the n rays of each region are given n distinct colors. This coloring is unique up to permuting the color set.*

Proof. Given a region Q, a coloring of the rays of Q, and a region R adjacent to Q, we define a coloring of the rays of R in the only natural way: The $n-1$ common rays are already colored, and the remaining ray of R is colored the same as the remaining ray of Q. Now, choose a base region B and start with a coloring of the n rays of B with n distinct colors. Let this coloring propagate to adjacent regions. In light of Lemma 9-1.12, every ray of $\mathcal{A}$ is assigned a color, so we need only rule out the possibility that some ray is assigned different colors based on how the coloring propagates from B to the ray. Specifically, we need to rule out the existence of a region R and two sequences of regions $B = Q_0, \ldots, Q_k = R$ with Q_{i-1} and Q_i adjacent for each $i = 1, \ldots, k$ and $B = R_0, \ldots, R_m = R$ with R_{i-1} and R_i adjacent for each $i = 1, \ldots, m$ such that the coloring of the rays of B propagates to two different colorings of the rays of R. If such a region and sequences exist, then we concatenate the sequences (reversing one of them) to obtain a sequence

$$B = Q_0, \ldots, Q_k = R_m, \ldots, R_0 = B$$

along which the coloring of B propagates to a different coloring of B.

Thus we complete the proof by verifying the following fact: Propagating the coloring of B along any sequence $B = R_0, \ldots, R_k = B$ of adjacent regions, we obtain the same coloring of B. Within this proof, we call a sequence $B = R_0, \ldots, R_k = B$ a *loop* and call a loop *good* if propagating the coloring along the loop does not change the coloring of B. We argue by a double induction that every loop is good, first by induction on the maximum value of $|S(R_i)|$ (the size of the separating set of R_i) on the loop and second, fixing this maximum, by induction on the number of times the maximum is attained. The base case of the induction is where this maximum is zero, so that $k = 0$ and the loop consists of a single region B.

If the maximum is positive, let i be an index such that $|S(R_i)|$ attains the maximum. Then $|S(R_{i-1})|$ and $|S(R_{i+1})|$ are both one less than the maximum. There are two cases: either $R_{i-1} = R_{i+1}$ or not. If $R_{i-1} = R_{i+1}$, then we consider the loop obtained from $B = R_0, \ldots, R_k = B$ by deleting R_i and R_{i+1}. This loop either has a lower maximum size of a separating set or attains the maximum fewer times. By induction, the shortened loop is good. In the original loop, propagating the coloring from R_{i-1} to R_i to $R_{i+1} = R_{i-1}$ returns the same coloring of R_{i-1}, and we conclude that the original loop $B = R_0, \ldots, R_k = B$ is good.

Now suppose $R_{i-1} \neq R_{i+1}$. Let H_- be the hyperplane containing the common facet of R_i and R_{i-1} and let H_+ be the hyperplane containing the common facet of R_i and R_{i+1}. Let $\mathcal{A}'$ be the rank-two subarrangement containing H_- and H_+. Lemma 9-1.27 says that $R' = R_{i-1} \wedge R_{i+1}$ has $S(R') = S(R_i) \setminus \mathcal{A}'$ and that the interval $[R', R_i]$ is the union of two chains of the same length, disjoint except at R' and R_i. Write $R' = Q_0 \prec \cdots \prec Q_\ell = R_i$ and $R' = Q'_0 \prec \cdots \prec Q'_\ell = R_i$ for these two chains, with $Q_{\ell-1} = R_{i-1}$ and $Q'_{\ell-1} = R_{i+1}$. Alter the loop $B = R_0, \ldots, R_k = B$ by deleting R_i and inserting in its place $Q_{\ell-2}, \ldots, Q_0 = Q'_0, \ldots, Q'_{\ell-2}$. The inserted sequence may consist of as few as one region, in which case the region is R'. The altered loop either has a lower maximum size of a separating set or attains the maximum fewer times, so by induction, it is good.

The proof can now be completed by checking that propagating the coloring along the original loop $B = R_0, \ldots, R_k = B$ yields the same coloring of B as propagating the coloring along the altered loop. Equivalently, propagating a coloring of R_i along the sequence $R_i = Q_\ell, \ldots, Q_0 = Q'_0, \ldots, Q'_\ell = R_i$ does not alter the coloring of R_i. By Lemma 9-1.27, the intersection of the regions in the latter sequence is an $(n-2)$-dimensional face F of $\mathcal{A}$. The coloring of the rays of F does not change along the sequence. The remaining two colors are switched at every step along the sequence, and thus are switched an even number of times. $\square$

When $\mathcal{A}$ is simplicial, the k-dimensional faces of a region are in bijection with the k-element subsets of the rays of the region. Thus we have the following immediate corollary to Theorem 9-3.15.

Corollary 9-3.16. *Suppose $\mathcal{A}$ is a simplicial arrangement in $\mathbb{R}^n$ and color the rays of $\mathcal{A}$ with n colors as in Theorem 9-3.15. Color each k-dimensional face F of $\mathcal{A}$ with the set of colors of the rays contained in F. Then the k-dimensional faces of each region are colored with distinct colors (that is, with distinct sets of colors). This coloring is unique up to permuting the color set.*

Suppose $\mathcal{A}$ is a simplicial hyperplane arrangement, with faces colored as in Corollary 9-3.16 using a color set S with $|S| = n$. In the adjacency graph $G(\mathcal{A})$ of $\mathcal{A}$, each edge corresponds to a facet (codimension-1 face) of $\mathcal{A}$, the intersection of the two regions connected by the edge in the adjacency graph. The facets of $\mathcal{A}$ are colored with sets of the form $S \setminus \{s\}$ for $s \in S$. Color each edge in $G(\mathcal{A})$ with the color s if the corresponding facet of $\mathcal{A}$ is colored $S \setminus \{s\}$. For each subset I of the colors, consider the graph $G_I(\mathcal{A})$ obtained from the adjacency graph by deleting all of the edges with colors in I. The following proposition shows how to recover the faces of a simplicial arrangement from its colored adjacency graph.

Proposition 9-3.17. *Suppose $\mathcal{A}$ is a simplicial hyperplane arrangement, colored as in Corollary 9-3.16. For each subset I of the color set and each component C of the graph $G_I(\mathcal{A})$, let F_C be the intersection of the regions of $\mathcal{A}$ corresponding to the vertices of C.*

(i) *F_C is an $|I|$-dimensional face of $\mathcal{A}$ colored I.*

(ii) *Every face of $\mathcal{A}$ colored I is F_C for a unique component C of $G_I(\mathcal{A})$.*

Since each face of $\mathcal{A}$ has a color, as I and C are allowed to vary, Proposition 9-3.17(ii) accounts for each face of $\mathcal{A}$ exactly once.

Proof. The vertices of $G_I(\mathcal{A})$ are regions of $\mathcal{A}$. Let R be a region in C. Since each ray of R is colored a distinct (singleton) color, and the color of a face F of R is the set of colors of the rays of R contained in F, there is an $|I|$-dimensional face F of R colored I. We will show that $F = F_C$. If $F = R$, then $C = \{R\}$ and we are done, so assume $F \subsetneq R$. The face F is the intersection of the facets of R containing it. These are exactly the facets of R colored $S \setminus \{s\}$

for $s \notin I$. The corresponding edges of $G(\mathcal{A})$ are colored with the colors s for $s \notin I$. Writing R_s for the region sharing with R a facet colored $S \setminus \{s\}$, we have $F = R \cap \bigcap_{s \notin I} R_s$, so $F_C \subseteq F$. On the other hand, all regions in C are connected to R by a path in $G(\mathcal{A})$ using only the colors in $S \setminus I$. Thus since the color of F is I, every region on this path contains F. We conclude that $F_C = F$, and we have established (i).

Furthermore, if F is a face of $\mathcal{A}$ colored I, then the construction above realizes F as F_C for a component C of $G_I(\mathcal{A})$, namely the component of $G_I(\mathcal{A})$ containing R. To prove uniqueness of C, we need to verify that if F is a face of R colored I and F is also a face of Q, then Q and R are in the same component of $G_I(\mathcal{A})$. Lemma 9-1.12 says that there exists a sequence of regions $Q = R_0, \ldots, R_k = R$ with R_{i-1} adjacent to R_i for $i = 1, \ldots, k$ and with the property that, moving from Q to R in the sequence, no hyperplane of $\mathcal{A}$ is crossed more than once. This sequence is a path in $G(\mathcal{A})$. If R' is a region containing F and H is a facet-defining hyperplane of R', then H contains F if and only if the color of the facet defined by H is $S \setminus \{s\}$ for $s \notin I$. Suppose an edge R_i, R_{i+1} in the path $Q = R_0, \ldots, R_k = R$ is colored with a color $s \in I$. Then either R_i or R_{i+1} is separated from F by the hyperplane H defining the common facet of R_i and R_{i+1}. Since F is a face of Q and of R, neither is separated from F by a hyperplane, and we see that the path crosses H twice. This contradiction shows that all of the edges in the path $Q = R_0, \ldots, R_k = R$ are colored with colors not in I. Thus Q and R are in the same connected component of $G_I(\mathcal{A})$. We have established (ii). $\square$

Proposition 9-3.17 can be rephrased as a method for determining whether two simplicial arrangements are combinatorially "the same."

Definition 9-3.18. The *face semilattice* of a hyperplane arrangement $\mathcal{A}$ is the set of faces of $\mathcal{A}$, partially ordered by containment. Corollary 9-1.11 implies that this is a meet-semilattice, with the meet being intersection. Two hyperplane arrangements are called *combinatorially isomorphic* if they have isomorphic face semilattices.

Proposition 9-3.19. *Suppose $\mathcal{A}$ and $\mathcal{A}'$ are simplicial arrangements with faces colored as in Corollary 9-3.16, using the same color set. Color each edge of $G(\mathcal{A})$ as described before Proposition 9-3.17, and color $G(\mathcal{A}')$ in the same way. Then $\mathcal{A}$ and $\mathcal{A}'$ are combinatorially isomorphic if and only if there exists an isomorphism from $G(\mathcal{A})$ to $G(\mathcal{A}')$ that preserves the colors (up to permuting the color set).*

Proof. An isomorphism that preserves colors also preserves the graphs $G_I(\mathcal{A})$ for each color set I, and thus the sets of regions defining faces as F_C. The isomorphism can be extended to an isomorphism of face semilattices by identifying each face F as F_C for C a component of the appropriate $G_I(\mathcal{A})$, and mapping F to $F_{C'}$ where C' is the corresponding component of $G_I(\mathcal{A}')$. Conversely, given an isomorphism of face lattices, the uniqueness in Corollary

9-3.16 implies that the isomorphism preserves face colors. The coloring on adjacency graphs is completely determined by the coloring on faces. □

Remark 9-3.20. A weaker notion of combinatorial isomorphism (isomorphism of adjacency graphs) is considered in Exercise 9.25. Isomorphism of adjacency graphs is a weaker condition because the adjacency graph is in essence the restriction of the face semilattice to faces of dimension n (regions) and faces of dimension $n-1$ (facets). Exercise 9.25 thus applies to say that weakly combinatorially isomorphic hyperplane arrangements have isomorphic posets of regions when the base regions are chosen to coincide under the isomorphism. For simplicial arrangements, the two notions coincide: When $\mathcal{A}$ is simplicial, the zonotope dual to $\mathcal{A}$ is simple, and its face lattice (and thus the face lattice of $\mathcal{A}$) is determined by the adjacency graph as explained in [263]. We will not need to appeal to this result of [263], so we omit the details here. Proposition 9-3.19, which appears to be a step in the same direction, is not as strong because it requires not only the adjacency graph but also a coloring of that graph's edges.

9-4. Biconvexity and rank-two biconvexity

In this section, we describe the connection between the poset of regions $\text{Pos}(\mathcal{A}, B)$ and several notions of combinatorial convexity of subsets of $\mathcal{A}$ with respect to B.

9-4.1 Convexity, biconvexity, and strong biconvexity

Definition 9-4.1. Let $\mathbf{b}$ be a vector in the interior of B, and for each $H \in \mathcal{A}$, let $\mathbf{n}_H$ be a normal vector to H such that $\langle \mathbf{b}, \mathbf{n}_H \rangle > 0$. A subset S of $\mathcal{A}$ is *convex with respect to* B if

$$\bigl(\text{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in S\}\bigr) \cap \{\mathbf{n}_H \mid H \in \mathcal{A}\} = \{\mathbf{n}_H \mid H \in S\}.$$

Here $\text{Span}_{\geq 0}$ denotes nonnegative linear span. Define a closure operator $S \mapsto \overline{S}$ on subsets $S \subseteq \mathcal{A}$ by

$$\overline{S} = \{H' \in \mathcal{A} \mid \mathbf{n}_{H'} \in \text{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in S\}\}.$$

Then S is convex if and only if it is closed in the sense that $\overline{S} = S$. Exercise 9.27 is to verify that $\overline{S}$ is the intersection of all convex sets containing S. The subset S is *biconvex with respect to* B if S and $\mathcal{A} \setminus S$ are both convex with respect to B. The subset S of $\mathcal{A}$ is *strongly biconvex* if

$$\bigl(\text{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in S\}\bigr) \cap \bigl(\text{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in \mathcal{A} \setminus S\}\bigr) = \{0\}.$$

The notion of convexity in Definition 9-4.1 corresponds to a well-established notion of convexity in finite sets of vectors. This is, for example, a linearization of the notion defined in Example 1 of [141, Section 3]. See also [70, Remark 5.3]. Exercise 9.28 establishes this linearized notion of convexity in general.

The following proposition is immediate from Definition 9-4.1.

Proposition 9-4.2. *Given a hyperplane arrangement $\mathcal{A}$ with base region B and a subset $S \subseteq \mathcal{A}$, the following implications hold for convexity with respect to B:*

$$S \text{ is strongly biconvex} \implies S \text{ is biconvex} \implies S \text{ is convex.}$$

We now show that strong biconvexity characterizes separating sets.

Theorem 9-4.3. *Given a hyperplane arrangement $\mathcal{A}$ with base region B, a subset $S \subseteq \mathcal{A}$ is the separating set of some region if and only if S is strongly biconvex with respect to B.*

Proof. Suppose $S = S(R)$ for a region R and let $\mathbf{r}$ be a vector in the interior of R. By Proposition 9-1.19, any nonzero vector $\mathbf{n} \in \mathrm{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in S\}$ has $\langle \mathbf{n}, \mathbf{r}\rangle < 0$, while any nonzero vector $\mathbf{n} \in \mathrm{Span}_{\geq 0}\{\mathbf{n}_H \mid H \in \mathcal{A} \setminus S\}$ has $\langle \mathbf{n}, \mathbf{r}\rangle > 0$. Thus S is strongly biconvex.

Conversely, if S is strongly biconvex, then a standard separation theorem from the theory of convexity (a special case of the Hahn-Banach Separation Theorem) implies that there exists a vector $\mathbf{x}$ such that $\langle \mathbf{n}_H, \mathbf{x}\rangle < 0$ for $H \in S$ and $\langle \mathbf{n}_H, \mathbf{x}\rangle \geq 0$ for $H \in \mathcal{A} \setminus S$. Let $\mathbf{b}$ be any vector in the interior of B. For small enough $\varepsilon > 0$, the vector $\mathbf{r} = \mathbf{x} + \varepsilon\mathbf{b}$ has $\langle \mathbf{n}_H, \mathbf{b}\rangle < 0$ for $H \in S$ and $\langle \mathbf{n}_H, \mathbf{b}\rangle > 0$ for $H \in \mathcal{A} \setminus S$. In particular, $\mathbf{r}$ is contained in a region R, and Proposition 9-1.19 says that $S(R) = S$. □

Convexity has a reformulation in terms of regions and halfspaces defined by hyperplanes in $\mathcal{A}$. Given vectors $\mathbf{n}_H$ as in Definition 9-4.1, write H^+ for the closed halfspace $\{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{x}, \mathbf{n}_H\rangle \geq 0\}$. This is a union of regions of $\mathcal{A}$. Let $\mathcal{R}_+(H)$ be the set of regions contained in H^+. In other words, $\mathcal{R}_+(H)$ is the set of regions whose separating set does not contain H. Exercise 9.29 verifies that for $S \subseteq \mathcal{A}$,

$$\overline{S} = \Big\{ H \in \mathcal{A} \mid \mathcal{R}_+(H) \supseteq \bigcap_{H' \in S} R_+(H')) \Big\}.$$

To understand this formulation, one should notice that $\bigcap_{H' \in S} R_+(H')$ is the set of regions that are **not** separated from B by any hyperplane in S. To form $\overline{S}$, we adjoin to S every hyperplane that we can adjoin without making $\bigcap_{H' \in S} R_+(H')$ any smaller.

9-4.2 Rank-two!biconvexity

Theorem 9-4.3 says that strong biconvexity characterizes separating sets of regions. We will see in this section that in the simplicial case, the weaker condition of biconvexity also characterizes separating sets of regions. In fact, an even weaker condition of rank-two biconvexity characterizes separating sets in the simplicial case. Furthermore, the join operation can be described in terms of the closure operation or the rank-two closure operation.

Definition 9-4.4. Given $(\mathcal{A}, B)$ and a rank-two subarrangement $\mathcal{A}'$, write B' for the $\mathcal{A}'$-region containing B. Say a set of hyperplanes $S \subseteq \mathcal{A}$ is *rank-two convex* with respect to B if, for every rank-two subarrangement $\mathcal{A}' \subseteq \mathcal{A}$, the intersection $S \cap \mathcal{A}'$ is convex with respect to B'. We define a *rank-two closure* operator which takes $U \subseteq \mathcal{A}$ to ${}^2\overline{U}^2$, defined to be the intersection of all rank-two convex sets in $\mathcal{A}$ containing U. Exercise 9.31 verifies that ${}^2\overline{U}^2$ is rank-two convex. A set $S \subseteq \mathcal{A}$ is *rank-two biconvex* if S and $\mathcal{A} \setminus S$ are both rank-two convex.

Theorem 9-4.5. *Suppose* $\mathrm{Pos}(\mathcal{A}, B)$ *is simplicial, and let S be a subset of $\mathcal{A}$. Then the following are equivalent:*

(i) *S is the separating set of some region.*

(ii) *S is strongly biconvex with respect to B.*

(iii) *S is biconvex with respect to B.*

(iv) *S is rank-two biconvex with respect to B.*

Without any hypotheses on $\mathcal{A}$, Theorem 9-4.3 says that conditions (i) and (ii) are equivalent, Proposition 9-4.2 says that (ii) implies (iii), and Exercise 9.32 shows that (iii) implies (iv). To prove Theorem 9-4.5, we will show that (iv) implies (i) under the hypothesis that $\mathcal{A}$ is simplicial.

Example 9-4.6. When $\mathcal{A}$ is not simplicial, a biconvex subset of $\mathcal{A}$ can fail to be the separating set of a region. In Figure 9-1.1 (Example 9-1.3), the set containing the hyperplanes numbered 1 and 3 is biconvex but no region has this separating set.

Example 9-4.7. On the other hand, the conclusion of Theorem 9-4.5 may hold when $\mathcal{A}$ is not simplicial. One example is the poset of regions in Example 9-3.5. In this example, there are 32 subsets of $\mathcal{A}$ and 18 regions. One can check that the 14 subsets that are not separating sets are not rank-two biconvex.

In the simplicial case, we can also make two definite statements about the join operation.

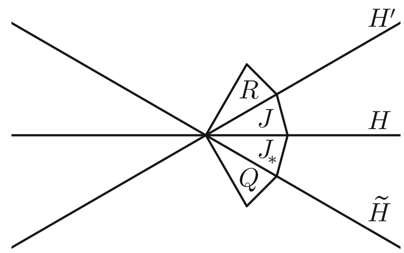

Figure 9-4.1: An illustration of the proof of Lemma 9-4.10

Theorem 9-4.8. *Suppose $\mathcal{A}$ is simplicial and let Q and R be regions. Then*

(i) *$Q \vee R$ is the unique region with separating set $\overline{S(Q) \cup S(R)}$.*

(ii) *$Q \vee R$ is also the unique region with separating set ${}^{2}\overline{S(Q) \cup S(R)}^{2}$.*

Using the self-duality of $\mathrm{Pos}(\mathcal{A}, B)$, Theorem 9-4.8 implies a similar description of the meet (Exercise 9.33).

We now prepare to prove Theorems 9-4.5 and 9-4.8.

Definition 9-4.9. The *depth* of a hyperplane $H \in \mathcal{A}$ with respect to B is the minimum, over regions R with $H \in S(R)$, of $|S(R)|$. For example, the set $\mathcal{B}(B)$ of facet-defining hyperplanes of B is the set of hyperplanes of depth 1.

Lemma 9-4.10. *Suppose $\mathcal{A}$ is simplicial and $H \in \mathcal{A} \setminus \mathcal{B}(B)$. Then there exists a rank-two subarrangement $\mathcal{A}'$ of $\mathcal{A}$ such that $H \in \mathcal{A}'$ and both basic hyperplanes of $\mathcal{A}'$ have depth strictly smaller than the depth of H.*

Proof. This proof is illustrated in Figure 9-4.1. Choose J such that $S(J)$ has minimal size among separating sets containing H. In particular, H is the unique lower hyperplane of J. By Lemma 9-1.17, the region J covers exactly one region J_*, which has $S(J_*) = S(J) \setminus \{H\}$. Since $H \notin \mathcal{B}(B)$, in particular J_* is not B, so J_* covers some region Q.

Let $\widetilde{H}$ be the hyperplane containing the common facet of Q and J_* and let $\mathcal{A}'$ be the rank-two subarrangement containing H and $\widetilde{H}$. Since J_* is a simplicial region, the intersection $F = Q \cap J_* \cap J$ (the intersection of the facets $Q \cap J_*$ and $J_* \cap J$ of J_*) is an $(n-2)$-dimensional face of J_*. (See Exercise 9.23.) Since F is a face of J_*, it is the nonnegative span of $n-2$ rays of J_*. Since $F \subseteq J_* \cap J$, all of these rays are in J as well, so F is a face of J. Specifically, F is the intersection of two facets of J, one being $J_* \cap J$ and the other $J \cap R$ for some R with $J \prec R$. The facet $J \cap J_*$ is defined by the hyperplane H. Write H' for the hyperplane defining the facet $J \cap R$. This hyperplane H' is also in $\mathcal{A}'$, since it contains F and therefore contains $H \cap \widetilde{H}$.

Now suppose $\mathcal{A}'$ is ordered $H_1, \ldots, H_k$ as in Lemma 9-1.24. Since $Q \prec J_* \prec J \prec R$, Lemma 9-1.24 implies that H is between $\widetilde{H}$ and H' in this total order. In particular, H is not basic in $\mathcal{A}'$. Lemma 9-1.25 says that the set of regions containing F is an interval in $\operatorname{Pos}(\mathcal{A}, B)$ isomorphic to the poset of regions $\operatorname{Pos}(\mathcal{A}', B')$, where B' is the $\mathcal{A}'$-region containing B. Furthermore, the regions in this interval have separating sets differing by hyperplanes in $\mathcal{A}'$. Lemma 9-1.24 implies that there are two distinct elements Q' and Q'' of this interval whose separating set contains exactly one hyperplane (a basic hyperplane) in $\mathcal{A}'$. Since $S(J)$ contains at least two hyperplanes (H and $\widetilde{H}$) in $\mathcal{A}'$, it is larger than the separating set of Q' and the separating set of Q''. In particular, the two basic hyperplanes of $\mathcal{A}'$ have depth strictly smaller than the depth of H. □

Lemma 9-4.11. *Suppose $\mathcal{A}$ is simplicial. Every nonempty rank-two biconvex subset of $\mathcal{A}$ contains at least one hyperplane in $\mathcal{B}(B)$.*

Proof. Suppose S is a biconvex subset of $\mathcal{A}$. We must show that S contains a hyperplane of depth 1. Suppose not, and take H to be a hyperplane of minimal depth in S. Lemma 9-4.10 says that H is contained in a rank-two subarrangement $\mathcal{A}'$ such that both basic hyperplanes of $\mathcal{A}'$ have strictly smaller depth. But $\mathcal{A} \setminus S$ is rank-two convex, and thus one of these two basic hyperplanes is in S, contradicting the minimality of H. □

Lemma 9-4.12. *Suppose $\mathcal{A}$ is a hyperplane arrangement, B is a base region, and H is a hyperplane in $\mathcal{B}(B)$. Let C be the region that shares with B the facet defined by H. Suppose a subset $S \subseteq \mathcal{A}$ contains H.*

(i) *If S is rank-two convex with respect to B then $S \setminus \{H\}$ is rank-two convex with respect to C.*

(ii) *If S is rank-two biconvex with respect to B then $S \setminus \{H\}$ is rank-two biconvex with respect to C.*

Proof. We argue the second assertion; the proof of the first assertion is similar, but simpler, and is left as Exercise 9.34. We work with normal vectors $\mathbf{n}_H$ as in Definition 9-4.1, but now we need to explicitly mention B in the notation. For each hyperplane $K \in \mathcal{A}$, choose a nonzero normal vector $\mathbf{n}_K^B$ with $\langle \mathbf{b}, \mathbf{n}_K^B \rangle > 0$. Let $\mathbf{c}$ be a vector in the interior of C. Setting $\mathbf{n}_K^C = \mathbf{n}_K^B$ for $K \neq H$ and setting $\mathbf{n}_H^C = -\mathbf{n}_H^B$, we see that $\langle \mathbf{c}, \mathbf{n}_K^C \rangle > 0$ for all $K \in \mathcal{A}$. Now let $\mathcal{A}'$ be any rank-two subarrangement. If $H \notin \mathcal{A}'$, then $S' \cap \mathcal{A}'$ is biconvex in $\mathcal{A}'$ with respect to C' (the $\mathcal{A}'$-region containing C) because $S' \cap \mathcal{A}' = S \cap \mathcal{A}'$ and because C' is also the $\mathcal{A}'$-region containing B. On the other hand, if $H \in \mathcal{A}'$, then H is basic in $\mathcal{A}'$ (with respect to B). In this case, $S' \cap \mathcal{A}'$ is still biconvex in $\mathcal{A}'$ with respect to C', as illustrated in Figure 9-4.2. In the figure, the vectors $\{\mathbf{n}_K^B \mid K \in S \cap \mathcal{A}'\}$ are shown by solid black arrows. The other vectors in $\{\mathbf{n}_K^B \mid K \in \mathcal{A}'\}$ are shown by dotted black arrows and the vector $\mathbf{n}_H^C = -\mathbf{n}_H^B$ is shown by a differently-dotted gray arrow. □

We are now prepared to prove the main results of this section. We begin with Theorem 9-4.5.

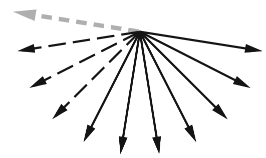

Figure 9-4.2: An illustration for the proof of Lemma 9-4.12

Proof of Theorem 9-4.5. As discussed above, it remains only to show that a rank-two biconvex set S in a simplicial arrangement $\mathcal{A}$ is necessarily the separating set of some region. We argue by induction on $|S|$, allowing the choice of base region to vary. The base case $|S| = 0$ is trivial. If S is rank-two biconvex with respect to B and $|S| > 0$, then Lemma 9-4.11 says that S contains some hyperplane H in $\mathcal{B}(B)$. Let C be the region that shares with B the facet defined by H. Lemma 9-4.12 says that the set $S' = S \setminus \{H\}$ is rank-two biconvex with respect to C. Thus by induction, S' is the separating set, with respect to C, of a region R of $\mathcal{A}$. The separating set of R with respect to B is $S' \cup \{H\} = S$. □

Remark 9-4.13. The hypothesis that $\mathcal{A}$ is simplicial is needed in the proof of Theorem 9-4.5 for several reasons. First, in the proof of Lemma 9-4.10 (which is the key to Lemma 9-4.11), the hypothesis that all regions are simplicial is used (specifically for the region J^*). Furthermore, the structure of the proof of Theorem 9-4.5 makes any generalization of hypotheses meaningless. The base region B varies in the induction, in such a way that every region serves as the base region at some point in the proof. In order for the hypotheses to hold as B varies, in particular, the poset of regions must be a lattice for each choice of B. Thus Corollary 9-3.10 says that $\mathcal{A}$ is simplicial.

Proof of Theorem 9-4.8. We prove the second assertion first, beginning with several claims.

Claim 1: *$S(Q \vee R) = {}^2\overline{S(Q) \cup S(R)}^2$ if and only if ${}^2\overline{S(Q) \cup S(R)}^2$ is rank-two biconvex. Proof:* If $S(Q \vee R) = {}^2\overline{S(Q) \cup S(R)}^2$, then Theorem 9-4.5 says that ${}^2\overline{S(Q) \cup S(R)}^2$ is rank-two biconvex. Conversely, if the set ${}^2\overline{S(Q) \cup S(R)}^2$ is rank-two biconvex, then Theorem 9-4.5 says that ${}^2\overline{S(Q) \cup S(R)}^2$ is the separating set of some region. But ${}^2\overline{S(Q) \cup S(R)}^2$ is contained in every rank-two convex set containing $S(Q)$ and $S(R)$, and thus (again by Theorem 9-4.5) contained in the separating set of every region above Q and R. Thus ${}^2\overline{S(Q) \cup S(R)}^2$ is $S(Q \vee R)$. □ (Claim 1)

Claim 2: *If Q and R are distinct regions covering B, then ${}^2\overline{S(Q) \cup S(R)}^2$ is rank-two biconvex. Proof:* In this case $S(Q) = \{H_1\}$ and $S(R) = \{H_2\}$ for distinct hyperplanes $H_1, H_2 \in \mathcal{B}(B)$. These hyperplanes are basic in the rank-two subarrangement $\mathcal{A}'$ containing them, so ${}^2\overline{S(Q) \cup S(R)}^2 = \mathcal{A}'$, and Lemma 9-1.26 says that $S(Q \vee R) = \mathcal{A}'$ as well. □ (Claim 2)

Suppose C is a region covering B with $C \le Q \wedge R$. Lemma 9-1.20 says that the interval $[C, -B]$ in $\mathrm{Pos}(\mathcal{A}, B)$ is isomorphic (by the identity map) to the interval $[C, -B]$ in $\mathrm{Pos}(\mathcal{A}, C)$. We continue to write $S(Q)$ for the separating set of Q with respect to B, and we now also write $S_C(Q)$ for the separating set of Q with respect to C. Let H be the hyperplane defining the common facet of B and C. Then any region X in $[C, -B]$ has $S_C(X) = S(X) \setminus \{H\}$. We also adopt the following notational convention for the remainder of the proof: If the subscript C appears anywhere beneath the closure marker ${}^{2}\overline{}^{2}$, the rank-two closure is taken with respect to C. Otherwise, the closure is taken with respect to B. Furthermore, all mentions of rank-two (bi)convexity are with respect to B, unless otherwise noted.

Claim 3: *If* ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} = S_C(Q \vee R)$, *then* ${}^{2}\overline{S(Q) \cup S(R)}^{2} = S(Q \vee R)$.

Proof: If ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2}$ equals $S_C(Q \vee R)$, then the set $S(Q \vee R)$ equals ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} \cup \{H\}$. In particular, ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} \cup \{H\}$ is rank-two biconvex by Theorem 9-4.5. Also, since ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} \cup \{H\}$ is a rank-two convex set containing $S(Q) \cup S(R)$, we have ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} \cup \{H\} \supseteq {}^{2}\overline{S(Q) \cup S(R)}^{2}$. If the containment is proper, then Lemma 9-4.12 implies that ${}^{2}\overline{S(Q) \cup S(R)}^{2} \setminus \{H\}$ is a rank-two convex set (with respect to C) containing $S_C(Q)$ and $S_C(R)$ but properly contained in ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2}$. This contradiction implies that ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2} \cup \{H\} = {}^{2}\overline{S(Q) \cup S(R)}^{2}$. □ (Claim 3)

We now prove the second assertion of the theorem by induction on the size of $S(Q \vee R)$, with the base region B varying. The proof is modeled after the proof of the BEZ lemma (Lemma 9-2.2).

If $Q \le R$ or $R \le Q$, then the result follows because $S(Q)$ and $S(R)$ are both rank-two biconvex by Theorem 9-4.5. Thus we assume that Q and R are incomparable, so that in particular both are strictly above B. Let C and D be regions of $\mathcal{A}$ such that $B \prec C \le Q$ and $B \prec D \le R$.

If $C = D$ then both Q and R lie in the interval $[C, -B]$ in $\mathrm{Pos}(\mathcal{A}, B)$. As mentioned above, the interval $[C, -B]$ in $\mathrm{Pos}(\mathcal{A}, B)$ is isomorphic (by the identity map) to the interval $[C, -B]$ in $\mathrm{Pos}(\mathcal{A}, C)$. In particular, $Q \vee R$ is the join both in $\mathrm{Pos}(\mathcal{A}, B)$ and in $\mathrm{Pos}(\mathcal{A}, C)$. Since $|S_C(Q \vee R)| = |S(Q \vee R)| - 1$, we can appeal to induction to conclude that ${}^{2}\overline{S_C(Q) \cup S_C(R)}^{2}$ is $S_C(Q \vee R)$. Claim 3 now says that ${}^{2}\overline{S(Q) \cup S(R)}^{2}$ is $S(Q \vee R)$.

If $C \ne D$ then define E to be the join $C \vee D$ in $\mathrm{Pos}(\mathcal{A}, B)$. Since $E \le R$, we see that $Q \vee E \le Q \vee R$. In particular $|S_C(Q \vee E)| \le |S_C(Q \vee R)| = |S(Q \vee R)| - 1$, so we can appeal to induction to see that ${}^{2}\overline{S_C(Q) \cup S_C(E)}^{2}$ is $S_C(Q \vee E)$. Claim 3 says that ${}^{2}\overline{S(Q) \cup S(E)}^{2}$ is $S(Q \vee E)$. Claim 2 says that $S(E) = {}^{2}\overline{S(C) \cap S(D)}^{2}$. Furthermore, ${}^{2}\overline{S(Q) \cup S(E)}^{2}$ is the smallest rank-two convex set both containing $S(Q)$ and containing the smallest rank-two convex set containing $S(C)$ and $S(D)$. Thus ${}^{2}\overline{S(Q) \cup S(E)}^{2} = {}^{2}\overline{S(Q) \cup S(C) \cup S(D)}^{2}$, which is ${}^{2}\overline{S(Q) \cup S(D)}^{2}$ because $C \le Q$.

But now $Q \vee E$ and R are both above D and $Q \vee E \vee R \le Q \vee R$, so by the same argument, ${}^2\overline{S(Q \vee (E)) \cup S(R)}^2$ equals $S(Q \vee (E) \vee R)$. But $Q \vee (E) \vee R$ equals $Q \vee R$. Also, ${}^2\overline{S(Q \vee (E)) \cup S(R)}^2$ is the smallest rank-two convex set both containing $S(R)$ and containing the smallest rank-two convex set containing $S(Q)$ and $S(D)$. Thus ${}^2\overline{S(Q \vee (E)) \cup S(R)}^2 = {}^2\overline{S(Q) \cup S(D) \cup S(R)}^2$, which equals ${}^2\overline{S(Q) \cup S(R)}^2$ because $D \leq R$. This completes the inductive argument.

We have proved the second assertion of the theorem. In particular, the set $S(Q \vee R)$ is contained in every rank-two convex set containing $S(Q) \cup S(R)$. But as verified in Exercise 9.32, every convex set is also rank-two convex, so $S(Q \vee R)$ is contained in every convex set containing $S(Q) \cup S(R)$. Theorem 9-4.5 implies in particular that $S(Q \vee R)$ is convex, so it is the intersection of all convex sets containing $S(Q) \cup S(R)$. But Exercise 9.27 verifies that that intersection is $\overline{S(Q) \cup S(R)}$. We have proved the first assertion of the theorem. □

Remark 9-4.14. The equivalence of conditions (i) and (iii) in Theorem 9-4.5 is proved without the simplicial hypothesis, but with the hypothesis that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice, as part of [70, Theorem 5.5]. The other part of [70, Theorem 5.5] proves assertion (i) of Theorem 9-4.8, also under the hypothesis that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. In [308, Theorem 5.1], assertion (ii) of Theorem 9-4.8 is proved under the weaker hypothesis that $(\mathcal{A}, B)$ is tight.

9-5. Lattice congruences for combinatorialists

This section presents basic notions and combinatorial tools surrounding lattice congruences of finite lattices, emphasizing order-theoretic characterizations of lattice-theoretic concepts. The "combinatorialists" targeted in the section title are those who may deal frequently with posets and lattices but have less contact with lattice theory *per se*, and in particular may not have thought very much about lattice homomorphisms and congruences. Anticipating that some of the targeted combinatorialists may read this section separately from the rest of the book, we give some basic definitions in this section that we have assumed earlier in the book.

Probably none of the results of this section are surprising to those familiar with lattice homomorphisms and congruences. Indeed, some have appeared in the literature — see the Notes at the end of this chapter — or have already appeared in this volume. However, this section may provide a point of view that is different from the traditional lattice-theoretic viewpoint.

Although here we highlight the case of finite lattices, there are several possible extensions to the infinite case. If one sacrifices the purely algebraic framework by passing to complete lattices and complete homomorphisms, then many statements for finite lattices generalize essentially verbatim. If, on the other hand, one avoids the notion of completeness, then the definition of a lattice as a triple $(L, \wedge, \vee)$ remains purely algebraic, but the order-theoretic statements become less satisfying. A third option is to consider only bounded homomorphisms and bounded congruences. We consider some of these generalizations beyond finite lattices in Exercises 9.40–9.49.

9-5.1 Homomorphisms and congruences

A *lattice homomorphism* is a map η from a lattice L_1 to a lattice L_2 such that $\eta(x \wedge y) = \eta(x) \wedge \eta(y)$ and $\eta(x \vee y) = \eta(x) \vee \eta(y)$ for all $x, y \in L_1$. An *isomorphism* of lattices is a bijective lattice homomorphism.

Given a map $\eta\colon L \to L'$ and a subset $S \subseteq L'$, the notation $\eta^{-1}(S)$ means $\{x \in L \mid \eta(x) \in S\}$.

Proposition 9-5.1. *Let L and L' be finite lattices. A surjective map $\eta\colon L \to L'$ is a lattice homomorphism if and only if the following two conditions hold:*

(i) *η is order-preserving.*

(ii) *For every interval $[x, y]$ in L', the set $\eta^{-1}([x, y])$ is an interval.*

Proof. Suppose η is a lattice homomorphism. If x and y are in L and $x \leq y$ then $x \wedge y = x$ and thus $\eta(x) \wedge \eta(y) = \eta(x \wedge y) = \eta(x)$, so $\eta(x) \leq \eta(y)$. Suppose $[x, y]$ is an interval in L'. If $\eta(a)$ and $\eta(b)$ are in $[x, y]$, then $\eta(a \vee b) = \eta(a) \vee \eta(b)$ is in $[x, y]$ and similarly $\eta(a \wedge b)$ is in $[x, y]$. Thus $\eta^{-1}([x, y])$ is contained in the interval $\left[\bigwedge \eta^{-1}\big([x, y]\big), \bigvee \eta^{-1}\big([x, y]\big)\right]$. If $\eta(a) \in [x, y]$ and $\eta(b) \in [x, y]$, then $\eta(a \vee b) = \eta(a) \vee \eta(b) \in [x, y]$. Therefore $\eta\big(\bigvee \eta^{-1}\big([x, y]\big)\big)$ is in $[x, y]$. Similarly, $\eta\big(\bigwedge \eta^{-1}\big([x, y]\big)\big)$ is in $[x, y]$. Furthermore, if $a \leq b \leq c$ and if $\eta(a)$ and $\eta(c)$ are in $[x, y]$, then since η is order-preserving, $\eta(b)$ is in $[x, y]$. Thus $\eta^{-1}([x, y])$ is the entire interval $\left[\bigwedge \eta^{-1}\big([x, y]\big), \bigvee \eta^{-1}\big([x, y]\big)\right]$. We have verified (i) and (ii).

Conversely, suppose η satisfies (i) and (ii) and let x and y be elements of L. By (i), $\eta(x \wedge y)$ is a lower bound for $\eta(x)$ and $\eta(y)$, so $\eta(x \wedge y) \leq \eta(x) \wedge \eta(y)$. By (ii), $\eta^{-1}([\eta(x) \wedge \eta(y), \eta(x) \vee \eta(y)])$ is an interval I in L. But $\eta(x)$ is in $[\eta(x) \wedge \eta(y), \eta(x) \vee \eta(y)]$, so $x \in I$, and similarly $y \in I$. Since I is an interval, also $x \wedge y$ is in I, and in particular $\eta(x \wedge y) \geq \eta(x) \wedge \eta(y)$. We have shown that $\eta(x \wedge y) = \eta(x) \wedge \eta(y)$, and the dual argument shows that $\eta(x \vee y) = \eta(x) \vee \eta(y)$. ☐

When η is a bijection, Proposition 9-5.1 reduces, with the help of Exercise 9.21(c), to the statement that a bijection between finite lattices is an isomorphism of lattices if and only if it is an isomorphism of posets. This is easy to prove for general lattices. (See, for example, LTF Lemma 4). Furthermore, if a lattice L and a poset P are isomorphic as posets then P is a lattice and L and P are isomorphic as lattices.

We now turn to the order-theoretic characterization of lattice congruences. A *congruence* on a lattice L is an equivalence relation $\boldsymbol{\alpha}$ on L such that if $x_1 \equiv x_2 \pmod{\boldsymbol{\alpha}}$ and $y_1 \equiv y_2 \pmod{\boldsymbol{\alpha}}$ then $x_1 \wedge y_1 \equiv x_2 \wedge y_2 \pmod{\boldsymbol{\alpha}}$ and

$x_1 \vee y_1 \equiv x_2 \vee y_2 \pmod{\boldsymbol{\alpha}}$. To check that a given $\boldsymbol{\alpha}$ is a congruence, it is enough to check that if $x_1 \equiv x_2 \pmod{\boldsymbol{\alpha}}$ then $x_1 \wedge y \equiv x_2 \wedge y \pmod{\boldsymbol{\alpha}}$ and $x_1 \vee y \equiv x_2 \vee y \pmod{\boldsymbol{\alpha}}$ for all y. (See Exercise 9.35.)

Proposition 9-5.2. *An equivalence relation $\boldsymbol{\alpha}$ on a finite lattice L is a lattice congruence if and only if the following three conditions hold:*

(i) *Each equivalence class is an interval in L.*

(ii) *The map $\pi_\downarrow^{\boldsymbol{\alpha}}$ mapping each element to the bottom element of its equivalence class is order-preserving.*

(iii) *The map $\pi^\uparrow_{\boldsymbol{\alpha}}$ mapping each element to the top element of its equivalence class is order-preserving.*

Proof. First, suppose $\boldsymbol{\alpha}$ is a congruence on L and let C be an $\boldsymbol{\alpha}$-class. Since C is finite, it has at least one minimal element, but if x_1 and x_2 are both minimal in C, then since $x_1 \equiv x_2 \pmod{\boldsymbol{\alpha}}$, we have $x_1 \wedge x_2 \equiv x_2 \wedge x_2 = x_2 \pmod{\boldsymbol{\alpha}}$, so that $x_1 \geq x_2$ and therefore $x_1 = x_2$ by the minimality of x_1. Thus C has a unique minimal element x and by the dual argument, it has a unique maximal element y. If $x \leq z \leq y$, then since $x \equiv y \pmod{\boldsymbol{\alpha}}$, we have $x = x \wedge z \equiv y \wedge z = z \pmod{\boldsymbol{\alpha}}$. Thus C is the entire interval $[x, y]$.

If $x \leq y$ then $x \wedge y = x$, and since $x \equiv \pi_\downarrow^{\boldsymbol{\alpha}} x \pmod{\boldsymbol{\alpha}}$ and $y \equiv \pi_\downarrow^{\boldsymbol{\alpha}} y \pmod{\boldsymbol{\alpha}}$, we have $x = x \wedge y \equiv \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y \pmod{\boldsymbol{\alpha}}$. Thus $\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y$ is in the $\boldsymbol{\alpha}$-class of x, so $\pi_\downarrow^{\boldsymbol{\alpha}} x \leq \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y \leq \pi_\downarrow^{\boldsymbol{\alpha}} y$. This is (ii), and the proof of (iii) is dual.

Conversely, suppose that (i) and (ii) hold for some equivalence relation $\boldsymbol{\alpha}$. For any $x, y \in L$, since $x \geq x \wedge y$ and $y \geq x \wedge y$, condition (ii) implies that $\pi_\downarrow^{\boldsymbol{\alpha}} x \geq \pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y)$ and $\pi_\downarrow^{\boldsymbol{\alpha}} y \geq \pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y)$. Thus

$$\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y \geq \pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y). \tag{9-5.1}$$

On the other hand, $x \geq \pi_\downarrow^{\boldsymbol{\alpha}} x$ and $y \geq \pi_\downarrow^{\boldsymbol{\alpha}} y$, so

$$x \wedge y \geq \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y. \tag{9-5.2}$$

Applying (ii) to (9-5.2) and combining it with (9-5.1), we obtain

$$\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y \geq \pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y) \geq \pi_\downarrow^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y). \tag{9-5.3}$$

But also $\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y$ and $\pi_\downarrow^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y)$ are in the same $\boldsymbol{\alpha}$-class, so (i) implies that $\pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y)$ is also in that $\boldsymbol{\alpha}$-class. Therefore also $x \wedge y$ is in that $\boldsymbol{\alpha}$-class, or in other words $x \wedge y \equiv \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y \pmod{\boldsymbol{\alpha}}$.

Now given $x_1 \equiv y_1 \pmod{\boldsymbol{\alpha}}$ and $x_2 \equiv y_2 \pmod{\boldsymbol{\alpha}}$, we have $\pi_\downarrow^{\boldsymbol{\alpha}} x_1 = \pi_\downarrow^{\boldsymbol{\alpha}} y_1$ and $\pi_\downarrow^{\boldsymbol{\alpha}} x_2 = \pi_\downarrow^{\boldsymbol{\alpha}} y_2$. Thus

$$x_1 \wedge x_2 \equiv \pi_\downarrow^{\boldsymbol{\alpha}} x_1 \wedge \pi_\downarrow^{\boldsymbol{\alpha}} x_2 = \pi_\downarrow^{\boldsymbol{\alpha}} y_1 \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y_2 \equiv y_1 \wedge y_2 \pmod{\boldsymbol{\alpha}}.$$

The proof is completed by arguing dually based on (i) and (iii). $\square$

Remark 9-5.3. The proof of Proposition 9-5.2 also shows that, in a lattice of arbitrary cardinality, an equivalence relation satisfying (i), (ii), and (iii) is a congruence. In addition, the proof shows that a congruence satisfying (i) also satisfies (ii) and (iii). Exercise 9.36 asks for an example of a congruence not satisfying (i). Congruences that satisfy (i) are called *bounded congruences*.

9-5.2 Quotient lattices

We now prove two order-theoretic characterizations of the quotient of a finite lattice modulo a congruence. The first and more obvious characterization describes the quotient as a partial order on congruence classes. See also Exercise 9.46.

Proposition 9-5.4. *If L is a lattice and $\boldsymbol{\alpha}$ is a congruence on L then the quotient lattice $L/\boldsymbol{\alpha}$ is the partially ordered set on $\boldsymbol{\alpha}$-classes with order relation described as follows: Two $\boldsymbol{\alpha}$-classes C_1 and C_2 have $C_1 \le C_2$ if and only if there exist $x \in C_1$ and $y \in C_2$ with $x \le y$. If L is finite, then distinct $\boldsymbol{\alpha}$-classes have $C_1 \prec C_2$ if and only if there exist $x \in C_1$ and $y \in C_2$ with $x \prec y$.*

Proof. We have $C_1 \le C_2$ if and only if $C_1 \wedge C_2 = C_1$. The latter condition is equivalent to the requirement that $x \wedge y \in C_1$ for any $x \in C_1$ and $y \in C_2$. Thus if $C_1 \le C_2$ and $x \in C_1$ and $y \in C_2$, then the elements $x \wedge y \in C_1$ and $y \in C_2$ satisfy $x \wedge y \le y$. Conversely, given $x \in C_1$ and $y \in C_2$ with $x \le y$, we have $x \wedge y = x \in C_1$.

Now suppose L is finite. If $C_1 \prec C_2$ then in particular, there exist $x \in C_1$ and $y \in C_2$ with $x \le y$. Any maximal chain from x to y is $x = x_0 \prec \cdots \prec x_k \prec y_0 \prec \cdots \prec y_\ell = y$ with each x_i in C_1 and each y_i in C_2. (If some other class appears in the chain, then we obtain a contradiction to $C_1 \prec C_2$.) Conversely, if $C_1 \ne C_2$ and there exist $x \in C_1$ and $y \in C_2$ with $x \prec y$, then in particular $C_1 < C_2$. If $C_1 \prec C_3 \le C_2$, then in particular there exist elements $x' \in C_1$ and z in C_3 with $x' < z$ and elements $z' \in C_3$ and $y' \in C_2$ with $z' \le y'$. Applying the order-preserving map $\pi_\downarrow^{\boldsymbol{\alpha}}$ to $x' < z$ and to $z' \le y'$, we obtain $\pi_\downarrow^{\boldsymbol{\alpha}} x \le \pi_\downarrow^{\boldsymbol{\alpha}} z \le \pi_\downarrow^{\boldsymbol{\alpha}} y$, because the map $\pi_\downarrow^{\boldsymbol{\alpha}}$ is constant on $\boldsymbol{\alpha}$-classes. Now $x = x \vee \pi_\downarrow^{\boldsymbol{\alpha}} x \le x \vee \pi_\downarrow^{\boldsymbol{\alpha}} z \le x \vee \pi_\downarrow^{\boldsymbol{\alpha}} y = y$ (with the latter equality holding because $x \prec y$). But $x \vee \pi_\downarrow^{\boldsymbol{\alpha}} x$ and $x \vee \pi_\downarrow^{\boldsymbol{\alpha}} z$ are not equal, because one is in C_1 and the other is in C_3. Because $x \prec y$, we conclude that $x \vee \pi_\downarrow^{\boldsymbol{\alpha}} z = x \vee \pi_\downarrow^{\boldsymbol{\alpha}} y$ and thus that $C_3 = C_2$. □

The second order-theoretic characterization of quotients uses Propositions 9-5.2 and 9-5.4 to give a more direct description of the quotient lattice as a partial order. Given a congruence $\boldsymbol{\alpha}$ on a finite lattice L, the notation $\pi_\downarrow^{\boldsymbol{\alpha}} L$ stands for $\{\pi_\downarrow^{\boldsymbol{\alpha}} x \mid x \in L\}$, the set of elements that are minimal in their equivalence class. The set $\pi_\downarrow^{\boldsymbol{\alpha}} L$ is partially ordered as an induced subposet of L.

Proposition 9-5.5. *If L is a finite lattice and $\boldsymbol{\alpha}$ is a congruence on L then $\pi_\downarrow^{\boldsymbol{\alpha}} L$ is a lattice, isomorphic to the quotient lattice $L/\boldsymbol{\alpha}$. The map $\pi_\downarrow^{\boldsymbol{\alpha}}$ is a lattice homomorphism from L to $\pi_\downarrow^{\boldsymbol{\alpha}} L$.*

Proof. We write $\overline{\pi}_\downarrow^{\boldsymbol{\alpha}} : (L/\boldsymbol{\alpha}) \to \pi_\downarrow^{\boldsymbol{\alpha}} L$ for the bijection taking $\boldsymbol{\alpha}$-classes to their bottom elements. Suppose C_1 and C_2 are $\boldsymbol{\alpha}$-classes. If $C_1 \leq C_2$ then by Proposition 9-5.4, there exist $x \in C_1$ and $y \in C_2$ with $x \leq y$. By Proposition 9-5.2, $\pi_\downarrow^{\boldsymbol{\alpha}} x \leq \pi_\downarrow^{\boldsymbol{\alpha}} y$, or in other words $\overline{\pi}_\downarrow^{\boldsymbol{\alpha}}(C_1) \leq \overline{\pi}_\downarrow^{\boldsymbol{\alpha}}(C_2)$. Conversely, if $\overline{\pi}_\downarrow(C_1) \leq \overline{\pi}_\downarrow(C_2)$ then $C_1 \leq C_2$ by Proposition 9-5.4. We have shown that $\overline{\pi}_\downarrow^{\boldsymbol{\alpha}}$ is an isomorphism.

The map $\pi_\downarrow^{\boldsymbol{\alpha}}$ from L to $\pi_\downarrow^{\boldsymbol{\alpha}} L$ coincides with the composition $\overline{\pi}_\downarrow^{\boldsymbol{\alpha}} \circ \psi$, where ψ is the natural homomorphism from L to $L/\boldsymbol{\alpha}$. (See, for example, LTF Theorem 16.) In particular $\pi_\downarrow^{\boldsymbol{\alpha}}$ is a lattice homomorphism to $\pi_\downarrow^{\boldsymbol{\alpha}} L$. □

Of course, the dual statement to Proposition 9-5.5, replacing $\pi_\downarrow^{\boldsymbol{\alpha}}$ by $\pi^\uparrow_{\boldsymbol{\alpha}}$, holds by the dual proof.

Remark 9-5.6. Proposition 9-5.2 may allow a combinatorialist to recognize situations where lattice-theoretic tools are applicable. Suppose L is a finite lattice, S is a set, and $\eta\colon L \to S$ is a surjective map such that the fiber $\eta^{-1}(x) = \{y \in L \mid \eta(y) = x\}$ of any $x \in S$ is an interval in L. In this case, one should check whether the fibers are the congruence classes of a lattice congruence on L, specifically by checking that the projection to bottom elements of fibers is order-preserving and that the projection to top elements of fibers is order-preserving. If the fibers define a congruence $\boldsymbol{\alpha}$, then the natural bijection from η-fibers to S allows the lattice structure on $L/\boldsymbol{\alpha}$ to be carried to S. That is, S admits a lattice structure isomorphic to $L/\boldsymbol{\alpha}$ and also, by Proposition 9-5.5, isomorphic to $\pi_\downarrow^{\boldsymbol{\alpha}} L$, the subposet of L induced by the elements that are minimal in their η-fibers. Furthermore, Proposition 9-5.5 implies that η is a lattice homomorphism from L to S. This kind of investigation led to the notion of Cambrian lattices discussed in Chapter 10; more details are found in [374].

Propositions 9-5.2 and 9-5.4 also lead to the following useful lemma.

Lemma 9-5.7. *Given a finite lattice L, a congruence $\boldsymbol{\alpha}$ on L and an interval $[x, y]$ in L, the intervals $[\pi_\downarrow^{\boldsymbol{\alpha}} x, \pi_\downarrow^{\boldsymbol{\alpha}} y]$ in $\pi_\downarrow^{\boldsymbol{\alpha}} L$ and $[x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha}]$ in $L/\boldsymbol{\alpha}$ are both isomorphic to the quotient of $[x, y]$ modulo the restriction of $\boldsymbol{\alpha}$ to $[x, y]$.*

Proof. The map $\overline{\pi}_\downarrow^{\boldsymbol{\alpha}}$ from the proof of Proposition 9-5.5 restricts to an isomorphism between the two intervals. Consider the restriction to $[x, y]$ of the natural homomorphism from L to $L/\boldsymbol{\alpha}$. Proposition 9-5.4 implies that the image of this restriction is contained in $[x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha}]$. We claim that the restriction is also surjective onto $[x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha}]$. That is, given $z/\boldsymbol{\alpha} \in [x/\boldsymbol{\alpha}), y/\boldsymbol{\alpha}]$, we claim that $z/\boldsymbol{\alpha} \cap [x, y] \neq \varnothing$. Proposition 9-5.4 says that there exist $x' \in x/\boldsymbol{\alpha}$ and $z' \in z/\boldsymbol{\alpha}$ with $x' \leq z'$ and that there exist $z'' \in z/\boldsymbol{\alpha}$ and $y'' \in y/\boldsymbol{\alpha}$ with

$z'' \le y''$. Proposition 9-5.2 implies that $\pi_\downarrow^{\boldsymbol{\alpha}} z = \pi_\downarrow^{\boldsymbol{\alpha}} z'' \le \pi_\downarrow^{\boldsymbol{\alpha}} y'' \le y$. Since also y is above x, we have $(\pi_\downarrow^{\boldsymbol{\alpha}} z) \vee x \in [x, y]$. Because $x' \le z'$, we have $x \le \pi^\uparrow_{\boldsymbol{\alpha}} x' \le \pi^\uparrow_{\boldsymbol{\alpha}} z' = \pi^\uparrow_{\boldsymbol{\alpha}} z$. Thus $\pi_\downarrow^{\boldsymbol{\alpha}} z \le (\pi_\downarrow^{\boldsymbol{\alpha}} z) \vee x \le \pi^\uparrow_{\boldsymbol{\alpha}} z$, and we conclude that $(\pi_\downarrow^{\boldsymbol{\alpha}} z) \vee x \in [\pi_\downarrow^{\boldsymbol{\alpha}} z, \pi^\uparrow_{\boldsymbol{\alpha}} z] = z/\boldsymbol{\alpha}$. We have proved the claim. Two elements of $[x, y]$ map to the same element of $[x/\boldsymbol{\alpha}), y/\boldsymbol{\alpha}]$ if and only if they are congruent mod $\boldsymbol{\alpha}$. Thus we have proved the lemma. $\square$

The following proposition is proved as Exercise 9.39.

Proposition 9-5.8. *If L is a finite lattice and $\boldsymbol{\alpha}$ is a congruence on L, then $\pi_\downarrow^{\boldsymbol{\alpha}} L$ is a join-subsemilattice of L, but can fail to be a sublattice of L.*

Remark 9-5.9. The last assertion of Proposition 9-5.5 (that $\pi_\downarrow^{\boldsymbol{\alpha}}$ is a lattice homomorphism from L to $\pi_\downarrow^{\boldsymbol{\alpha}} L$) may lead the unwary to write down a statement like "$\pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge y) = \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge \pi_\downarrow^{\boldsymbol{\alpha}} y$." In general, that statement is, at best, ambiguous and, at worst, incorrect. A correct statement is $\pi_\downarrow^{\boldsymbol{\alpha}}(x \wedge_L y) = \pi_\downarrow^{\boldsymbol{\alpha}} x \wedge_{\pi_\downarrow^{\boldsymbol{\alpha}} L} \pi_\downarrow^{\boldsymbol{\alpha}} y$, where we use subscripts to distinguish the meet in L from the meet in $\pi_\downarrow^{\boldsymbol{\alpha}} L$. Since $\pi_\downarrow^{\boldsymbol{\alpha}} L$ is a join-subsemilattice of L by Proposition 9-5.8, the statement $\pi_\downarrow^{\boldsymbol{\alpha}}(x \vee y) = \pi_\downarrow^{\boldsymbol{\alpha}} x \vee \pi_\downarrow^{\boldsymbol{\alpha}} y$ is correct and unambiguous. When $\pi_\downarrow^{\boldsymbol{\alpha}} L$ is a sublattice of L, the offending statement above is also correct.

It is straightforward to describe the cover relations in $\pi_\downarrow^{\boldsymbol{\alpha}} L$.

Proposition 9-5.10. *Suppose L is a finite lattice and $\boldsymbol{\alpha}$ is a congruence on L. For each $y \in \pi_\downarrow^{\boldsymbol{\alpha}} L$, the map $\pi_\downarrow^{\boldsymbol{\alpha}}$ restricts to a bijection between elements of L covered by y in L and elements of $\pi_\downarrow^{\boldsymbol{\alpha}} L$ covered by y in $\pi_\downarrow^{\boldsymbol{\alpha}} L$.*

The proposition amounts to two assertions: First, if $y \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ and x is covered by y in L, then $\pi_\downarrow^{\boldsymbol{\alpha}} x$ is covered by y in $\pi_\downarrow^{\boldsymbol{\alpha}} L$. Second, if x is covered by y in $\pi_\downarrow^{\boldsymbol{\alpha}} L$, then there exists a unique element x' covered by y in L such that $\pi_\downarrow^{\boldsymbol{\alpha}} x' = x$.

Proof. We use Proposition 9-5.2 throughout the proof. Suppose $y \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ and x is covered by y in L. Then $\pi_\downarrow^{\boldsymbol{\alpha}} x \le x \prec y$. Suppose there is some element $z \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ such that $\pi_\downarrow^{\boldsymbol{\alpha}} x < z < y$. If $z \le x$, so that $\pi_\downarrow^{\boldsymbol{\alpha}} x < z \le x$, then $\pi_\downarrow^{\boldsymbol{\alpha}} z = \pi_\downarrow^{\boldsymbol{\alpha}} x$, and since $\pi_\downarrow^{\boldsymbol{\alpha}} z = z$, this is a contradiction. If $z \not\le x$, then since $x \prec y$ and y is an upper bound for x and z, we see that $x \vee_L z = y$. Propositions 9-5.5 and 9-5.8 imply that $\pi_\downarrow^{\boldsymbol{\alpha}} x \vee \pi_\downarrow^{\boldsymbol{\alpha}} z = \pi_\downarrow^{\boldsymbol{\alpha}} y$. But $\pi_\downarrow^{\boldsymbol{\alpha}} z = z$ and $\pi_\downarrow^{\boldsymbol{\alpha}} y = y$, so $\pi_\downarrow^{\boldsymbol{\alpha}} x \vee z = y$, contradicting the supposition that $\pi_\downarrow^{\boldsymbol{\alpha}} x < z < y$.

Now suppose $x, y \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ and x is covered by y in $\pi_\downarrow^{\boldsymbol{\alpha}} L$. In particular $x < y$ in L. Let x' be an element of L with $x \le x' \prec y$ in L. Then $x = \pi_\downarrow^{\boldsymbol{\alpha}} x \le \pi_\downarrow^{\boldsymbol{\alpha}} x' < \pi_\downarrow^{\boldsymbol{\alpha}} y = y$. Since x is covered by y in $\pi_\downarrow^{\boldsymbol{\alpha}} L$, we see that $\pi_\downarrow^{\boldsymbol{\alpha}} x' = x$. If there are two distinct elements x' and x'' covered by y in L with $\pi_\downarrow^{\boldsymbol{\alpha}} x' = x$ and $\pi_\downarrow^{\boldsymbol{\alpha}} x'' = x$, then $x' \vee x'' = y$. Applying $\pi_\downarrow^{\boldsymbol{\alpha}}$ to $x' \vee x'' = y$ and appealing again to Propositions 9-5.5 and 9-5.8, we obtain the contradiction $x = y$. $\square$

9-5.3 Join-irreducible elements and congruences

An element j of a finite lattice L is *join-irreducible* if it cannot be written as $j = \bigvee U$ for some $U \subset L$ with $j \notin U$. Equivalently, j is join-irreducible if it covers exactly one element. We write j_* for the unique element covered by j. A congruence $\boldsymbol{\alpha}$ *contracts* a join-irreducible element j if $j \equiv j_* \pmod{\boldsymbol{\alpha}}$.

Proposition 9-5.11. *Let $\boldsymbol{\alpha}$ be a congruence on a finite lattice L. An element $j \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ is join-irreducible as an element of $\pi_\downarrow^{\boldsymbol{\alpha}} L$ if and only if it is join-irreducible as an element of L. The join-irreducible elements of $\pi_\downarrow^{\boldsymbol{\alpha}} L$ are exactly the join-irreducible elements of L that are not contracted by $\boldsymbol{\alpha}$.*

Proof. The first statement is an immediate consequence of Proposition 9-5.10. The second statement follows because a join-irreducible element j of L is in $\pi_\downarrow^{\boldsymbol{\alpha}} L$ if and only if it is not contracted by $\boldsymbol{\alpha}$. $\square$

The notion of contracting join-irreducible elements leads to another characterization of the quotient of a finite lattice L modulo a congruence $\boldsymbol{\alpha}$.

Proposition 9-5.12. *Let $\boldsymbol{\alpha}$ be a congruence on a finite lattice L. For $x \in L$, let $D^{\boldsymbol{\alpha}} x$ be the set of join-irreducible elements $j \leq x$ that are not contracted by $\boldsymbol{\alpha}$. Then $x, y \in L$ have $x \equiv y \pmod{\boldsymbol{\alpha}}$ if and only if $D^{\boldsymbol{\alpha}}(x) = D^{\boldsymbol{\alpha}}(y)$. The quotient $L/\boldsymbol{\alpha}$ is isomorphic to the set $\{D^{\boldsymbol{\alpha}} x \mid x \in L\}$ partially ordered by containment.*

Proof. Suppose $x \in L$. Since $\pi_\downarrow^{\boldsymbol{\alpha}} x \leq x$, it is immediate that $D^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} x) \subseteq D^{\boldsymbol{\alpha}} x$. If, on the other hand, $j \in D^{\boldsymbol{\alpha}} x$, then $j = \pi_\downarrow^{\boldsymbol{\alpha}} j \leq \pi_\downarrow^{\boldsymbol{\alpha}} x$, so $j \in D^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} x)$. Thus $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} x)$ for all $x \in L$, and we conclude that $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}} y$ whenever $x \equiv y \pmod{\boldsymbol{\alpha}}$.

To prove the converse, first consider the case where $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}} y$ for elements $x, y \in L$ with $x \prec y$. Choose some element j which is minimal among elements below y but not below x. If j covers two distinct elements y and y', then both are below x by our choice of j. But then since x is an upper bound for both y and y', the join $y \vee y' = j$ is below x, and this is a contradiction. Thus j is join-irreducible and covers a unique element j_*. Since $x \prec y$, we have $j \vee x = y$. Since $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}} y$, we see that $j_* \equiv j \pmod{\boldsymbol{\alpha}}$ and therefore $j_* \vee x \equiv j \vee x \pmod{\boldsymbol{\alpha}}$, or in other words $x \equiv y \pmod{\boldsymbol{\alpha}}$. Next, if $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}} y$ for $x \leq y$, then repeating the above argument along a maximal chain from x to y, we conclude that $x \equiv y \pmod{\boldsymbol{\alpha}}$. Finally, for general x and y, the set $D^{\boldsymbol{\alpha}}(x \wedge y)$ is the set of uncontracted join-irreducible elements below x and below y, so $D^{\boldsymbol{\alpha}}(x \wedge y) = D^{\boldsymbol{\alpha}} x \cap D^{\boldsymbol{\alpha}} y$. Thus if $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}} y$, then $D^{\boldsymbol{\alpha}} x = D^{\boldsymbol{\alpha}}(x \wedge y) = D^{\boldsymbol{\alpha}} y$. Since $x \wedge y \leq x$, we see that $x \equiv x \wedge y \pmod{\boldsymbol{\alpha}}$ and similarly $y \equiv x \wedge y \pmod{\boldsymbol{\alpha}}$, so $x \equiv y \pmod{\boldsymbol{\alpha}}$.

By Proposition 9-5.5, to prove the second statement, it is enough to show that for $x, y \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ we have $x \leq y$ if and only if $D^{\boldsymbol{\alpha}} x \subseteq D^{\boldsymbol{\alpha}} y$. The "only if" direction is immediate. Conversely, suppose $D^{\boldsymbol{\alpha}} x \subseteq D^{\boldsymbol{\alpha}} y$. As before,

$D^{\boldsymbol{\alpha}}(x \wedge y) = D^{\boldsymbol{\alpha}}x \cap D^{\boldsymbol{\alpha}}y$, which equals $D^{\boldsymbol{\alpha}}x$ because $D^{\boldsymbol{\alpha}}x \subseteq D^{\boldsymbol{\alpha}}y$. Thus $x \wedge y \equiv x \pmod{\boldsymbol{\alpha}}$, but since $x \in \pi_{\downarrow}^{\boldsymbol{\alpha}} L$, we conclude that $x \wedge y = x$, so that $x \leq y$. □

The *congruence lattice* $\mathrm{Con}\, L$ is the set of all congruences of L, partially ordered by containment of equivalence relations, or equivalently as a subposet of the lattice of partitions of L.

The *partition lattice* of a set S is the set of equivalence relations, partially ordered by containment. The intersection of equivalence relations is an equivalence relation, and thus intersection of relations is the meet in the partition lattice. Since there is a unique maximal equivalence relation (with exactly one class), the dual to Lemma 9-2.1 implies that the partitions of L form a lattice. Alternately, one can easily construct the join directly: It is the operation of taking the union and then transitive closure of relations.

The following proposition is a combination of part of LTF Theorem 12 and LTF Exercise I.3.60. See also Corollary 9-5.17.

Proposition 9-5.13. *If L is a lattice, then* $\mathrm{Con}\, L$ *is a lattice, specifically a sublattice of the partition lattice of L. The meet in* $\mathrm{Con}\, L$ *is the intersection of relations, and the join is the transitive closure of the union of relations.*

Proof. It is easy to see that the intersection of two congruence relations is a congruence relation. It is only slightly harder to see that $\mathrm{Con}\, L$ is also closed under taking union and then transitive closure. Given two congruences $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathrm{Con}\, L$, write $\boldsymbol{\gamma}$ for the transitive closure of their union. Then $x \equiv y \pmod{\boldsymbol{\gamma}}$ if and only if there is a sequence $x = x_0, \ldots, x_k = y$ such that for each $i = 1, \ldots, k$, either $x_{i-1} \equiv x_i \pmod{\boldsymbol{\alpha}}$ or $x_{i-1} \equiv x_i \pmod{\boldsymbol{\beta}}$. For any $z \in L$, since $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are congruences, either $x_{i-1} \wedge z \equiv x_i \wedge z \pmod{\boldsymbol{\alpha}}$ or $x_{i-1} \wedge z \equiv x_i \wedge z \pmod{\boldsymbol{\beta}}$ for all $i = 1, \ldots, k$, and we conclude that $x \wedge z \equiv y \wedge z \pmod{\boldsymbol{\gamma}}$. Similarly, $x \vee z \equiv y \vee z \pmod{\boldsymbol{\gamma}}$, and we conclude by Exercise 9.35 that $\boldsymbol{\gamma}$ is a congruence. □

By Proposition 9-5.13, for any elements a and b of a lattice L, there is a unique smallest congruence relation on L with $a \equiv b$ (the meet of all congruences with $a \equiv b$). We write $\mathrm{con}(a, b)$ for this congruence. If j is a join-irreducible element, then we write $\mathrm{con}(j)$ for $\mathrm{con}(j_*, j)$. This is the unique smallest congruence contracting j. We now characterize join-irreducible congruences on finite lattices.

Proposition 9-5.14. *If L is a finite lattice and $\boldsymbol{\alpha} \in \mathrm{Con}\, L$, then the following are equivalent.*

(i) $\boldsymbol{\alpha}$ *is join-irreducible in* $\mathrm{Con}\, L$.

(ii) $\boldsymbol{\alpha} = \mathrm{con}(a, b)$ *for some covering pair* $a \lessdot b$.

(iii) $\boldsymbol{\alpha} = \mathrm{con}(j)$ *for some join-irreducible element* j.

Proof. As an immediate consequence of Proposition 9-5.12, $\boldsymbol{\alpha}$ is the join $\bigvee \mathrm{con}(j)$ over all join-irreducible elements j contracted by $\boldsymbol{\alpha}$. Thus if $\boldsymbol{\alpha}$ is join-irreducible in $\mathrm{Con}\, L$, then it is of the form $\mathrm{con}(j)$. That is, (i) implies (iii). Keeping in mind that $\mathrm{con}(j)$ means $\mathrm{con}(j_*, j)$, we see that (iii) implies (ii).

Suppose $a \prec b$ and suppose $\boldsymbol{\beta} \vee \boldsymbol{\gamma} = \mathrm{con}(a, b)$ in $\mathrm{Con}\, L$. By Proposition 9-5.13, there is a sequence of elements $a = x_0, \ldots, x_k = b$ such that for each $i = 1, \ldots, k$, either $x_{i-1} \equiv x_i \pmod{\boldsymbol{\beta}}$ or $x_{i-1} \equiv x_i \pmod{\boldsymbol{\gamma}}$. Defining $y_i = (x_i \vee a) \wedge b$ for each i, we obtain a sequence $a = y_0, \ldots, y_k = b$ such that $y_{i-1} \equiv y_i$ (mod $\boldsymbol{\beta}$ or $\boldsymbol{\gamma}$) for each i. Furthermore, each $x_i \vee a$ is above a, and therefore each y_i is between a and b. But $a \prec b$ so each y_i is either a or b, and we conclude that either $a \equiv b \pmod{\boldsymbol{\beta}}$ or $a \equiv b \pmod{\boldsymbol{\gamma}}$. Thus either $\boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ equals $\mathrm{con}(a, b)$. We see that $\mathrm{con}(a, b)$ is join-irreducible in $\mathrm{Con}\, L$. That is, (ii) implies (i). ☐

When L is infinite, each $\mathrm{con}(a, b)$ for $a \prec b$ is join-irreducible, but there may be join-irreducible congruences not of this form. See LTF Section III.1.4.

9-5.4 Forcing among edges and join-irreducible elements

We call a cover relation in a finite lattice an *edge*, because it is an edge in the Hasse diagram. A congruence $\boldsymbol{\alpha}$ *contracts* an edge $a \prec b$ if $a \equiv b \pmod{\boldsymbol{\alpha}}$. This use of the term "contract" is not at odds with our earlier use: a congruence $\boldsymbol{\alpha}$ contracts a join-irreducible element j if and only if it contracts the edge $j_* \prec j$. The following fact is an immediate consequence of Proposition 9-5.14 (and indeed was stated in part in the proof of Proposition 9-5.14).

Corollary 9-5.15. *A congruence $\boldsymbol{\alpha}$ on a finite lattice L is determined by the set of join-irreducible elements it contracts or by the set of edges it contracts. Specifically, $\boldsymbol{\alpha}$ is the join $\bigvee \mathrm{con}(j)$ over all join-irreducible elements j contracted by $\boldsymbol{\alpha}$. Also, $\boldsymbol{\alpha}$ is the join $\bigvee \mathrm{con}(a, b)$ over all edges $a \prec b$ it contracts.*

We say that an edge $a \prec b$ *forces* an edge $c \prec d$ if $\mathrm{con}(a, b) \geq \mathrm{con}(c, d)$. Equivalently, $a \prec b$ forces $c \prec d$ if every congruence contracting $a \prec b$ also contracts $c \prec d$, or in other words, if $c \equiv d \pmod{\mathrm{con}(a, b)}$. One approach to understanding congruences on a finite lattice is to determine the forcing relation on edges. We will see in Section 9-6 that the forcing relation on $\mathrm{Pos}(\mathcal{A}, B)$ is given by simple local rules when $\mathcal{A}$ is tight with respect to B.

For join-irreducible j, we continue to write $\mathrm{con}(j)$ for $\mathrm{con}(j_*, j)$, the smallest congruence contracting j. Given join-irreducible elements j and j' of L, we say j *forces* j' and write $j \to j'$ if $\mathrm{con}(j) \geq \mathrm{con}(j')$ in $\mathrm{Con}\, L$. The reflexive-transitive closure of the forcing relation is a pre-order on the join-irreducible elements of L, taking the convention that $\to$ corresponds to $\geq$. Setting two join-irreducible elements j and j' to be equivalent if and only if $\mathrm{con}(j) = \mathrm{con}(j')$, the forcing relation defines a partial order on equivalence classes of join-irreducible elements by the usual construction of a partial order from a pre-order. (See

LTF Section I.1.2.) We write $[j]$ for the equivalence class of a join-irreducible element. Call this partial order the *forcing order* on (equivalence classes of) join-irreducible elements. The notation $\mathrm{Con}_{\mathrm{Ji}}\, L$ stands for the subposet of $\mathrm{Con}\, L$ induced by join-irreducible congruences. In light of Proposition 9-5.14, the following proposition is simply a rephrasing of the definition of the forcing order.

Proposition 9-5.16. *The map* $[j] \mapsto \mathrm{con}(j)$ *is an isomorphism from the forcing order to* $\mathrm{Con}_{\mathrm{Ji}}\, L$. *Thus the congruence lattice* $\mathrm{Con}\, L$ *is isomorphic to the poset of down-sets in the forcing order.*

The poset of downsets in an arbitrary poset is distributive.[5] Thus Proposition 9-5.16 implies the following corollary. The same result holds for arbitrary lattices. (See LTF Theorem 149.)

Corollary 9-5.17. *The congruence lattice* $\mathrm{Con}\, L$ *of a finite lattice* L *is distributive.*

For emphasis, we also mention the following immediate corollary to Proposition 9-5.16.

Corollary 9-5.18. *Suppose* L *is a finite lattice and* $\boldsymbol{\alpha}$ *and* $\boldsymbol{\beta}$ *are congruences on* L *such that* J *is the set of join-irreducible elements contracted by* $\boldsymbol{\alpha}$ *and* K *is the set of join-irreducible elements contracted by* $\boldsymbol{\beta}$. *Then* $\boldsymbol{\alpha} \wedge \boldsymbol{\beta}$ *contracts exactly the join-irreducible elements* $J \cap K$ *and* $\boldsymbol{\alpha} \vee \boldsymbol{\beta}$ *contracts exactly the join-irreducible elements* $J \cup K$.

When the forcing relation is acyclic, the forcing order is a partial order on join-irreducible elements of L, rather than on equivalence classes of join-irreducible elements. In this case, the map $j \mapsto \mathrm{con}(j)$ is a bijection from join-irreducible elements of L to join-irreducible elements of $\mathrm{Con}\, L$. Lattices with this property and the dual property[6] are called *congruence uniform* by Day [116]. (Unfortunately, this terminology conflicts with the terminology of LTF Section IV.4.5.) This very natural combinatorial condition turns out to coincide with another very natural combinatorial condition and with a less combinatorial condition of great lattice-theoretic interest. *Doubling* an interval in a lattice means replacing the interval by its product with a 2-element chain and defining order relations between the doubled interval and the rest of the lattice in a natural way as explained in Section 3-2.7. A lattice that is the quotient of a finitely generated free lattice modulo a bounded congruence in the sense of Remark 9-5.3 is often called a *bounded lattice*. However, since

[5]Indeed, a finite lattice is distributive if and only if it is the poset of downsets in some poset. (See LTF Theorem 107 and Corollary 108.) This result is sometimes called the Fundamental Theorem of Finite Distributive Lattices [421, Section 3.4].

[6]The dual property is: The map $m \mapsto \mathrm{con}(m, m^*)$ is a bijection from meet-irreducible elements of L to join-irreducible elements of $\mathrm{Con}\, L$. Here m^* is the unique element covering m.

this definition conflicts with a very common usage of the term "bounded" to refer to a poset with a unique minimal and a unique maximal element, we will avoid using the term "bounded lattice" here. In the following theorem, the equivalence of (i) and (iii) is obtained by combining Lemma 3-2.33 and Corollary 3-2.35 and the equivalence of (ii) and (iii) is Theorem 3-2.40. The assertion about semidistributivity is Lemma 3-2.18. We mark Theorem 9-5.19 with a diamond here because Theorem 3-2.40 is also a diamond theorem.

◇ **Theorem 9-5.19.** *For a finite lattice L, the following are equivalent:*

(i) *L is congruence uniform (in the sense of Day).*

(ii) *L is obtained from a one-element lattice by a sequence of doublings of intervals.*

(iii) *L is the quotient of a finitely generated free lattice modulo a bounded congruence.*

If these conditions hold, then L is semidistributive.

In a finite congruence uniform lattice, for each edge $a \prec b$, there is a unique join-irreducible element j such that $\mathrm{con}(a, b) = \mathrm{con}(j_*, j)$. The map sending $a \prec b$ to this j and the dual map to meet-irreducible elements are made explicit in the following proposition, which is proved as Exercise 9.48.

Proposition 9-5.20. *Let L be a finite congruence uniform lattice and let $a \prec b$ be a cover relation in L.*

(i) *The unique join-irreducible element j of L with $\mathrm{con}(a, b) = \mathrm{con}(j_*, j)$ is $j = \bigwedge\{x \in L \mid x \leq b,\, x \not\leq a\}$. Furthermore, $j \leq b$ but $j \not\leq a$.*

(ii) *The unique meet-irreducible element m of L with $\mathrm{con}(a, b) = \mathrm{con}(m, m^*)$ is $m = \bigvee\{x \in L \mid x \geq a,\, x \not\geq b\}$. Furthermore, $m \geq a$ but $m \not\geq b$.*

In particular, if j is a join-irreducible element and m is a meet-irreducible element with $\mathrm{con}(j_*, j) = \mathrm{con}(m, m^*)$, then $j = \bigwedge\{x \in L \mid x \leq m^*,\, x \not\leq m\}$ and $m = \bigvee\{x \in L \mid x \geq j_*,\, x \not\geq j\}$.

The bijection between meet-irreducible elements and join-irreducible elements in Proposition 9-5.20 coincides with the map κ which appears in Theorem 3-1.4 in the more general context of join-semidistributive lattices. (In the more general context, the map need not be a bijection.)

We briefly discuss a weaker condition than congruence uniformity. Recall from Section 3-2.7 that a lattice is called *congruence normal* if whenever j is join-irreducible and m is meet-irreducible and $\mathrm{con}(j_*, j) = \mathrm{con}(m, m^*)$ then $j \not\leq m$. Theorem 3-2.39 says that a finite lattice is congruence normal if and only if it is obtained from a one-element lattice by a sequence of doublings of convex sets. Theorem 3-2.41 says that a finite lattice is congruence uniform if and only if it is both congruence normal and semidistributive. We quote a result that characterizes congruence normality of a finite lattice L combinatorially in terms of edge labelings.

A *join-fundamental pair* of chains is a pair $C = x_0 \prec x_1 \prec \cdots \prec x_k$ and $D = y_0 \prec y_1 \prec \cdots \prec y_\ell$ with $x_0 = y_0$ and $x_k = y_\ell = x_1 \vee y_1$. Each chain is *unrefinable*, meaning that adjacent elements of the chain are related by cover relations in L. Dually, we can define *meet-fundamental pairs*. A *CN-labeling* of a finite lattice L is a map γ from the set of edges of L to some poset P, satisfying the following properties for each join-fundamental $C = x_0 \prec x_1 \prec \cdots \prec x_k$ and $D = y_0 \prec y_1 \prec \cdots \prec y_\ell$, and satisfying the dual properties for each meet-fundamental pair of chains.

(i) $\gamma(x_0, x_1) = \gamma(y_{\ell-1}, y_\ell)$.

(ii) If $1 < i < k$, then $\gamma(x_{i-1}, x_i) < \gamma(x_0, x_1)$ and $\gamma(x_{i-1}, x_i) < \gamma(x_{k-1}, x_k)$.

(iii) The labels $\gamma(x_{i-1}, x_i)$ for $1 \leq i \leq k$ are all distinct.

$\diamondsuit$ **Theorem 9-5.21.** *A finite lattice is congruence normal if and only if it admits a CN-labeling. In this case, the map taking an edge $x \prec y$ to $\mathrm{con}(x, y) \in \mathrm{Con}_{\mathrm{Ji}}(L)$ is a CN-labeling.*

9-5.5 Congruences on quotients

Since a quotient $L/\boldsymbol{\alpha}$ is isomorphic to $\pi_\downarrow^{\boldsymbol{\alpha}} L$, we will describe congruences on both representations of the quotient, beginning with $L/\boldsymbol{\alpha}$.

First, we prove one of the standard Isomorphism Theorems[7] for finite lattices. For a congruence $\boldsymbol{\alpha}$ on L, the notation $x/\boldsymbol{\alpha}$ stands for the $\boldsymbol{\alpha}$-class of $x \in L$. Given a lattice L and congruences $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ on L such that $\boldsymbol{\alpha} \leq \boldsymbol{\beta}$, define a relation $\boldsymbol{\beta}/\boldsymbol{\alpha}$ on $L/\boldsymbol{\alpha}$ by setting $x/\boldsymbol{\alpha} \equiv y/\boldsymbol{\alpha} \pmod{\boldsymbol{\beta}/\boldsymbol{\alpha}}$ if and only if $x \equiv y \pmod{\boldsymbol{\beta}}$.

Theorem 9-5.22. *Let L be a finite lattice and let $\boldsymbol{\alpha}$ be a congruence on L. Then the map $\boldsymbol{\beta} \mapsto \boldsymbol{\beta}/\boldsymbol{\alpha}$ is an isomorphism from $\{\boldsymbol{\beta} \in \mathrm{Con}\, L \mid \boldsymbol{\beta} \geq \boldsymbol{\alpha}\}$ to $\mathrm{Con}(L/\boldsymbol{\alpha})$.*

Proof. Suppose $\boldsymbol{\beta} \geq \boldsymbol{\alpha}$. Then each $\boldsymbol{\beta}$-class C is an interval $[x, y]$ and is a union of $\boldsymbol{\alpha}$-classes. The set of $\boldsymbol{\alpha}$-classes contained in C is a $\boldsymbol{\beta}/\boldsymbol{\alpha}$-class and equals the interval $[x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha}]$ in $L/\boldsymbol{\alpha}$. We see that $\pi_\downarrow^{\boldsymbol{\beta}/\boldsymbol{\alpha}}(x/\boldsymbol{\alpha}) = (\pi_\downarrow^{\boldsymbol{\beta}} x)/\boldsymbol{\alpha}$. Suppose C_1 and C_2 are $\boldsymbol{\alpha}$-classes with $C_1 \leq C_2$, so that by Proposition 9-5.4, there exist $x \in C_1$ and $y \in C_2$ with $x \leq y$. Then $\pi_\downarrow^{\boldsymbol{\beta}} x \leq \pi_\downarrow^{\boldsymbol{\beta}} y$, so $\pi_\downarrow^{\boldsymbol{\beta}/\boldsymbol{\alpha}}(x/\boldsymbol{\alpha}) = (\pi_\downarrow^{\boldsymbol{\beta}} x)/\boldsymbol{\alpha} \leq (\pi_\downarrow^{\boldsymbol{\beta}} y)/\boldsymbol{\alpha} = \pi_\downarrow^{\boldsymbol{\beta}/\boldsymbol{\alpha}}(y/\boldsymbol{\alpha})$ by Proposition 9-5.4. The

[7] Although the theorems are standard, their numbering is apparently not. This theorem appears (in a more general setting) as LTF Lemma 220, where it is called the Second Isomorphism Theorem. Some other references call this the Third Isomorphism Theorem.

dual argument shows that $\pi^{\uparrow}_{\boldsymbol{\beta}/\boldsymbol{\alpha}}$ is also order-preserving, so $\boldsymbol{\beta}/\boldsymbol{\alpha}$ is a congruence by Proposition 9-5.2.

It is immediate for $\boldsymbol{\beta} \geq \boldsymbol{\alpha}$ and $\boldsymbol{\gamma} \geq \boldsymbol{\alpha}$ that $\boldsymbol{\beta} \leq \boldsymbol{\gamma}$ if and only if $\boldsymbol{\beta}/\boldsymbol{\alpha} \leq \boldsymbol{\gamma}/\boldsymbol{\alpha}$, so the map $\boldsymbol{\beta} \mapsto \boldsymbol{\beta}/\boldsymbol{\alpha}$ is an isomorphism from $\{\boldsymbol{\beta} \in \operatorname{Con} L \mid \boldsymbol{\beta} \geq \boldsymbol{\alpha}\}$ to its image, which is an induced subposet of $\operatorname{Con}(L/\boldsymbol{\alpha})$. We need to show that this image is all of $\operatorname{Con}(L/\boldsymbol{\alpha})$. If $\boldsymbol{\beta}'$ is a congruence on $L/\boldsymbol{\alpha}$, let $\boldsymbol{\beta}$ be the equivalence relation on L with $x \equiv y \pmod{\boldsymbol{\beta}}$ if and only if $x/\boldsymbol{\alpha} \equiv y/\boldsymbol{\alpha} \pmod{\boldsymbol{\beta}'}$. To see that $\boldsymbol{\beta}$ is a congruence, let C be a $\boldsymbol{\beta}$-class and let C' be the corresponding $\boldsymbol{\beta}'$-class. By Proposition 9-5.4, C is the interval between the bottom element of the bottom $\boldsymbol{\alpha}$-class in C' and the top element of the top $\boldsymbol{\alpha}$-class in C'. The projection $\pi_{\downarrow}^{\boldsymbol{\beta}}$ takes an element $x \in L$ to the bottom element of $\pi_{\downarrow}^{\boldsymbol{\beta}'}(x/\boldsymbol{\alpha})$. If $x \leq y$ then $x/\boldsymbol{\alpha} \leq y/\boldsymbol{\alpha}$ by Proposition 9-5.4, and then $\pi_{\downarrow}^{\boldsymbol{\beta}'}(x/\boldsymbol{\alpha}) \leq \pi_{\downarrow}^{\boldsymbol{\beta}'}(y/\boldsymbol{\alpha})$. Then Proposition 9-5.4 says that some element of the $\boldsymbol{\alpha}$-class $\pi_{\downarrow}^{\boldsymbol{\beta}'}(x/\boldsymbol{\alpha})$ is less than or equal to some element of $\pi_{\downarrow}^{\boldsymbol{\beta}'}(y/\boldsymbol{\alpha})$, but then since $\pi_{\downarrow}^{\boldsymbol{\alpha}}$ is order-preserving, we see that the bottom element of $\pi_{\downarrow}^{\boldsymbol{\beta}'}(x/\boldsymbol{\alpha})$ is less than or equal to the bottom element of $\pi_{\downarrow}^{\boldsymbol{\beta}'}(y/\boldsymbol{\alpha})$. We have shown that $\pi_{\downarrow}^{\boldsymbol{\beta}}$ is order-preserving, and the dual argument shows that $\pi^{\uparrow}_{\boldsymbol{\beta}}$ is order-preserving, so $\boldsymbol{\beta}$ is a congruence by Proposition 9-5.2. Therefore $\boldsymbol{\beta}' = \boldsymbol{\beta}/\boldsymbol{\alpha}$, and the proof is complete. □

Given a congruence $\boldsymbol{\alpha}$ on L and a subset U of L, we write $\boldsymbol{\alpha}|_U$ for the restriction of the equivalence relation $\boldsymbol{\alpha}$ to U. The following rephrasing of Theorem 9-5.22 is immediate in light of Proposition 9-5.5.

Corollary 9-5.23. *Let L be a finite lattice and let $\boldsymbol{\alpha}$ be a congruence on L. Then the map $\boldsymbol{\beta} \mapsto \boldsymbol{\beta}|_{\pi_{\downarrow}^{\boldsymbol{\alpha}} L}$ is an isomorphism from $\{\boldsymbol{\beta} \in \operatorname{Con} L \mid \boldsymbol{\beta} \geq \boldsymbol{\alpha}\}$ to $\operatorname{Con}(\pi_{\downarrow}^{\boldsymbol{\alpha}} L)$.*

We can also relate join-irreducible congruences on a quotient of L to join-irreducible congruences on L. Recall from Proposition 9-5.10 that the cover relations in $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$ are exactly $\pi_{\downarrow}^{\boldsymbol{\alpha}} x \prec y$ such that $y \in \pi_{\downarrow}^{\boldsymbol{\alpha}} L$ and $x \prec y$ in L.

Proposition 9-5.24. *Let L be a finite lattice and let $\boldsymbol{\alpha}$ be a congruence on L. Suppose $y \in \pi_{\downarrow}^{\boldsymbol{\alpha}} L$ and suppose $x \prec y$ in L. Then the congruence $\operatorname{con}(x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha})$ on $L/\boldsymbol{\alpha}$ is $(\operatorname{con}(x,y) \vee \boldsymbol{\alpha})/\boldsymbol{\alpha}$. Equivalently, the congruence $\operatorname{con}(\pi_{\downarrow}^{\boldsymbol{\alpha}} x, y)$ on $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$ is the restriction of $\operatorname{con}(x,y) \vee \boldsymbol{\alpha}$ to $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$.*

Proof. Theorem 9-5.22 implies that $\operatorname{con}(x/\boldsymbol{\alpha}, y/\boldsymbol{\alpha})$ is $\boldsymbol{\beta}/\boldsymbol{\alpha}$ such that $\boldsymbol{\beta}$ is the smallest congruence greater than or equal to $\boldsymbol{\alpha}$ such that $x \equiv y \pmod{\boldsymbol{\beta}}$. In other words, $\boldsymbol{\beta} = \operatorname{con}(x,y) \vee \boldsymbol{\alpha}$. The second statement is equivalent in light of Proposition 9-5.5. □

By Proposition 9-5.14, the congruences on $L/\boldsymbol{\alpha}$ described in Proposition 9-5.24 are exactly the join-irreducible congruences. Thus we can relate $\operatorname{Con}_{\mathrm{Ji}}(L/\boldsymbol{\alpha})$ to $\operatorname{Con}_{\mathrm{Ji}}(L)$.

Proposition 9-5.25. *Suppose L is a finite lattice and $\boldsymbol{\alpha}$ is a congruence on L. Then the map $[j] \mapsto (\mathrm{con}(j) \vee \boldsymbol{\alpha})/\boldsymbol{\alpha}$ is an isomorphism from the forcing order, restricted to equivalence classes of join-irreducible elements not contracted by $\boldsymbol{\alpha}$, to $\mathrm{Con}_{\mathrm{Ji}}(L/\boldsymbol{\alpha})$. Thus the congruence lattice $\mathrm{Con}(L/\boldsymbol{\alpha})$ is isomorphic to the poset of down-sets in the restriction of the forcing order to classes of join-irreducible elements not contracted by $\boldsymbol{\alpha}$.*

If we realize the quotient as $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$ instead, then the isomorphism takes $[j]$ to the restriction of $\mathrm{con}(j) \vee \boldsymbol{\alpha}$ to $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$.

Proof. Proposition 9-5.16 implies that the map $[j] \mapsto \mathrm{con}(j)$ restricts to an isomorphism from the forcing order, restricted to equivalence classes of join-irreducible elements not contracted by $\boldsymbol{\alpha}$, to its image inside $\mathrm{Con}_{\mathrm{Ji}} L$. The congruence $\boldsymbol{\alpha}$ corresponds to the downset $(\downarrow \boldsymbol{\alpha}) \cap \mathrm{Con}_{\mathrm{Ji}} L$ in $\mathrm{Con}_{\mathrm{Ji}} L$, and $\mathrm{con}(j) \vee \boldsymbol{\alpha}$ corresponds to the downset generated by $\mathrm{con}(j)$ and $(\downarrow \boldsymbol{\alpha}) \cap \mathrm{Con}_{\mathrm{Ji}} L$. Thus the map $[j] \mapsto (\mathrm{con}(j) \vee \boldsymbol{\alpha})$ is an isomorphism to the set of congruences in $\uparrow \boldsymbol{\alpha} \subseteq \mathrm{Con}\, L$ that cover exactly one other congruence in $\uparrow \boldsymbol{\alpha}$. Applying Theorem 9-5.22, we see that $[j] \mapsto (\mathrm{con}(j) \vee \boldsymbol{\alpha})/\boldsymbol{\alpha}$ is an isomorphism to $\mathrm{Con}_{\mathrm{Ji}}(L/\boldsymbol{\alpha})$. □

9-5.6 Semidistributive lattices

Recall that semidistributivity was considered at length in Chapters 3–6 and very briefly in Section 9-2.2. The following appears as part of Corollary 3-1.22, but here we give a proof using the tools of this section.

Proposition 9-5.26. *If L is a finite semidistributive lattice and $\boldsymbol{\alpha}$ is a congruence on L, then $L/\boldsymbol{\alpha}$ (or equivalently $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$) is semidistributive. The same is true for join- or meet-semidistributivity separately.*

Proof. Suppose L is join-semidistributive and suppose $x, y, z \in \pi_{\downarrow}^{\boldsymbol{\alpha}} L$ have $x \vee y = x \vee z$ in $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$. Because $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$ is a join-subsemilattice of L by Proposition 9-5.8, we have $x \vee y = x \vee z$ in L as well. Since L is join-semidistributive, $x \vee (y \wedge z) = x \vee y$ in L. Because x, y, and z are in $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$, Proposition 9-5.5 implies that $x \vee (y \wedge z) = x \vee y$ in $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$. We see that $\pi_{\downarrow}^{\boldsymbol{\alpha}} L$ is join-semidistributive, or in other words, by Proposition 9-5.5, that $L/\boldsymbol{\alpha}$ is join-semidistributive.

The assertion for meet-semidistributivity holds by the dual proof (using the analogous results on $\pi^{\uparrow}_{\boldsymbol{\alpha}}$). □

Remark 9-5.27. The finiteness hypothesis in Proposition 9-5.26 is indispensable: Quotients of infinite semidistributive lattices need not be semidistributive. Indeed, free lattices are semidistributive and every lattice is a quotient of some free lattice. See, however, Exercise 9.49.

Definition 9-5.28. Suppose L is a finite lattice. A *join representation* for $x \in L$ is an identity of the form $x = \bigvee U$ for some $U \subseteq L$. A join representation

$x = \bigvee U$ is *irredundant* if there does not exist a proper subset $U' \subsetneq U$ such that $x = \bigvee U'$. If $x = \bigvee U$ is irredundant, then in particular U is an antichain in L. For subsets U and V of L, say $U \ll V$ if, for every $u \in U$, there exists a $v \in V$ with $u \leq v$. In other words, $U \ll V$ if and only if $\operatorname{down}(U) \subseteq \operatorname{down}(V)$. This relation is called *join-refinement* because if $U \ll V$ then $\bigvee U \leq \bigvee V$. It restricts to a partial order on antichains. (See Exercise 9.50.) The expression $x = \bigvee U$ is called the *canonical join representation* of x if it is irredundant and if any other join representation $x = \bigvee V$ has $U \ll V$. The elements of U are called the *canonical joinands* of x. Exercise 9.51 shows that $x = \bigvee U$ is the canonical join representation of x if and only if U is the unique minimal (in the sense of $\ll$) antichain joining to x. Exercise 9.52 shows that if x has a canonical join representation, then each canonical joinand of x is join-irreducible. (Exercise 3.1(a) establishes something even stronger.) Exercise 9.53 shows that x has exactly one canonical joinand (which equals x itself) if and only if x is join-irreducible.

Recall that Theorem 3-1.4 states, among other assertions, that a finite lattice L is join-semidistributive if and only if every element of L has a canonical join representation.

Proposition 9-5.29. *Suppose L is a finite join-semidistributive lattice and $\boldsymbol{\alpha}$ is a congruence on L. Then an element $x \in L$ is in $\pi_\downarrow^{\boldsymbol{\alpha}} L$ if and only if none of its canonical joinands is contracted by $\boldsymbol{\alpha}$. If $x \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ then its canonical join representation in $\pi_\downarrow^{\boldsymbol{\alpha}} L$ coincides with its canonical join representation in L.*

Proof. Let J be the set of canonical joinands of x, so that $x = \bigvee J$ is the canonical join representation of x. By Exercise 9.52, each $j \in J$ is join-irreducible, and we write j_* for the unique element covered by j.

Suppose $\boldsymbol{\alpha}$ contracts some $j \in J$. Writing x' for $j_* \vee \bigvee (J \setminus \{j\})$, we have $x \equiv x' \pmod{\boldsymbol{\alpha}}$ because $x = j \vee \bigvee (J \setminus \{j\})$ and $j \equiv j_* \pmod{\boldsymbol{\alpha}}$. But since $\bigvee J$ is irredundant, $J \not\ll (j_* \cup J \setminus \{j\})$, and then since $x = \bigvee J$ is the canonical join representation of x, we see that $x' < x$. We conclude that $x \notin \pi_\downarrow^{\boldsymbol{\alpha}} L$. Conversely, suppose $x \notin \pi_\downarrow^{\boldsymbol{\alpha}} L$, so that there exists $x' < x$ with $x \equiv x' \pmod{\boldsymbol{\alpha}}$. Because $x = \bigvee J$ and because $x' < x$, there exists $j \in J$ such that $x' \not\geq j$. Since $x \equiv x' \pmod{\boldsymbol{\alpha}}$, we have $j = j \wedge x \equiv j \wedge x' \pmod{\boldsymbol{\alpha}}$, but $j \wedge x' \leq j_*$, so $j \equiv j_* \pmod{\boldsymbol{\alpha}}$.

When $x \in \pi_\downarrow^{\boldsymbol{\alpha}} L$, we have showed that no element of J is contracted. Proposition 9-5.8 implies that $x = \bigvee J$ is also the canonical join representation of x in $\pi_\downarrow^{\boldsymbol{\alpha}} L$. $\square$

The proof of Theorem 3-1.4 includes an explicit construction of canonical join representations in finite (join-)semidistributive lattices. We pause to point out a construction of canonical join and meet representations in the special case of congruence uniform lattices, using Proposition 9-5.20. In the spirit of that proposition, given a cover relation $a \prec b$ in a finite congruence uniform lattice L, define $j_{a \prec b}$ to be $\bigwedge\{x \in L \mid x \leq b,\, x \not\leq a\}$ and define $m_{a \prec b}$ to be $\bigvee\{x \in L \mid x \geq a,\, x \not\geq b\}$. The following proposition is proved as Exercise 9.54.

Proposition 9-5.30. *If L is a finite congruence uniform lattice and $x \in L$, then the canonical join representation of x is $x = \bigvee\{j_{w \prec x} \mid w \prec x\}$ and the canonical meet representation of x is $x = \bigwedge\{m_{x \prec w} \mid x \prec w\}$.*

9-6. Polygonal lattices

We now introduce a lattice property called polygonality and discuss its consequences. We also show that polygonality is inherited by quotients. Our motivation for introducing this property will be made clear in Theorem 9-6.10, which asserts that tightness and polygonality coincide for posets of regions.

Definition 9-6.1. A *polygon* in a lattice is an interval $[x, y]$ that is the union of two finite maximal chains from x to y, with these chains disjoint except at x and y. A given lattice may have many polygons or none. A lattice L is called *polygonal* if the following two dual conditions hold:

(i) If distinct elements y_1 and y_2 both cover an element x, then $[x, y_1 \vee y_2]$ is a polygon.

(ii) If an element y covers distinct elements x_1 and x_2, then $[x_1 \wedge x_2, y]$ is a polygon.

Less formally, a polygonal lattice is a lattice that has as many polygons as possible.

Example 9-6.2. Recall that Example 9-3.7 features a lattice of regions $\mathrm{Pos}(\mathcal{A}, B)$ such that $(\mathcal{A}, B)$ is not tight. The regions R_1 and R_2 with separating sets $\{1, 2\}$ and $\{1, 4\}$ both cover the region Q with separating set $\{1\}$, but $[Q, R_1 \vee R_2]$ is not a polygon. Thus $\mathrm{Pos}(\mathcal{A}, B)$ is not polygonal.

Maximal chains in a polygonal lattice are related by local changes in polygons, as we now describe. Suppose L is a polygonal lattice and $x \leq y$ in L. Distinct maximal chains $x = x_0 \prec \cdots \prec x_k = y$ and $x = y_0 \prec \cdots \prec y_\ell = y$ in the interval $[x, y]$ are related by a *polygon move* if there exist i, j with $0 \leq i < j \leq k$ such that the interval $[x_i, x_j]$ is a polygon in L and such that the two chains coincide except in that interval. That is, the two chains differ only in that one chain covers one side of the polygon while the other chain covers the other side.

Lemma 9-6.3. *Suppose L is a finite polygonal lattice and $x \leq y$ in L. Then any two maximal chains in $[x, y]$ are related by a sequence of polygon moves.*

Proof. It is immediate that any interval in a polygonal lattice is itself a polygonal lattice. Thus we may as well argue in the case where $x = 0$. We argue by induction on the *height* $h(y)$ of y in L, the length (the number of edges) of the longest chain from 0 to y. If $h(y) \leq 1$, then there is a unique maximal chain in $[0, y]$, so suppose $h(y) > 1$. Given two distinct maximal chains in $[0, y]$, write the (weakly) longer one as $0 = x_0 \prec \cdots \prec x_k = y$ and write the shorter one as $0 = x'_{k-m} \prec \cdots \prec x'_m = y$. The argument we will make is illustrated in Figure 9-6.1.

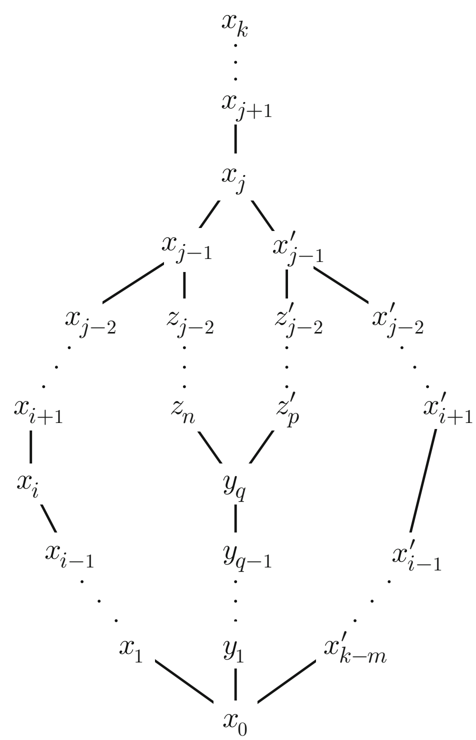

Figure 9-6.1: An illustration of the proof of Lemma 9-6.3

There exists a smallest integer j with $1 < j \leq k$ such that $x_\ell = x'_\ell$ for all $\ell = j, \ldots, k$. In particular, x_{j-1} and x'_{j-1} are distinct elements covered by x_j. Since L is polygonal, writing y for the meet $x_{j-1} \wedge x'_{j-1}$, the interval $[y, x_j]$ is a polygon. Choose any maximal chain $0 \prec y_1 \prec \cdots \prec y_q = y$ in the interval $[0, y]$. Let C be the maximal chain obtained by concatenating $0 \prec y_1 \prec \cdots \prec y_q$ with the side of the polygon containing x_{j-1} and then $x_j \prec \cdots \prec x_k$. This is the chain containing elements of the form z_i in Figure 9-6.1. Let C' agree with C except that we take the side of the polygon containing x'_{j-1}. This is the chain with elements z'_i in Figure 9-6.1. By induction, the chains $0 = x_0 \prec \cdots \prec x_{j-1}$ and $0 \prec y_1 \prec \cdots \prec y_q \prec z_n \prec \cdots \prec z_{j-1} = x_{j-1}$ are related by a sequence of polygon moves. The corresponding polygon moves also relate the chains

$0 = x_0 \prec \cdots \prec x_k = y$ and C. Similarly, the chains $0 = x'_{k-m} \prec \cdots \prec x'_m = y$ and C' are related by a sequence of polygon moves. The chains C and C' are, by construction, related by a single polygon move. Thus the original chains $0 = x_0 \prec \cdots \prec x_k = y$ and $0 = x'_{k-m} \prec \cdots \prec x'_m = y$ are related by a sequence of polygon moves. $\square$

Exercise 9.57 gives another proof of Lemma 9-6.3 and appears to generalize the lemma to lattices with 0 and 1 having no infinite chains. The next result, proved as Exercise 9.58, shows that such a generalization is meaningless.

Proposition 9-6.4. *If L is a polygonal lattice having* 0 *and* 1 *and having no infinite chains, then L is finite.*

9-6.1 Congruences on polygonal lattices

Recall from Section 9-5.4 that an edge $a \prec b$ forces an edge $c \prec d$ if every congruence contracting $a \prec b$ also contracts $c \prec d$, or in other words, if $c \equiv d \pmod{\mathrm{con}(a, b)}$. If a lattice L is itself a polygon $[x, y]$, then edge forcing on L is entirely straightforward. There are two edges incident to x that we call *bottom edges* of the interval and two edges incident to y that we call *top edges* of the interval. The remaining edges in the interval are called *side edges*. In Exercise 9.59, we verify that the only forcing relations are as follows: Each bottom edge forces the opposite top edge (the top edge in the other chain) and also forces all side edges. Each top edge forces the opposite bottom edge (the bottom edge in the other chain) and also forces all side edges.

Accordingly, given edges $a \prec b$ and $c \prec d$, we say $a \prec b$ *forces* $c \prec d$ *in a polygon* if there is some polygon in L containing $a \prec b$ and $c \prec d$ such that one of the following holds:

(i) $a \prec b$ is a bottom edge of the polygon and $c \prec d$ is the opposite top edge.

(ii) $a \prec b$ is a bottom edge of the polygon and $c \prec d$ is a side edge.

(iii) $a \prec b$ is a top edge of the polygon and $c \prec d$ is the opposite bottom edge.

(iv) $a \prec b$ is a top edge of the polygon and $c \prec d$ is a side edge.

Figure 9-6.2 illustrates cases (i) and (ii) of forcing in a polygon. Contracted edges are indicated by gray shading. The picture shows that a bottom edge of a polygon forces the opposite top edge and all side edges. The other forcing relations in a polygon are dual.

The following theorem holds more generally for polygonal lattices without infinite bounded chains, but we prove it here only for finite polygonal lattices. See the Notes to this chapter.

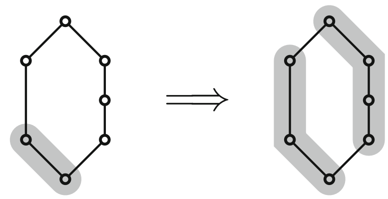

Figure 9-6.2: Forcing in a polygon

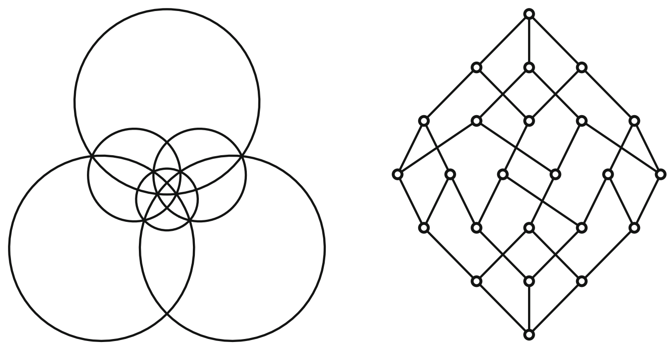

Figure 9-6.3: A simplicial arrangement and its poset of regions

Theorem 9-6.5. *If L is a finite polygonal lattice, then the edge forcing relation on L is the transitive closure of forcing in polygons. That is, given an edge $a \prec b$ that forces another edge $c \prec d$, there exists a sequence of edges*

$$(a \prec b) = (a_0 \prec b_0), (a_1 \prec b_1), \ldots, (a_k \prec b_k) = (c \prec d)$$

such that $a_{i-1} \prec b_{i-1}$ forces $a_i \prec b_i$ in a polygon for all $i = 1, \ldots, k$.

Example 9-6.6. Figure 9-6.3 shows a simplicial arrangement $\mathcal{A}$ and its poset of regions, taking B to correspond to the small triangle inside all of the circles shown. One can verify directly (or, later, as a result of Theorem 9-6.10) that this poset is a polygonal lattice. Figure 9-6.4 shows the steps in applying Theorem 9-6.5 to find the smallest congruence contracting two given edges. As before, contracted edges are shaded. The given edges are shaded in the top-left picture of the figure and the smallest congruence is shown in the bottom-right picture. We will see in Example 10-3.2 that $\mathcal{A}$ is a Coxeter arrangement and in Example 10-6.3 that the congruence shown is a Cambrian congruence.

Remark 9-6.7. Arbitrary finite lattices have a property weaker than, but similar in spirit to, the conclusion of Theorem 9-6.5. Specifically, one can replace forcing in polygonal intervals with forcing in sublattices that are isomorphic to polygons. The main result of [203] is that in any finite lattice, the edge

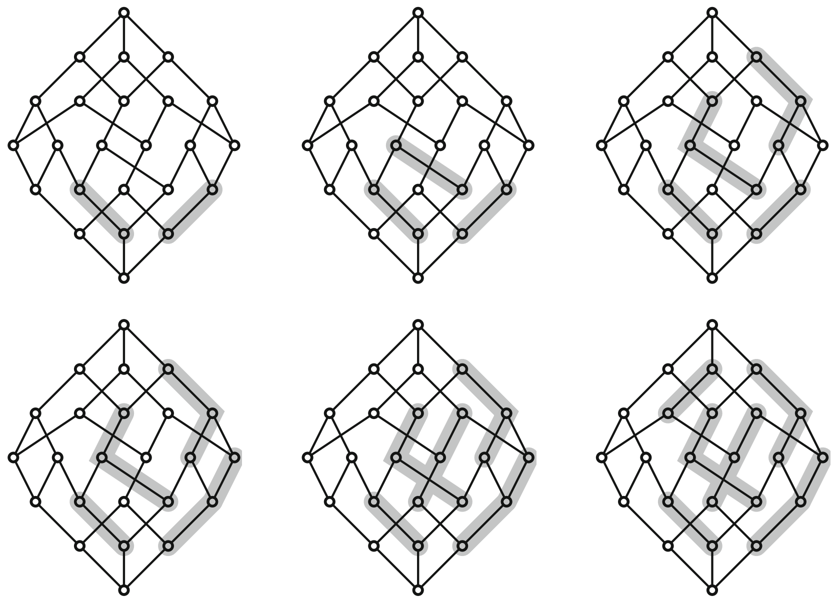

Figure 9-6.4: Applying Theorem 9-6.5

forcing relation is the transitive closure of forcing in such “polygon-sublattices.” Thus we can understand Theorem 9-6.5 to say that, in a polygonal lattice, congruences can be understood in a more “local” way than in general lattices.

The proof of Theorem 9-6.5 rests on the following proposition.

Proposition 9-6.8. *Suppose L is a finite polygonal lattice and $\mathcal{E}$ is a collection of edges of L that is closed under forcing in polygons.*

(i) *The relation $\boldsymbol{\theta}_{\mathcal{E}}$ generated by $\mathcal{E}$ (by reflexive-transitive closure) is a congruence relation.*

(ii) *An edge $a \prec b$ has $a \equiv b \pmod{\boldsymbol{\theta}_{\mathcal{E}}}$ if and only if $a \prec b$ is in $\mathcal{E}$.*

Proof. We verify the first assertion using Proposition 9-5.2. Let C be a $\boldsymbol{\theta}_{\mathcal{E}}$-class and suppose x and y are both in C, so that there exists a path $x = x_0, \ldots, x_k = y$ with every edge (x_{i-1}, x_i) in $\mathcal{E}$. A *local maximum* in the path is x_i with $i \in \{1, \ldots, k-1\}$ and $x_{i-1} \prec x_i \succ x_{i+1}$. We claim that there is a path from x to y consisting of edges in $\mathcal{E}$ and having no local maxima.

Again, we write $h(x)$ for the *height* of x in L, the length of the longest chain from 0 to x. If the path $x = x_0, \ldots, x_k = y$ has a local maximum, then choose x_i to maximize $h(x_i)$ among local maxima. If $x_{i-1} = x_{i+1}$, then we modify the path by deleting x_i and x_{i+1} to create a new path with all edges in $\mathcal{E}$. Otherwise, since L is polygonal, the interval $[x_{i-1} \wedge x_{i+1}, x_i]$ is a polygon. Since the top edges of this polygon are in $\mathcal{E}$ and since $\mathcal{E}$ is closed under forcing in polygons, all edges of the polygon are in $\mathcal{E}$. Replacing x_{i-1}, x_i, x_{i+1} by the path along the bottom of the polygon, we obtain a new path from x to y with all edges in $\mathcal{E}$, and in this path, the maximum of $h(x)$ among local maxima x is either lower or is attained fewer times than in the original path. Repeating, we eventually construct a path with no local maxima, thus proving the claim.

Now suppose that x and y are elements of C that are minimal in L among elements of C. By the claim, there is a path from x to y with all edges in $\mathcal{E}$ and having no local maxima. Since x and y are both minimal, we conclude that $x = y$. The dual argument (using the dual claim) shows that C contains only one maximal element. Write a for the minimal element of C and b for the maximal element.

We next show that C is the entire interval $[a, b]$, and indeed that every edge in $[a, b]$ is in $\mathcal{E}$. By the claim (taking a and b for x and y), there exists a path $a = x_0, \ldots, x_k = b$ with all edges in $\mathcal{E}$, having no local maxima. This path is therefore a maximal chain $a = x_0 \prec \cdots \prec x_k = b$ in $[a, b]$. If $a = y_0 \prec \cdots \prec y_m = b$ is a maximal chain related to $a = x_0 \prec \cdots \prec x_k = b$ by a polygon move, then since $\mathcal{E}$ is closed under forcing in polygons, $a = y_0 \prec \cdots \prec y_m = b$ also has all edges in $\mathcal{E}$. Lemma 9-6.3 says that any maximal chain in $[a, b]$ can be obtained from $a = x_0 \prec \cdots \prec x_k = b$ by a sequence of polygon moves. We conclude that every maximal chain in $[a, b]$ has all edges in $\mathcal{E}$, and thus that every edge in $[a, b]$ is in $\mathcal{E}$. In particular, C is all of $[a, b]$.

We have verified condition (i) of Proposition 9-5.2 for $\boldsymbol{\theta}_{\mathcal{E}}$. We write $\pi_\downarrow^{\mathcal{E}}$ for the map sending each element of L to the bottom element of its $\boldsymbol{\theta}_{\mathcal{E}}$-class. To verify condition (ii), it is enough (Exercise 9.21) to consider elements x and y of L with $x \prec y$ and show that $\pi_\downarrow^{\mathcal{E}} x \le \pi_\downarrow^{\mathcal{E}} y$. If x and y are equivalent mod $\boldsymbol{\theta}_{\mathcal{E}}$, then $\pi_\downarrow^{\mathcal{E}} x = \pi_\downarrow^{\mathcal{E}} y$, so assume they are not equivalent. We argue by induction on the length of a maximal chain from $\pi_\downarrow^{\mathcal{E}} y$ to y. If $y = \pi_\downarrow^{\mathcal{E}} y$, then $\pi_\downarrow^{\mathcal{E}} x \le x \prec \pi_\downarrow^{\mathcal{E}} y$. Otherwise, there exists $z \prec y$ with $z \equiv y$ in $\boldsymbol{\theta}_{\mathcal{E}}$. Since L is polygonal, the interval $[x \wedge z, y]$ is a polygon consisting of chains $x \wedge z = x_0 \prec \cdots \prec x_k = y$ (with $x_{k-1} = x$) and $x \wedge z = z_0 \prec \cdots \prec z_\ell = y$ (with $z_{\ell-1} = z$). Since $\mathcal{E}$ is closed under forcing in polygons, we have $z_1 \equiv y$ and $x \wedge z \equiv x$ in $\boldsymbol{\theta}_{\mathcal{E}}$. By induction, $\pi_\downarrow^{\mathcal{E}}(x \wedge z) \le \pi_\downarrow^{\mathcal{E}} z_1$ but these equal $\pi_\downarrow^{\mathcal{E}} x$ and $\pi_\downarrow^{\mathcal{E}} y$ respectively, so $\pi_\downarrow^{\mathcal{E}} x \le \pi_\downarrow^{\mathcal{E}} y$ and we have proved condition (ii). Condition (iii) is proved by the dual argument. We have verified all three conditions of Proposition 9-5.2, so $\boldsymbol{\theta}_{\mathcal{E}}$ is a lattice congruence. We have also verified that an edge $a \prec b$ is contracted by $\boldsymbol{\theta}_{\mathcal{E}}$ if and only if $a \prec b$ is in $\mathcal{E}$. □

Proof of Theorem 9-6.5. As discussed above (and verified in Exercise 9.59), if $a \prec b$ forces $c \prec d$ in a polygon, then $a \prec b$ forces $c \prec d$ in the usual sense. Thus the transitive closure of forcing in polygons is contained in the edge forcing relation. Given an edge $a \prec b$, let $\mathcal{E}$ be the smallest set of edges that contains $a \prec b$ and is closed under forcing in polygons. If an edge $c \prec d$ is not contained in $\mathcal{E}$, then Proposition 9-6.8 implies that $\boldsymbol{\theta}_{\mathcal{E}}$ is a congruence and $c \not\equiv d$ (mod $\boldsymbol{\theta}_{\mathcal{E}}$). Since there is a congruence contracting $a \prec b$ but not

$c \prec d$, by definition $a \prec b$ does not force $c \prec d$. Thus the edge forcing relation is contained in the transitive closure of forcing in polygons. $\square$

9-6.2 Quotients of polygonal lattices

We now show that polygonality is inherited by quotient lattices.

Proposition 9-6.9. *Suppose L is a finite polygonal lattice and $\boldsymbol{\alpha}$ is a congruence on L. Then $L/\boldsymbol{\alpha}$ is polygonal.*

Proof. Suppose $a_1/\boldsymbol{\alpha}$, $a_2/\boldsymbol{\alpha}$, and $b/\boldsymbol{\alpha}$ are distinct $\boldsymbol{\alpha}$-classes such that $a_1/\boldsymbol{\alpha} \prec b/\boldsymbol{\alpha}$ and $a_2/\boldsymbol{\alpha} \prec b/\boldsymbol{\alpha}$. We want to show that $[(a_1/\boldsymbol{\alpha}) \wedge (a_2/\boldsymbol{\alpha}), b/\boldsymbol{\alpha}]$ is a polygon. Propositions 9-5.5 and 9-5.10 imply that we can choose $b \in \pi_\downarrow^{\boldsymbol{\alpha}} L$ and $a_1, a_2 \in L$ such that $a_1 \prec b$ and $a_2 \prec b$ in L. Then $[a_1 \wedge a_2, b]$ (the meet and the interval in L) is a polygon because L is a polygonal lattice. The top and bottom edges of this interval are not contracted, because $a_1/\boldsymbol{\alpha}$, $a_2/\boldsymbol{\alpha}$, and $b/\boldsymbol{\alpha}$ are distinct $\boldsymbol{\alpha}$-classes. Thus $\boldsymbol{\alpha}$ only contracts side edges of the interval. In particular, the quotient of $[a_1 \wedge a_2, b]$ mod the restriction of $\boldsymbol{\alpha}$ is a polygon. Lemma 9-5.7 says that this quotient is isomorphic to the interval $[(a_1 \wedge a_2)/\boldsymbol{\alpha}), b/\boldsymbol{\alpha}]$ in $L/\boldsymbol{\alpha}$. The latter equals $[(a_1/\boldsymbol{\alpha}) \wedge (a_2/\boldsymbol{\alpha}), b/\boldsymbol{\alpha}]$, which is therefore a polygon. We have established the second condition of Definition 9-6.1 and the first condition holds by the dual argument. $\square$

Proposition 9-6.9 says that the class of finite polygonal lattices is closed under passing to quotients. Exercise 9.62 verifies that the class is closed under finite products. However, the class is not closed under passing to sublattices. (See Exercise 9.63 and/or 9.64.)

9-6.3 Polygonality and tightness

Theorem 9-6.10. *The poset of regions* $\mathrm{Pos}(\mathcal{A}, B)$ *is a polygonal lattice if and only if $\mathcal{A}$ is tight with respect to B.*

Proof. If $\mathcal{A}$ is tight with respect to B, then Lemmas 9-1.26(v) and 9-1.27(v) combine with Theorem 9-3.2 to imply that $\mathrm{Pos}(\mathcal{A}, B)$ is a polygonal lattice.

Conversely, suppose $\mathrm{Pos}(\mathcal{A}, B)$ is a polygonal lattice. Let R be a region and let F_1 and F_2 be lower facets of R, with facet-defining hyperplanes H_1 and H_2 respectively. Let Q_1 and Q_2 be the regions covered by R sharing the facets F_1 and F_2 respectively with R. Write Q for $Q_1 \wedge Q_2$. Since $\mathrm{Pos}(\mathcal{A}, B)$ is polygonal, the interval $[Q, R]$ is a polygon. By Lemma 9-1.20 the interval $[Q, -B]$ in $\mathrm{Pos}(\mathcal{A}, B)$ is isomorphic, by the identity map, to the interval $[Q, -B]$ in $\mathrm{Pos}(\mathcal{A}, Q)$. Since $[Q, R]$ is a polygon, the region R has exactly two lower facets with respect to Q, namely F_1 and F_2, and thus Lemma 9-3.13 implies that $F_1 \cap F_2$ is $(n-2)$-dimensional. $\square$

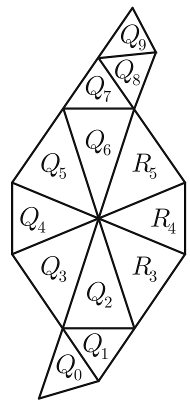

Figure 9-6.5: Two maximal chains related by a rank-two move

Remark 9-6.11. Looking at Theorems 9-3.8 and 9-6.10 together, one might be tempted to guess that semidistributivity implies polygonality and/or *vice versa* in a general finite lattice. Exercises 9.60 and 9.61 ask for counterexamples to both directions of implication. The hint to Exercise 9.60 suggests that the relationship between polygonality and congruence uniformity is also interesting. The relationship between polygonality and congruence normality is addressed in Exercises 9.55 and 9.56.

Combining Lemma 9-6.3 and Theorem 9-6.10, we obtain a useful result on maximal chains in intervals in $\mathrm{Pos}(\mathcal{A}, B)$ when $(\mathcal{A}, B)$ is tight. Two distinct maximal chains $Q = Q_0 \prec \cdots \prec Q_k = R$ and $Q = R_0 \prec \cdots \prec R_k = R$ in the interval $[Q, R]$ are related by a *rank-two move* if there exist a rank-two subarrangement $\mathcal{A}'$ of $\mathcal{A}$ and indices i and j with $1 \leq i < j \leq k$ such that

(i) $Q_m = R_m$ for $m \in \{0, \ldots, i\} \cup \{j, \ldots, k\}$.

(ii) $\mathcal{A}' \cap S(Q_i) = \varnothing$.

(iii) $S(Q_j) = S(Q_i) \cup \mathcal{A}'$.

In the last two requirements, it is useful to remember that $Q_i = R_i$ and $Q_j = R_j$. These requirements immediately imply that for $m \in \{i+1, \ldots, j\}$, the hyperplane defining the common facet of Q_{m-1} and Q_m is in $\mathcal{A}'$ and the hyperplane defining the common facet of R_{m-1} and R_m is in $\mathcal{A}'$. Informally, the two chains are the same except that they go opposite ways around $\mathcal{A}'$ as illustrated schematically in Figure 9-6.5. The picture should be understood in the context of the stereographic-projection pictures of rank-three arrangements that have appeared earlier.

It is immediate that the rank-two moves in $\mathrm{Pos}(\mathcal{A}, B)$ are exactly the polygon moves. Thus the following lemma is Lemma 9-6.3 in the special case of tight lattice of regions.

Lemma 9-6.12. *Suppose $\mathcal{A}$ is tight with respect to B and suppose $Q \leq R$ in* $\mathrm{Pos}(\mathcal{A}, B)$*. Then any two maximal chains in $[Q, R]$ are related by a sequence of rank-two moves.*

We stated and proved Lemma 9-6.12 here with the hypothesis of tightness, but it holds without that hypothesis. See the Notes at the end of the chapter.

9-7. Shards

In this section, we define certain closed polyhedral cones called shards, obtained by decomposing the hyperplanes in an arrangement. The name is meant to suggest breaking the hyperplane, like a pane of glass, into pieces. We will mostly consider shards in the case where $(\mathcal{A}, B)$ is tight, although the same or analogous results may hold whenever $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice.

To define shards, we first define a binary relation on $\mathcal{A}$ called cutting.

Definition 9-7.1. Recall from Definition 9-1.23 the notion of a rank-two subarrangement $\mathcal{A}'$ of $\mathcal{A}$ and the notion of the basic hyperplanes in $\mathcal{A}'$. Given two distinct hyperplanes H_1 and H_2, let $\mathcal{A}'$ be the rank-two subarrangement containing them. We say H_1 *cuts* H_2 if H_1 is basic in $\mathcal{A}'$ and H_2 is not basic in $\mathcal{A}'$. This definition depends on the choice of B *via* the definition of basic hyperplanes in $\mathcal{A}'$. The cutting relation is irreflexive and may fail to be transitive. It is, however, antisymmetric, because by definition if H_1 cuts H_2 then H_2 does not cut H_1.

Definition 9-7.2. Let H be a hyperplane in $\mathcal{A}$. For each hyperplane H' that cuts H, the intersection $H \cap H'$ is an $(n-2)$-dimensional subspace contained in H. In particular, removing $H' \cap H$ from H "cuts" H into two pieces along the intersection. The *shards* in H are the closures of the connected components of $H \setminus \bigcup(H' \cap H)$, where the union is taken over all H' that cut H. Informally, each of the hyperplanes H' slices H, leaving a collection of pieces whose closures are the shards in H. We use the term "shards of $\mathcal{A}$" to refer to the set of all shards in all hyperplanes of $\mathcal{A}$. To emphasize the dependence on B, we may call these the shards of $\mathcal{A}$ with respect to B, or the shards of $(\mathcal{A}, B)$. Each shard Σ belongs to a unique hyperplane in $\mathcal{A}$, and we write H_Σ for this hyperplane. Since each shard is a subset of a hyperplane in $\mathcal{A}$ and is cut out by hyperplanes in $\mathcal{A}$, it is a union of facets of $\mathcal{A}$.

Example 9-7.3. When the rank of $\mathcal{A}$ is 2, there is only one rank-two subarrangement, $\mathcal{A}$ itself. The two basic hyperplanes are the facet-defining hyperplanes of B. The shards are illustrated in Figure 9-7.1 for an arrangement with 6 hyperplanes. Hyperplanes in $\mathbb{R}^2$ are lines. Two of the shards are lines, and the remaining shards are halflines (rays) containing the origin. However, in the figure, the halflines are offset from the origin for clarity. (Otherwise, a picture of all of the shards in $\mathcal{A}$ would be indistinguishable from a picture of all of the hyperplanes in $\mathcal{A}$.)

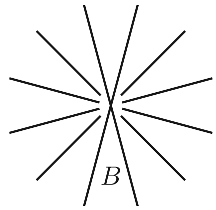

Figure 9-7.1: The shards in a hyperplane arrangement of rank 2

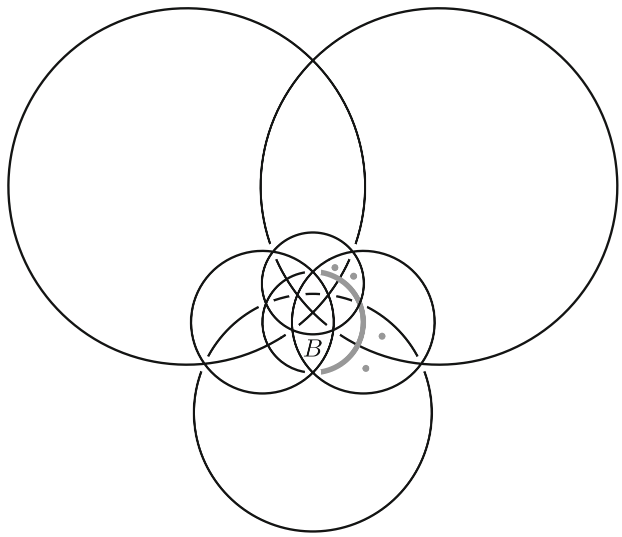

Figure 9-7.2: The shards in a hyperplane arrangement of rank 3

Example 9-7.4. Figure 9-7.2 shows the shards in the arrangement from Example 9-1.7 (Figure 9-1.2) and a particular choice of base region B. Recall that that picture shows a stereographic projection of the intersection of the arrangement with a sphere about the origin. Thus lines in $\mathbb{R}^3$ become pairs of antipodal points in the picture. In particular, when the hyperplanes are cut into shards at rank-two subarrangements, the picture looks locally like a rank-two picture as in Example 9-7.3. In each of these local pictures, we have offset the shards just as in Example 9-7.3 (Figure 9-7.1) for the same reason. The gray dots and the gray color of one shard is explained in Example 9-7.6.

9-7.1 Shards and join-irreducible elements

Definition 9-7.5. Suppose Σ is a shard. An *upper region* of Σ is a region R that intersects Σ in dimension $n-1$ and has $H_\Sigma \in S(R)$. A *lower region* of Σ is a region R that intersects Σ in dimension $n-1$ and has $H_\Sigma \notin S(R)$. In particular, if R is an upper region of Σ, then H_Σ is a lower hyperplane of R in the sense of Definition 9-1.16, and if R is a lower region of Σ, then H_Σ is an upper hyperplane of R. If R is an upper or lower region of Σ, then H_Σ defines a facet of R, and this facet equals $\Sigma \cap R$, so that we might reasonably say that Σ defines a facet of R. Write $\mathrm{Upper}(\Sigma)$ for the set of upper regions of Σ and $\mathrm{Lower}(\Sigma)$ for the set of lower regions of Σ. We will think of $\mathrm{Upper}(\Sigma)$ as partially ordered by the order induced by $\mathrm{Pos}(\mathcal{A}, B)$.

Example 9-7.6. The shard shaded gray in Figure 9-7.2 has four upper regions, and those four regions are marked with gray dots.

Lemma 9-7.7. *The partial order* $\mathrm{Upper}(\Sigma)$ *is connected.*

Proof. Let Q and R be regions in $\mathrm{Upper}(\Sigma)$. Consider the collection $\tilde{\mathcal{A}} = \{H \cap H_\Sigma \mid H \in \mathcal{A} \setminus \{H_\Sigma\}\}$ as an arrangement of hyperplanes in the vector space H_Σ. The $(n-1)$-dimensional cones $Q \cap H_\Sigma$ and $R \cap H_\Sigma$ are regions of $\tilde{\mathcal{A}}$. Lemma 9-1.12 says that there exists a sequence of $\tilde{\mathcal{A}}$-regions $Q \cap H_\Sigma = F_0, \ldots, F_k = R \cap H_\Sigma$ with F_{i-1} adjacent to F_i for $i = 1, \ldots, k$ such that, moving from $Q \cap H_\Sigma$ to $R \cap H_\Sigma$ in the sequence, no hyperplane of $\tilde{\mathcal{A}}$ is crossed more than once. The shard Σ is a closed polyhedral cone in H_Σ whose facet-defining hyperplanes are certain hyperplanes in $\tilde{\mathcal{A}}$. The sequence $F_0, \ldots, F_k$ does not cross any of these facet-defining hyperplanes more than once, and since it starts in Σ and ends in Σ, the entire sequence is in Σ.

The sequence $F_0, \ldots, F_k$ of $\tilde{\mathcal{A}}$-regions corresponds to a sequence $Q = R_0, \ldots, R_k = R$ of regions in $\mathrm{Upper}(\Sigma)$ with $R_{i-1} \cap R_i \cap H_\Sigma$ having dimension $n-2$ for each i from 1 to k. For some i, let $\mathcal{A}'$ be the rank-two subarrangement consisting of hyperplanes containing $R_{i-1} \cap R_i \cap H_\Sigma$. Besides H_Σ, the rank-two subarrangement $\mathcal{A}'$ contains another facet-defining hyperplane of R_{i-1} and another facet-defining hyperplane of R_i. (The latter two hyperplanes might coincide.) Since H_Σ is not cut along the intersection of the hyperplanes in $\mathcal{A}'$, we see that H_Σ is basic in $\mathcal{A}'$. The set $R_{i-1} \cap R_i \cap H_\Sigma$ is an $(n-2)$-dimensional face of $\mathcal{A}$, so Lemma 9-1.25 implies that the separating sets of R_{i-1} and R_i differ only by hyperplanes in $\mathcal{A}'$. Since both separating sets contain H_Σ, Lemma 9-1.24 implies that either $R_{i-1} \leq R_i$ or $R_{i-1} \geq R_i$. We conclude that Q and R are in the same connected component of the partial order induced on $\mathrm{Upper}(\Sigma)$. □

Proposition 9-7.8. *Suppose $\mathcal{A}$ is a tight arrangement with respect to B, and let Σ be a shard of $(\mathcal{A}, B)$. The set* $\mathrm{Upper}(\Sigma)$ *has a unique minimal region J_Σ, which is also the unique join-irreducible region of* $\mathrm{Pos}(\mathcal{A}, B)$ *in* $\mathrm{Upper}(\Sigma)$*. Every join-irreducible region of* $\mathrm{Pos}(\mathcal{A}, B)$ *is J_Σ for some unique Σ.*

Proof. First, suppose J is a minimal region of $\mathrm{Upper}(\Sigma)$. Then J covers the region J_* that shares with J the facet $J \cap \Sigma$. If J is not join-irreducible, then J also covers some region R with $H_\Sigma \in S(R)$. Let $\mathcal{A}'$ be the rank-two subarrangement containing H_Σ and the hyperplane H defining the common facet of J and R. Since $\mathcal{A}$ is tight with respect to B, we apply Lemma 9-1.27 to conclude that the region $J_* \wedge R$ has separating set $S(J) \setminus \mathcal{A}'$ and that there is a region Q covering $J_* \wedge R$ and strictly below J, having $H_\Sigma \in S(Q)$, and sharing with $J_* \wedge R$ a facet defined by H_Σ. The hyperplane H_Σ is basic in $\mathcal{A}'$, so H_Σ is not cut at the intersection $H \cap H_\Sigma$. We see that $Q \in \text{Upper}(\Sigma)$, contradicting the minimality of J, since $Q < J$. We conclude from this contradiction that J is join-irreducible in $\text{Pos}(\mathcal{A}, B)$. The unique region covered by J is J_*.

Suppose now that there is some region $R \in \text{Upper}(\Sigma)$ that is not above J. By Lemma 9-7.7, there exists a sequence $J = R_0, \dots, R_k = R$ of regions in $\text{Upper}(\Sigma)$ such that either $R_{i-1} \le R_i$ or $R_{i-1} \ge R_i$ for each i from 1 to k. We can assume $R_i \ne J$ for all $i > 0$. Since J is minimal, $J = R_0 < R_1$. Let R_i be the first region in the sequence that is not above J. In particular, $R_{i-1} > J$ and $R_i < R_{i-1}$. Let Q_{i-1} be the region in $\text{Lower}(\Sigma)$ that shares a facet with R_{i-1}. Since $S(Q_{i-1}) = S(R_{i-1}) \setminus \{H_\Sigma\}$, we see that $J \vee Q_{i-1} = R_{i-1}$ and $R_i \vee Q_{i-1} = R_{i-1}$. By Theorem 9-3.8, $\text{Pos}(\mathcal{A}, B)$ is semidistributive, so $(J \wedge R_i) \vee Q_{i-1} = R_{i-1}$. Since $R_i \not\ge J$, we have $J \wedge R_i < J$, so that $J \wedge R_i \le J_*$. Therefore, $H_\Sigma \notin S(J \wedge R_i)$. Since also $S(Q_{i-1}) = S(R_{i-1}) \setminus \{H_\Sigma\}$, we see that $(J \wedge R_i) \vee Q_{i-1} \le Q_{i-1} < R_{i-1}$. This contradiction shows that every region of $\text{Upper}(\Sigma)$ is above J, so that J is the unique minimal region.

Now, every region in $\text{Upper}(\Sigma)$ covers a region in $\text{Lower}(\Sigma)$. If J' is a join-irreducible region in $\text{Upper}(\Sigma)$, then the only region covered by J' is in $\text{Lower}(\Sigma)$. In particular every region R strictly below J' has $H_\Sigma \notin S(R)$, and thus $R \notin \text{Upper}(\Sigma)$. We see that J' is minimal in $\text{Upper}(\Sigma)$, so that $J' = J$.

Finally, any join-irreducible region J is in $\text{Upper}(\Sigma)$, where Σ is the unique shard separating J from the unique region J_* covered by J. Thus $H_\Sigma \notin S(J_*)$, so J must be J_Σ. Since every $J_{\Sigma'}$ has its unique lower facet contained in Σ, we see that J does not equal $J_{\Sigma'}$ for any $\Sigma' \ne \Sigma$. □

An interval $[a, b]$ is *prime* if and only if it has exactly two elements. Equivalently, $a \prec b$. Prime intervals $[a, b]$ and $[c, d]$ are *perspective* if either $a \wedge d = c$ and $a \vee d = b$ or $b \wedge c = a$ and $b \vee c = d$. Two intervals are *projective* if they are related in the transitive closure of the perspectivity relation. See LTF Section I.3.5. The notion of projectivity give more insight into the lattice-theoretic significance of shards. If $Q \prec R$ in $\text{Pos}(\mathcal{A}, B)$, then Q and R share a common facet, and that common facet is contained in some shard. We write $\Sigma(Q, R)$ for this shard.

Proposition 9-7.9. *Suppose* $(\mathcal{A}, B)$ *is tight and let* $[Q, R]$ *and* $[Q', R']$ *be prime intervals in* $\text{Pos}(\mathcal{A}, B)$. *Then* $[Q, R]$ *and* $[Q', R']$ *are projective if and only if* $\Sigma(Q, R) = \Sigma(Q', R')$.

Proof. Write Σ for $\Sigma(Q, R)$. Proposition 9-7.8 says that $J_\Sigma \le R$. Let J_* be the unique region covered by J_Σ. Since $S(J_*) = S(J_\Sigma) \setminus \{H_\Sigma\}$ and $S(Q) = S(R) \setminus \{H_\Sigma\}$, we see that $Q \wedge J_\Sigma = J_*$ and $Q \vee J_\Sigma = R$. Thus $[J_*, J_\Sigma]$ and $[Q, R]$ are perspective. Therefore if $\Sigma(Q, R) = \Sigma(Q', R')$, then $[Q, R]$ and $[Q', R']$ are projective.

To prove the converse, we may as well take $[Q, R]$ and $[Q', R']$ to be distinct perspective intervals, and without loss of generality, $Q \wedge R' = Q'$ and $Q \vee R' = R$. Let H be the hyperplane defining the common facet of Q and R, and let H' be the hyperplane defining the common facet of Q' and R'. In particular $S(R) = S(Q) \cup \{H\}$ and $S(R') = S(Q') \cup \{H'\}$. Since $Q \wedge R' = Q'$, we have $Q' \leq Q$ and $R' \not\leq Q$, and therefore $H' \notin S(Q)$. Since $Q \vee R' = R$, we have $R' \leq R$, and therefore $H' \in S(R)$. We conclude that $H = H'$.

Write J for $J_{\Sigma(Q,R)}$ and J' for $J_{\Sigma(Q',R')}$. Then $J \vee Q = R$ and also $Q \leq J' \vee Q \leq R' \vee Q = R$, but since $H \in S(J')$, we see that $J' \vee Q = R$. Theorem 9-3.8 says that $\mathrm{Pos}(\mathcal{A}, B)$ is semidistributive, so $(J \wedge J') \vee Q = R$. If $\Sigma(Q, R) \neq \Sigma(Q', R')$, then $J \neq J'$, so $J \wedge J'$ is strictly below J or J' or both. But every element strictly below J does not have H in its separating set, and similarly for J', so $H \notin S(J \wedge J')$. Therefore since $J \wedge J' \leq R$, also $J \wedge J' \leq Q$, contradicting the fact that $(J \wedge J') \vee Q = R$. This contradiction implies that $\Sigma(Q, R) = \Sigma(Q', R')$. □

9-7.2 Shards and canonical join representations

Definition 9-7.10. A *lower shard* of a region R is a shard Σ such that R is an upper region of Σ in the sense of Definition 9-7.5. Write $\Lambda(R)$ for the set of lower shards of R. The lower facets of R are the sets $\Sigma \cap R$ for $\Sigma \in \Lambda(R)$.

Theorem 9-7.11. *Suppose $\mathcal{A}$ is tight with respect to B. Then the canonical join representation of a region R is $R = \bigvee\{J_\Sigma \mid \Sigma \in \Lambda(R)\}$.*

Our proof of Theorem 9-7.11 uses Theorem 9-3.8 (connecting tightness to semidistributivity) but does not rely on Theorem 3-1.4 (connecting semidistributivity to existence of canonical join representations). We begin with the following lemma, which strengthens Proposition 9-7.8 and also provides more insight into the statement of Theorem 9-7.11.

Lemma 9-7.12. *Suppose $\mathcal{A}$ is tight with respect to B, let R be a region, and let Σ be a lower shard of R. Then J_Σ is the unique minimal element of the set $\{Q \leq R \mid H_\Sigma \in S(Q)\}$.*

Proof. Proposition 9-7.8 implies that $J_\Sigma \leq R$ and $H_\Sigma \in S(J_\Sigma)$. The separating set of the unique region covered by J_Σ does not contain H_Σ, so J_Σ is minimal in $\{Q \leq R \mid H_\Sigma \in S(Q)\}$. If J' is minimal in $\{Q \leq R \mid H_\Sigma \in S(Q)\}$, then J' covers a unique element J'_* with $S(J'_*) = S(J') \setminus \{H_\Sigma\}$. Let Q_0 be the region with $Q_0 \prec R$ and $S(Q_0) = S(R) \setminus \{H_\Sigma\}$. Then $Q_0 \wedge J' = J'_*$ and $Q_0 \vee J' = R$. By Proposition 9-7.9, $\Sigma(J'_*, J) = \Sigma(Q_0, R) = \Sigma$, so $J' = J_\Sigma$. □

Proof of Theorem 9-7.11. Given a lower hyperplane H of R, there is a unique lower shard of R containing the facet $R \cap H$ of R, and given a lower shard Σ, the hyperplane H_Σ is a lower hyperplane. Thus Lemmas 9-3.14 and 9-7.12 together imply that $R = \bigvee\{J_\Sigma \mid \Sigma \in \Lambda(R)\}$ is a join representation of R.

Let Σ_1 and Σ_2 be distinct shards in $\Lambda(R)$. Let Q_1 be the region sharing with R a facet defined by H_{Σ_1} and let Q_2 be the region sharing with R a facet defined by H_{Σ_2}. Since $\mathcal{A}$ is tight with respect to B, Lemma 9-1.27 says that $S(Q_1 \wedge Q_2) = S(R) \setminus \mathcal{A}'$, where $\mathcal{A}'$ is the rank-two subarrangement containing H_{Σ_1} and H_{Σ_2}. Furthermore, in the interval $[Q_1 \wedge Q_2, R]$, there is a region with separating set $(S(R) \setminus \mathcal{A}') \cup \{H_{\Sigma_1}\}$. This region is in $\{Q \leq R \mid H_{\Sigma_1} \in S(Q)\}$ but does not have H_{Σ_2} in its separating set. Lemma 9-7.12 implies that J_{Σ_1} also does not have H_{Σ_2} in its separating set. By Lemma 9-3.14, we conclude that no proper subset of $\{J_\Sigma \mid \Sigma \in \Lambda(R)\}$ joins to R. Thus $R = \bigvee\{J_\Sigma \mid \Sigma \in \Lambda(R)\}$ is an irredundant join representation of R.

Finally, for any $\Sigma \in \Lambda(R)$ and any join representation $R = \bigvee U$, Lemma 9-3.14 implies that there exists a region $Q \in U$ with $H_\Sigma \in S(Q)$. Thus $Q \geq J_\Sigma$ by Lemma 9-7.12. We have shown that $\{J_\Sigma \mid \Sigma \in \Lambda(R)\} \ll U$ for any join representation $R = \bigvee U$. □

9-7.3 Shards and congruences

We have seen in Theorem 9-6.10 that when $(\mathcal{A}, B)$ is tight, then $\mathrm{Pos}(\mathcal{A}, B)$ is polygonal. Applying Theorem 9-6.5, we obtain a description of the congruence lattice $\mathrm{Con}\,\mathrm{Pos}(\mathcal{A}, B)$ in the tight case *via* forcing in polygons. We now use the forcing-in-polygons description to describe $\mathrm{Con}\,\mathrm{Pos}(\mathcal{A}, B)$ in terms of shards and incidences among shards. The description in terms of shards is more compact than the description in terms of polygons and refers directly to the geometry of $(\mathcal{A}, B)$. As a corollary (Corollary 9-7.22) to the description, we characterize the posets of regions which are congruence uniform lattices (in the sense of Day), or equivalently the posets of regions which are quotients of finitely generated free lattices modulo bounded congruences.

Recall that for elements a and b of a lattice L, the notation $\mathrm{con}(a, b)$ denotes the smallest congruence relation on L with $a \equiv b$. It is an easy exercise (or a consequence of the much stronger LTF Theorem 230) that given projective intervals $[a, b]$ and $[c, d]$, a congruence $\boldsymbol{\alpha}$ has $a \equiv b \pmod{\boldsymbol{\alpha}}$ if and only if $c \equiv d$ $\pmod{\boldsymbol{\alpha}}$. In other words, $\mathrm{con}(a, b) = \mathrm{con}(c, d)$. Thus Proposition 9-7.9 has the following immediate consequence.

Proposition 9-7.13. *Suppose $(\mathcal{A}, B)$ is tight. For regions $Q \prec R$, the congruence $\mathrm{con}(Q, R)$ on $\mathrm{Pos}(\mathcal{A}, B)$ depends only on $\Sigma(Q, R)$.*

Definition 9-7.14. Given a shard Σ, there exist adjacent regions Q and R such that $Q \cap R \subseteq \Sigma$. We write $\mathrm{con}(\Sigma)$ to mean $\mathrm{con}(Q, R)$. Proposition 9-7.13 says that $\mathrm{con}(\Sigma)$ is well-defined. Let $\boldsymbol{\alpha}$ be a congruence on a lattice of regions $\mathrm{Pos}(\mathcal{A}, B)$. Say $\boldsymbol{\alpha}$ *removes* the shard $\Sigma(Q, R)$ if $Q \equiv R \pmod{\boldsymbol{\alpha}}$. Proposition 9-7.13 also implies that the notion of removing shards is well-defined. Equivalently, $\boldsymbol{\alpha}$ removes a shard Σ if $\mathrm{con}(\Sigma) \leq \boldsymbol{\alpha}$ in $\mathrm{Con}\,\mathrm{Pos}(\mathcal{A}, B)$.

The term "removing" shards will make more sense in Section 9-8 when we discuss the geometry of quotients of $\mathrm{Pos}(\mathcal{A}, B)$. We will see that such a quotient is a partial order on the cones cut out by the unremoved shards.

Propositions 9-5.14 and 9-7.13 together imply that the join-irreducible congruences of a tight lattice of regions are exactly the congruences $\mathrm{con}(\Sigma)$. (Possibly some join-irreducible congruence may be $\mathrm{con}(\Sigma)$ for more than one shard Σ.) Since a congruence is determined by which join-irreducible congruences it contracts, Proposition 9-7.13 implies the following fact.

Proposition 9-7.15. *Suppose* $\mathrm{Pos}(\mathcal{A}, B)$ *is a lattice. Then a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$ *is determined uniquely by the set of shards it removes.*

To characterize the congruences of $\mathrm{Pos}(\mathcal{A}, B)$, we must thus determine which collections of shards can be removed. We give a complete determination when $(\mathcal{A}, B)$ is tight in terms of a directed graph on shards that we now define.

Definition 9-7.16. The *shard digraph* is the directed graph whose vertices are the shards of $(\mathcal{A}, B)$ with a directed edge $\Sigma_1 \to \Sigma_2$ if and only if H_{Σ_1} cuts H_{Σ_2} and $\Sigma_1 \cap \Sigma_2$ has dimension $n-2$. The reflexive-transitive closure of the shard digraph is a pre-order on the shards of $\mathcal{A}$, again taking the convention that $\to$ corresponds to $\geq$. We set two shards Σ and Σ' to be equivalent if and only if there is a directed path from Σ_1 to Σ_2 in the shard digraph and a directed path from Σ_2 to Σ_1 in the shard digraph. The shard digraph defines a partial order on equivalence classes of shards by the usual construction. (See LTF Section I.1.2 or compare Section 9-5.4.) We call this poset the *shard poset* even though it is a partial order on equivalence classes rather than on shards. Say a set Δ of shards is *closed under arrows* if it satisfies the following condition: If $\Sigma_1 \in \Delta$ and $\Sigma_1 \to \Sigma_2$, then $\Sigma_2 \in \Delta$. Essentially, Δ is a down-set in the shard digraph (keeping in mind that the shard digraph may have directed cycles).

Theorem 9-7.17. *Suppose* $\mathcal{A}$ *is tight with respect to* B*, and suppose* Σ_1 *and* Σ_2 *are shards. Then* $\mathrm{con}(\Sigma_1) \geq \mathrm{con}(\Sigma_2)$ *if and only if there is a directed path in the shard digraph from* Σ_1 *to* Σ_2*.*

Proof. For one direction, it is enough to show that if $\Sigma_1 \to \Sigma_2$ then $\mathrm{con}(\Sigma_1) \geq \mathrm{con}(\Sigma_2)$. Suppose $\Sigma_1 \to \Sigma_2$ and suppose $\boldsymbol{\alpha}$ is a congruence removing Σ_1. Write H_1 for H_{Σ_1} and H_2 for H_{Σ_2}. The intersection $\Sigma_1 \cap \Sigma_2$ is an $(n-2)$-dimensional closed polyhedral cone contained in the $(n-2)$-dimensional subspace $H_1 \cap H_2$. Since each shard is a union of facets of $\mathcal{A}$, the intersection $\Sigma_1 \cap \Sigma_2$ is a union of faces of $\mathcal{A}$, and at least one of these must be $(n-2)$-dimensional. Let F be an $(n-2)$-dimensional face of $\mathcal{A}$ in $\Sigma_1 \cap \Sigma_2$. Let $\mathcal{A}'$ be the rank-two subarrangement of $\mathcal{A}$ consisting of hyperplanes containing F. Since $\Sigma_1 \to \Sigma_2$, H_1 is basic in $\mathcal{A}'$ but H_2 is not. Lemma 9-1.25 says that the regions containing F constitute an interval $[Q, R]$ in $\mathrm{Pos}(\mathcal{A}, B)$ isomorphic to the poset $\mathrm{Pos}(\mathcal{A}', B')$ where B' is the $\mathcal{A}'$-region containing B. The left picture of Figure 9-7.3 illustrates the situation. The picture represents the shards containing F, with F itself

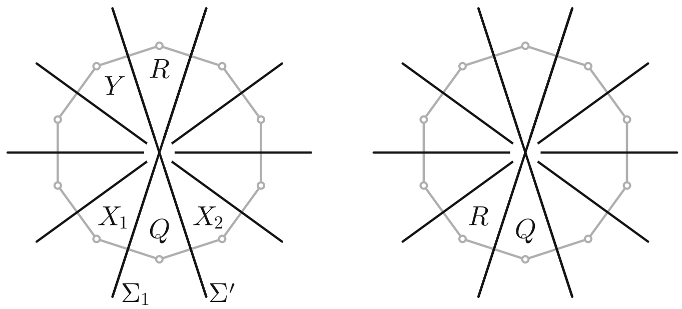

Figure 9-7.3: Illustrations for the proof of Theorem 9-7.17

represented by the origin. The Hasse diagram of $[Q,R]$ is superimposed in gray. All of the shards contain F but in the picture certain shards are offset from the center to indicate that they do not continue through F. We adopt the labeling of regions shown in the figure. Since Σ_1 is removed by $\boldsymbol{\alpha}$, we have $Q \equiv X_1 \pmod{\boldsymbol{\alpha}}$. Since edge forcing in polygons implies edge forcing (this is the easy direction of Theorem 9-6.5), we see that $Q \equiv Y \pmod{\boldsymbol{\alpha}}$ and $X_2 \equiv R \pmod{\boldsymbol{\alpha}}$. Thus $\boldsymbol{\alpha}$ removes all of the shards pictured, except possibly the shard labeled Σ' in the figure. One of the removed shards is Σ_2.

Now suppose $\mathrm{con}(\Sigma_1) \geq \mathrm{con}(\Sigma_2)$. Let $Q_1 \prec R_1$ and $Q_2 \prec R_2$ be edges in $\mathrm{Pos}(\mathcal{A}, B)$ such that $\Sigma_1 = \Sigma(Q_1, R_1)$ and $\Sigma_2 = \Sigma(Q_2, R_2)$. Then $\mathrm{con}(Q_1, R_1) = \mathrm{con}(\Sigma_1) \geq \mathrm{con}(\Sigma_2) = \mathrm{con}(Q_2, R_2)$, so Theorem 9-6.5 implies that there is sequence of edges, starting at $Q_1 \prec R_1$ and ending at $Q_2 \prec R_2$ such that each edge in the sequence forces the following edge *in a polygon* in $\mathrm{Pos}(\mathcal{A}, B)$. Thus we can complete the proof by showing that whenever an edge $Q \prec R$ forces an edge $Q' \prec R'$ in a polygon, either $\Sigma(Q,R) = \Sigma(Q',R')$ or $\Sigma(Q,R) \to \Sigma(Q',R')$. Since $(\mathcal{A}, B)$ is tight, the polygon is an interval consisting of all regions containing some codimension-2 face of $\mathcal{A}$. The situation is illustrated in the right picture of Figure 9-7.3, which coincides with the left picture except for labels on regions and shards. (The case pictured is where $Q \prec R$ is a bottom edge of the polygon. If $Q \prec R$ is a top edge in the polygon, then we should draw a similar picture, but upside-down.) The picture shows the polygon in gray and indicates the shards associated to each edge in the polygon. If the edge $Q' \prec R'$ is the right-top edge of the polygon in the figure, then $\Sigma(Q,R) = \Sigma(Q',R')$. Otherwise, since $Q \prec R$ forces $Q' \prec R'$ in the polygon, $Q' \prec R'$ is one of the side edges, in which case $\Sigma(Q,R) \to \Sigma(Q',R')$. $\square$

The following rephrasing of Theorem 9-7.17 is immediate.

Theorem 9-7.18. *Suppose $\mathcal{A}$ is tight with respect to B and suppose Δ is a set of shards. Then there exists a congruence on $\mathrm{Pos}(\mathcal{A}, B)$ removing exactly the shards in Δ if and only if Δ is closed under arrows.*

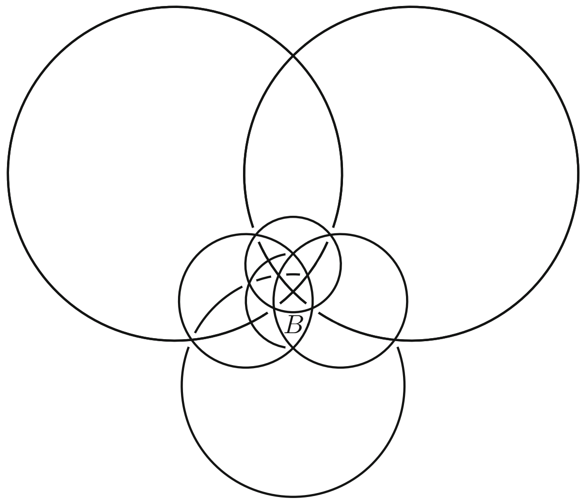

Figure 9-7.4: The shards not removed by a certain lattice congruence

Theorem 9-7.17 implies in particular that $\mathrm{con}(\Sigma_1) = \mathrm{con}(\Sigma_2)$ if and only if Σ_1 and Σ_2 are equivalent in the sense of Definition 9-7.16. In particular, the map $\Sigma \mapsto \mathrm{con}(\Sigma)$ factors through a map $\overline{\mathrm{con}}$ from equivalence classes to congruences. Thus we have a further rephrasing of Theorem 9-7.17.

Theorem 9-7.19. *Suppose $\mathcal{A}$ is tight with respect to B. The map $\overline{\mathrm{con}}$ is an isomorphism from the shard poset to the poset $\mathrm{Con}_{\mathrm{Ji}} \mathrm{Pos}(\mathcal{A}, B)$ of join-irreducible congruences of $\mathrm{Pos}(\mathcal{A}, B)$.*

Example 9-7.20. This is a continuation of Example 9-7.4. Figure 9-7.4 shows the shards not removed by a certain congruence $\boldsymbol{\alpha}$. Specifically, $\boldsymbol{\alpha}$ is the smallest congruence in $\mathrm{Con}\,\mathrm{Pos}(\mathcal{A}, B)$ that removes the shard shaded gray in Figure 9-7.2. This congruence removes exactly two shards besides the gray-shaded shard.

Example 9-7.21. In Theorem 9-7.18, the hypothesis that $(\mathcal{A}, B)$ is tight cannot be weakened to the hypothesis that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice. To see why, we continue Example 9-3.7, which exhibited a non-tight pair $(\mathcal{A}, B)$ such that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice (Figure 9-3.2). In this arrangement, all rank-two subarrangements contain exactly two hyperplanes, so each hyperplane is a shard and the shard digraph has no arrows. Let Σ be the hyperplane numbered 1. Then $\{\Sigma\}$ is a set of shards that is closed under arrows. Suppose there is a congruence $\boldsymbol{\alpha}$ removing only Σ. Then $Q \equiv R \pmod{\boldsymbol{\alpha}}$ if and only if $Q = R$ or $S(Q)$ and $S(R)$ differ only by the hyperplane labeled 1. Naming regions by their labels in Figure 9-3.2, we have $2 \equiv 12 \pmod{\boldsymbol{\alpha}}$. But $2 \vee 24 = 24$ and $12 \vee 24 = 1234$ and these two regions are not equivalent modulo $\boldsymbol{\alpha}$, contradicting the supposition that $\boldsymbol{\alpha}$ is a congruence.

In the following corollary to Theorem 9-7.19, we mean (as usual) congruence uniform in the sense of Day.

Corollary 9-7.22. *The poset of regions* $\mathrm{Pos}(\mathcal{A}, B)$ *is a congruence uniform lattice (or equivalently it is the quotient of a finitely generated free lattice modulo a bounded congruence) if and only if* $\mathcal{A}$ *is tight with respect to* B *and the shard digraph is acyclic.*

Proof. A finite congruence uniform lattice is in particular a semidistributive lattice by Theorem 9-5.19. Theorem 9-3.8 says that $\mathrm{Pos}(\mathcal{A}, B)$ is a semidistributive lattice if and only if $\mathcal{A}$ is tight with respect to B. Thus it remains to show, for $(\mathcal{A}, B)$ tight, that $\mathrm{Pos}(\mathcal{A}, B)$ is congruence uniform if and only if the shard digraph is acyclic.

The shard digraph is acyclic if and only if the shard poset is a partial order on shards (rather than on equivalence classes). By Theorem 9-7.19, this is equivalent to the statement that the map $\Sigma \mapsto \mathrm{con}(\Sigma)$ is a bijection from the set of shards to the set of join-irreducible congruences. By Proposition 9-7.8, the map $J \mapsto \Sigma_J$ is a bijection from join-irreducible elements of $\mathrm{Pos}(\mathcal{A}, B)$ to shards, and by Proposition 9-7.13 we see that $\mathrm{con}(J) = \mathrm{con}(\Sigma_J)$. Thus $\Sigma \mapsto \mathrm{con}(\Sigma)$ is a bijection if and only if $J \mapsto \mathrm{con}(J)$ is a bijection from join-irreducible elements of $\mathrm{Pos}(\mathcal{A}, B)$ to join-irreducible congruences. Since $\mathrm{Pos}(\mathcal{A}, B)$ is self-dual (Exercise 9.4), the latter condition holds if and only if its dual condition holds. $\square$

Exercise 9.69 exhibits a poset of regions $\mathrm{Pos}(\mathcal{A}, B)$ which is a semidistributive lattice but not congruence uniform.

It is instructive to see how Corollary 9-7.22 follows from the other equivalent characterization of congruence uniformity in Theorem 9-5.19. If the shard poset is a partial order on shards (rather than on equivalence classes), then any linear extension $\Sigma_1 \prec \Sigma_2 \prec \cdots \prec \Sigma_k$ of the shard poset gives rise to a maximal chain $\boldsymbol{\alpha}_0 \prec \boldsymbol{\alpha}_1 \prec \cdots \prec \boldsymbol{\alpha}_k$ of congruences in $\mathrm{Con}\,\mathrm{Pos}(\mathcal{A}, B)$ such that each $\boldsymbol{\alpha}_i$ removes the shards $\Sigma_1, \ldots, \Sigma_i$ but not the shards $\Sigma_{i+1}, \ldots, \Sigma_k$. One can check that $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}_i$ is obtained from $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}_{i+1}$ by doubling an interval. We leave the details to Exercise 9.70.

9-7.4 The shard intersection order

Although our main motivation for considering shards is to understand lattice congruences on $\mathrm{Pos}(\mathcal{A}, B)$, we pause to mention a construction that deserves more lattice-theoretic attention. Let $\Psi(\mathcal{A}, B)$ be the set of *shard intersections*: sets that arise as intersections of collections of shards. By convention, the intersection of an empty set of shards is interpreted to be $\mathbb{R}^n$ and is an element of $\Psi(\mathcal{A}, B)$. The results quoted here were proved for shard intersections in simplicial arrangements. We see no obvious obstacle to generalizing them to tight arrangements, but this work has not been carried out. Define a map

$\psi\colon \mathcal{R}(\mathcal{A}) \to \Psi(\mathcal{A}, B)$ sending a region R to the intersection of its lower shards. Define a map $\rho\colon \Psi(\mathcal{A}, B) \to \mathcal{R}(\mathcal{A})$ sending a shard intersection Γ to $\bigvee_{\Sigma \supseteq \Gamma} J_\Sigma$. The join takes place in $\mathrm{Pos}(\mathcal{A}, B)$ and is indexed by all shards containing Γ. The following theorem is proved as [372, Proposition 4.7(i)].

$\diamondsuit$ **Theorem 9-7.23.** *Suppose $\mathcal{A}$ is simplicial. Then the map ψ is a bijection from $\mathcal{R}(\mathcal{A})$ to $\Psi(\mathcal{A}, B)$ with inverse ρ.*

Theorem 9-7.23 is particularly interesting because $\Psi(\mathcal{A}, B)$ admits a natural lattice structure, which we call the *shard intersection order*. As a partial order on shard intersections, it is reverse containment. That is, $\Gamma_1 \leq \Gamma_2$ if and only if $\Gamma_1 \supseteq \Gamma_2$. Thus the join operation is intersection. Since $\mathbb{R}^n$ is a unique minimal element of the shard intersection order, Lemma 9-2.1 implies that the shard intersection order is a lattice. By Theorem 9-7.23, we interpret the shard intersection order as an alternate lattice structure on regions. As a partial order on regions, the shard intersection order is weaker than the poset of regions. (That is, ρ is order-preserving [372, Proposition 4.7(ii)].)

Precisely how much weaker the shard intersection is can be seen in the following theorem. The result is proved in [427] for the case of Coxeter arrangements (see Section 10-2), but the argument given there works for arbitrary simplicial arrangements. Recall the notation $\mathcal{L}(R)$ for the set of lower hyperplanes of R in the sense of Definition 9-1.16.

$\diamondsuit$ **Theorem 9-7.24.** *Suppose $\mathcal{A}$ simplicial. The regions Q and R have $Q \leq R$ in the shard intersection order if and only if $Q \leq R$ in the poset of regions and $\bigcap_{H \in \mathcal{L}(Q)} H \subseteq \bigcap_{H \in \mathcal{L}(R)} H$.*

We conclude by quoting a result on the shard intersection order that suggests generalizations. For any R in $\mathrm{Pos}(\mathcal{A}, B)$, define $\mathcal{J}(R)$ to be the set $\{J_{\Sigma(M,N)} \mid (\bigwedge_{P \prec R} P) \leq M \prec N \leq R\}$. The following is [372, Proposition 5.7].

$\diamondsuit$ **Theorem 9-7.25.** *Suppose $\mathcal{A}$ simplicial. The regions Q and R have $Q \leq R$ in the shard intersection order if and only if $\mathcal{J}(Q) \subseteq \mathcal{J}(R)$.*

Theorem 9-7.25 suggests a generalization of the shard intersection order beyond posets of regions. For simplicity's sake, suppose L is a congruence uniform lattice (in the sense of Day), so the map $j \mapsto \mathrm{con}(j)$ is a bijection from join-irreducible elements of L to join-irreducible congruences on L. Equivalently, by Theorem 9-5.19, L is the quotient of a finitely generated free lattice modulo a bounded congruence. In this case, for any x in L, one can define $\mathcal{J}(x) = \{\mathrm{con}(a, b) \mid (\bigwedge_{w \prec x} w) \leq a \prec b \leq x\}$, and then put an alternate partial order on L by setting $x \leq' y$ if and only if $\mathcal{J}(x) \subseteq \mathcal{J}(y)$. Exercise 9.73 is to show that the relation $\leq'$ is a partial order when L is congruence uniform, but can fail to be antisymmetric otherwise. As a consequence of Proposition 9-7.13, this alternate partial order coincides with the shard intersection order in the case where $\mathrm{Pos}(\mathcal{A}, B)$ is congruence uniform (the case characterized in Corollary 9-7.22).

The fact that the shard intersections for a simplicial arrangement form a lattice is immediate; what is hard to prove is Theorem 9-7.23, which says that the shard intersection order is in fact a partial order on the regions. For general congruence uniform L, the alternate partial order is defined *a priori* as a partial order on the set L, but without the geometric information provided by shards, the lattice property is not obvious. It would be interesting to know for which congruence uniform L this alternate partial order is a lattice. (See Problem 9.5.) Furthermore, does the alternate partial order provide any insight into the structure of congruence uniform lattices? The case where L is the weak order on a finite Coxeter group (see Chapter 10) is encouraging in this regard, as the shard intersection order in this case appears to capture fundamental combinatorics of Coxeter groups. The case where L is a Cambrian lattice is even more encouraging. Here the alternate partial order is, surprisingly, the corresponding noncrossing partition lattice. (See Theorem 10-6.34.)

9-8. Quotients of posets of regions

In this section, we describe the properties of quotients of lattices $\mathrm{Pos}(\mathcal{A}, B)$ of regions, focusing, as in Section 9-7, on the case where $\mathcal{A}$ is tight with respect to B. Some of these properties follow immediately from the corresponding properties of $\mathrm{Pos}(\mathcal{A}, B)$. More interestingly, we discuss how the interplay between geometry and lattice theory is inherited by quotients.

First, quotients of finite lattices inherit the properties of semidistributivity and polygonality, as established in Propositions 9-5.26 and 9-6.9. Thus Theorems 9-3.8 and 9-6.10 imply the following theorem.

Theorem 9-8.1. *Suppose $\mathcal{A}$ is tight with respect to B and suppose $\boldsymbol{\alpha}$ is a congruence on $\mathrm{Pos}(\mathcal{A}, B)$. Then $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ is semidistributive and polygonal.*

Quotients of finite congruence uniform lattices also inherit the property of congruence uniformity. (See the Notes to this section.) However, we say something stronger about congruence uniformity of quotients of $\mathrm{Pos}(\mathcal{A}, B)$ later as Corollary 9-8.20, without relying on Corollary 9-7.22.

9-8.1 The geometric viewpoint

When $(\mathcal{A}, B)$ is tight, a quotient $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ can be realized as a partial order on n-dimensional cones in $\mathbb{R}^n$.

Definition 9-8.2. Given $\mathcal{A}$ tight with respect to B, a congruence $\boldsymbol{\alpha}$ on $\mathrm{Pos}(\mathcal{A}, B)$, and an $\boldsymbol{\alpha}$-class C, the union $\bigcup_{R\in C} R$ of the regions in C will be called an $\boldsymbol{\alpha}$-*cone*. The following proposition justifies the term.

Proposition 9-8.3. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on $\mathrm{Pos}(\mathcal{A}, B)$.*

(i) *Each* $\boldsymbol{\alpha}$*-cone is a closed polyhedral cone.*

(ii) *The* $\boldsymbol{\alpha}$*-cones are the closures of the connected components of* $\mathbb{R}^n \setminus (\bigcup \Sigma)$*, where the union runs over all shards of* $\mathcal{A}$ *not removed by* $\boldsymbol{\alpha}$*.*

(iii) *The interior of each* $\boldsymbol{\alpha}$*-cone is a connected component of* $\mathbb{R}^n \setminus (\bigcup \Sigma)$*.*

(iv) *Given an* $\boldsymbol{\alpha}$*-cone* U*, each facet of* U *is* $U \cap \Sigma$ *for a unique shard* Σ *not removed by* $\boldsymbol{\alpha}$*.*

Proof. Proposition 9-5.2 says that each $\boldsymbol{\alpha}$-class is an interval in $\mathrm{Pos}(\mathcal{A}, B)$. Exercise 9.12 is to show that the union $\bigcup_{R\in I} R$ over any interval I is a closed polyhedral cone. This proves the first assertion of the proposition.

Each shard is a closed set, and there are finitely many shards. Thus each connected component of $\mathbb{R}^n \setminus (\bigcup \Sigma)$ is open and in particular full-dimensional. Since the shards not removed by $\boldsymbol{\alpha}$ are pieces of the hyperplanes that define the regions of $\mathcal{A}$, the closure of each component of $\mathbb{R}^n \setminus (\bigcup \Sigma)$ is a union of regions. By Proposition 9-5.15, two regions Q and R are in the same $\boldsymbol{\alpha}$-class if and only if there exists a sequence of adjacent regions $Q = R_0, \ldots, R_k = R$ such that $R_{i-1} \equiv R_i \pmod{\boldsymbol{\alpha}}$ for all $i = 1, \ldots, k$. Equivalently, each shard $\Sigma(R_{i-1}, R_i)$ or $\Sigma(R_i, R_{i-1})$ is removed by $\boldsymbol{\alpha}$. We see that each $\boldsymbol{\alpha}$-cone U is contained in the closure of some component of $\mathbb{R}^n \setminus (\bigcup \Sigma)$. On the other hand, U has facets that are unions of facets of $\mathcal{A}$ (facets of regions of $\mathcal{A}$). Each facet of $\mathcal{A}$ contained in a facet of U is necessarily contained in a shard that is not removed by $\boldsymbol{\alpha}$. Thus the boundary of U is covered by unremoved shards, and we conclude that U is the closure of an entire connected component of $\mathbb{R}^n \setminus (\bigcup \Sigma)$. We have proved the second assertion of the proposition.

Since each $\boldsymbol{\alpha}$-cone U is the closure of a connected component V of the complement $\mathbb{R}^n \setminus (\bigcup \Sigma)$ and since V is open, $V \subseteq \operatorname{int} U$. If $\operatorname{int} U \not\subseteq V$, then some unremoved shard Σ intersects the interior of U. Therefore there are adjacent regions R_1 and R_2 in C whose common facet is contained in Σ. Since $R_1 \equiv R_2 \pmod{\boldsymbol{\alpha}}$, the shard $\Sigma = \Sigma(R_1, R_2)$ is removed by $\boldsymbol{\alpha}$. This contradiction implies that $V = \operatorname{int} U$. This is the third assertion of the proposition.

Finally, given an $\boldsymbol{\alpha}$-cone U and a facet F of U, there is some region $R \subseteq U$ such that $R \cap F$ is a facet of R. Proposition 9-1.8 says that there exists an adjacent region Q with $Q \cap R \subseteq F$. Necessarily, $Q \not\subseteq U$. We claim that $F = U \cap \Sigma(Q, R)$. Write H for $H_{\Sigma(Q,R)}$, so that $F = U \cap H$. Thus $F \supseteq U \cap \Sigma(Q, R)$. If $F \neq U \cap \Sigma(Q, R)$, then there are regions Q' and R' with $Q' \cap R' \subseteq F$ and $\Sigma(Q', R') \neq \Sigma(Q, R)$. A line segment from the relative interior of $Q \cap R$ to the relative interior of $Q' \cap R'$ exits $\Sigma(Q, R)$ at some point in the relative interior of F. That exit point is in the boundary of Σ, so it is contained in some hyperplane H' that cuts H, and thus also contained in some shard Σ' with $\Sigma' \to \Sigma$. But since H does not cut H', the shard Σ' extends on both sides of H, and in particular intersects the interior of U, contradicting the third assertion of the proposition. We conclude that $F = U \cap \Sigma(Q, R)$. Since $\Sigma(Q, R)$ is the unique shard containing $Q \cap R$, it is also the unique shard with $F = U \cap \Sigma(Q, R)$. □

Example 9-8.4. An example of Proposition 9-8.3 and of later results in this section is provided by Figure 9-7.4 (Example 9-7.20).

Proposition 9-8.3 is our first indication that lattice quotients of tight lattices $\mathrm{Pos}(\mathcal{A}, B)$ of regions have a geometric structure that echoes the geometric definition of the poset of regions. Just as the hyperplanes of $\mathcal{A}$ cut $\mathbb{R}^n$ into open convex sets whose closures are the regions, the unremoved shards of a congruence $\boldsymbol{\alpha}$ on $\mathrm{Pos}(\mathcal{A}, B)$ cut $\mathbb{R}^n$ into open convex sets whose closures are the $\boldsymbol{\alpha}$-cones. Since $\boldsymbol{\alpha}$-cones are in particular in bijection with the $\boldsymbol{\alpha}$-congruence classes, in what follows, the quotient $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ will be realized as a lattice on the set of $\boldsymbol{\alpha}$-cones. To understand this realization of the quotient, we first prove versions of Propositions 9-1.8 and 9-1.15 and Lemma 9-1.18 for quotients.

Proposition 9-8.5. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on $\mathrm{Pos}(\mathcal{A}, B)$. If U and V are distinct $\boldsymbol{\alpha}$-cones and $U \cap V$ is $(n-1)$-dimensional, then $U \cap V$ is a facet of U and a facet of V. If U is an $\boldsymbol{\alpha}$-cone, then every facet of U is shared by a unique other $\boldsymbol{\alpha}$-cone V.*

Proof. Since U is some union of regions, and V is some union of regions, the intersection $U \cap V$ is a union of intersections $Q \cap R$ of regions with $Q \subseteq U$ and $R \subseteq V$. If $U \cap V$ is $(n-1)$-dimensional, then there is at least one such intersection $Q \cap R$ that is $(n-1)$-dimensional, so that in particular Q and R are adjacent regions. Let H be the hyperplane defining their common facet. Since U and V are distinct $\boldsymbol{\alpha}$-cones and both are closed polyhedral cones by Proposition 9-8.3(i), the hyperplane H defines a facet $H \cap U$ of U and a facet $H \cap V$ of V. We need to show that $H \cap U = H \cap V$.

The set $H \cap U$ is a union of facets $H \cap Q'$ of certain regions Q' of $\mathcal{A}$ contained in U. These facets are the regions of a hyperplane arrangement in H, namely $\tilde{\mathcal{A}} = \{H' \cap H \mid H' \in \mathcal{A} \setminus \{H\}\}$. In light of Lemma 9-1.12 and since $Q \cap R$ is a $\tilde{\mathcal{A}}$-region, to prove that $H \cap U \subseteq V$, it is enough to show that whenever F and G are adjacent $\tilde{\mathcal{A}}$-regions such that $F \subseteq U \cap V$ and $G \subseteq U$, then also $G \subseteq V$. Consider the set of shards containing the $(n-2)$-dimensional set $F \cap G$, and consider the rank-two subarrangement $\mathcal{A}'$ consisting of hyperplanes of $\mathcal{A}$ containing $F \cap G$. By Proposition 9-8.3, one of the shards (call it Σ) contains F and G, because they are both contained in the same facet of U. Therefore H equals H_Σ and is basic in $\mathcal{A}$. By the same proposition, no unremoved shard intersects the interior of U, so no other unremoved shard containing F and G continues through its intersection with Σ. In particular, the shard Σ' containing $F \cap G$ and contained in the other basic hyperplane of $\mathcal{A}$ is removed by $\boldsymbol{\alpha}$. But Σ' arrows every other shard (besides Σ) containing $F \cap G$, so Theorem 9-7.18 implies that Σ is the only unremoved shard containing $F \cap G$. Thus by Proposition 9-8.3 again, since F is in V, also G is in V. We have shown that $H \cap U \subseteq H \cap V$, and by symmetry we conclude that $H \cap U = H \cap V$.

We have proved the first assertion of the proposition. For the second assertion, suppose F is a facet of an $\boldsymbol{\alpha}$-cone U and let R be a region contained in U such that $R \cap F$ is $(n-1)$-dimensional. Proposition 9-1.8 says that there exists a unique region Q sharing the facet $R \cap F$ with R. The $\boldsymbol{\alpha}$-cone V containing Q is the unique $\boldsymbol{\alpha}$-cone sharing the facet F with U. $\square$

Since $\boldsymbol{\alpha}$-cones are in bijection with $\boldsymbol{\alpha}$-classes, we may think of the quotient $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ as a partial order on the $\boldsymbol{\alpha}$-cones. Just as for regions, we will say that two $\boldsymbol{\alpha}$-cones are *adjacent* if they share a facet in common. Just as we defined the adjacency graph of $\mathcal{A}$, we can define the adjacency graph on $\boldsymbol{\alpha}$-cones. The following proposition implies that the undirected Hasse diagram of $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ is isomorphic to the adjacency graph on $\boldsymbol{\alpha}$-cones.

Proposition 9-8.6. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *Then $U \prec V$ is a cover relation in the quotient* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *if and only if U and V are adjacent and the hyperplane defining their common facet separates V from B.*

Proof. Proposition 9-5.4 says that $U \prec V$ if and only if there exist regions $Q \subseteq U$ and $R \subseteq V$ with $Q \prec R$ in $\mathrm{Pos}(\mathcal{A}, B)$. By Proposition 9-1.15, this is if and only if there exist Q in U and R in V that are adjacent and have $S(R) = S(Q) \cup \{H\}$, where H defines the common facet of Q and R. If such Q and R exist, then Proposition 9-8.5 says that H also defines a common facet of U and V and separates V from B. Conversely, if U and V are adjacent and the hyperplane H defining their common facet separates V from B, then as argued in the first paragraph of the proof of Proposition 9-8.5, there exist adjacent regions $Q \subseteq U$ and $R \subseteq V$. The common facet of Q and R is defined by H, and $S(R) = S(Q) \cup \{H\}$. $\square$

The following lemma is immediate by Proposition 9-8.6 and Lemma 9-1.18.

Lemma 9-8.7. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *Let $\mathbf{b}$ be a vector in the interior of B. Suppose U and V are adjacent $\boldsymbol{\alpha}$-cones and let $\mathbf{n}$ be a normal vector to their shared facet with* $\langle \mathbf{x}, \mathbf{n} \rangle > 0$ *for all $\mathbf{x}$ in the interior of U. Then $U \prec V$ if and only if* $\langle \mathbf{b}, \mathbf{n} \rangle > 0$.

Proposition 9-8.5 and Theorem 9-1.10 together imply the following fact.

Corollary 9-8.8. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *Then the $\boldsymbol{\alpha}$-cones are the maximal cones of a fan.*

Some fans have a special property of being *normal fans* of polytopes. (See [464, Section 7.1] for definitions.) The vertex-edge graph of the polytope is the adjacency graph on maximal cones of the fan, and each edge of the polytope is normal to the corresponding codimension-1 cone in the fan. The following theorem is immediate from Lemma 9-8.7 and Corollary 9-8.8.

Theorem 9-8.9. *Suppose $\mathcal{A}$ is tight with respect to B and $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$*. If the fan defined by $\boldsymbol{\alpha}$-cones is the normal fan of a polytope P, then the undirected Hasse diagram of the quotient* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *is isomorphic to the vertex-edge graph of P. If $X \prec Y$ in* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$*, then the isomorphism maps X to a vertex $\mathbf{x}$ and Y to a vertex $\mathbf{y}$ such that $\langle \mathbf{b}, \mathbf{x} \rangle > \langle \mathbf{b}, \mathbf{y} \rangle$, where $\mathbf{b}$ is any vector in the interior of B.*

It is helpful to pass between two points of view on a quotient of $\mathrm{Pos}(\mathcal{A}, B)$ modulo $\boldsymbol{\alpha}$. We continue to write $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ for the partial order on $\boldsymbol{\alpha}$-cones obtained by identifying $\boldsymbol{\alpha}$-cones with $\boldsymbol{\alpha}$-classes. We will write $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$ for the subposet of $\mathrm{Pos}(\mathcal{A}, B)$ induced by the regions $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$. Recall that Proposition 9-5.5 states that these two posets are isomorphic. To more easily pass between the two points of view, we extend $\pi_\downarrow^{\boldsymbol{\alpha}}$ to a map from $\boldsymbol{\alpha}$-cones to $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$ in the obvious way: an $\boldsymbol{\alpha}$-cone maps to the minimal region (in $\mathrm{Pos}(\mathcal{A}, B)$) that it contains. We also define a map $\mathrm{Cone}^{\boldsymbol{\alpha}}$ from $\mathrm{Pos}(\mathcal{A}, B)$ to $\boldsymbol{\alpha}$-cones taking a region to the union of the regions in its $\boldsymbol{\alpha}$-class. The map $\mathrm{Cone}^{\boldsymbol{\alpha}}$ restricts to an isomorphism from $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$ to $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ with inverse $\pi_\downarrow^{\boldsymbol{\alpha}}$.

9-8.2 Canonical join representations

Definition 9-8.10. Suppose Σ is a shard not removed by $\boldsymbol{\alpha}$. An *upper $\boldsymbol{\alpha}$-cone* of Σ is an $\boldsymbol{\alpha}$-cone U such that $U \cap \Sigma$ has dimension $(n-1)$ and such that U is greater, in the quotient, than the unique $\boldsymbol{\alpha}$-cone (Proposition 9-8.5) sharing the facet $U \cap \Sigma$ with U. Write $\mathrm{Upper}_{\boldsymbol{\alpha}}(\Sigma)$ for the set of upper $\boldsymbol{\alpha}$-cones of Σ. (Compare Definition 9-7.5.)

Proposition 9-8.11. *Suppose $\mathcal{A}$ is a tight arrangement with respect to B and suppose $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$*. The join-irreducible regions of* $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$ *are exactly the regions J_Σ such that Σ is not removed by $\boldsymbol{\alpha}$. If Σ is not removed by $\boldsymbol{\alpha}$, then the set* $\mathrm{Upper}_{\boldsymbol{\alpha}}(\Sigma)$ *has a unique minimal element* $\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma)$*, which is also the unique join-irreducible element of* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *in* $\mathrm{Upper}_{\boldsymbol{\alpha}}(\Sigma)$*. Every join-irreducible element of* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *is* $\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma)$ *for some unique Σ not removed by $\boldsymbol{\alpha}$.*

Proof. By Definition 9-7.14, a shard Σ is removed by $\boldsymbol{\alpha}$ if and only if J_Σ is not contracted by $\boldsymbol{\alpha}$. Thus Proposition 9-5.11 immediately implies the first assertion. Passing between two points of view on the quotient, we see that the join-irreducible elements of $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ are exactly the cones $\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma)$ for shards Σ not removed by $\boldsymbol{\alpha}$. It remains only to show that $\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma)$ is the unique minimal element of $\mathrm{Upper}_{\boldsymbol{\alpha}}(\Sigma)$. But every cone in $\mathrm{Upper}_{\boldsymbol{\alpha}}(\Sigma)$ is

$\mathrm{Cone}^{\boldsymbol{\alpha}}(R)$ for some region $R \in \mathrm{Upper}(\Sigma)$. We have $J_\Sigma \le R$ by Proposition 9-7.8, and so

$$\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma) = \mathrm{Cone}^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} J_\Sigma) \le \mathrm{Cone}^{\boldsymbol{\alpha}}(\pi_\downarrow^{\boldsymbol{\alpha}} R) = \mathrm{Cone}^{\boldsymbol{\alpha}}(R)$$

because $\pi_\downarrow^{\boldsymbol{\alpha}}$ is order-preserving on $\mathrm{Pos}(\mathcal{A}, B)$ and because $\mathrm{Cone}^{\boldsymbol{\alpha}}$ restricts to an isomorphism from $\pi_\downarrow^{\boldsymbol{\alpha}}(\mathrm{Pos}(\mathcal{A}, B))$ to $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$. $\square$

Definition 9-8.12. A *lower shard* of an $\boldsymbol{\alpha}$-cone U is a shard Σ such that U is an upper $\boldsymbol{\alpha}$-cone of Σ in the sense of Definition 9-8.10. No lower shard of U is removed by $\boldsymbol{\alpha}$. Write $\Lambda(U)$ for the set of lower shards of U. (Compare Definition 9-7.10.)

The following proposition is an immediate consequence of Propositions 9-1.15, 9-5.10 and 9-8.6.

Proposition 9-8.13. *Suppose* $(\mathcal{A}, B)$ *is tight,* $\boldsymbol{\alpha}$ *is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$*, and* U *is a* $\boldsymbol{\alpha}$*-cone. Then the lower shards of the* $\boldsymbol{\alpha}$*-cone* U *are exactly the lower shards of the region* $\pi_\downarrow^{\boldsymbol{\alpha}} U$*.*

The following generalization of Theorem 9-7.11 follows immediately from Proposition 9-5.29, Theorem 9-7.11, and Proposition 9-8.13.

Theorem 9-8.14. *Suppose* $\mathcal{A}$ *is tight with respect to* B *and suppose* $\boldsymbol{\alpha}$ *is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$*. Then the canonical join representation of an* $\boldsymbol{\alpha}$*-cone* U *is* $U = \bigvee\{\mathrm{Cone}^{\boldsymbol{\alpha}}(J_\Sigma) \mid \Sigma \in \Lambda(U)\}$*.*

Remark 9-8.15. The shard intersection order described just after Theorem 9-7.23 interacts nicely with congruences. Given a simplicial arrangement $\mathcal{A}$, a base region B, and a congruence $\boldsymbol{\alpha}$ on $\mathrm{Pos}(\mathcal{A}, B)$, one may naturally consider the subset of the shard intersection order consisting of intersections of shards not removed by $\boldsymbol{\alpha}$. This turns out to be a join-subsemilattice of the shard intersection order. Furthermore, the map ϕ restricts to a bijection from the quotient $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ to the join-subsemilattice, and thus we obtain an alternate lattice structure on the quotient. Theorem 9-7.25 also generalizes, as do many of the properties of the shard intersection order, as explained in [372, Section 7].

9-8.3 Congruences on quotients

Definition 9-8.16. Suppose U and V are $\boldsymbol{\alpha}$-cones and suppose $U \prec V$ in $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$. Proposition 9-7.13 and Theorem 9-5.22 imply that the congruence $\mathrm{con}(U, V)$ on $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ depends only on the shard $\Sigma(U, V)$ defining the common facet of U and V. We write $\mathrm{con}_{\boldsymbol{\alpha}}(\Sigma)$ for the congruence $\mathrm{con}(U, V)$ on $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$. A congruence $\boldsymbol{\beta}$ on $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *removes* the shard $\Sigma(U, V)$ if $U \equiv V \pmod{\boldsymbol{\beta}}$.

The following generalizations of Theorems 9-7.17 and 9-7.18 follow from the original theorems and Proposition 9-5.25.

Theorem 9-8.17. *Suppose $\mathcal{A}$ is tight with respect to B, and suppose $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *Suppose Σ_1 and Σ_2 are shards not removed by $\boldsymbol{\alpha}$. Then* $\mathrm{con}_{\boldsymbol{\alpha}}(\Sigma_1) \geq \mathrm{con}_{\boldsymbol{\alpha}}(\Sigma_2)$ *if and only if there is a directed path in the shard digraph from Σ_1 to Σ_2.*

Theorem 9-8.18. *Suppose $\mathcal{A}$ is tight with respect to B, suppose $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$, *and suppose Δ is a set of shards not removed by $\boldsymbol{\alpha}$. Then there exists a congruence on* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *removing exactly the shards in Δ if and only if Δ is closed under arrows among shards not removed by $\boldsymbol{\alpha}$.*

For shards Σ_1 and Σ_2 not removed by $\boldsymbol{\alpha}$, Theorem 9-7.17 says that $\mathrm{con}(\Sigma_1)$ and $\mathrm{con}(\Sigma_2)$ on $\mathrm{Pos}(\mathcal{A}, B)$ are equal if and only if Σ_1 and Σ_2 are equivalent in the sense of Definition 9-7.16. As in the case of congruences on $\mathrm{Pos}(\mathcal{A}, B)$, the map $\Sigma \mapsto \mathrm{con}_{\boldsymbol{\alpha}}(\Sigma)$ from unremoved shards to join-irreducible congruences on $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ factors through a map $\overline{\mathrm{con}}_{\boldsymbol{\alpha}}$ from equivalence classes to congruences. Thus we have the following generalization of Theorem 9-7.19.

Theorem 9-8.19. *Suppose $\mathcal{A}$ is tight with respect to B, and suppose $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *The map $\overline{\mathrm{con}}_{\boldsymbol{\alpha}}$ is an isomorphism from the shard poset, restricted to shards not removed by $\boldsymbol{\alpha}$, to the poset* $\mathrm{Con}_{\mathrm{Ji}}(\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha})$ *of join-irreducible congruences of* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$.

Finally, a version of Corollary 9-7.22 follows from Theorem 9-8.19 just as Corollary 9-7.22 follows from Theorem 9-7.19.

Corollary 9-8.20. *Suppose $\mathcal{A}$ is tight with respect to B and suppose $\boldsymbol{\alpha}$ is a congruence on* $\mathrm{Pos}(\mathcal{A}, B)$. *Then* $\mathrm{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$ *is a congruence uniform lattice (or equivalently it is the quotient of a finitely generated free lattice modulo a bounded congruence) if and only if the restriction of the shard digraph to shards not removed by $\boldsymbol{\alpha}$ is acyclic.*

9-9. Exercises

Basic notions

9.1. Given a non-redundant expression $\bigcap_{\mathbf{n} \in N} H_{\mathbf{n}}^-$ for an n-dimensional cone R in $\mathbb{R}^n$, show that for each $\mathbf{n} \in N$, the intersection $R \cap H_{\mathbf{n}}$ (a facet of R) is $(n-1)$-dimensional.

9.2. Given a full-dimensional closed polyhedral cone R in $\mathbb{R}^n$ and a vector $\mathbf{n}$ such that $H_{\mathbf{n}}^- \supseteq R$ and $H_{\mathbf{n}} \cap R$ is $(n-1)$-dimensional, show that $H_{\mathbf{n}} \cap R$ is a facet of R. That is, show that the halfspace $H_{\mathbf{n}}^-$ appears in every non-redundant expression for R as a finite intersection of closed halfspaces.

9.3. Verify the antisymmetric property of the order (the poset of regions) defined by $Q \leq R$ if and only if $S(Q) \subseteq S(R)$. Specifically, prove that $R \mapsto S(R)$ is an injective map from $\mathcal{R}(\mathcal{A})$ to subsets of $\mathcal{A}$. (Put another way, a region is uniquely determined by its separating set.)

9.4. Show that the poset of regions is self-dual (meaning isomorphic to its dual). The isomorphism takes a region R to its antipodal region $-R = \{-\mathbf{x} \mid \mathbf{x} \in R\}$.

9.5. A finite lattice L is *orthocomplemented* if there exists an order-reversing involution (an *orthocomplementation*) $x \mapsto x^\perp$ on L with $x^\perp \wedge x = 0$ and $x^\perp \vee x = 1$ for all $x \in L$. In a non-lattice with a smallest element 0 and largest element 1, one can interpret the statement $x^\perp \wedge x = 0$ to mean "the element 0 is the greatest lower bound of x and $x^\perp$" and interpret $x^\perp \vee x = 1$ similarly, and thus define the notion of an orthocomplemented poset. Show that $\mathrm{Pos}(\mathcal{A}, B)$ is orthocomplemented.

9.6. Show that if $|\mathcal{A}| > 1$ then the number of maximal chains in $\mathrm{Pos}(\mathcal{A}, B)$ is even. (*Use Exercise 9.4.*)

9.7. Prove Proposition 9-1.15.

9.8. Show that the function $h(R) = |S(R)|$ has the property that $Q \prec R$ if and only if $Q \leq R$ and $h(R) = h(Q) + 1$. In other words, h is a rank function for $\mathrm{Pos}(\mathcal{A}, B)$, which is therefore graded.

9.9. Prove Proposition 9-1.19.

9.10. Prove Lemma 9-1.24. (*One way to do this: Consider the obvious map from $\mathcal{R}(\mathcal{A})$ to $\mathcal{R}(\mathcal{A}')$. What are the possible separating sets of regions of $\mathcal{A}'$ with respect to B', the region of $\mathcal{A}'$ containing B? Another way: Use Proposition 9-1.19.*)

9.11. Let $Q \in \mathrm{Pos}(\mathcal{A}, B)$ and let X_1 and X_2 be distinct regions covering Q in $\mathrm{Pos}(\mathcal{A}, B)$. Let R be a minimal upper bound of $\{X_1, X_2\}$ in $\mathrm{Pos}(\mathcal{A}, B)$. By Lemma 9-1.17, there are upper hyperplanes H_1 and H_2 of Q such that $S(X_i) = S(Q) \cup \{H_i\}$ for each $i \in \{1, 2\}$. Prove that there exist regions Y_1 and Y_2 covered by R with $S(Y_i) = S(R) \setminus \{H_i\}$ for each $i \in \{1, 2\}$, and furthermore, Y_1 and Y_2 are the only regions covered by R in $[Q, R]$. (Possibly Q has other upper covers besides X_1 and X_2. See for example Figure 9-3.2.)

9.12. Suppose I is an interval in $\mathrm{Pos}(\mathcal{A}, B)$. Show that the union $\bigcup_{R \in I} R$ of regions in I is a closed polyhedral cone. Give a precise description of the cone in terms of the vectors $\mathbf{n}_H$ defined in Proposition 9-1.19.

Lattice-theoretic shortcuts

9.13. Show that a poset is well-founded if and only if it satisfies the Descending Chain Condition.

9.14. Find a well-founded partially ordered set satisfying condition (iii) of Lemma 9-2.3 but not conditions (i) and (ii).

9.15. Recall the definition of completeness from the last paragraph of Section 9-2.1. Show that a well-founded meet-semilattice is a complete meet-semilattice. Conclude that the condition "P is a complete meet-semilattice" can be added to the list of equivalent conditions in Lemma 9-2.3. (*Use the implication* (i) $\Longrightarrow$ (ii) *in Lemma 9-2.3 and adapt the argument that* (ii) $\Longrightarrow$ (i).)

9.16. Lemma 9-2.4 says that a lower finite join-semilattice P is a lattice. Show that this lattice is a complete meet-semilattice but need not be a complete join-semilattice. (*Use Exercise 9.15.*)

9.17. Under the hypotheses of Lemma 9-2.5, show that P is a complete meet-semilattice. (*Use Exercise 9.15.*)

9.18. Prove Lemma 9-2.7.

9.19. Find a counterexample to the following false BEZ Lemma for distributivity (an incorrect dualization of Lemma 9-2.7): Suppose L is a finite lattice such that the distributive law $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$ holds whenever y and z cover a common element. Then L is distributive.

9.20. Since Lemma 9-2.8 is only about joins, one might consider weakening the hypotheses of Lemma 9-2.8 to only require L to be a lower finite join-semilattice. Show that the lemma fails with this weaker hypothesis. (*To find a counterexample faster, recall Lemma 9-2.4.*)

9.21. Call a poset P *interval-finite* if every interval is finite. Recall that P is *lower finite* if the downset $\downarrow x$ of every element $x \in P$ is finite.

(a) Show that a lower finite lattice is interval-finite.

(b) Show that in a interval-finite poset, the order relation $\leq$ is the reflexive-transitive closure of the cover relation $\prec$.

(c) If P is interval-finite, show that a map η from P to a poset P' is order-preserving if and only if $x \prec y \implies \eta(x) \leq \eta(y)$.

9.22. Let P be a poset with 0 and let I be a down-set in P that is a lower finite meet-semilattice. Write $\vee_I$ for the join in I if it exists and similarly $\vee_P$. Suppose I and P satisfy the following condition.

> If $x, y \in I$ cover a common element and $x \vee_I y$ exists then $x \vee_P y$ exists and $x \vee_I y = x \vee_P y$.

Prove that for any $x, y \in I$, if $x \vee_I y$ exists then $x \vee_P y$ exists and $x \vee_I y = x \vee_P y$.

Tight posets of regions

9.23. Let C be a simplicial cone in $\mathbb{R}^n$. Show that every pair of facets of C intersects in a face of C of dimension $n-2$.

9.24. Show that if R has no lower hyperplanes with respect to B, then $R = B$.

9.25. Recall that the *adjacency graph* of $\mathcal{A}$ is the graph $G(\mathcal{A})$ whose vertices are regions and whose edges are pairs of adjacent regions. Two hyperplane arrangements are called *weakly combinatorially isomorphic* if they have isomorphic adjacency graphs. (See Remark 9-3.20.)

(a) Suppose $\mathcal{A}$ and $\mathcal{A}'$ are hyperplane arrangements and B and B' are respective base regions. Show that if $\mathrm{Pos}(\mathcal{A}, B)$ is isomorphic to $\mathrm{Pos}(\mathcal{A}', B')$, then $\mathcal{A}$ and $\mathcal{A}'$ are weakly combinatorially isomorphic.

(b) Suppose $\mathcal{A}$ and $\mathcal{A}'$ are weakly combinatorially isomorphic and write $R \mapsto R'$ for the combinatorial isomorphism. Show that for any base region B of $\mathcal{A}$, the map $R \mapsto R'$ is also an isomorphism from $\mathrm{Pos}(\mathcal{A}, B)$ to $\mathrm{Pos}(\mathcal{A}', B')$.

Biconvexity and rank-two biconvexity

9.26. Show that the intersection of convex subsets of $\mathcal{A}$ (with respect to B) is a convex subset of $\mathcal{A}$ (with respect to B).

9.27. Show that for any $S \subseteq \mathcal{A}$, the closure $\overline{S}$ of S is the intersection of all convex sets containing S.

9.28. Recall from Definition 8-3.1 that a *closure operator* on a set Ω is an isotone and idempotent map φ from the powerset of Ω to itself that is *extensive* (i.e., $X \subseteq \varphi(X)$ for all $X \subseteq \Omega$) and that has $\varphi(\varnothing) = \varnothing$. Recall also that if φ is a closure operator on Ω, then (Ω, φ) is a *convex geometry* if $\varphi(X \cup \{p\}) = \varphi(X \cup \{q\})$ and $p \neq q$ imply together that $p \in \varphi(X)$ (equivalently, $q \in \varphi(X)$), for all $p, q \in P$ and all $X \subseteq \Omega$.

Pick any $\Omega \subseteq \mathbb{R}^n \setminus \{0\}$ such that no two distinct elements of Ω are collinear (i.e., related by scaling). Denote by $\mathrm{cone}(X)$ the convex cone generated by X, for any $X \subseteq \mathbb{R}^n$, and set $\varphi(X) = \mathrm{cone}(X) \cap \Omega$. Prove that if $\mathrm{cone}(\Omega) \cap (-\mathrm{cone}(\Omega)) = \{0\}$, then (Ω, φ) is a convex geometry. Find an example where (Ω, φ) is not a convex geometry.

9.29. Show that a set $S \subseteq \mathcal{A}$ of hyperplanes is convex with respect to B in the sense of Definition 9-4.1 if and only if the set $\{\mathbf{n}_B(H) \mid H \in S\}$ of vectors is convex with respect to $\{\mathbf{n}_B(H) \mid H \in \mathcal{A}\}$.

9.30. Recall that $\mathcal{B}(R)$ is the set of boundary hyperplanes of a region R, the hyperplanes defining facets of R. Show that $\overline{\mathcal{B}(B)}$ (in the sense of Definition 9-4.1) is $\mathcal{A}$.

9.31. Show that the intersection of any collection of rank-two convex subsets of $\mathcal{A}$ is a rank-two convex subset of $\mathcal{A}$.

9.32. Suppose a set $S \subseteq \mathcal{A}$ of hyperplanes is convex with respect to B. Show that S is rank-two convex with respect to B.

9.33. Suppose $\mathcal{A}$ is simplicial and let B be a base region. Formulate two descriptions of the meet in $\mathrm{Pos}(\mathcal{A}, B)$, in terms of convexity and in terms of rank-two convexity. (*Use Theorem 9-4.8 and Exercise 9.4 but phrase the description without reference to antipodal regions.*)

9.34. Prove the first assertion of Lemma 9-4.12.

Lattice congruences for combinatorialists

9.35. Suppose $\boldsymbol{\alpha}$ is an equivalence relation on a lattice L such that, for all $x_1, x_2, y \in L$, if $x_1 \equiv x_2 \pmod{\boldsymbol{\alpha}}$ then $x_1 \wedge y \equiv x_2 \wedge y \pmod{\boldsymbol{\alpha}}$ and $x_1 \vee y \equiv x_2 \vee y \pmod{\boldsymbol{\alpha}}$. Show that $\boldsymbol{\alpha}$ is a congruence.

9.36. Find a lattice L and a congruence on L that fails to have the property that every congruence class is an interval. (In light of Proposition 9-5.2, L must be infinite.) Among such examples, find an example that minimizes the number of congruence classes.

9.37. Suppose L is a lattice and suppose $\pi^\uparrow \colon L \to L$ and $\pi_\downarrow \colon L \to L$ are *order-preserving* maps satisfying

(i) $\pi_\downarrow(x) \le x \le \pi^\uparrow(x)$ for every $x \in L$,

(ii) $\pi^\uparrow \circ \pi^\uparrow = \pi^\uparrow$, $\pi^\uparrow \circ \pi_\downarrow = \pi^\uparrow$, $\pi_\downarrow \circ \pi_\downarrow = \pi_\downarrow$, and $\pi_\downarrow \circ \pi^\uparrow = \pi_\downarrow$.

Then the fibers of $\pi^\uparrow$ determine the same equivalence relation on L as the fibers of $\pi_\downarrow$ and this equivalence relation is a lattice congruence with projection maps $\pi^\uparrow$ and $\pi_\downarrow$.

9.38. Suppose L is a finite lattice and $\boldsymbol{\alpha}$ is a congruence on L. Show that, for any $x \in L$, the element $\pi_\downarrow^{\boldsymbol{\alpha}} x$ is the unique maximal element of $(\downarrow x) \cap \pi_\downarrow^{\boldsymbol{\alpha}} L$.

9.39. Prove Proposition 9-5.8.

9.40. Prove the following assertion for an arbitrary lattice L and deduce Proposition 9-5.2 in the special case where L is finite. An equivalence relation $\boldsymbol{\alpha}$ is a lattice congruence if and only if

(i) Each equivalence class is a convex sublattice of L.

(ii) If $x \le y$ and $x \equiv x' \pmod{\boldsymbol{\alpha}}$ then there exists y' such that $x' \le y'$ and $y \equiv y' \pmod{\boldsymbol{\alpha}}$.

(iii) If $x \le y$ and $y \equiv y' \pmod{\boldsymbol{\alpha}}$ then there exists x' such that $x' \le y'$ and $x \equiv x' \pmod{\boldsymbol{\alpha}}$.

9.41. Prove the following assertion for an arbitrary lattice L and deduce Proposition 9-5.2 in the special case where L is finite. An equivalence relation $\boldsymbol{\alpha}$ is a lattice congruence if and only if

(i) Each equivalence class is a convex sublattice of L.

(ii) If $w \leq x$ and $w \leq y$ and $w \equiv x \pmod{\boldsymbol{\alpha}}$ then there exists z such that $x \leq z$ and $y \leq z$ and $y \equiv z \pmod{\boldsymbol{\alpha}}$.

(iii) If $x \leq z$ and $y \leq z$ and $y \equiv z \pmod{\boldsymbol{\alpha}}$ then there exists w such that $w \leq x$ and $w \leq y$ and $w \equiv x \pmod{\boldsymbol{\alpha}}$.

9.42. Recall the definition of a complete lattice from the last paragraph of Section 9-2.1. A *complete congruence* on a complete lattice L is a congruence $\boldsymbol{\alpha}$ such that if $x_i \equiv y_i \pmod{\boldsymbol{\alpha}}$ for all $i \in I$, then $\bigvee\{x_i \mid i \in I\} \equiv \bigvee\{y_i \mid i \in I\} \pmod{\boldsymbol{\alpha}}$ and similarly for meets. (See LTF Section IV.4.10.) A map $\eta\colon L \to L'$ is a *complete homomorphism* if $\eta\left(\bigvee S\right) = \bigvee\{\eta(x) \mid x \in S\}$, and similarly for meets, for any set $S \subseteq L$. Show that Propositions 9-5.1, 9-5.2 and 9-5.5 hold if the phrase "finite lattice" is replaced with "complete lattice" throughout and the adjective "complete" is added to the phrases "(lattice) congruence" and "(lattice) homomorphism" throughout.

9.43. Do Propositions 9-5.1, 9-5.2, and/or 9-5.5 hold when the phrase "finite lattice" is replaced by "interval-finite lattice" throughout? (See Exercise 9.21 for the definition of interval-finite.)

9.44. Say a surjective homomorphism η between lattices is *bounded* if each of its fibers $\eta^{-1}(x)$ is an interval. Show that Proposition 9-5.1 holds for arbitrary lattices L and L' if the adjective "bounded" is placed before the phrase "lattice homomorphism."

9.45. Say a congruence on a lattice is *bounded* if every congruence class is an interval. Show that Proposition 9-5.2 holds for an arbitrary lattice L if the adjective "bounded" is placed before the phrase "lattice congruence."

9.46. Show that the assertion about cover relations in Proposition 9-5.4 holds in one direction for arbitrary lattices: If distinct $\boldsymbol{\alpha}$-classes C_1 and C_2 have $x \in C_1$ and $y \in C_2$ with $x \prec y$ then $C_1 \prec C_2$. Give a counterexample to the converse in general, but show that the converse hold when L is interval-finite. (See Exercise 9.21.)

9.47. Prove Proposition 9-5.8 in the more general setting of an arbitrary lattice L and a bounded congruence $\boldsymbol{\alpha}$.

9.48. Prove Proposition 9-5.20.

9.49. Show that Proposition 9-5.26 holds if L is an arbitrary semidistributive lattice and $\boldsymbol{\alpha}$ is a bounded congruence.

9.50. Show that the relation $\ll$ of Definition 9-5.28 restricts to a partial order on antichains.

9.51. Show that $x = \bigvee U$ is the canonical join representation of x if and only if U is the unique minimal (in the sense of $\ll$) antichain joining to x.

9.52. Suppose $x = \bigvee U$ is the canonical join representation of x in L. Show that each element of U is join-irreducible in L.

9.53. Let x be an element of a finite lattice. Show that x has canonical join representation $\bigvee\{x\}$ if and only if x is join-irreducible.

9.54. Prove Proposition 9-5.30.

Polygonal lattices

9.55. Suppose L is a finite polygonal lattice. Use Theorem 9-5.21 to show that L is congruence normal if and only if, for every polygon $[x, y]$ and each maximal chain $x = z_0 \prec \cdots \prec z_k = y$ in $[x, y]$, the congruences $\mathrm{con}(z_{i-1}, z_i)$ are all distinct.

9.56. Suppose L is a finite polygonal lattice with the property that, for every maximal chain $0 = x_0 \prec \cdots \prec x_k = 1$, the congruences $\mathrm{con}(x_{i-1}, x_i)$ are all distinct. Show that L is congruence normal.

9.57. Denote by $\Gamma(P)$ the set of all maximal chains in a poset P. For $X, Y \in \Gamma(P)$, say $X \asymp_P Y$ holds if there exist $a, b \in X \cap Y$ such that $a \leq b$, $X \cap {\downarrow} a = Y \cap {\downarrow} a$, $X \cap {\uparrow} b = Y \cap {\uparrow} b$, and

$$\text{(9-9.1)}\quad x \wedge y = a \text{ and } x \vee y = b \quad \text{within } [a, b]\,, \quad \text{whenever} \\ (x, y) \in X \times Y\,, \quad a < x < b\,, \text{ and } a < y < b\,.$$

The transitive closure $\equiv_P$ of $\asymp_P$ is an equivalence relation on $\Gamma(P)$.

(a) Suppose $u \leq v$ in P, $U \in \Gamma({\downarrow} u)$, $V \in \Gamma({\uparrow} v)$, and $X, Y \in \Gamma([u, v])$. Show that if $X \asymp_{[u,v]} Y$, then $U \cup X \cup V \asymp_P U \cup Y \cup V$. Show also that if $X \equiv_{[u,v]} Y$, then $U \cup X \cup V \equiv_P U \cup Y \cup V$.

(b) Show that part (a) of the exercise fails if the phrase "within $[a, b]$" is deleted from (9-9.1).

(c) Show that if P is a bounded poset with no infinite chains and $X, Y \in \Gamma(P)$, then $X \equiv_P Y$. (*Since P has no infinite chains, its poset of closed intervals under containment is well-founded. Thus we can argue by induction this poset of closed intervals and use part* (a).)

(d) Use part (c) to give an alternate proof of Lemma 9-6.3.

9.58. This exercise proves Proposition 9-6.4. Suppose L is a polygonal lattice having 0 and 1 and having no infinite chains. We want to show that L is finite.

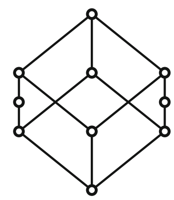

Figure 9-9.1: A polygonal lattice for Exercise 9.63

(a) Show that it is enough to consider the special case where L is infinite but every proper interval in L is finite. In parts (b) and (c), we reach a contradiction in this special case.

(b) Recall that an *atom* of L is an element covering 0. Show that no two atoms a and b of L have $a \vee b = 1$.

(c) Given an atom a, show that there exists an atom b such that $a \vee b = 1$. (*Consider the set* $\bigcup_{x \in \uparrow a \setminus \{1\}} \downarrow x$.)

9.59. Suppose L is a polygon $[x, y]$. Verify the following two assertions: Each bottom edge forces the opposite top edge (the top edge in the other chain) and also forces all side edges. Each top edge forces the opposite bottom edge (the bottom edge in the other chain) and also forces all side edges.

9.60. Find a small example of a lattice that is semidistributive but not polygonal. (*There is a seven-element congruence uniform example.*)

9.61. Find a small example of a lattice that is polygonal but not semidistributive. (*Again, seven elements is enough.*)

9.62. Prove that the class of finite polygonal lattices is closed under finite products.

9.63. Find a finite polygonal lattice with a non-polygonal sublattice. (*See the lattice of Figure 9-9.1.*)

9.64. Find a congruence uniform polygonal finite lattice with a sublattice that is not polygonal. (*There is a seven-element congruence uniform example for Exercise 9.60 that is not subdirectly irreducible, but rather embeds into* $\mathsf{N}_5 \times \mathsf{N}_5$*, which is congruence uniform and polygonal. Here* N_5 *is the* 5*-element non-modular lattice.*)

9.65. Let L be a meet semidistributive lattice and let x_1 and x_2 be distinct elements of L that both cover their meet $a = x_1 \wedge x_2$.

(a) Let y_i be a lower cover of $b = x_1 \vee x_2$ in $[x_i, b]$, for each $i \in \{1, 2\}$. Prove that $y_1 \wedge y_2 = a$.

(b) Prove that y_1 and y_2 are the only lower covers of b in $[a, b]$.

(c) Prove that if L is semidistributive, then x_1 and x_2 are the only upper covers of a in $[a, b]$.

How does this exercise relate to Exercise 9.60?

9.66. Taking $(\mathcal{A}, B)$ as in Example 9-3.5 (Figure 9-3.1), use Theorems 9-6.10 and 9-6.5 to compute the following congruences on $\text{Pos}(\mathcal{A}, B)$, perhaps by shading the contracted edges on a copy of the diagram of $\text{Pos}(\mathcal{A}, B)$.

(a) $\text{con}(5, 25)$

(b) $\text{con}(5, 235)$

(c) $\text{con}(5, 245)$

(Here elements are represented by their separating sets, written without commas or set braces.) For each congruence $\boldsymbol{\alpha}$ listed above, use Proposition 9-5.5 to draw a representation of the quotient lattice $\text{Pos}(\mathcal{A}, B)/\boldsymbol{\alpha}$. (Cf. Exercise 9.67.)

Shards

9.67. Consider $(\mathcal{A}, B)$ as in Example 9-3.5 (Figure 9-3.1).

(a) Find the shards (perhaps drawing the shards on a copy of the figure), find J_Σ for each shard, and find the shard digraph.

(b) Find $\text{Con}_{\text{Ji}} \text{Pos}(\mathcal{A}, B)$ (the partial order on join-irreducible congruences of $\text{Pos}(\mathcal{A}, B)$. Represent $\text{Con}_{\text{Ji}} \text{Pos}(\mathcal{A}, B)$ as a partial order on join-irreducible elements of $\text{Pos}(\mathcal{A}, B)$.

9.68. Show that the poset of regions $\text{Pos}(\mathcal{A}, B)$ is a distributive lattice if and only if a set of vectors, one normal to each hyperplane in $\mathcal{A}$, is linearly independent. Describe the poset of regions precisely in this case.

9.69. Consider the hyperplane arrangement $\mathcal{A}$ shown in Figure 9-9.2 and take B to be the region marked with a dot in the picture. Note that five of the hyperplanes shown intersect “at infinity” in this stereographic projection. Show that $\text{Pos}(\mathcal{A}, B)$ is a semidistributive lattice but is not congruence uniform (equivalently, not a bounded homomorphic image of a finitely generated free lattice).

9.70. Show directly (that is, not using Theorem 9-5.19 and Corollary 9-7.22) that when $\mathcal{A}$ is tight with respect to B and the shard digraph is acyclic, then $\text{Pos}(\mathcal{A}, B)$ can be obtained from a one-element lattice by a sequence of doublings of intervals. (For the definition of doubling, see Section 3-2.7.)

9.71. Suppose $\text{Pos}(\mathcal{A}, B)$ is a lattice, suppose $H \in \mathcal{A}$, let $\mathcal{A}' = \mathcal{A} \setminus \{H\}$, and for any $\mathcal{A}$-region R, let R' be the $\mathcal{A}'$-region containing R. Show that $\text{Pos}(\mathcal{A}, B)$ is obtained from $\text{Pos}(\mathcal{A}', B')$ by doubling a convex set if and only if the map $R \mapsto R'$ is a homomorphism from $\text{Pos}(\mathcal{A}, B)$ to $\text{Pos}(\mathcal{A}', B')$. (See Exercise 9.70.)

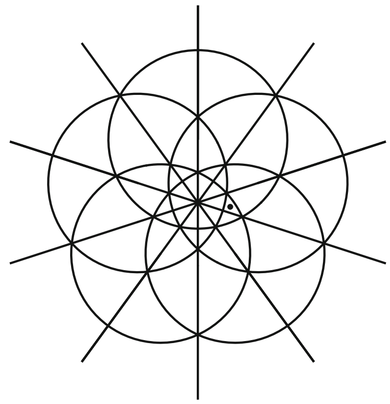

Figure 9-9.2: A hyperplane arrangement and base region such that $\mathrm{Pos}(\mathcal{A}, B)$ is semidistributive but not congruence uniform

9.72. Given a total order $H_1, \ldots, H_k$ on $\mathcal{A}$, define $\mathcal{A}' = \mathcal{A} \setminus \{H_k\}$, and define B' to be the $\mathcal{A}'$-region containing B. Then $H_1, \ldots, H_k$ is a *quotient order* (with respect to B) if $\mathrm{Pos}(\mathcal{A}, B)$ is obtained from $\mathrm{Pos}(\mathcal{A}', B')$ by doubling a convex set and $H_1, \ldots, H_{k-1}$ is a quotient order on $\mathcal{A}'$ (with respect to B'). Show that $\mathcal{A}$ has a quotient order with respect to B if and only if the cutting relation (Definition 9-7.1) is acyclic.

9.73. Show that the alternate order $\leq'$, defined in Section 9-7.4, is a partial order when L is congruence uniform. Give an example of a lattice for which $\leq'$ fails to be antisymmetric.

Quotients of posets of regions

9.74. Consider the hyperplane arrangement $\mathcal{A}$ shown in Figure 9-9.2 and take B to be the region marked with a dot in the picture. Exercise 9.69 shows that $\mathrm{Pos}(\mathcal{A}, B)$ is a semidistributive lattice but is not congruence uniform. Find a minimal lattice quotient of $\mathrm{Pos}(\mathcal{A}, B)$ that is semidistributive but not congruence uniform.

9-10. Notes

Basic notions

Edelman first defined the poset of regions in [139], but he credits Purdy [360] with studying the same poset in $\mathbb{R}^2$. The results in Section 9-1.4 are due to Björner, Edelman, and Ziegler [139, 70], although for convenience we have added some more detailed statements. We give here a few specific citations beyond what appears in the text. Theorem 9-1.21 is [70, Theorem 3.1].

Exercise 9.4 is [139, Proposition 2.1]. Exercise 9.6 is part of [139, Theorem 2.2]. Proposition 9-1.15 is part of the proof of [139, Proposition 1.1], which is a more general version of Exercise 9.8. Lemma 9-1.26 is essentially a special case of [139, Lemma 2.3]. Exercise 9.11 was suggested by Friedrich Wehrung.

Cordovil [102] later generalized the definition and basic results on posets of regions to the setting of oriented matroids. Using the framework of [70, Section 6], we expect that most of the results of this chapter can be extended to the generality of oriented matroids, but we have seen no reason to do so here. One can imagine a possible reason to generalize these results to oriented matroids: One might contemplate a statement along the lines of "A lattice has Property X if and only if it is isomorphic to $\mathrm{Pos}(\mathcal{A}, B)$ for $(\mathcal{A}, B)$ having Property Y." But a true statement along these lines would almost certainly involve oriented matroids rather than posets of regions.

Lattice-theoretic shortcuts

The most basic shortcuts (Lemmas 9-2.1, 9-2.3, and 9-2.4) for proving the (semi-)lattice property are well known. The BEZ Lemma (Lemma 9-2.2) is [70, Lemma 2.1] and is named for the authors Björner, Edelman, and Ziegler. (This name for the lemma has apparently not appeared in the literature before.) The other BEZ-type lemmas have not appeared in the literature, except that Lemma 9-2.9 recently appeared as [310, Lemma 2.2.1] and Lemma 9-2.6 follows from another BEZ-type lemma [310, Lemma 2.2.2]. Other BEZ-type lemmas include [380, Lemma 2.6] and Exercise 9.22. The main result of [202] is similar in spirit to the BEZ lemmas. The proof of Theorem 9-4.8 is modeled after the proof of the BEZ lemma. The author gratefully acknowledges extensive conversations with David Speyer on the topic of BEZ-type lemmas, including some of the lemmas proved here.

Tight posets of regions

The notion of tightness and Theorems 9-3.2 and 9-3.8 and Proposition 9-3.11 are new, but the definition (under a different name) and results were given independently by McConville [309] while this chapter was being written. The result in [309] goes further than Theorem 9-3.8 by establishing that $\mathrm{Pos}(\mathcal{A}, B)$ is semidistributive if and only if it is *crosscut simplicial*. Corollary 9-3.4 is [70, Theorem 3.4]. Corollary 9-3.9 appeared as [364, Theorem 3].

Theorem 9-3.15 is part of a result of Edelman and Reiner [143, Theorem 3.3]. The theorem there is equivalent but phrased differently (as alluded to at the beginning of Section 9-3.2): It states that the simplicial complex defined by the regions of $\mathcal{A}$ is *balanced*. The proof in [143] relies on [123, Proposition 1.12], which is a special case of Lemma 9-6.12, which is in turn a special case of Lemma 9-6.3. The proof given here uses the lattice property directly without appealing to [123, Proposition 1.12].

Biconvexity and rank-two biconvexity

Theorem 9-4.3 is equivalent, by a standard affinization argument, to [397, Lemma 6.1], and the argument given here is taken from [397]. Theorems 9-4.5 and 9-4.8 are inspired by [70, Theorem 5.5], which has the weaker hypothesis that $\mathrm{Pos}(\mathcal{A}, B)$ is a lattice but characterizes regions and the join only in terms of convexity, not in terms of rank-two convexity.

As far as the author is aware, the rank-two convexity statement in Theorems 9-4.5 has not appeared in print except here and in [308, Theorem 5.4]. Similarly, the rank-two convexity statement in Theorem 9-4.8 has not appeared in print except here and in [308, Theorem 5.1]. The latter is stronger than Theorem 9-4.8, in that it has the weaker hypothesis of tightness, rather than simpliciality. For rank-two convexity in the special case where $\mathcal{A}$ is a finite Coxeter arrangement, see the Notes to Chapter 10.

Exercise 9.29 is taken from [70, Remark 5.3], and Exercise 9.28 is a linearization of Example 1 of [141, Section 3].

Lattice congruences for combinatorialists

As mentioned in the text, probably none of the results proved in this section are surprising to lattice-theorists. Some of them are standard. For example, Proposition 9-5.2 has appeared in several of the author's papers (with the proof omitted). It follows fairly easily from Dorfer [130, Corollary 3.4] or from Chajda and Snášel [95], as verified in Exercises 9.40 and 9.41. The exercises themselves (aside from deducing Proposition 9-5.2) are results of [130] and [95] respectively. See also Kolibiar [281] and Grätzer [202, Lemma 2]. Some of the other results in the section have also appeared in the author's earlier papers. Proposition 9-5.12 is [170, Lemma 2.32]. Proposition 9-5.14 is [170, Theorem 2.30], which is attributed to Dilworth [127]. The Fundamental Theorem of Distributive Lattices (used in the proof of Corollary 9-5.17) is due to Birkhoff [63, Theorem 17.3]. Theorem 9-5.19 is due to Day ([113, Lemma 4.2] and [113, Theorem 5.1]). See also [170, Theorem 2.20] and [170, Corollary 2.43] and more generally all of [170, Section II.3]. Exercise 9.37 was suggested by Vincent Pilaud and Aram Dermenjian.

Theorem 9-5.21 is [364, Theorem 4], which was inspired by the notion of $\mathcal{HH}$-lattices in [91]. Unfortunately, the definition of CN-labelings is misstated in [364]. Specifically, the dual requirement on meet-fundamental pairs is omitted. Since the lattices considered in [364] are posets of regions, which are self dual (Exercise 9.4), and since the labelings used there respect that duality, the results of [364] are not affected by the misstatement.

Polygonal lattices

The results of this section have not, to our knowledge, appeared in the literature. However, the $\mathcal{HH}$-*property*, defined in [91], is stronger than polygonality. A related notion is in [146]. See the Notes to Chapter 10. Exercises 9.57, 9.58, 9.62, 9.63, 9.64, and 9.65 were suggested by Friedrich Wehrung. Proposition 9-6.4 and Theorem 9-6.5 are both due to Wehrung, in response to questions posed in an early version of this chapter. Wehrung's proof of Theorem 9-6.5 is based on ideas from [445] and applies more generally to polygonal lattices without infinite bounded chains. The proof given here is quite different and is based on an earlier argument for tight lattices of regions.

Lemma 9-6.12 is due to Deligne [123, Proposition 1.12] in the case of simplicial arrangements. Edelman proved that the simplicial case of Lemma 9-6.12 follows from the fact that the poset of regions of a simplicial arrangement is a lattice. Edelman's proof is unpublished but was communicated to the author by Vic Reiner, and the proof of Lemma 9-6.3 given here is Edelman's proof, generalized to polygonal lattices. In fact, Lemma 9-6.12 holds without the hypothesis of tightness. This was proved by Salvetti [391, Lemma 11] and also by Cordovil and Moreira [103, Theorem 2.4] in the more general setting of oriented matroids. The proof also appears, with definitions and notation more in the style of the current chapter, as [373, Theorem 3.4]. For even stronger and more general results of this kind, see Athanasiadis, Edelman, and Reiner [40] and Athanasiadis and Santos [42].

Shards

Shards were first defined in [364] in order to prove versions of Theorem 9-7.19 and Corollary 9-7.22 for simplicial arrangements. They appeared again in [365] as part of the proof of an upper bound on the order dimension of $\mathrm{Pos}(\mathcal{A}, B)$. (For more on order dimension, see Trotter [436].) In [364, 365], closures were not taken in the definition of shards, but this difference is inconsequential except for changing the phrasing of some definitions and results. More detailed studies of shards were carried out in [366] (expanding on Theorem 9-7.19) and in [372] (studying the shard intersection order mentioned in connection with Theorem 9-7.23). Most of the results of Section 9-7 appear in those references, under the stronger hypothesis of simpliciality (rather than tightness). In particular, Theorem 9-7.11 appeared for simplicial arrangements as [372, Theorem 3.6]. (See also [381, Theorem 8.1].) Theorems 9-7.18, 9-7.17, and 9-7.19 generalize the second assertion of [364, Theorem 25] to tight arrangements and remove the requirement that the shard digraph be acyclic. Exercises 9.71 and 9.69 are results from [364], while Exercise 9.72 generalizes [364, Theorem 26]. Exercise 9.73 was suggested by Friedrich Wehrung. Corollary 9-7.22 generalizes the first assertion of [364, Theorem 25] to tight arrangements. The shard intersection order on permutations was also studied by Bancroft [48, 49] and the shard intersection order on the classical Coxeter groups was studied further by Petersen [348]. Theorem 9-7.24 is [427, Proposition 4.47].

Quotients of posets of regions

The results of this section overlap slightly with [367, Theorem 5.1], but most of the contents of [367, Theorem 5.1] are not in the section and most of the contents of this section are not in [367, Theorem 5.1]. Other than this small overlap, the results of this section are new. The fact that quotients of finite congruence uniform lattices inherit the property of congruence uniformity is proved as [170, Corollary 2.17].

9-11. Problems

Problem 9.1. Find a necessary and sufficient local condition on $(\mathcal{A}, B)$ for $\mathrm{Pos}(\mathcal{A}, B)$ to be a lattice. Specifically, the condition should be based on local configurations of hyperplanes/regions. *This problem is posed in [70], where the sufficient condition of simpliciality is discussed. A weaker sufficient condition is tightness* (*see Theorem 9-3.2*)*, but tightness is not necessary, as illustrated by Example 9-3.7.*

Problem 9.2. Characterize the lattices $\mathrm{Pos}(\mathcal{A}, B)$ for $(\mathcal{A}, B)$ tight and/or, more generally, the lattices that appear as $\mathrm{Pos}(\mathcal{A}, B)$. For example, develop the equational theory of such lattices. (Is it decidable? Is there a lattice identity that holds in all lattices $\mathrm{Pos}(\mathcal{A}, B)$, or in such lattices with $(\mathcal{A}, B)$ tight, that does not hold for all lattices? Compare [398].) What other "non-equational" properties do these lattices have? (See also Problem 10.1 in Section 10-10.)

Problem 9.3. Find the order dimension of the poset of regions of a simplicial hyperplane arrangement. More generally, find the order dimension of $\mathrm{Pos}(\mathcal{A}, B)$ when $\mathcal{A}$ is tight with respect to B. *This problem is considered in the special case of the weak order on a Coxeter group in [365]. The problem is solved there for some Coxeter groups, including the symmetric group. In the latter case, the problem was solved earlier by Flath [154]. There are no known counterexamples to the guess that the order dimension is the rank* (*the dimension of the linear span of normal vectors to the hyperplanes in* $\mathcal{A}$)*. The rank n is an obvious lower bound, as it is easy to find in* $\mathrm{Pos}(\mathcal{A}, B)$ *a subposet isomorphic to the standard example [436, Example 5.1] of a poset of dimension n.* (See also Problem 10.2 in Section 10-10.)

Problem 9.4. Let L be a lattice and let $\boldsymbol{\alpha}$ be an arbitrary *equivalence* relation on L (not necessarily a congruence). Define a relation "$\leq$" on $\boldsymbol{\alpha}$-classes by setting $C_1 \leq C_2$ if and only if there exist $x \in C_1$ and $y \in C_2$ with $x \leq y$. (Compare Proposition 9-5.4.) When is the set $L/\boldsymbol{\alpha}$ of equivalence classes, endowed with the relation $\leq$, a lattice? A necessary condition for $L/\boldsymbol{\alpha}$ to be a poset is that each $\boldsymbol{\alpha}$-class be order-convex (that is, closed under taking intervals). One might specialize this problem to special classes of lattices L or generalize it to allow L to be a poset. *This problem was posed by Christian Stump.*

Problem 9.5. The paragraphs after Theorem 9-7.25 define an "alternate" partial order on a congruence uniform lattice L. This alternate partial order is important when L is the weak order on a finite Coxeter group or when L is a Cambrian lattice. (In the latter case, the alternate partial order is the noncrossing partition lattice. See Theorem 10-6.34.) For which congruence uniform L is the alternate partial order a lattice?

Problem 9.6. Can every finite, congruence uniform lattice be embedded into a finite, congruence uniform, polygonal lattice? Can it be embedded into a congruence uniform lattice of regions for some tight (simplicial) arrangement? *This problem is suggested by Exercise 9.64 and was posed by Friedrich Wehrung.*

Chapter

10

Finite Coxeter Groups and the Weak Order

N. Reading

In this chapter, we develop the basic theory of finite Coxeter groups, drawing on results already proved for posets of regions. There are two main points to this chapter: First, to show how the geometry and lattice theory of hyperplane arrangements underlies the theory of finite Coxeter groups, and second, to point out the weak orders on finite Coxeter groups as an important class of lattice-theoretic examples. A broader class of examples is obtained as lattice quotients of weak orders. Several examples of such quotients are given in Sections 10-6 and 10-7.

10-1. Coxeter groups and the weak order

A *Coxeter group* is a group presented by generators and relations of a very specific form. There is a finite[1] set S of generators, and for each pair s, t of distinct generators, we choose a quantity $m(s, t)$, which must be either an

[1]Some authors allow S to be infinite, but we have no need to do so.

G. Grätzer, F. Wehrung (eds.), *Lattice Theory: Special Topics and Applications*,
DOI 10.1007/978-3-319-44236-5_10

integer ≥ 2 or ∞. We require that $m(s,t) = m(t,s)$ for each pair s,t. The Coxeter group is the group W given by the presentation

$$W = \left\langle S \;\middle|\; s^2 = 1 \,\forall s \in S,\ (st)^{m(s,t)} = 1 \,\forall s \neq t \in S \right\rangle \tag{10-1.1}$$

Here 1 is the identity element. If $m(s,t) = \infty$, then the notation $(st)^{m(s,t)} = 1$ means that no relation of the form $(st)^k = 1$ is imposed. We call S the set of *defining generators* of W. The cardinality of S is called the *rank* of W.

For readers not familiar with groups presented by generators and relations, we provide a brief explanation. We restrict our attention to the presentation (10-1.1), in order to avoid some complications that can arise for general group presentations. A formal treatment of generators and relations is found in most modern algebra texts. Least formally, the group W is the largest group generated by the symbols S such that all the *defining relations* $s^2 = 1$ and $(st)^{m(s,t)} = 1$ hold. Here, "largest" means that any other group generated by S and satisfying the defining relations is a quotient (i.e., a group-homomorphic image) of W.

Suppose we are given a group $\overline{W}$ and a map $s \mapsto \bar{s}$ from S into $\overline{W}$ such that the image $\bar{S}$ of the map generates $\overline{W}$. Suppose also that the defining relations of W hold in $\overline{W}$ with respect to the given map $s \mapsto \bar{s}$. In other words, suppose that for each $s \in S$ the element $\bar{s}^2$ equals the identity in $\overline{W}$ and for each $s \neq t$ in S, the element $(\bar{s}\,\bar{t}\,)^{m(s,t)}$ equals the identity in $\overline{W}$. There are some immediate consequences of the fact that the defining relations hold. First, the fact that $\overline{W}$ is generated by $\bar{S}$ means that each element of $\overline{W}$ is a product $a_1 \cdots a_k$ where each a_i either is in $\bar{S}$ or is the inverse of an element of $\bar{S}$. But because the defining relations hold, in particular each element of $\bar{S}$ is its own inverse, so $a_1 \cdots a_k$ is a sequence of elements of $\bar{S}$. We call such a sequence a *word* in the alphabet $\bar{S}$. Second, given a word $a_1 \cdots a_k$, if we identify a sequence of adjacent entries in $a_1 \cdots a_k$ that is identical to some $\bar{s}\bar{s}$ or to some $\bar{s}\,\bar{t}\,\bar{s}\,\bar{t} \cdots$ with $2m(s,t)$ letters, then we can delete that subsequence without changing the product of the word. Third, if we insert the word $\bar{s}\bar{s}$ or the $2m(s,t)$-letter word $\bar{s}\,\bar{t}\,\bar{s}\,\bar{t} \cdots$ between two adjacent letters of a word $a_1 \cdots a_k$, then the product of the word is unchanged. Now suppose that the word $a_1 \cdots a_k$ is a *relation* in $\overline{W}$, meaning that its product is the identity. We say that the relation $a_1 \cdots a_k$ is a *consequence of the defining relations* if it can by transformed into the empty word by a sequence of insertions or deletions of $\bar{s}\bar{s}$ or $\bar{s}\,\bar{t}\,\bar{s}\,\bar{t} \cdots$ (with $2m(s,t)$ letters).

We can now define W more formally: A group $\overline{W}$ is isomorphic to W if and only if there is a map $s \mapsto \bar{s}$ from S to $\overline{W}$ such that the image of S generates $\overline{W}$, such that the defining relations of W hold in $\overline{W}$ with respect to $s \mapsto \bar{s}$, and such that every relation in $\overline{W}$ is a consequence of the defining relations.

Most formally, W is the quotient of the free group generated by S, modulo the smallest group congruence having s^2 congruent to 1 for all $s \in S$ and $(st)^{m(s,t)}$ congruent to 1 for all $s \neq t$ in S. In fact, the term "group congruence" is not typically used, although the term makes sense in the context of universal algebra. Rather, the usual phrasing is that W is the free group on S modulo the smallest normal subgroup containing s^2 for all $s \in S$ and containing $(st)^{m(s,t)}$ for all $s \neq t$ in S. The normal subgroup is the congruence class of 1, and this normal subgroup determines the congruence completely.

In general, groups presented by generators and relations are difficult to understand. For example, the Word Problem (deciding whether two words describe the same element) is known to be recursively unsolvable for some such groups, and the Finiteness Problem for such groups (deciding whether the group is finite) is known to be recursively unsolvable.

Some group presentations are more tractable than others, however, and Coxeter groups are a particularly well-behaved case. But *a priori*, we can't rule out certain "bad" behavior on the part of the Coxeter group W given by (10-1.1). For example, we can't show that a given generator s is not the identity element. That is, we can't rule out, *a priori*, the possibility that the relation $s = 1$ is not a consequence of the defining relations. Similarly, the relation $(st)^{m(s,t)} = 1$ implies that the order of the element st is a divisor of $m(s,t)$, but we can't rule out the possibility that the order of st is a *proper* divisor of $m(s,t)$. We can't even rule out the possibility that the order of st is 1, or in other words that $s = t$ in W for distinct elements $s, t \in S$. We rule out all these bad behaviors for finite Coxeter groups later in Proposition 10-2.17, using the geometry of hyperplane arrangements.

Given any group presented by generators and relations, there is a natural partial order called the *prefix order*. In the case of Coxeter groups, the prefix order is called the *weak order* and has particularly nice properties. Let W be the Coxeter group defined in (10-1.1). A finite sequence $a_1 \cdots a_k$ of generators in S whose product is $w \in W$ is called a *word for* w of length k. The *length* $\ell(w)$ of an element w is the minimal k such that there exists a word $a_1 \cdots a_k$ for w. A word of this minimal length is called a *reduced word*. Given v and w in W, we say $v \leq w$ in the weak order[2] if there exists a reduced word $a_1 \cdots a_k$ for w and an index i such that $a_1 \cdots a_i$ is a word for v. Informally, we might say that v is a *prefix* of w. Equivalently, $v \leq w$ if $\ell(v) + \ell(v^{-1}w) = \ell(w)$. The cover relations of the weak order are given by $w \prec ws$ for any $w \in W$ and $s \in S$ such that $\ell(ws) > \ell(w)$.

Example 10-1.1. When $S = \{s_1, s_2\}$ (i.e., when the rank of W is 2), the number of elements of W is $2m(s_1, s_2)$. The rank-two Coxeter groups for $m(s_1, s_2) \in \{2, 3, 4, 5\}$ are shown arranged in the weak order in Figure 10-1.1.

The first goal of this chapter is to show that the weak order on a finite Coxeter group is a lattice, and to establish some of its lattice-theoretic properties. We do this by showing that finite Coxeter groups are essentially the same

[2]This is the right weak order. There is also a left weak order, isomorphic but not identical, with left and right switched in the definition. Thus left weak order is the *postfix order*.

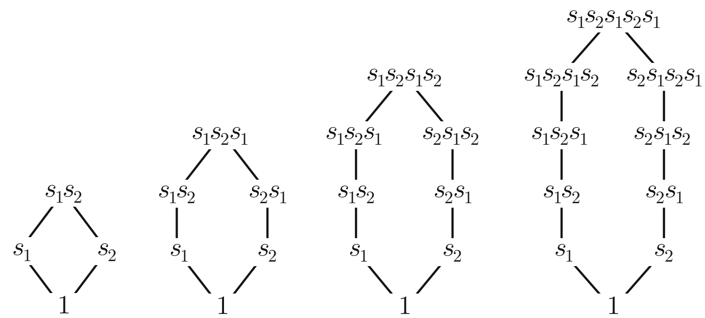

Figure 10-1.1: The weak order on some Coxeter groups of rank 2

as finite reflection groups (finite groups generated by reflections), showing that each finite reflection group defines a hyperplane arrangement, and showing that the poset of regions of that arrangement is isomorphic to the weak order on the Coxeter group. Similar results exist for infinite Coxeter groups, but to realize infinite Coxeter groups as reflection groups, one must leave behind the comfort of Euclidean geometry. An indication that finiteness and Euclidean geometry go hand-in-hand is found in Proposition 10-2.7.

10-2. Finite reflection groups

A Euclidean (orthogonal) reflection in $\mathbb{R}^n$ is a linear transformation that fixes an $(n-1)$-dimensional subspace H and negates vectors in the 1-dimensional subspace orthogonal to H. Orthogonality is defined in terms of the usual inner product on $\mathbb{R}^n$. A *(finite, real, Euclidean) reflection group* [3] is a finite group generated by Euclidean reflections in $\mathbb{R}^n$. In this section, we consider reflection groups and the hyperplane arrangements that arise from reflection groups. Specifically, let W be a reflection group and let T be the set of reflections in W. By hypothesis, the group W is generated by reflections, but T may be larger than the generating set because additional elements may act as reflections. For each reflection $t \in T$, let H_t be the hyperplane fixed by t (the *reflecting hyperplane* of t). Foreshadowing the connection to finite Coxeter groups, the arrangement $\mathcal{A} = \{H_t \mid t \in T\}$ is called a *Coxeter arrangement*. The group W is uniquely determined by $\mathcal{A}$.

[3] Other versions of reflection groups replace the real numbers with another field and/or relax the requirement of finiteness (thus necessarily deleting the adjective Euclidean). We will see in Proposition 10-2.7 that a slightly broader definition, keeping the adjectives "finite" and "real" but not requiring *a priori* that the reflections preserve a Euclidean metric, is essentially equivalent to this definition of finite reflection groups.

10-2.1 Coxeter arrangements are simplicial

The first goal of this section is to prove the following theorem, which will let us apply results from Chapter 9 on simplicial (or tight) arrangements to Coxeter arrangements.

Theorem 10-2.1. *Every essential Coxeter arrangement is simplicial.*

Recall that an essential hyperplane arrangement is an arrangement $\mathcal{A}$ such that $\bigcap_{H\in\mathcal{A}} H$ is the origin. We emphasize that for every non-essential Coxeter arrangement, an essential Coxeter arrangement can be obtained in the quotient vector space $\mathbb{R}^n/\bigcap_{H\in\mathcal{A}} H$ by taking the quotient modulo $\bigcap_{H\in\mathcal{A}} H$ of each hyperplane in $\mathcal{A}$. The quotient arrangement is a Coxeter arrangement for a finite reflection group in $\mathbb{R}^n/\bigcap_{H\in\mathcal{A}} H$ that is isomorphic to the original finite reflection group.

Of key importance in what follows will be the action of a reflection group W on the hyperplanes and regions of the Coxeter arrangement for W. We write wH for the image of a hyperplane H under the action of $w \in W$ and write wR for the image of a region R under the action of w.

Theorem 10-2.1 is immediate from the following two propositions.

Proposition 10-2.2. *Let $\mathcal{A}$ be a Coxeter arrangement defined by a finite reflection group W. Then W acts transitively on the set of regions.*

We will see soon, as an easy consequence of Theorem 10-2.5, that W acts *simply* transitively on regions.

Proof. The action of W permutes the hyperplanes of $\mathcal{A}$, because if $w \in W$ and $t \in T$, then wtw^{-1} is a reflection with $H_{wtw^{-1}} = wH_t$. (This is verified in Exercise 10.1.) Thus also W permutes the set of regions. Suppose Q and R are regions of $\mathcal{A}$. Lemma 9-1.12 constructs a sequence of regions $Q = R_0, \dots, R_k = R$ with R_{i-1} adjacent to R_i for $i = 1, \dots, k$. The common facet of each pair R_{i-1}, R_i is contained in some hyperplane of $\mathcal{A}$, which is a reflecting hyperplane for some reflection $t_i \in T$. Since W permutes the regions, we see that $R_i = t_i R_{i-1}$. Thus $R = t_k t_{k-1} \cdots t_1 Q$. □

Proposition 10-2.3. *Every essential central arrangement has at least one simplicial region.*

Proof. A *line* of $\mathcal{A}$ is a 1-dimensional linear subspace of $\mathbb{R}^n$ that is the intersection of some collection of hyperplanes in $\mathcal{A}$. We argue, by induction on n, an assertion that is stronger than the proposition: If $\mathcal{A}$ is essential and H_0 is a hyperplane containing no line of $\mathcal{A}$, then there exists a simplicial region R of $\mathcal{A}$ with $H_0 \cap R = \{\mathbf{0}\}$. (Such a hyperplane H_0 is necessarily not in $\mathcal{A}$.)

Let H_0 be a hyperplane containing no line of $\mathcal{A}$. Let v be a nonzero normal vector to H_0 and write H_1 for the set $v + H_0$. Since H_0 contains no line of $\mathcal{A}$, every line of $\mathcal{A}$ intersects H_1 in exactly one point.

Choose a hyperplane $H \in \mathcal{A}$. Exercise 10.2 verifies that there exists a line of $\mathcal{A}$ not contained in H. Among all lines of $\mathcal{A}$ not contained in H, choose a line ℓ of $\mathcal{A}$ to minimize the distance from the point $\mathbf{p} = \ell \cap H_1$ to the set $H \cap H_1$. Each hyperplane H' of $\mathcal{A}$ containing ℓ intersects H in a linear subspace of dimension $n-2$. Thus the set $\mathcal{A}' = \{H' \cap H \mid \ell \subseteq H' \in \mathcal{A}\}$ is a central hyperplane arrangement in $H \cong \mathbb{R}^{n-1}$. Furthermore, $\mathcal{A}'$ is essential because the intersection of all of its hyperplanes is $\left(\bigcap_{\ell \subseteq H' \in \mathcal{A}} H'\right) \cap H = \ell \cap H = \{\mathbf{0}\}$. Every line of $\mathcal{A}'$ is also a line of $\mathcal{A}$, so the hyperplane $H_0 \cap H$ contains no line of $\mathcal{A}'$.

By induction on n, there is a simplicial region R' of $\mathcal{A}'$ whose intersection with $H_0 \cap H$ is $\{\mathbf{0}\}$. Since $R' \subseteq H$, we know that $H_0 \cap R'$ is also $\{\mathbf{0}\}$. Thus (up to passing from R' to $-R'$) we can take R' to be on the same side of H_0 as H_1. Let R be the nonnegative linear span of $R' \cup \mathbf{p}$. This is a simplicial cone whose facet-defining hyperplanes are all in $\mathcal{A}$. (One of the facet-defining hyperplanes is H. The others are the hyperplanes H' containing ℓ such that $H' \cap H$ is a facet-defining hyperplane of R' as a subset of H.) Since R' is on on the same side of H_0 as H_1, also R is on the same side of H_0 as H_1.

We complete the proof by showing that R is a region of $\mathcal{A}$. Since the facet-defining hyperplanes of R are in $\mathcal{A}$, this amounts to showing that R is not the union of more than one region of $\mathcal{A}$. The cone R is the nonnegative linear span of a set of n vectors. We may as well take all of these vectors to lie in H_1, so that one of them is $\mathbf{p}$, and write $\mathbf{p}_1, \ldots, \mathbf{p}_{n-1}$ for the others. The region R' of $\mathcal{A}'$ is the nonnegative linear span of $\mathbf{p}_1, \ldots, \mathbf{p}_{n-1}$.

Now suppose for the sake of contradiction that R is the union of more than one region. Then there exists a hyperplane $\widetilde{H} \in \mathcal{A}$ with some of the points $\mathbf{p}, \mathbf{p}_1, \ldots, \mathbf{p}_{n-1}$ strictly on one side of H and some strictly on the other side. If $\mathbf{p}$ is contained in $\widetilde{H}$ and some pair of points $\mathbf{p}_i$ and $\mathbf{p}_j$ are on opposite sides of $\widetilde{H}$, then $\widetilde{H} \cap H$ is in $\mathcal{A}'$, and we obtain a contradiction to the fact that R' is a region of $\mathcal{A}'$. Otherwise, for some i from 1 to $n-1$, the points $\mathbf{p}$ and $\mathbf{p}_i$ are on opposite sides of $\widetilde{H}$. Now $\mathbf{p}$ is $\ell \cap H_1$ and also $\mathbf{p}_i$ is $\ell' \cap H_1$ for some line ℓ' of $\mathcal{A}'$. Since $\mathcal{A}' = \{H' \cap H \mid \ell \subseteq H' \in \mathcal{A}\}$, there is a set $\mathcal{A}''$ of hyperplanes, all containing ℓ, such that the intersection U of $\mathcal{A}''$ has $H \cap U = \ell'$. Thus U is the span of ℓ and ℓ'. But $U \cap \widetilde{H}$ is a line of $\mathcal{A}$, and the point $(U \cap \widetilde{H}) \cap H_1$ is in the interior of the line segment $\overline{\mathbf{p}\mathbf{p}_i}$. This point is strictly closer to $H \cap H_1$ than $\mathbf{p}$, contradicting our choice of $\mathbf{p}$. $\square$

We have proved Theorem 10-2.1. Recall that Theorem 9-3.15 and Corollary 9-3.16 define colorings of the rays (and more generally of the faces) of a simplicial arrangement. In the case of Coxeter arrangements, the colorings are compatible with the action of the associated reflection group.

Proposition 10-2.4. *If $\mathcal{A}$ is an essential Coxeter arrangement defined by a finite reflection group W, then the ray coloring of Theorem 9-3.15 (and more generally the face coloring of Corollary 9-3.16) is preserved by the action of W.*

Proof. It is enough to prove the statement about the ray coloring. Since W is generated by reflections, it is enough to show that each reflection preserves the ray coloring. Let t be a reflection and let R be a region. Since W permutes the regions, there is a region tR.

By Lemma 9-1.12, there is a sequence $R = R_0, \ldots, R_k = tR$ with R_{i-1} adjacent to R_i for $i = 1, \ldots, k$. Since R and tR are on opposite sides of H_t, there is some i such that R_0 and R_i are on the same side of H_t but R_{i+1} is on the opposite side of H_t. In particular, R_i shares with tR_i a facet contained in H_t. The rest of the proof considers only the subsequence $R = R_0, \ldots, R_i$. We show by induction on j, that t takes the ray coloring of R_{i-j} to the ray coloring of tR_{i-j}. The base case of the induction, where $j = 0$, is immediate because R_i and tR_i share a facet contained in H_t. If $j > 0$, then by induction, t takes the ray coloring of R_{i-j+1} to the ray coloring of tR_{i-j+1}. Since t also takes the unique ray of R_{i-j} not contained in R_{i-j+1} to the unique ray of tR_{i-j} not contained in tR_{i-j+1}, we conclude that t also takes the ray coloring of R_{i-j} to the ray coloring of tR_{i-j}. This fact, for $j = i$, completes the proof. □

Not only do reflection groups provide examples of simplicial arrangements, but also simplicial arrangements provide insight into reflection groups. This insight is more fully realized in Section 10-2.3. For now, we point out a fundamental result: the correspondence between regions and group elements.

Theorem 10-2.5. *Let $\mathcal{A}$ be a Coxeter arrangement defined by a finite reflection group W and let B be any region of $\mathcal{A}$. Then the map $w \mapsto wB$ is a bijection from W to regions of $\mathcal{A}$.*

Proof. It is enough to prove the proposition in the case where $\mathcal{A}$ is essential. The set wB is a region by Proposition 10-2.2, which furthermore implies that the map $w \mapsto wB$ is surjective. Suppose two elements v and w of W have $vB = wB = R$ for some region R. Proposition 10-2.4 implies that each ray of B is taken to the same ray of R by v and by w. Take a nonzero vector $\mathbf{x}_\rho$ in each ray ρ of B. If $v(\mathbf{x}_\rho) \neq w(\mathbf{x}_\rho)$, then $w^{-1}v(\mathbf{x}_\rho)$ is a positive, non-unit scalar multiple of $\mathbf{x}_\rho$. In particular, $w^{-1}v$ is of infinite order, contradicting the fact that W is finite. Therefore, $v(\mathbf{x}_\rho) = w(\mathbf{x}_\rho)$ for each ρ. Since the $\mathbf{x}_\rho$ are a basis for $\mathbb{R}^n$ and the maps v and w are linear, we conclude that $v = w$. □

10-2.2 Generalized reflection groups

Our main focus will be finite reflection groups in a Euclidean vector space. We pause briefly to consider more general reflections and in particular to show that we can safely relax the definition of finite reflection groups without losing the results of this section. Dropping the relationship with a Euclidean metric, we define a *generalized reflection* in $\mathbb{R}^n$ to be a linear map having an $(n-1)$-dimensional 1-eigenspace (i.e., having a fixed hyperplane) and an eigenvalue -1. A *generalized reflection group* is a group generated by generalized reflections, without any *a priori* requirement that the reflections preserve the Euclidean metric.

In dealing with generalized reflections, it is useful to introduce a symmetric bilinear form to take the place of the usual Euclidean inner product. This is a map from pairs of vectors in $\mathbb{R}^n$ to real numbers that is linear in each entry and symmetric in exchanging the vectors. We will continue to use $\langle\,\cdot\,,\,\cdot\,\rangle$ for the usual Euclidean inner product and refer to other symmetric bilinear forms by letters such as f. Vectors $\mathbf{x}$ and $\mathbf{y}$ are *orthogonal* with respect to the form f if $f(\mathbf{x},\mathbf{y}) = 0$. A linear transformation w is an *isometry* of the form f if $f(w\mathbf{x}, w\mathbf{y}) = f(\mathbf{x},\mathbf{y})$ for all vectors $\mathbf{x}$ and $\mathbf{y}$. A form is *preserved* by a group of transformations if every element of the group is an isometry of the form. The following basic facts are left as Exercise 10.3.

Proposition 10-2.6. *Let f be a symmetric bilinear form.*

(i) *If $f(\mathbf{x},\mathbf{x}) \neq 0$, then the map $r_{\mathbf{x},f}\colon \mathbb{R}^n \to \mathbb{R}^n$ with $r_{\mathbf{x},f}(\mathbf{y}) = \mathbf{y} - 2\frac{f(\mathbf{x},\mathbf{y})}{f(\mathbf{x},\mathbf{x})}\mathbf{x}$ is an isometry of f.*

(ii) *If $f(\mathbf{x},\mathbf{x}) \neq 0$, then $\mathbf{x}_f^{\perp} = \{\mathbf{y} \in \mathbb{R}^n \mid f(\mathbf{x},\mathbf{y}) = 0\}$ is a hyperplane.*

(iii) *If $f(\mathbf{x},\mathbf{x}) \neq 0$, then the map $r_{\mathbf{x},f}$ is a reflection with fixed hyperplane $\mathbf{x}_f^{\perp}$ and (-1)-eigenspace $\mathbb{R}\mathbf{x}$.*

(iv) *A reflection is an isometry of f if and only if its fixed space is orthogonal, with respect to f, to its (-1)-eigenspace.*

(v) *Suppose r is a reflection and $\mathbf{x}$ is an (-1)-eigenvector of r. If $f(\mathbf{x},\mathbf{x}) \neq 0$, then r is an isometry of f if and only if $r = r_{\mathbf{x},f}$.*

Given a symmetric bilinear form f, there is a matrix M such that $f(\mathbf{x},\mathbf{y}) = \mathbf{x}^T M\mathbf{y}$, where $\mathbf{x}$ and $\mathbf{y}$ are column vectors and $\mathbf{x}^T$ is the transpose of $\mathbf{x}$. The matrix M is real and symmetric, so it has a basis $\mathbf{v}_1,\ldots,\mathbf{v}_n$ of eigenvectors with real eigenvalues $\lambda_1,\ldots,\lambda_n$, and furthermore $\mathbf{v}_1,\ldots,\mathbf{v}_n$ can be taken to be an orthonormal basis with respect to the usual Euclidean inner product.[4] Orthogonality of the eigenvectors with respect to the usual Euclidean product implies orthogonality of the eigenvectors with respect to the form f, because $f(\mathbf{v}_i,\mathbf{v}_j) = \mathbf{v}_i^T M\mathbf{v}_j = \mathbf{v}_i^T\lambda_j\mathbf{v}_j = \lambda_j\langle\mathbf{v}_i,\mathbf{v}_j\rangle$.

The form f is called *Euclidean* or *positive definite* if $f(\mathbf{x},\mathbf{x}) > 0$ for all nonzero vectors $\mathbf{x}$. The form f is Euclidean if and only if all of the eigenvalues of M are positive, in which case we replace each $\mathbf{v}_i$ with $\mathbf{v}_i/\sqrt{\lambda_i}$ to obtain a basis on which f is described by the identity matrix. Thus a Euclidean form f is essentially the same as the usual Euclidean inner product; the only difference is the choice of basis.

[4]This standard result, or its generalization to Hermitian matrices, is known as the Spectral Theorem and is found in most linear algebra textbooks.

Given a symmetric bilinear form f, a *reflection group with respect to* f is a group generated by reflections $r_{\mathbf{x},f}$. The main point of our discussion of generalized reflection groups is the following proposition, which says that, up to a change of basis, every finite generalized reflection group is a (Euclidean) reflection group.

Proposition 10-2.7. *Any generalized finite reflection group W in $\mathbb{R}^n$ is a reflection group with respect to some Euclidean form f.*

For the proof, we use the following standard trick.

Proposition 10-2.8. *Given any finite group of linear transformations, there is a Euclidean symmetric bilinear form preserved by the group.*

Proof. Define a bilinear form f by $f(\mathbf{x},\mathbf{y}) = \sum_{w\in W} \langle w\mathbf{x}, w\mathbf{y}\rangle$. The form f is symmetric and Euclidean (positive definite) because the usual Euclidean inner product is. We see that f is preserved by W as well, because for any $u \in W$,

$$f(u\mathbf{x}, u\mathbf{y}) = \sum_{w\in W} \langle wu\mathbf{x}, wu\mathbf{y}\rangle = \sum_{v\in W} \langle v\mathbf{x}, v\mathbf{y}\rangle = f(\mathbf{x},\mathbf{y}). \qquad \square$$

Proof of Proposition 10-2.7. By Proposition 10-2.8, W preserves a Euclidean form f, and so W is generated by reflections that are isometries of f. $\square$

10-2.3 Finite Coxeter groups and finite reflection groups

In this section, we show that finite reflection groups and finite Coxeter groups are essentially the same (in particular, justifying the name "Coxeter arrangement" for the set of reflecting hyperplanes for a finite reflection group). Specifically, we prove the following theorems.

Theorem 10-2.9. *Let W be a finite reflection group with Coxeter arrangement $\mathcal{A}$. Choose any region B and let S be the set of reflections in facet-defining hyperplanes of B. For each pair $s \neq t$ in S, define $m(s,t)$ to be π divided by the angle between H_s and H_t. Then W is the Coxeter group with the presentation* (10-1.1).

Theorem 10-2.10. *Every finite Coxeter group is isomorphic to some finite reflection group. The isomorphism can be chosen so that it restricts to a bijection from the set S of defining generators of W to the set of reflections in the facet-defining hyperplanes of some region B of the associated Coxeter arrangement. Given $s \neq t$ in S, the angle between the reflecting hyperplanes corresponding to s and t is $\frac{\pi}{m(s,t)}$.*

Theorem 10-2.9 has been stated in a slightly informal way. In the reflection group, the elements of S are certain linear transformations. In the presentation (10-1.1), those same elements play the role of formal symbols that generate a free group. The quotient of that free group, modulo some normal subgroup N, is the Coxeter group. More formally, Theorem 10-2.9 asserts not that the two groups are identical, but rather that S generates the reflection group and that the map taking each $s \in S$ to the coset sN extends to an isomorphism from the reflection group to the Coxeter group.

We now prepare to prove Theorem 10-2.9. The key geometric insight is the following well known fact which is proved as Exercise 10.4.

Proposition 10-2.11. *Suppose s and t are Euclidean reflections whose reflecting hyperplanes H_s and H_t meet at an angle θ. Then the composition st is a rotation through an angle 2θ fixing the subspace $H_s \cap H_t$.*

The other key to the proof is the facet coloring of $\mathcal{A}$. Corollary 9-3.16 colors the facets of $\mathcal{A}$ with n distinct colors, each of which is an $(n-1)$-element subset of $\{1, \ldots, n\}$. We write $\langle i \rangle$ for the $(n-1)$-element subset $\{1, \ldots, n\} \setminus \{i\}$. Choose any base region B, and let s_i denote the reflection in the hyperplane defining the facet of B that is colored $\langle i \rangle$.

Lemma 10-2.12. *Let $w \in W$ and let $i \in \{1, \ldots, n\}$. Then $(ws_i)B$ is the region that shares a facet colored $\langle i \rangle$ with wB.*

Proof. The regions B and s_iB are adjacent, sharing a facet colored $\langle i \rangle$. Thus the regions wB and $w(s_iB) = (ws_i)B$ are adjacent, and Proposition 10-2.4 says that their common facet is colored $\langle i \rangle$. □

Lemma 9-1.12 says that any region R is connected to B by a sequence $B = R_0, \ldots, R_k = R$ with R_{i-1} adjacent to R_i for $i = 1, \ldots, k$. For each i from 1 to k, set a_i equal to the generator s_j such that $\langle j \rangle$ is the color of the facet shared by R_{i-1} and R_i. Iterating Lemma 10-2.12, we see that R is $(a_1 \cdots a_k)B$. Theorem 10-2.5 now implies that every element of W can be expressed as a product $a_1 \cdots a_k$ of reflections in facet-defining hyperplanes of B. We have established one part of the proof of Theorem 10-2.9, which we record as the following proposition.

Proposition 10-2.13. *Let W be a finite reflection group and let B be any region of the corresponding Coxeter arrangement. Then W is generated by the set S of reflections in the facet-defining hyperplanes of B.*

Another part of the proof of Theorem 10-2.9 is the following proposition:

Proposition 10-2.14. *For S and $m(s,t)$ as defined in Theorem 10-2.9, each $m(s,t)$ is an integer greater than or equal to 2. The relation $s^2 = 1$ holds for every $s \in S$ and the relation $(st)^{m(s,t)} = 1$ holds for every $s \neq t \in S$.*

Proof. Every $s \in S$ is a reflection, so s^2 is the identity. Let i and j be distinct numbers in $\{1, \ldots, n\}$, so that s_i and s_j are distinct elements of S. The intersection of the facets of B colored $\langle i \rangle$ and $\langle j \rangle$ is an $(n-2)$-dimensional face

F colored with the set $\{1,\ldots,n\} \setminus \{i,j\}$. The reflections s_i and s_j each fix F. Since W is finite, there is some smallest number k such that the alternating product $s_i s_j s_i s_j \cdots$ of k generators is the identity. Then k cannot be odd, because the determinant of a reflection is -1, so $k = 2m$ for some integer $m \geq 2$. By Theorem 10-2.5, k is the smallest number such that that alternating product, applied to B, returns B.

Lemma 10-2.12 implies that $B, s_iB, s_is_jB, \ldots, (s_is_j)^mB$ is a sequence of adjacent regions, each containing F, with the last term $(s_is_j)^mB$ equaling B. Each region in the sequence has two facets whose intersection is F, and the angle between these facets is the same for each region in the sequence. We conclude that each angle is $\frac{\pi}{m}$. But this angle is the angle between the reflecting hyperplanes H_{s_i} and H_{s_j}, so $m = m(s_i, s_j)$. $\square$

We now complete the proof of Theorem 10-2.9.

Proof of Theorem 10-2.9. It remains to show that every relation among the generators S in the reflection group W is a consequence of the defining relations of the Coxeter group. Let $a_1 \cdots a_k$ be such a relation. That is, the product $a_1 \cdots a_k$ is the identity, or equivalently, $(a_1 \cdots a_k)B = B$. Define $R_i = (a_1 \cdots a_i)B$ for each i from 1 to k, so that $B = R_0, R_1, \ldots, R_k = B$ is a sequence of adjacent regions by Lemma 10-2.12.

In the proof of Theorem 9-3.15, we showed inductively how to reduce $B = R_0, R_1, \ldots, R_k = B$ to the singleton sequence B by a series of moves of two types. The first type of move found an index i with $R_{i-1} = R_{i+1}$ and deleted R_i. This means that $a_i = a_{i+1}$, and the new sequence of regions corresponds to the word obtained from $a_1 \cdots a_k$ by deleting a_i and a_{i+1}.

The second type of move found an index with $R_{i-1} \neq R_{i+1}$. Let c be the color of the facet shared by R_{i-1} and R_i and let d be the color of the facet shared by R_i and R_{i+1}. Thus $a_i = s_c$ and $a_{i+1} = s_d$. Let F be the intersection of those two facets. The second type of move replaced R_{i-1}, R_i, R_{i+1} with a new subsequence, consisting of all of the regions containing F except R_i. Since F is colored by the set $\{1,\ldots,n\} \setminus \{c,d\}$, the corresponding change in the word $a_1 \cdots a_k$ is to replace a_ia_{i+1} with an alternating word $a_{i+1}a_i \cdots a_{i+1}a_i$. In light of Proposition 10-2.14 and Lemma 10-2.12, the alternating word has $2m(a_i, a_{i+1}) - 2$ entries. Replacing a_ia_{i+1} with the alternating word $a_{i+1}a_i \cdots a_{i+1}a_i$ can be carried out in three steps. First, insert a different alternating word $a_ia_{i+1} \cdots a_ia_{i+1}$ with $m(a_i, a_{i+1})$ elements between the original a_i and a_{i+1}, then remove a_ia_i near the beginning of the inserted word, and finally remove $a_{i+1}a_{i+1}$ near the end of the inserted word.

We have seen that both kinds of moves on the sequence of regions correspond to changes in the word $a_1 \cdots a_k$ that are immediate consequences of the defining relations of the Coxeter group. Since there is a sequence of such moves changing $a_1 \cdots a_k$ to the empty word, we conclude that the relation $a_1, \ldots, a_k$ is a consequence of the defining relations. $\square$

We now turn to the proof of Theorem 10-2.10. The difficult part of the proof (aside from Theorem 10-2.9, which is also used) is the following proposition.

Proposition 10-2.15. *Let B be a simplicial cone whose facet-defining hyperplanes are $H_1, \dots, H_n$. Suppose, for any distinct indices i and j from 1 to n, that the facets defined by H_i and H_j meet at an angle $\frac{\pi}{m(i,j)}$ for some integer $m(i,j) \geq 2$. If the group W generated by Euclidean reflections in $H_1, \dots, H_n$ is finite, then B is a region in the Coxeter arrangement $\mathcal{A}$ associated to W.*

Example 10-2.16. To clarify Proposition 10-2.15, consider a simple example. Let L_1 and L_2 be lines through the origin in $\mathbb{R}^2$ meeting at an angle $\frac{\pi}{3}$. The reflection group generated by Euclidean reflections in these lines has six elements, three of which are reflections. The third reflecting line is the unique line meeting both L_1 and L_2 at an angle $\frac{\pi}{3}$. The associated Coxeter arrangement defines six regions, each of which is a sector with angle $\frac{\pi}{3}$. A sector with angle $\frac{\pi}{3}$ defined by L_1 and L_2 satisfies the hypotheses of the proposition and is one of the regions in the Coxeter arrangement. A sector with angle $\frac{2\pi}{3}$ defined by L_1 and L_2 does not satisfy the hypotheses of the proposition and is the union of two regions in the Coxeter arrangement.

Proof of Proposition 10-2.15. For each i, let s_i be the reflection in the hyperplane H_i. Suppose the group W generated by the s_i is finite. Since B is a closed polyhedral cone defined by reflecting hyperplanes for reflections in W, it is a union of regions of $\mathcal{A}$. If B is not a single region of $\mathcal{A}$, then there exist regions Q and R contained in B, with Q adjacent to R. The reflection t fixing the common facet of Q and R has the property that $R = tQ$. This reflection is in W, by the definition of $\mathcal{A}$, and since W is generated by $\{s_1, \dots, s_n\}$, there is a word $a_1 \cdots a_k$ for t with each a_i in $\{s_1, \dots, s_n\}$.

We wish to make an argument similar to the proofs of Theorems 9-3.15 and 10-2.9, inductively applying a sequence of moves to $a_1 \cdots a_k$, with every move preserving the property that the product of the word is t. But we must approach the induction differently in this case: Since we are trying to *prove* that B is a region, we don't know *a priori* that words in $\{s_1, \dots, s_n\}$ correspond to sequences of adjacent region. In particular, the moves may not interact well with separating sets.

Choose a unit vector $\mathbf{v}_0$ in the relative interior of the common facet of Q and R. In particular, $\mathbf{v}_0$ is in the interior of B. For i from 1 to k, let B_i be the cone $(a_1 \cdots a_i)B$ and let $\mathbf{v}_i$ be $(a_1 \cdots a_i)\mathbf{v}_0$. Since $\mathbf{v}_0$ is in the interior of B, each $\mathbf{v}_i$ is in the interior of B_i. Let $d(a_1 \cdots a_k)$ be the maximum, over $i \in \{1, \dots, k\}$, of the quantity $1 - \langle \mathbf{v}_i, \mathbf{v}_0 \rangle$. We will say $d(a_1 \cdots a_k)$ is *attained at i* if $d(a_1 \cdots a_k) = 1 - \langle \mathbf{v}_i, \mathbf{v}_0 \rangle$. Given another word $a'_1 \cdots a'_\ell$, we say that $a'_1 \cdots a'_\ell$ is *closer to* $\mathbf{v}_0$ than $a_1 \cdots a_k$ if either $d(a'_1 \cdots a'_\ell) < d(a_1 \cdots a_k)$ or $d(a'_1 \cdots a'_\ell) = d(a_1 \cdots a_k)$ but the maximum is attained fewer times on $a'_1 \cdots a'_\ell$ than on $a_1 \cdots a_k$.

We now show that whenever $d(a_1 \cdots a_k) > 0$, there exists another word $a'_1 \cdots a'_\ell$ for t that is closer to $\mathbf{v}_0$ than $a_1 \cdots a_k$. Suppose $d(a_1 \cdots a_k)$ is attained at i. Since $d(a_1 \cdots a_k) > 0$ and $\mathbf{v}_k = (a_1 \cdots a_k)\mathbf{v}_0 = t\mathbf{v}_0 = \mathbf{v}_0$, we know that $i < k$. We consider two cases depending on whether $B_{i-1} = B_{i+1}$.

If $B_{i-1} = B_{i+1}$, then $a_i = a_{i+1}$ and we delete a_i and a_{i+1} from $a_1 \cdots a_k$. The shortened word has the same product as $a_1 \cdots a_k$, and it is closer to $\mathbf{v}_0$ than $a_1 \cdots a_k$.

If $B_{i-1} \neq B_{i+1}$, then consider the hyperplanes H, defining the common facet of B_{i-1} and B_i, and H', defining the common facet of B_i and B_{i+1}. Since B_i is the image of B under a Euclidean isometry, the angle between H and H' is $\frac{\pi}{m}$ for some integer m. By Proposition 10-2.11, the element $a_i a_{i+1}$ is a rotation, fixing $H \cap H'$, about an angle $\frac{2\pi}{m}$. In particular, $(a_i a_{i+1})^m$ is the identity, and since each reflection is its own inverse, $a_i a_{i+1}$ equals $(a_{i+1}a_i)^{m-1}$. Thus, we replace the letters $a_i a_{i+1}$ in $a_1 \cdots a_k$ with an alternating word $a_{i+1}a_i a_{i+1} a_i \cdots$ of length $2m-2$ to obtain a new word whose product is still t. We must verify that the new word is closer to $\mathbf{v}_0$.

The images of the vector $\mathbf{v}_i$ under the group generated by a_i and a_{i+1} form a convex polygon with $2n$ vertices in a plane orthogonal to $H \cap H'$. The neighbors of $\mathbf{v}_i$ in this polygon are $\mathbf{v}_{i-1}$ and $\mathbf{v}_{i+1}$. The vectors $\mathbf{v}_0$ and $a_i\mathbf{v}_0$ are related by the reflection a_i in the hyperplane defining the common facet of B and a_iB. Applying the transformation $a_1 \cdots a_{i-1}$ throughout, we see that $\mathbf{v}_{i-1} = (a_1 \cdots a_{i-1})\mathbf{v}_0$ is related to $\mathbf{v}_i = (a_1 \cdots a_i)\mathbf{v}_0$ by reflection in H. In particular, the vector $\mathbf{v}_i - \mathbf{v}_{i-1}$ is orthogonal to H. The reflection fixing H is in W (it can be written $a_1 a_2 \cdots a_i \cdots a_2 a_1$), so H is in $\mathcal{A}$.

If $H \neq H_t$, then since $\mathbf{v}_0$ is in the relative interior of the common facet of Q and R and H_t is the hyperplane containing that facet, $\mathbf{v}_0$ is not in H, and thus $\mathbf{v}_0$ is not orthogonal to $\mathbf{v}_i - \mathbf{v}_{i-1}$. That is, $\langle \mathbf{v}_i - \mathbf{v}_{i-1}, \mathbf{v}_0 \rangle \neq 0$, and therefore $\langle \mathbf{v}_{i-1}, \mathbf{v}_0 \rangle$ and $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle$ are not equal except when $H = H_t$. Similarly, $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle$ and $\langle \mathbf{v}_{i+1}, \mathbf{v}_0 \rangle$ are not equal when $H' \neq H_t$. Since $H \neq H'$, the quantities $1 - \langle \mathbf{v}_{i-1}, \mathbf{v}_0 \rangle$ and $1 - \langle \mathbf{v}_{i+1}, \mathbf{v}_0 \rangle$ are not both equal to $1 - \langle \mathbf{v}_i, \mathbf{v}_0 \rangle$, and so at least one is strictly less. Therefore, because $\langle \cdot, \mathbf{v}_0 \rangle$ is a linear map, the maximum of the function $1 - \langle \cdot, \mathbf{v}_0 \rangle$ on the polygon is attained either only at the vertex $\mathbf{v}_i$ or only on an edge of the polygon incident to $\mathbf{v}_i$. The effect of replacing $a_i a_{i+1}$ with $a_{i+1} a_i a_{i+1} a_i \cdots$ is to replace $\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}$ with the other sequence of adjacent vertices of the polygon that starts at $\mathbf{v}_{i-1}$ and goes to $\mathbf{v}_{i+1}$. This replacement removes a vertex where the maximum $d(a_1 \cdots a_k)$ is attained and inserts a number of vertices $\mathbf{v}'$ with $1 - \langle \mathbf{v}', \mathbf{v}_0 \rangle < d(a_1 \cdots a_k)$. We conclude that the new word is closer to $\mathbf{v}_0$ than $a_1 \cdots a_k$.

In either case, we have replaced $a_1 \cdots a_k$ with a word strictly closer to $\mathbf{v}_0$. Each $\mathbf{v}'$ in the W-orbit of $\mathbf{v}_0$ is a unit vector, so $\langle \mathbf{v}', \mathbf{v}_0 \rangle \leq 1$, with equality if and only if $\mathbf{v}' = \mathbf{v}_0$. There are finitely many vectors in the orbit, so there are finitely many inner products. Thus if we continue to find words strictly closer to $\mathbf{v}_0$, we eventually find a word $a'_1 \cdots a'_\ell$ with $d(a'_1 \cdots a'_\ell) = 0$. In that case, $\mathbf{v}_0$ is fixed by every element $a'_1 \cdots a'_i$ for $1 \leq i \leq \ell$. Since $\mathbf{v}_0$ is in the relative interior of the common facet of Q and R, Theorem 10-2.5 implies that an element fixing $\mathbf{v}_0$ either fixes Q or maps Q to R, and thus that every element fixing $\mathbf{v}_0$ is either the identity or the reflection t. In particular, either $\ell = 0$ or $a_1' = t$. If $\ell = 0$, then $Q = R$, and if $a_1' = t$, then since a_1' is a reflection in one of the facet-defining hyperplanes of B, the regions Q and R are not both in B. In either case, we have a contradiction. We conclude that B contains only one region of $\mathcal{A}$, which region is therefore B. □

Proof of Theorem 10-2.10. Let W be the group presented by (10-1.1), and suppose W is finite. Let the elements of S be $s_1, \ldots, s_n$ and let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be a basis for $\mathbb{R}^n$. We define a symmetric bilinear form g on $\mathbb{R}^n$ by setting $g(\mathbf{v}_i, \mathbf{v}_i) = 1$ for all i and $g(\mathbf{v}_i, \mathbf{v}_j) = -\cos\left(\frac{\pi}{m(s_i,s_j)}\right)$ for all $i \neq j$. Define r_i to be the reflection $r_{\mathbf{v}_i,g}$ orthogonal to $\mathbf{v}_i$ with respect to g. Let W' be the group of transformations generated by $r_1, \ldots, r_n$.

We now show that $(r_i r_j)^{m(s_i,s_j)} = 1$, for all $i \neq j$. The element $r_i r_j$ fixes the $(n-2)$-dimensional intersection of the fixed space of r_i and the fixed space of r_j. This fixed space, together with the vectors $\mathbf{v}_i$ and $\mathbf{v}_j$, span $\mathbb{R}^n$, so it is enough to show that $(r_i r_j)^{m(s_i,s_j)}$ also fixes the plane P_{ij} spanned by the vectors $\mathbf{v}_i$ and $\mathbf{v}_j$. Writing $c = \cos\frac{\pi}{m(s_i,s_j)}$ and $s = \sin\frac{\pi}{m(s_i,s_j)}$, the vector $\mathbf{x} = \frac{1}{s}(c\mathbf{v}_i + \mathbf{v}_j)$ is in P_{ij} and is orthogonal to $\mathbf{v}_i$ with respect to g. Using Proposition 10-2.6, we see that $r_i r_j$ fixes P_{ij} as a set, and we compute the matrix for the action of $r_i r_j$ on P_{ij}, in terms of the basis $\{\mathbf{v}_i, \mathbf{x}\}$ to be $\begin{bmatrix} 2c^2-1 & -2cs \\ 2cs & 2c^2-1 \end{bmatrix}$. Using double-angle formulas, we see that this is the rotation matrix $\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ for rotation about an angle $\theta = \frac{2\pi}{m(s_i,s_j)}$. In particular, $(r_i r_j)^{m(s_i,s_j)}$ fixes P_{ij} pointwise.

We have verified that W' satisfies the defining relations of the Coxeter group W. In particular, W' is a homomorphic image of W, which is finite. Thus W' is finite, and Proposition 10-2.8 says that there exists a Euclidean form f preserved by W'. In particular, the reflections r_i are isometries of f, and therefore of the form $r_{\mathbf{x}_i,f}$ for some vectors $\mathbf{x}_i$ by Proposition 10-2.6(v). Since the r_i generate W', it is a reflection group with respect to f. Thus, up to a change of basis, we can apply all of our results about finite reflection groups to W'.

By Proposition 10-2.11, each $r_i r_j$ acts by a rotation about some angle with respect to f. Since $r_i r_j$ acts on P_{ij} by rotation about an angle $\frac{2\pi}{m(s_i,s_j)}$ with respect to g, it must act by rotation about the same angle with respect to f. By Proposition 10-2.11 again, the reflecting hyperplanes for r_i and r_j meet at the angle $\frac{\pi}{m(s_i,s_j)}$. Choose B to be a closed polyhedral cone defined by the reflecting hyperplanes for $r_1, \ldots, r_n$ such that the facets of B meet at internal angles $\frac{\pi}{m(s_i,s_j)}$ (rather than $\pi - \frac{\pi}{m(s_i,s_j)}$). Proposition 10-2.15 says that B is a region in the Coxeter arrangement for W'. Theorem 10-2.9 now implies that W' is isomorphic to W. □

As a result of Theorem 10-2.10, we are able to rule out the conceivable bad behaviors of Coxeter groups that we mentioned in Section 10-1. With the tools we have developed here, we can only prove the following proposition for finite Coxeter groups, but it is true in general.

Proposition 10-2.17. *Let W be a finite Coxeter group given by the presentation* (10-1.1). *Then no element of S is the identity element in W. Furthermore, for distinct generators $s, t \in S$, the order of the element st is $m(s,t)$, and in particular s and t are distinct elements of W.*

Proof. By Theorem 10-2.10, we identify W with a reflection group W' and identify S with the set of reflections in some region B of the Coxeter arrangement $\mathcal{A}$ associated to W'. Let s and t be distinct elements of S. Lemma 10-2.12, with $w = 1$, says in particular that $sB \neq B$, so Theorem 10-2.5 implies that $s \neq 1$. Since Theorem 10-2.10 gives the angle between the reflecting hyperplanes H_s and H_t as $\frac{\pi}{m(s,t)}$, the composition st is a rotation through an angle $\frac{2\pi}{m(s,t)}$ by Proposition 10-2.11. Thus the order of st is $m(s,t)$. Since by definition $m(s,t) > 1$, we have $st \neq 1$, so that $s \neq t$. □

10-2.4 The classification of finite Coxeter groups

The ideas used in Section 10-2.3 are also central to the problem of classifying finite Coxeter groups (and the related Cartan-Killing classification in Lie theory). We state two theorems that we do not prove here. The first is the main idea behind the classification.

◇ **Theorem 10-2.18.** *Let W be the Coxeter group presented by* (10-1.1). *Let M be the matrix $[m_{ij}]$ with $m_{ii} = 1$ for all i and $m_{ij} = -\cos\left(\frac{\pi}{m(s_i,s_j)}\right)$ for all $i \neq j$. Then W is finite if and only if M is positive definite.*

A symmetric matrix M is positive definite if and only if the associated form $(\mathbf{x}, \mathbf{y}) \mapsto \mathbf{x}^T M \mathbf{y}$ is positive definite. Equivalently, the eigenvalues of M are all positive. The form associated to the matrix M from Theorem 10-2.18 appeared under the name g in the proof of Theorem 10-2.10. Thus Theorem 10-2.18 implies that g was already Euclidean.

To state the actual classification result, it is useful to introduce a standard combinatorial shorthand for describing Coxeter groups. The *Coxeter diagram* of a Coxeter group W is a graph whose vertices are the defining generators S, with an edge connecting vertices r and s if and only if $m(r,s) \geq 3$. If $m(r,s) = 3$, then the edge is left unlabeled, and if $m(r,s) > 3$ then the edge is labeled by $m(r,s)$. A Coxeter group is *irreducible* if its diagram is a connected graph. Every Coxeter group is a direct product of irreducible Coxeter groups. (See Exercise 10.6.) The following theorem completely classifies finite Coxeter groups (and thus in light of Theorems 10-2.9 and 10-2.10, finite reflection groups).

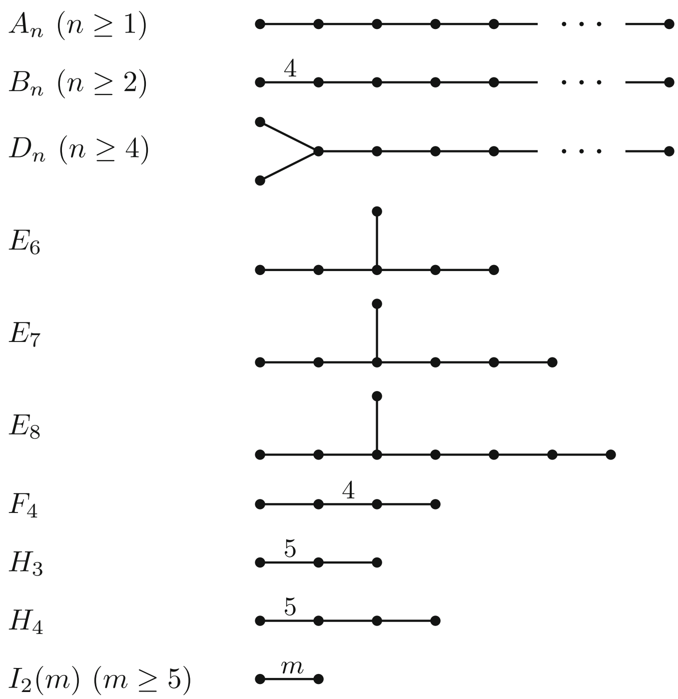

Figure 10-2.1: Coxeter diagrams of finite irreducible Coxeter groups

♢ **Theorem 10-2.19.** *An irreducible Coxeter group W is finite if and only if its diagram is on the list shown in Figure 10-2.1.*

The subscripts in Figure 10-2.1 show the rank of W (the size of S). The labels to the left of the diagrams are called the *types* of Coxeter groups. The finite irreducible Coxeter groups consist of four infinite families A_n, B_n, D_n, and $I_2(m)$ and 6 exceptional groups. Some of these are familiar. For example, the Coxeter group of type A_n is the symmetric group of permutations of $n+1$ symbols, which is isomorphic to the symmetry group of a regular n-dimensional simplex. The group of type B_n is the symmetry group of an n-dimensional cube, and the groups of type F_4, H_3, H_4 and $I_2(m)$ are also the symmetry groups of regular polytopes. All of the finite Coxeter groups except H_3, H_4 and most cases of $I_2(m)$ appear as Weyl groups of semi-simple Lie groups/Lie algebras. The group $I_2(6)$ is a Weyl group, named G_2 in the Cartan-Killing classification of root systems. There is no Coxeter group called C_n because the root systems of types B_n and C_n both define Coxeter groups of type B_n.

Example 10-2.20. As an example of Theorem 10-2.19, we list all of the finite Coxeter groups W of rank 3. If the diagram of W has three connected components, then W is of type $A_1 \times A_1 \times A_1$. If the diagram has two connected components, then W is of type $I_2(m) \times A_1$ for any $m \geq 3$. (If $m = 3$ or 4 then $I_2(m)$ is referred to as A_2 or B_2 in the classification.) If W is irreducible, then

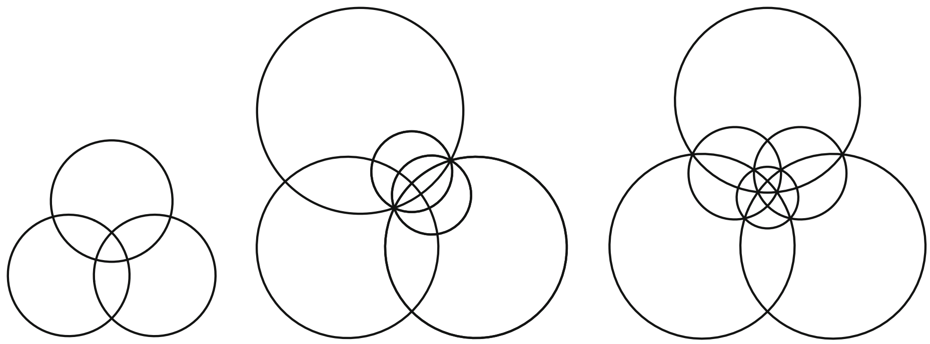

$A_1 \times A_1 \times A_1$ $B_2 \times A_1$ A_3

Figure 10-2.2: Some rank-3 Coxeter arrangements

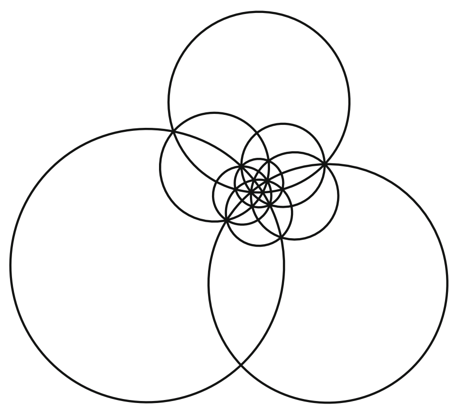

B_3

Figure 10-2.3: Another rank-3 Coxeter arrangement

W is of type A_3, B_3, or H_3. The rank-3 Coxeter arrangements are shown in Figures 10-2.2, 10-2.3, and 10-2.4. (We have chosen $I_2(4) \times A_1$ (i.e., $B_2 \times A_1$) to represent the infinite family $I_2(m) \times A_1$.)

10-2.5 Detecting Coxeter arrangements combinatorially

Corollary 9-3.16 and Theorem 10-2.10 provide a way to detect whether a given hyperplane arrangement is combinatorially isomorphic to a Coxeter arrangement, in the sense of Definition 9-3.18.

Theorem 10-2.21. *An essential hyperplane arrangement is combinatorially isomorphic to a Coxeter arrangement if and only if it is simplicial and the*

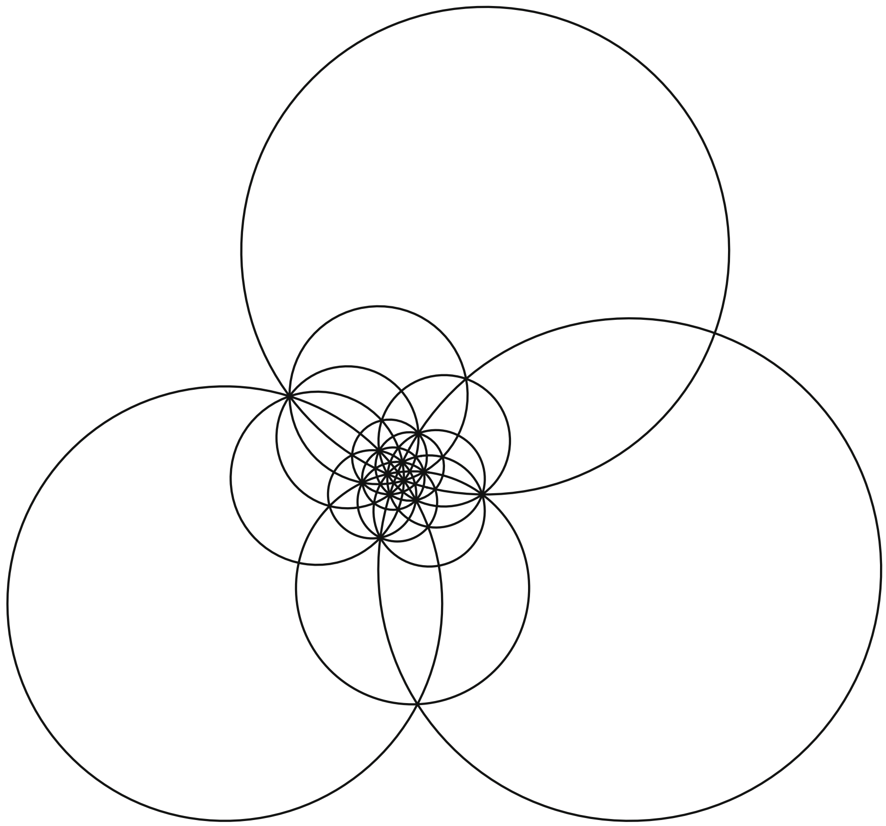

H_3

Figure 10-2.4: One more rank-3 Coxeter arrangement

number of regions containing a given codimension-2 face F depends only on the color of F in the coloring defined in Corollary 9-3.16.

Proof. We first prove the easy direction. Suppose $\mathcal{A}$ is a Coxeter arrangement with base region B and consider any codimension-2 face F of $\mathcal{A}$. Then there exists a region R having F as a face, and Theorem 10-2.5 says that $R = wB$ for some $w \in W$. Proposition 10-2.4 implies that F is the image under w of the face F' of B having the same color as F. Since w is an isometry mapping regions to regions, we see that F is contained in the same number of regions as F'. We conclude that the number of regions containing a given codimension-2 face F depends only of the color of F. Also $\mathcal{A}$ is simplicial by Theorem 10-2.1. These properties are therefore also true of any arrangement $\mathcal{A}'$ that is combinatorially isomorphic to $\mathcal{A}$.

Conversely, suppose that $\mathcal{A}$ is simplicial and that for every codimension-2 face F of $\mathcal{A}$, the number of regions of $\mathcal{A}$ containing F depends only on the color of F. Choose a base region B of $\mathcal{A}$. We will define a group W whose elements are the regions of $\mathcal{A}$ and whose product is defined as follows: Given a region Q, Lemma 9-1.12 says that there is a sequence of regions $B = Q_0, \ldots, Q_k = Q$ with Q_{i-1} adjacent to Q_i for $i = 1, \ldots, k$, not crossing any hyperplane twice. This defines a sequence $c_1, \ldots, c_k$ of colors, where c_i is the color of the facet shared by Q_{i-1} and Q_i. (Each c_j is of the form $\langle i \rangle$ for some $i \in \{i, \ldots, n\}$.) Similarly, given another region R, we have a sequence of regions $B = R_0, \ldots, R_\ell = R$, not crossing any hyperplane twice, which defines a sequence of colors $d_1, \ldots, d_\ell$. Concatenating the color sequences, we obtain a sequence $c_1, \cdots, c_k, d_1, \ldots, d_\ell$ which in turn defines a sequence of regions starting at B. We define the product QR to the be the region at the end of the sequence given by the colors $c_1, \cdots, c_k, d_1, \ldots, d_\ell$.

To see that QR is well-defined, note first that it is formally independent of the path chosen from B to Q. Since the sequence $B = R_0, \ldots, R_\ell = R$ does not cross any hyperplane twice, it is a maximal chain in the interval $[B, R]$ in $\mathrm{Pos}(\mathcal{A}, B)$. Lemma 9-6.12 says that any two such choices of path from B to R are related by a sequence of rank-two moves. Each rank-two move involves all of the regions containing some codimension-2 face of F of $\mathcal{A}$. If there are $2k$ regions containing F, then the rank-two move alters the sequence of colors by changing some subsequence $ababab\ldots$ of length to a sequence $bababa\ldots$ of length k, where a and b are facet-colors. These alterations to the sequence $d_1, \ldots, d_\ell$ do not change the region at the end of the sequence $c_1, \cdots, c_k, d_1, \ldots, d_\ell$, because for every codimension-2 face F of $\mathcal{A}$, the number of regions of $\mathcal{A}$ containing F depends only on the color of F. Thus QR is well-defined.

The product is associative because concatenation is associative. The element B is the identity element. Given R, we choose a sequence of adjacent regions from B to R, interpret it as a sequence of colors and then reverse the sequence of colors. The reversed sequence defines a path from B to some region R', and R' is the inverse of R for this product. Thus we have defined a group, which we call W. The regions adjacent to B correspond to the color sequences of length 1, so we identify these regions with the set of facet-colors, which we call S. The group W is generated by S, and we now show that W is a Coxeter group.

First, each element of S is an involution. If r and s are elements of S, let F be the intersection of the facets of B colored r and s and define $m(r, s)$ to be half the number of regions containing F. In particular, $(rs)^{m(r,s)}$ is the identity in W. We have showed that W satisfies the defining relations of the Coxeter group with generators S and these choices of $m(r, s)$.

The argument that every relation in W is a consequence of the defining relations of the Coxeter group now proceeds exactly as in the proof of Theorem 10-2.9, except that instead of appealing to Proposition 10-2.14 and Lemma 10-2.12 to know that the alternating sequence inserted has the right number of entries, we appeal to the hypothesis that for every codimension-2 face F of

$\mathcal{A}$, the number of regions of $\mathcal{A}$ containing F depends only on the color of F.

By Theorem 10-2.10, there is a Coxeter arrangement $\mathcal{A}'$ for a finite reflection group isomorphic to W, with the isomorphism taking S to the set of reflections in the facet-defining hyperplanes of some region B' of $\mathcal{A}'$. Recall that S was defined as a set of facet-colors of $\mathcal{A}$. Assign the corresponding colors to the facets of B'. Each facet color is $\langle i \rangle$, so we assign the color i to the ray of B opposite that facet. Complete this to a coloring of the rays, and then the faces, of $\mathcal{A}$ as in Theorem 9-3.15 and Corollary 9-3.16. Since the elements of W are the regions of $\mathcal{A}$, Theorem 10-2.5 says that the elements of W are in bijection with the regions of $\mathcal{A}'$. Given two adjacent regions Q and R of $\mathcal{A}$, let s be the color of their common facet. The corresponding elements of W are related by multiplication on the right by s. If Q' and R' are the regions of $\mathcal{A}'$ corresponding to Q and R, then Lemma 10-2.12 says that Q' and R' are adjacent and that s is the color of their common facet. The argument reverses, and we see that the adjacency graphs of $\mathcal{A}$ and $\mathcal{A}'$ are isomorphic, by an isomorphism that preserves colors of edges. Proposition 9-3.19 now says that $\mathcal{A}$ and $\mathcal{A}'$ are combinatorially isomorphic. $\square$

10-3. The weak order and the poset of regions

Having established that every finite Coxeter group is isomorphic to some finite reflection group, and *vice versa*, we now also make the connection between the weak order on the Coxeter group and the poset of regions of the associated Coxeter arrangement. In the process, we derive some well known combinatorial facts about reduced words in Coxeter groups. The latter facts are necessarily stated and proved here only for finite Coxeter groups, but hold more generally for arbitrary Coxeter groups.

In what follows, when we consider a finite Coxeter group W with defining generators S, we will assume that a specific representation has been chosen for W as a finite reflection group and moreover we will identify W with that reflection group. We will write $\mathcal{A}$ for the associated Coxeter arrangement, and we will assume that a base region B has been chosen such that the defining generators S are identified with the reflections in the facet-defining hyperplanes of B. Theorem 10-2.10 validates these assumptions. We will refer to $\mathcal{A}$ as the Coxeter arrangement for W and B as the base region for W.

10-3.1 The isomorphism

Theorem 10-3.1. *Suppose W is a finite Coxeter group with Coxeter arrangement $\mathcal{A}$ and base region B. Then the map $w \mapsto wB$ is an isomorphism from the weak order on W to the poset of regions* $\mathrm{Pos}(\mathcal{A}, B)$.

Example 10-3.2. Theorem 10-3.1 allows us to continue Example 10-2.20 by showing the weak order on each Coxeter group of rank 3. These are shown in Figures 10-3.1 and 10-3.2. One can also construct these weak orders directly

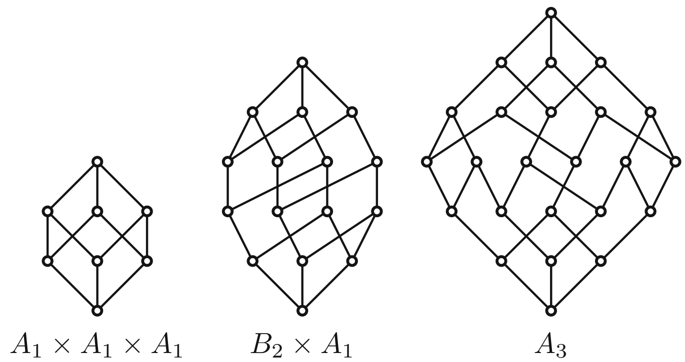

Figure 10-3.1: The weak order on some rank-3 Coxeter groups

using the combinatorics of reduced words, but that will be easier once we prove Theorem 10-4.1. (See Exercise 10.22.)

We now prepare to prove Theorem 10-3.1. Earlier, we defined T to be the set of reflections in the reflection group W. We now show that T is $\{wsw^{-1} \mid w \in W,\, s \in S\}$.

Proposition 10-3.3. *An element of W is a reflection if and only if it is conjugate to an element of S.*

Proof. If $t = wsw^{-1}$ for some $s \in S$ and $w \in W$, then t is a reflection with $H_t = wH_s$, as verified in Exercise 10.1. Conversely, suppose t is a reflection in W. Find two regions Q and R of $\mathcal{A}$ sharing a facet defined by H_t. Then $Q = wB$ for some $w \in W$. The region $w^{-1}R$ is adjacent to B, and thus it is sB for some $s \in S$. Therefore $R = wsB$. But also $R = tQ = twB$, so Theorem 10-2.5 implies that $ws = tw$, so that $t = wsw^{-1}$. $\square$

A reflection $t \in T$ is an *inversion* [5] of an element $w \in W$ if $\ell(tw) < \ell(w)$. Write $\mathrm{inv}(w)$ for the set of inversions of w. Given a word $a_1 \cdots a_k$, each index $i \in \{1, \ldots, l\}$ defines a *left reflection* t_i given by the palindrome $a_1 \cdots a_i \cdots a_1$. (Since t_i is conjugate to $a_i \in S$, it is a reflection by Proposition 10-3.3.) Write $T(a_1 \cdots a_k)$ for the set $\{t_i \mid i = 1 \ldots k\}$ of left reflections of $a_1 \cdots a_k$. The significance of left reflections derives from the following observation: Writing $R_i = a_1 \cdots a_i B$ for each i from 0 to k, the left reflection $t_i = a_1 \cdots a_i \cdots a_1$ is the reflection taking R_{i-1} to R_i. Equivalently, the reflecting hyperplane for t_i defines the common facet of R_{i-1} and R_i. The observation, which becomes clear upon glancing back at the proof of Proposition 10-3.3, leads to the proof of the following proposition.

[5] These might be called *left inversions*. Just as there is a left weak order, there are also right inversions.

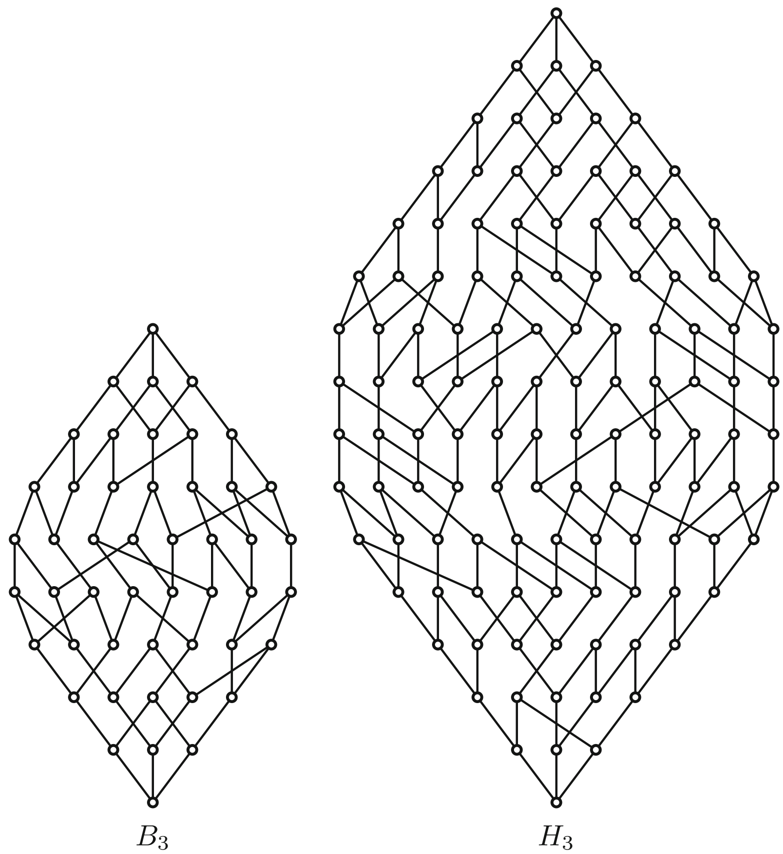

Figure 10-3.2: The weak order on some other rank-3 Coxeter groups

Proposition 10-3.4. *Let* $a_1 \cdots a_k$ *be a word for* w. *Then* $a_1 \cdots a_k$ *is reduced if and only if it has* k *distinct left reflections, in which case* $\mathrm{inv}(w) = T(a_1 \cdots a_k)$.

Proof. Suppose first that for some i and j with $1 \le i < j \le n$, the left reflections t_i and t_j coincide. Then

$$w = t_i t_j w = (a_1 \cdots a_i \cdots a_1)(a_1 \cdots a_j \cdots a_1)(a_1 \cdots a_k).$$

By repeatedly deleting pairs of coinciding adjacent letters, we can transform this expression for w into a word for w with $k-2$ letters, namely the word obtained from $a_1 \cdots a_k$ by deleting a_i and a_j. We conclude that $a_1 \cdots a_k$ is not reduced. Thus if $a_1 \cdots a_k$ is reduced then it has k distinct left reflections.

On the other hand, suppose that the word $a_1 \cdots a_k$ has k distinct left reflections $t_1, \ldots, t_k$. Write $R_i = a_1 \cdots a_i B$ for each i from 0 to k. Then there are k distinct hyperplanes H_{t_i}, each separating R_{i-1} from R_i. Moving from $B = R_0$ to R_1 and so on to $wB = R_k$, we cross each of these hyperplanes exactly once, and cross no other hyperplanes $\mathcal{A}$. Thus $\{H_t \mid t \in T(a_1 \cdots a_k)\}$ is the separating set of wB. If $a_1 \cdots a_k$ is not reduced, then there is a strictly shorter word $a_1' \cdots a_m'$ for w. Writing $R_i' = a_1' \cdots a_i' B$ for each i from 0 to m, we see that as we move from $B = R_0'$ to $wB = R_m'$, we cross only $m < k$ hyperplanes in $\mathcal{A}$. This contradicts what we just established, that $S(wB)$ contains k distinct hyperplanes, so we conclude that $a_1 \cdots a_k$ is reduced.

For each $t_i = a_1 \cdots a_i \cdots a_1 \in T(a_1 \cdots a_k)$, we have

$$t_i w = (a_1 \cdots a_i \cdots a_1)(a_1 \cdots a_k).$$

Thus a word for $t_i w$ is obtained by deleting the letter a_i from $a_1 \cdots a_k$. Since $\ell(w) = k$, we conclude that $\ell(t_i w) < \ell(w)$, so that $t_i \in \operatorname{inv}(w)$. On the other hand, suppose $t \in \operatorname{inv}(w)$. Since wB and twB are on opposite sides of H_t, either $H_t \in S(wB)$ or $H_t \in S(twB)$. If $H_t \in S(twB)$, then write a reduced word $a_1' \cdots a_m'$ for tw. As we argued above, $a_1' \cdots a_m'$ has m distinct reflections, so as we also argued above, $\{H_u \mid u \in T(a_1' \cdots a_m')\}$ is the separating set of twB. In particular, t equals some $t_i' = a_1' \cdots a_i' \cdots a_1' \in T(a_1' \cdots a_m')$. But then, arguing as in the beginning of this paragraph, we conclude that $\ell(ttw) < \ell(tw)$, or in other words $\ell(tw) > \ell(w)$, contradicting the supposition that t is an inversion of w. We conclude that $H_t \in S(wB)$. The latter equals $\{H_t \mid t \in T(a_1 \cdots a_k)\}$, so $t \in T(a_1 \cdots a_k)$. □

The proof of Proposition 10-3.4 also established the following statement.

Proposition 10-3.5. *If $w \in W$, then $S(wB) = \{H_t \mid t \in \operatorname{inv}(w)\}$.*

Proposition 10-3.5 enables us to prove Theorem 10-3.1 and a useful corollary.

Proof of Theorem 10-3.1. The map $w \mapsto wB$ is a bijection by Proposition 10-2.4. The cover relations in $\operatorname{Pos}(\mathcal{A}, B)$ are $Q \prec R$ if and only if Q and R are adjacent and $|S(R)| > |S(Q)|$. The cover relations in the weak order on W are $w \prec ws$ for $w \in W$ and $s \in S$ such that $\ell(ws) > \ell(w)$. We may as well assume $\mathcal{A}$ is essential, so that it is simplicial by Theorem 10-2.1. Thus there are n regions adjacent to a given region Q, and there are n elements ws for every element $w \in W$. Lemma 10-2.12 implies that the ordered pairs (w, ws) of elements of W are in bijection with the ordered pairs of adjacent regions of $\mathcal{A}$. Proposition 10-3.5 implies in particular that the conditions $\ell(ws) > \ell(w)$ and $|S(wsB)| > |S(wB)|$ are equivalent for any $w \in W$ and $s \in S$. □

Corollary 10-3.6. *Suppose W is a finite Coxeter group and let v and w be elements of W. Then $v \leq w$ in the weak order if and only if $\operatorname{inv}(v) \subseteq \operatorname{inv}(w)$.*

Proof. By Theorem 10-3.1, $v \leq w$ if and only if $S(vB) \subseteq S(wB)$. By Proposition 10-3.5, $S(vB) = \{H_t \mid t \in \operatorname{inv}(v)\}$ and $S(wB) = \{H_t \mid t \in \operatorname{inv}(w)\}$. □

10-3.2 Properties of the weak order

Theorem 10-3.1 let us establish some key lattice-theoretic properties of the weak order.

Theorem 10-3.7. *Let W be a finite Coxeter group. The weak order on W is a semidistributive, congruence uniform, polygonal lattice.*

Proof. All of the assertions about the weak order on W are immediate from Theorems 10-2.1 and 10-3.1, Proposition 9-3.3 and Theorems 9-3.8, and 9-6.10, except for the assertion that the weak order is congruence uniform. To obtain that assertion, by Corollary 9-7.22, we must verify that the shard digraph of $(\mathcal{A}, B)$ is acyclic. The key point will be Proposition 10-2.2, which says in particular that any reflection in W takes regions of $\mathcal{A}$ to regions of $\mathcal{A}$.

For any shard arrow $\Sigma_1 \to \Sigma_2$, Definition 9-7.16 says that the hyperplane H_{Σ_1} cuts the hyperplane H_{Σ_2} in the sense of Definition 9-7.2. Thus to show that the shard digraph is acyclic, it suffices to show that the cutting relation is acyclic. To do this, we recall from Definition 9-4.9 the definition of the depth of hyperplanes in $\mathcal{A}$ and will use the fact that $\mathcal{A}$ is a Coxeter arrangement to show that $\mathrm{depth}(H_1) < \mathrm{depth}(H)$ whenever H_1 cuts H.

Suppose H_1 cuts H and let $\mathcal{A}'$ be the rank-two subarrangement containing H_1 and H, so that H_1 is basic in $\mathcal{A}'$ and H is not. Let H_2 be the other basic hyperplane in $\mathcal{A}'$. Choose any region R with $|S(R)| = \mathrm{depth}(H)$ and $H \in S(R)$. Let $B = R_0 \prec R_1 \prec \cdots \prec R_{\mathrm{depth}(H)} = R$ be any maximal chain in $[B, R]$. Since H_1 and H_2 are basic in $\mathcal{A}'$ and H is not, and since $H \in S(R)$, Lemma 9-1.24 implies that at least one of the hyperplanes H_1 and H_2 is in $S(R)$. Thus there exists j with $1 \le j < \mathrm{depth}(H)$ such that R_{j-1} and R_j share a facet defined by H_1 or H_2. (We know that $j < \mathrm{depth}(H)$ because $R_{\mathrm{depth}(H)-1}$ and $R_{\mathrm{depth}(H)}$ share a facet defined by H.) If they share a facet defined by H_1, then R_j is a region with $H_1 \in S(R_j)$, so $\mathrm{depth}(H_1) \le |S(R_j)| = j < \mathrm{depth}(H)$.

If R_{j-1} and R_j share a facet defined by H_2, then let t be the reflection whose fixed hyperplane is H_2, define $Q = tR$, and define $Q_i = tR_i$ for $i = 0, \ldots, \mathrm{depth}(H)$. By Proposition 10-2.2, each Q_i is a region. Then $Q_j = R_{j-1}$, so

$$B = R_0, \ldots, R_{j-1}, Q_{j+1}, \ldots, Q_{\mathrm{depth}(H)} = Q$$

is a sequence of adjacent regions of $\mathcal{A}$. Now $H_2 \notin S(Q)$. If also $H_1 \notin S(Q)$, then Q is contained in B', the unique region of $\mathcal{A}'$ containing B. But in this case $R = tQ$ is in the $\mathcal{A}'$-region separated from B' by H_2, and in particular $S(R) \cap \mathcal{A}' = \{H_2\}$. This contradicts the fact that $H \in S(R)$, and we conclude that $H_1 \in S(Q)$. Since $R_0, \ldots, R_{j-1}, Q_{j+1}, \ldots, Q_{\mathrm{depth}(H)}$ is a sequence of adjacent regions, and since separating sets of adjacent regions differ by exactly one hyperplane, we have $|S(Q)| \le \mathrm{depth}(H) - 1$, so $\mathrm{depth}(H_1) \le \mathrm{depth}(H) - 1$.

In either case, we have seen that $\mathrm{depth}(H_1) < \mathrm{depth}(H)$, and we conclude that the shard digraph is acyclic. ∎

We can also describe canonical join representations in the weak order explicitly. Since the weak order on W is semidistributive, every element has a canonical join representation, and we can use the characterization of canonical join representations developed in Section 9-7.2.

Definition 10-3.8. If $v \prec w$ in the weak order on W, then Lemma 10-2.12 says that the associated regions vB and wB are adjacent, sharing a facet defined by a hyperplane H in the Coxeter arrangement $\mathcal{A}$. This hyperplane is H_t for some reflection t in W. The reflection t is called a *cover reflection* of w. The reflection t maps wB to vB, so $twB = vB$ and thus $tw = v$ by Theorem 10-2.5. But also $v = ws$ for some $s \in S$, so $tw = ws$ and thus $t = wsw^{-1}$. Conversely, if $\ell(ws) < \ell(w)$ for some $s \in S$, then the element ws is covered by w and furthermore $ws = (wsw^{-1})w$. Writing $\mathrm{cov}(w)$ for the set of cover reflections of w, we see that

$$\mathrm{cov}(w) = \{wsw^{-1} \mid s \in S,\, \ell(ws) < \ell(w)\} = \{t \in T \mid \ell(tw) = \ell(w) - 1\}.$$

The following is a combination of Theorem 9-7.11 and Lemma 9-7.12, rewritten in the language of Coxeter groups. It follows immediately from those two results together with Theorem 10-3.1 and Proposition 10-3.5.

Theorem 10-3.9. *Suppose W is a finite Coxeter group and $w \in W$. For each $t \in \mathrm{cov}(w)$, there is a unique minimal element j_t in $\{v \mid v \leq w,\, t \in \mathrm{inv}(v)\}$. The canonical join representation of w is $w = \bigvee\{j_t \mid t \in \mathrm{cov}(w)\}$.*

We emphasize also that Theorems 9-7.17, 9-7.18, and 9-7.19, which describe congruences in terms of shards, apply to Coxeter groups by Theorems 10-2.1 and 10-2.10 and Proposition 9-3.3. We will not restate these results separately for Coxeter groups.

The key properties of the weak order are inherited by lattice quotients. Specifically, the following results holds as a corollary of Theorem 10-3.7. (For the congruence uniformity, one can use the fact that congruence uniformity of finite lattices is inherited by quotients or argue using Theorem 9-8.1 and Corollary 9-8.20 and the acyclicity established in the proof of Theorem 10-3.7.)

Corollary 10-3.10. *Let W be a finite Coxeter group. Any lattice quotient of the weak order on W is semidistributive, congruence uniform, and polygonal.*

10-3.3 Combinatorial consequences

The ideas in the proof of Proposition 10-3.4 are sufficient to prove the following properties of finite Coxeter groups (Propositions 10-3.11 through 10-3.15). These hold more generally for not-necessarily-finite Coxeter groups, although our treatment here can only establish the finite case. We leave the few remaining details to Exercises 10.9, 10.10, 10.11, 10.12 and 10.13. The first two of these properties have been given names in the literature.

Proposition 10-3.11 (The Deletion Property). *Suppose W is a finite Coxeter group and let $a_1 \cdots a_k$ be a non-reduced word for some $w \in W$. Then there exist two distinct indices i and j in $\{1,\ldots,k\}$ such that, when a_i and a_j are both deleted from $a_1 \cdots a_k$, the result is another word for w.*

Proposition 10-3.12 (The Exchange Property). *Suppose W is a finite Coxeter group with defining generators S and let $a_1 \cdots a_k$ be a reduced word for some $w \in W$. If $\ell(ws) < \ell(w)$ for some $s \in S$, then a reduced word for ws can be obtained by deleting one letter from $a_1 \cdots a_k$.*

Proposition 10-3.13. *Let W be a finite Coxeter group and let $w \in W$. Then $\ell(w) = |\operatorname{inv}(w)|$.*

Proposition 10-3.14. *Let W be a finite Coxeter group with defining generators S, let $w \in W$, and let $s \in S$. Then $\ell(ws) \neq \ell(w)$.*

Since the weak order on a finite Coxeter group W is a finite lattice, it has a unique maximal element, traditionally called w_0. In light of the following proposition, w_0 is usually called the *longest element* of W. Some additional properties of w_0 are gathered in Exercises 10.14–10.18.

Proposition 10-3.15. *Suppose W is a finite Coxeter element. The maximal element w_0 of the weak order on W has length $\ell(w_0) = |T|$, the number of reflections in W. Every other element of W has strictly shorter length.*

10-3.4 Root systems and convexity

Since the weak order on a finite Coxeter group is isomorphic to a poset of regions, the convexity results of Section 9-4 apply. These results are most naturally explained in the context of root systems.

Given a finite reflection group W with Coxeter arrangement $\mathcal{A}$, a *root system* Φ associated to W is a collection of vectors[6] consisting of exactly two distinct (nonzero) normal vectors to each hyperplane in $\mathcal{A}$, such that the action of W permutes Φ. (That is, $W\Phi = \Phi$.) We already know that W permutes $\mathcal{A}$, so the additional requirement that W permutes Φ is purely a question of choosing the correct scaling of the normal vectors. Since in particular, $t\Phi = \Phi$ for each reflection t in W, the two normal vectors to each hyperplane are of the form $\pm\beta$ for some β.

It is easy to construct a root system Φ for W by taking the two unit normal vectors to each hyperplane. More generally, start with Φ empty. Choose some $H \in \mathcal{A}$ and any nonzero normal vector β to H, and take the orbit of β under W. Since the action of W preserves the (Euclidean) length of vectors, the orbit contains either two or zero normal vectors to each hyperplane in $\mathcal{A}$. If some hyperplane H' in $\mathcal{A}$ has no normal vectors in the orbit, choose a

[6]To conform to the tradition of naming roots by Greek letters, we break with our convention of naming vectors by bold roman letters.

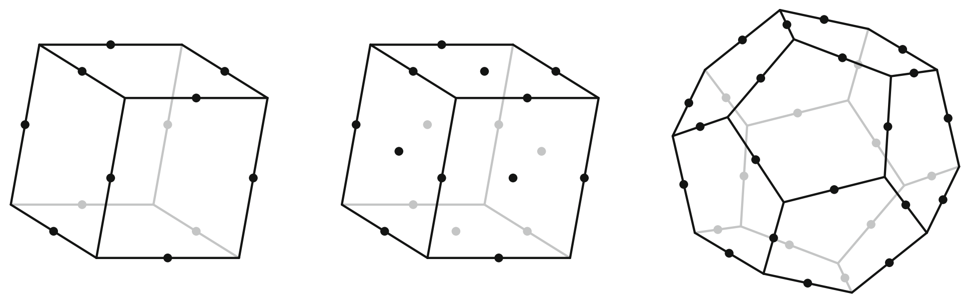

Figure 10-3.3: Root systems of types A_3, B_3, and H_3

nonzero normal vector β' to H' and take the union of the orbits of β and β'. Keep adjoining additional orbits until there are two normal vectors to every hyperplane in $\mathcal{A}$. (In practice, only one or two orbits are necessary when W is irreducible, but that is less obvious. It can be seen by considering the cases in Theorem 10-2.19.)

Given a choice of base region B, the *positive roots* Φ^+ are the roots $\beta \in \Phi$ such that $\langle \mathbf{b}, \beta \rangle > 0$ for any vector $\mathbf{b}$ in the interior of B. It is immediate that each root is either positive or is $-\beta$ for some positive root, so that $\Phi = \Phi^+ \cup (-\Phi^+)$. The map taking a positive root to its normal hyperplane is a bijection from Φ^+ to $\mathcal{A}$. But $\mathcal{A}$ is also in bijection with the reflections in W, and we write β_t for the positive root associated to a reflection t. Of special importance are the *simple roots*, which are the positive roots associated to the reflections S. (The associated hyperplanes contain the facets of B.) We write α_s for the simple root associated to $s \in S$.

Explicit constructions of root systems for all of the irreducible finite Coxeter groups are found in [241, Section 2.10, 2.13]. Here, we give some low-rank examples and describe root systems for Coxeter groups of types A_n, B_n, and D_n.

Example 10-3.16. A root system for a Coxeter group of rank 1 (that is, of type A_1) consists of two opposite nonzero vectors.

Example 10-3.17. A root system for a Coxeter group of rank two (type $I_2(m)$ for $m \geq 2$) consists of the vertices of a regular $2m$-gon centered at the origin. (Alternately, and more in keeping with the next few examples, the root system can be taken to be the midpoints of edges of a regular $2m$-gon.)

Example 10-3.18. Root systems for the irreducible rank-three Coxeter groups are shown in Figure 10-3.3. A root system of type A_3 can be obtained as the set of midpoints of edges of a cube centered at the origin. Adding in also the centers of the square faces, we obtain a root system of type B_3. For type H_3, we take the midpoints of edges of a dodecahedron centered at the origin.

Example 10-3.19. The standard choice of root system for the Coxeter group of type A_{n-1} is $\{\mathbf{e}_i - \mathbf{e}_j \mid i \neq j\} \subseteq \mathbb{R}^n$, where the $\mathbf{e}_i$ are the standard unit basis vectors. This root system is contained in the subspace of $\mathbb{R}^n$ consisting of vectors whose coordinates sum to zero, and thus the corresponding Coxeter arrangement is not essential.

Example 10-3.20. A standard choice of root system for the Coxeter group of type B_n is $\{\pm\mathbf{e}_i \pm \mathbf{e}_j \mid 1 \leq i < j \leq n\} \cup \{\pm\mathbf{e}_i \mid 1 \leq i \leq n\} \subseteq \mathbb{R}^n$, generalizing the B_3 root system in Example 10-3.18. Another standard choice replaces the roots $\pm\mathbf{e}_i$ by $\pm 2\mathbf{e}_i$. The latter root system is said to be of type C_n, but both root systems are associated to the same Coxeter group of type B_n.

Example 10-3.21. A standard choice of root system for the Coxeter group of type D_n is $\{\pm\mathbf{e}_i \pm \mathbf{e}_j \mid 1 \leq i < j \leq n\} \subseteq \mathbb{R}^n$, generalizing the A_3 root system in Example 10-3.18. (A Coxeter group of type D_3 is excluded from the list in Figure 10-2.1 precisely because it would coincide with A_3.)

Inversion sets of elements can be characterized in terms of roots. We leave the proof of the following proposition to Exercise 10.19.

Proposition 10-3.22. *Suppose W is a finite Coxeter group and suppose Φ is a root system for W. Each $w \in W$ has* $\mathrm{inv}(w) = \{t \in T \mid w\beta_t \in -\Phi^+\}$.

The positive roots can serve as the vectors $\mathbf{n}_H$ that appear in Definition 9-4.1. We restate Definitions 9-4.1 and 9-4.4 for root systems.

Definition 10-3.23. A set $\Psi \subseteq \Phi^+$ of positive roots is *convex* if Ψ equals the intersection of Φ^+ with the nonnegative real span of Ψ. The set Ψ is *biconvex* if both Ψ and its complement $\Phi^+ \setminus \Psi$ are convex. The set Ψ is *rank-two convex* if for every 2-dimensional linear subspace X of $\mathbb{R}^n$, the set $\Psi \cap X$ is convex in $\Phi^+ \cap X$. The set Ψ is *rank-two biconvex* if both Ψ and $\Phi^+ \setminus \Psi$ are rank-two convex.

The *closure* operator on positive roots takes a set $\Psi \subseteq \Phi^+$ to $\overline{\Psi}$, defined to be the intersection of all convex sets containing Ψ. The *rank-two closure* operator takes $\Psi \subseteq \Phi^+$ to ${}^2\overline{\Psi}^2$, defined to be the intersection of all rank-two convex sets containing Ψ.

Exercise 9.29, Theorems 10-2.1 and 10-2.10, and Proposition 10-3.5 imply that Theorems 9-4.5 and 9-4.8 can be restated as follows for finite Coxeter groups.

Theorem 10-3.24. *Suppose W is a finite Coxeter group and let Ψ be a subset of the positive roots Φ^+. Then the following are equivalent:*

(i) *Ψ is $\{\beta_t \mid t \in \mathrm{inv}(w)\}$ for some element w of W.*

(ii) *Ψ is biconvex in Φ^+.*

(iii) Ψ *is rank-two biconvex in* Φ^+.

Theorem 10-3.25. *Suppose* W *is a finite Coxeter group and let* v *and* w *be elements of* W. *Then*

(i) $v \vee w$ *is the unique element of* W *with inversion set* $\overline{\mathrm{inv}(v) \cup \mathrm{inv}(w)}$.

(ii) $v \vee w$ *is the unique element of* W *with inversion set* ${}^{2}\overline{\mathrm{inv}(v) \cup \mathrm{inv}(w)}^{2}$.

10-4. The Word Problem for finite Coxeter groups

The *Word Problem* for a finite Coxeter group is to give an algorithm that takes two words in the generators S and determines if the two words represent the same element of the group. This is equivalent to the problem of giving an algorithm that decides if a given word represents the identity element. We present two solutions to the Word Problem for finite Coxeter groups. One is a combinatorial algorithm and the other is geometric. However, the proof of the combinatorial algorithm relies heavily on the geometric setup that we have already developed. Both solutions can be extended to infinite Coxeter groups, with the same proofs, once the geometric setup has been developed without restricting to finite type.

The combinatorial solution is due to Tits [435]. The proof given here is in the same spirit as the proof in [435], but restricted to the finite case. The proof is also closely related to the proof of Theorems 9-3.15 and 10-2.9.

Suppose $a_1 \cdots a_k$ is a word in the defining generators S of a Coxeter group W. If $a_i = a_{i+1}$ for some i, then the word $a_1 \cdots a_{i-1} a_{i+2} \cdots a_n$ represents the same element of W, since in this case $a_i a_{i+1} = 1$. The operation of passing from $a_1 \cdots a_k$ to $a_1 \cdots a_{i-1} a_{i+2} \cdots a_n$, when $a_i = a_{i+1}$ is called a *nil move*. Similarly, if some subsequence $a_i \cdots a_j$ of $a_1 \cdots a_n$ is an alternating sequence $stst \cdots$ with exactly $m(s,t)$ letters, then we can replace the sequence with the alternating sequence $tsts \cdots$ with $m(s,t)$ letters and obtain a word for the same element. This is because the relation $(st)^{m(s,t)} = 1$ is equivalent to the relation that the two alternating sequences $stst \cdots$ and $tsts \cdots$ with $m(s,t)$ letters are equal in W. The operation of replacing one of these alternating sequences with the other is called a *braid move*. The following is Tits' solution to the Word Problem for Coxeter groups, stated in the case of finite W.

Theorem 10-4.1. *Suppose* W *is a finite Coxeter group with defining generators* S. *Then*

(i) *Any word in the generators* S *can be transformed to a reduced word for the same element by a sequence of changes consisting of braid moves and nil moves.*

(ii) *Any two reduced words for the same element are related by a sequence of braid moves.*

Since S is finite by hypothesis, the set of words of a given length using the letters S is also finite. Thus, given a word $a_1 \cdots a_k$, there are only finitely many other words that can be obtained by sequences of braid moves. Thus it is possible to check whether some word obtained from $a_1 \cdots a_k$ by braid moves admits a nil move. If so, then we perform the nil move and continue. If not, then $a_1 \cdots a_k$ is reduced. Since nil moves always decrease the number of letters in the word, the process eventually terminates with a reduced word. To solve the Word Problem using Theorem 10-4.1, one uses braid moves and nil moves to reduce the two given words. Since all reduced words for the same element are related by a sequence of braid moves, one can then tell whether the two given words represent the same element.

We continue to represent W as a reflection group with Coxeter arrangement $\mathcal{A}$ and base region B as in Theorem 10-2.10. The key to proving Theorem 10-4.1 is to use Theorem 10-3.1 to realize reduced words as maximal chains in lower intervals in $\mathrm{Pos}(\mathcal{A}, B)$ and appeal to Lemma 9-6.12. Let $a_1, \ldots, a_k$ be a reduced word in the generators S representing an element w, and let R be the region wB. As in the proof of Theorem 10-2.9, define $R_i = (a_1 \cdots a_i)B$ for each i from 1 to k, so that $B = R_0 \prec R_1 \prec \cdots \prec R_k = R$ is a maximal chain in the interval $[B, R]$. The proof of the following easy lemma is left as Exercise 10.21.

Lemma 10-4.2. *Under the correspondence given above between reduced words for w and maximal chains in $[B, R]$, rank-two moves on maximal chains correspond to braid moves on reduced words.*

Proof of Theorem 10-4.1. Let $a_1, \ldots, a_k$ and $a'_1 \cdots a'_k$ be reduced words in the generators S representing the same element w and let $B = R_0 \prec R_1 \prec \cdots \prec R_k = R$ and $B = Q_0 \prec Q_1 \prec \cdots \prec Q_k = R$ be the corresponding maximal chains. Lemma 9-6.12 says that these two maximal chains are related by a sequence of rank-two moves. Lemma 10-4.2 thus implies that $a_1, \ldots, a_k$ and $a'_1 \cdots a'_k$ are related by a sequence of braid moves.

Now suppose $a_1 \cdots a_k$ is a non-reduced word. Then there exists a largest index j with $1 \leq j < k$ such that $a_1 \cdots a_j$ is reduced. If w is the element represented by $a_1 \cdots a_j$, then $\ell(wa_{j+1}) = j - 1$. (The length is less than $j+1$ because $a_1 \cdots a_{j+1}$ is not reduced, but then the length is less than j by Proposition 10-3.14. If the length is less than $j-1$, then $a_1 \cdots a_j$ is not a reduced word for w.) Let $a'_1 \cdots a'_{j-1}$ be any reduced word for wa_{j+1}, so that $a'_1 \cdots a'_{j-1} a_{j+1}$ is a reduced word for w. The second assertion of the theorem says that $a_1 \cdots a_j$ can be transformed to $a'_1 \cdots a'_{j-1} a_{j+1}$ by a sequence of braid moves. The same braid moves transform $a_1 \cdots a_k$ into $a'_1 \cdots a'_{j-1} a_{j+1} a_{j+1} \cdots a_k$, which is transformed by a nil move to the word $a'_1 \cdots a'_{j-1} a_{j+2} \cdots a_k$. We repeat this procedure until the resulting word is reduced, and we have established the first assertion. $\square$

We now give the second, more geometric solution to the Word Problem.

The following theorem holds for all Coxeter groups, although we prove it only in the finite case. The theorem uses notation introduced in Proposition 10-2.6.

Theorem 10-4.3. *Let W be a finite Coxeter group with defining generators $S = \{s_1, \ldots, s_n\}$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be a basis for an n-dimensional real vector space. Define a symmetric bilinear form f with $f(\mathbf{v}_i, \mathbf{v}_i) = 1$ for all i and $f(\mathbf{v}_i, \mathbf{v}_j) = -\cos\left(\frac{\pi}{m(s_i,s_j)}\right)$ for all $i \neq j$. Define r_i to be $r_{\mathbf{v}_i,f}$, the reflection orthogonal to $\mathbf{v}_i$ with respect to f, and let $r(s_i)$ stand for r_i. Let $\mathbf{v}$ be a vector in $\mathbb{R}^n$ with $f(\mathbf{v}, \mathbf{v}_i) > 0$ for all i. Two words $a_1 \cdots a_j$ and $a'_1 \cdots a'_k$ in S stand for the same element of W if and only if $r(a_1) \cdots r(a_j)\mathbf{v} = r(a'_1) \cdots r(a'_k)\mathbf{v}$.*

Proof. Let W' be the group generated by the r_i. In the proof of Theorem 10-2.10, we showed that W is isomorphic to W', that the isomorphism restricts to a map $s_i \mapsto r(s_i)$, and that there is some Euclidean bilinear form with respect to which W' is a finite reflection group. We take $\mathcal{A}$ to be the Coxeter arrangement associated to W'. The set $B = \{\mathbf{x} \in V \mid f(\mathbf{x}, \mathbf{v}_i) \geq 0\,\forall i\}$ is a region of $\mathcal{A}$. Theorem 10-2.5 says that the map $w \mapsto wB$ is a bijection from W to the regions of $\mathcal{A}$. Since $\mathbf{v}$ is in the interior of B and the interiors of regions are disjoint by construction, two words in S define the same element if and only if the corresponding words in the r_i map $\mathbf{v}$ to the same vector. □

Theorem 10-4.3 provides an algorithm that is computationally more efficient than using braid moves and nil moves. However, there are computational issues arising from the fact that $-\cos\left(\frac{\pi}{m(s_i,s_j)}\right)$ need not be an integer in general. For most finite Coxeter groups (called *crystallographic* finite Coxeter groups), the definition of each r_i can be modified so that r_i is given by an integer matrix. The only noncrystallographic finite Coxeter groups are H_3, H_4, and $I_2(m)$ with $m \in \{5, 7, 8, \ldots\}$. (See Figure 10-2.1.) For $I_2(m)$, the word problem is easy, and there are reasonable ways to deal with H_3 and H_4 (including the "Non-Crystallographic Kludge" of [426, Section 7] or working symbolically over the field extension $\mathbb{Q}[\sqrt{5}]$).

10-5. Coxeter groups of type A

For a specific Coxeter group W, one can obtain a combinatorial model of W by making a good choice of the vectors $\mathbf{v}_i$ and $\mathbf{v}$ in Theorem 10-4.3. (In particular, it is helpful to choose the $\mathbf{v}_i$ such that the bilinear form defined in the theorem is the usual Euclidean inner product.) In this section, we describe such a model in detail for Coxeter groups of type A_{n-1}, and also describe many lattice-theoretic properties of the weak order in terms of the model. A similar construction yields combinatorial models for the finite Coxeter groups of types B_n and D_n as well. For details (without the lattice theory), see [69, Chapter 8].

Suppose W is of type A_{n-1} and suppose we label the diagram for W (see Figure 10-2.1) linearly s_1 through s_{n-1}. Define $\mathbf{v}_i$ to be $\mathbf{e}_{i+1} - \mathbf{e}_i$ for $i = 1, \ldots, n-1$, where the $\mathbf{e}_i$ are the standard unit basis vectors. (This is a set of simple roots for the root system described in Example 10-3.19.) These vectors are a basis for an $(n-1)$-dimensional subspace of $\mathbb{R}^n$. The symmetric bilinear form defined in Theorem 10-4.3 is the restriction of the usual Euclidean form on $\mathbb{R}^n$ to this subspace. Setting $\mathbf{v}$ to be the vector $(1, 2, \ldots, n) \in \mathbb{R}^n$, we have satisfied the hypotheses of Theorem 10-4.3. The action of $r(s_i)$ on a vector is to interchange the entries in positions i and $i+1$. Thus the action of W on $\mathbf{v}$ produces all vectors whose entries are a permutation of the entries of $\mathbf{v}$, and Theorem 10-4.3 implies that these permutations are in bijection with the elements of W. The associated Coxeter arrangement consists of all hyperplanes normal to vectors $\mathbf{e}_j - \mathbf{e}_i$ for all $1 \leq i < j \leq n$. For the rest of the section, we fix this representation of W as the reflection group generated by the reflections s_i orthogonal to $\mathbf{e}_{i+1} - \mathbf{e}_i$ for $i = 1, \ldots, n-1$.

Let $\mathfrak{S}_n$ be the group of permutations of $\{1, \ldots, n\}$. This is the group of bijections from $\{1, \ldots, n\}$ to itself, with product given by composition. We will write a permutation π in *one-line notation* as $\pi_1\pi_2 \cdots \pi_n$, meaning that π maps 1 to π_1, maps 2 to π_2, and so forth. The *inversion set* of π is $\mathrm{inv}(\pi) = \{(\pi_i, \pi_j) \mid i < j, \pi_i > \pi_j\}$. We follow the usual convention of composing permutations from right to left, so that, for example the product $213 \cdot 132$ is 231 (not 312). To fit this convention into the geometric description of the paragraph above, we identify each element w of W with the permutation π such that $w(\pi_1, \ldots, \pi_n) = \mathbf{v}$. Thus $(\pi_1, \ldots, \pi_n) = w^{-1}\mathbf{v}$ and $w\mathbf{v}$ is the vector with 1 in position π_1, with 2 in position π_2, etc.

Under this identification of W with $\mathfrak{S}_n$, the generator s_i corresponds to the transposition $(i\ i+1)$. Each transposition $(i\ j)$ in $\mathfrak{S}_n$ is identified with the reflection t_{ij} orthogonal to $\mathbf{e}_j - \mathbf{e}_i$. This is a reflection in W, and these are all of the reflections in W. The following proposition, which is verified in Exercise 10.24, shows that inversion sets of permutations (as defined above) correspond to inversion sets, in the sense of Coxeter groups.

Proposition 10-5.1. *If π is a permutation in $\mathfrak{S}_n$, then the inversion set of π (as an element of W) is the set of all transpositions $(\pi_i\ \pi_j)$ such that (π_i, π_j) in the set* $\mathrm{inv}(\pi)$ *defined above.*

Example 10-5.2. Figure 10-5.1 shows a Coxeter arrangement for a Coxeter group W of type A_3 with the regions labeled with the permutations in $\mathfrak{S}_4$. Figure 10-5.2 shows the weak order on W written in terms of permutations. (Compare Figures 10-2.2 and 10-3.1.)

It will be useful to have a few facts about multiplying permutations by transpositions, proved as Exercise 10.25.

Proposition 10-5.3. *Suppose $\pi = \pi_1 \cdots \pi_n \in \mathfrak{S}_n$ and $s_i = (i\ i+1) \in \mathfrak{S}_n$.*

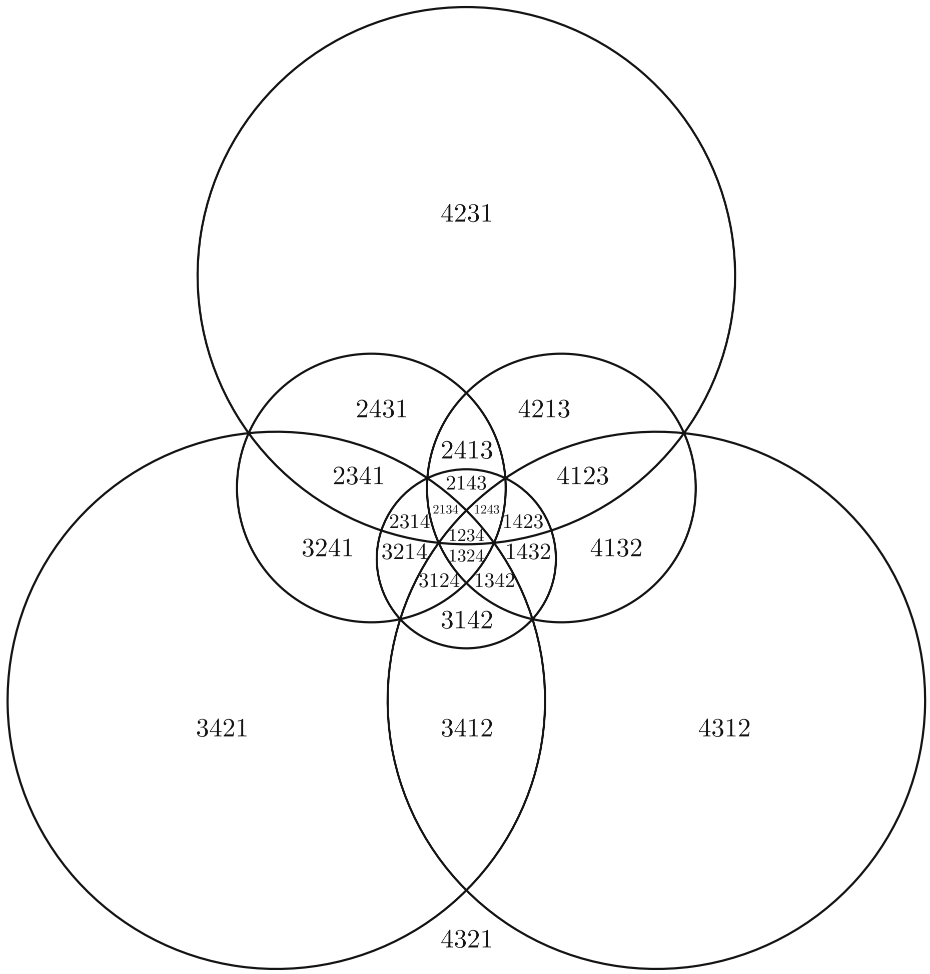

Figure 10-5.1: Regions labeled by permutations in $\mathfrak{S}_4$

(i) *$s_i\pi$ is obtained from π by swapping the values i and $i+1$ in the sequence $\pi_1 \cdots \pi_n$.*

(ii) *$(i\ j)\pi$ is obtained from π by swapping the values i and j in $\pi_1 \cdots \pi_n$.*

(iii) *πs_i is obtained from π by swapping the entries in positions i and $i+1$ (the values π_i and π_{i+1}) in the sequence $\pi_1 \cdots \pi_n$.*

(iv) *$\pi(i\ j)$ is obtained from π by swapping the entries in positions i and j (the values π_i and π_j) in the sequence $\pi_1 \cdots \pi_n$.*

For example, $s_2(35124) = 25134 = 35124 \cdot (1\ 4)$ and $(35124)s_2 = 31524 = (1\ 5) \cdot 35124$.

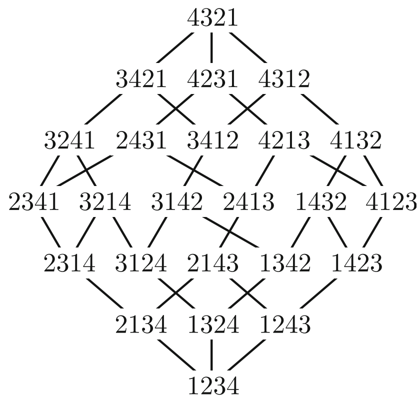

Figure 10-5.2: The weak order on permutations in $\mathfrak{S}_4$

Theorem 10-3.1 implies that the weak order on W is isomorphic to $\operatorname{Pos}(\mathcal{A}, B)$, where $\mathcal{A}$ is the set of hyperplanes orthogonal to vectors of the form $\mathbf{e}_j - \mathbf{e}_i$ for $1 \leq i < j \leq n$ and B is the region $\{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{x}, \mathbf{v}_i \rangle \geq 0 \; \forall \, i\}$. The following proposition is verified as Exercise 10.26.

Proposition 10-5.4. *The cover relations in the weak order on W correspond to the relations $\sigma \prec \pi$ such that, for some i, the permutations σ and π differ only in positions i and $i+1$, with $\pi_i > \pi_{i+1}$.*

A pair (π_i, π_{i+1}) such that $\pi_i > \pi_{i+1}$ is called a *descent* of π. Propositions 10-5.3 and 10-5.4 imply that the cover reflections of π are the transpositions $(\pi_i \; \pi_{i+1})$ such that (π_i, π_{i+1}) are the descents of π. The following is an immediate consequence of Proposition 10-5.4.

Proposition 10-5.5. *A permutation π is join-irreducible in the weak order if and only if π has a unique descent.*

Given $1 \leq a < b \leq n$ and a set $R \subseteq \{a+1, \ldots, b-1\}$, let R^c be $\{a+1, \ldots, b-1\} \setminus R$. We write $\tau(b, a, R)$ for the permutation in $\mathfrak{S}_n$ given by $1, \ldots, a-1$, then the values in R^c in increasing order, then b, then a, then the values in R in increasing order, then the values $b+1, \ldots, n$. Any of the sequences $a+1, \ldots, b-1$ and/or $1, \ldots, a-1$ and/or $b+1, \ldots, n$ may be empty. The construction of $\tau(b, a, R)$ depends on having specified n. For example, in $\mathfrak{S}_9$, the permutation $\tau(8, 3, \{5, 6\})$ is 124783569. Proposition 10-5.5 implies that $\tau(b, a, R)$ is join-irreducible (with descent (b, a)) and that every join-irreducible permutation is of this form. As a fairly easy consequence of Theorem 10-3.9, we have the following theorem. The details are left to Exercise 10.27.

Theorem 10-5.6. *The canonical join representation of a permutation $\pi \in \mathfrak{S}_n$ is*

$$\pi = \bigvee_{\pi_i > \pi_{i+1}} \tau(\pi_i, \pi_{i+1}, R_i),$$

where R_i is the set $\{\pi_j \mid i+1 < j$ and $\pi_{i+1} < \pi_j < \pi_i\}$.

Example 10-5.7. Let $\pi = 395284176 \in \mathfrak{S}_9$. Then Theorem 10-5.6 says that

$$\pi = 123495678 \vee 135246789 \vee 12358467 \vee 234156789 \vee 123457689$$

is the canonical join representation of π. This can also be checked directly from Theorem 10-3.9. One first verifies that the cover reflections of π are $(9\ 5)$, $(5\ 2)$, $(8\ 4)$, $(4\ 1)$, and $(7\ 6)$. One then checks that the smallest element below π having $(9\ 5)$ as an inversion is 123495678, that the smallest element below π having $(5\ 2)$ as an inversion is 135246789, etc.

Proposition 9-7.8 says that the join-irreducible elements are in bijection with the shards of $\mathcal{A}$ with respect to B. Exercises 10.30 and 10.31 are to describe the cutting relation on hyperplanes and prove the following fact.

Proposition 10-5.8. *The shard $\Sigma(b, a, R)$ associated to the join-irreducible permutation $\tau(b, a, R)$ is*

$$\{\mathbf{x} \in \mathbb{R}^n \mid x_a = x_b \textit{ and } x_a \le x_i \text{ for all } i \in R \textit{ and } x_a \ge x_i \text{ for all } i \in R^c\}.$$

We can also describe the shard digraph (Definition 9-7.16), and thus the forcing relation on join-irreducible elements.

Proposition 10-5.9. *The shard $\Sigma(b, a, R)$ arrows the shard $\Sigma(b', a', R')$ if and only if $a' = a < b < b'$ and $R = \{i \in R' \mid i < b\}$, or $a' < a < b = b'$ and $R = \{i \in R' \mid a < i\}$. A join-irreducible element $\tau(b, a, R)$ forces a join-irreducible element $\tau(b', a', R')$ if and only if $a' \le a < b \le b'$ and $R = \{i \in R' \mid a < i < b\}$.*

Proof. For each $a' < b'$, let $H = (\mathbf{e}_{b'} - \mathbf{e}_{a'})^\perp$. Exercise 10.30 says that a hyperplane cuts H if and only if it is of the form $(\mathbf{e}_{b'} - \mathbf{e}_c)^\perp$ or $(\mathbf{e}_c - \mathbf{e}_{a'})^\perp$ for some c with $a' < c < b'$. Choose such a c and write $H_1 = (\mathbf{e}_{b'} - \mathbf{e}_c)^\perp$ and $H_2 = (\mathbf{e}_c - \mathbf{e}_{a'})^\perp$. The shards in H correspond to the subsets $R \subseteq \{a'+1, \ldots, b'-1\}$. Choose such a subset R and write Σ for $\Sigma(b', a', R)$. The set $I = \Sigma \cap H_1 \cap H_2$ consists of vectors $\mathbf{x}$ with $x_{a'} = x_c = x_{b'}$, with $x_{a'} \le x_i$ for all $i \in R$, and with $x_{a'} \ge x_i$ for all $i \in R^c$. In particular, I is $(n-2)$-dimensional. Setting $\Sigma_1 = \Sigma(c, a', R \cap \{a'+1, \ldots, c-1\})$ and $\Sigma_2 = \Sigma(b', c, R \cap \{c+1, \ldots, b'-1\})$, we see that I equals $\Sigma(b', a', R) \cap \Sigma_1$ and $\Sigma(b', a', R) \cap \Sigma_2$. Thus Σ_1 and Σ_2 both arrow Σ, and these are the only shards that arrow Σ and intersect Σ in $H_1 \cap H_2$. Letting c and R vary, we obtain the first assertion of the proposition. Taking the transitive closure of the shard arrows, the second assertion follows by Proposition 10-5.8. $\square$

Example 10-5.10. Figure 10-5.3 shows the forcing order on join-irreducible permutations in $\mathfrak{S}_4$.

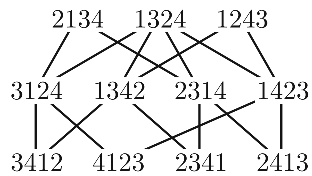

Figure 10-5.3: The forcing order on join-irreducible permutations in $\mathfrak{S}_4$

The forcing order on join-irreducible permutations also has a useful combinatorial description in terms of certain "arcs." Canonical join representations then correspond to certain "noncrossing arc diagrams." See in particular [376, Corollary 3.5] and [376, Theorem 4.4].

Proposition 10-5.9 lets us describe quite precisely the lattice quotients of the weak order on permutations. Given $1 \le a < b \le n$ and $R \subseteq \{a+1,\dots,b-1\}$, take $R^c = \{a+1,\dots,b-1\} \setminus R$ as before. Let π be a permutation. Say π has a (b,a,R)*-pattern* if (π_i,π_{i+1}) is a descent of π with $\pi_i \ge b$ and $\pi_{i+1} \le a$ such that all of the elements of R^c appear in π *before* the descent (π_i,π_{i+1}) while all of the elements of R appear in π *after* (π_i,π_{i+1}). Say that π *avoids* (b,a,R) if it has no (b,a,R)-patterns.

Theorem 10-5.11. *Given a set $\{\tau(b_i,a_i,R_i) \mid i \in I\}$ of join-irreducible permutations in $\mathfrak{S}_n$, let $\boldsymbol{\alpha}$ be the smallest congruence on the weak order on $\mathfrak{S}_n$ contracting each of the given join-irreducible permutations. The quotient of the weak order on $\mathfrak{S}_n$ modulo $\boldsymbol{\alpha}$ is isomorphic to the restriction of the weak order to permutations that avoid (b_i,a_i,R_i), for every $i \in I$.*

Proof. If π has a (b_i,a_i,R_i)-pattern involving the descent (π_i,π_{i+1}), then Proposition 10-5.9 says that the join-irreducible permutation $\tau(\pi_i,\pi_{i+1},R'_i)$ with $R'_i = \{\pi_j \mid j > i+1,\, \pi_{i+1} < \pi_j < \pi_i\}$ is forced by $\tau(b_i,a_i,R_i)$. Thus Theorem 10-5.6 and Proposition 9-5.29 imply that $\pi \notin \pi_\downarrow^{\boldsymbol{\alpha}}\mathfrak{S}_n$. If, for every $i \in I$, the permutation π has no (b_i,a_i,R_i)-pattern, then Proposition 10-5.9 and Theorem 10-5.6 imply that the canonical joinands of π are not contracted by $\boldsymbol{\alpha}$. Proposition 9-5.29 then says that $\pi \in \pi_\downarrow^{\boldsymbol{\alpha}}\mathfrak{S}_n$. Thus $\pi \in \pi_\downarrow^{\boldsymbol{\alpha}}\mathfrak{S}_n$ is exactly the set of permutations described by the avoidance condition of the theorem. Proposition 9-5.5 completes the proof. $\square$

10-6. Cambrian lattices and sortable elements

In this section, we define a family of congruences on the weak orders on finite Coxeter groups called the Cambrian congruences. The quotient of the weak order modulo a Cambrian congruence is called a Cambrian lattice. We also give a combinatorial description of the bottom elements of Cambrian congruence classes (the sortable elements). As we will see in Section 10-6.3, Cambrian lattices and sortable elements are very closely related to cluster algebras of finite type and to Coxeter-Catalan combinatorics, particularly generalized associahedra and noncrossing partitions.

10-6.1 Cambrian congruences

Let W be a Coxeter group with defining generators S. Exercise 10.32 is the statement that the join of r and s in the weak order is the element with two distinct reduced words $rsrs\cdots$ and $srsr\cdots$, each of length $m(r,s)$. (Compare Figure 10-1.1.) The interval below $r \vee s$ is a polygon. More specifically, the polygon has $2m(r,s)$ vertices, so the interval has $2m(r,s)-4$ side edges. On one side, the side edges are $r \prec rs \prec rsr \prec \cdots$ with the last cover relation in the list relating an element of length $m(r,s)-2$ to an element of length $m(r,s)-1$. On the other side, the side edges are $s \prec sr \prec srs$, etc. Each of these side edges is of the form $j_* \prec j$ where j is join-irreducible and j_* is the unique element covered by j. Taken together, these polygons contain the same information as the Coxeter diagram of W (defined earlier in connection with Theorem 10-2.19). That is, if the polygon $[1, r\vee s]$ is a square, then $m(r,s)=2$, so no edge exists between r and s in the Coxeter diagram. Similarly, hexagons correspond to unlabeled edges, octagons correspond to edges labeled 4, etc.

An *orientation* of a graph is a directed graph obtained by replacing each edge r, s of the graph by an arrow, either $r \to s$ or $s \to r$. There is a Cambrian lattice for each *oriented Coxeter diagram*, or in other words each orientation of the Coxeter diagram. The symbol c will represent an oriented Coxeter diagram. (Later, we will also think of c as a particular element of W called a *Coxeter element*. See Definition 10-6.11.) Given an orientation c, the opposite orientation (reversing all arrows) is denoted c^{-1}.

Given an oriented Coxeter diagram c for a finite Coxeter group W, the *c-Cambrian congruence* $\boldsymbol{\theta}_c$ is the smallest congruence contracting all of the side edges of the form $s \prec sr \prec srs \prec \cdots$ for each directed edge $r \to s$. (Here, as above, the last cover relation in the list relates an element of length $m(r,s)-2$ to an element of length $m(r,s)-1$.) The *c-Cambrian lattice* is the quotient of the weak order modulo the c-Cambrian congruence $\boldsymbol{\theta}_c$. The *c-Cambrian fan* is the fan consisting of the $\boldsymbol{\theta}_c$-cones and their faces. (See Definition 9-1.9.) Corollary 10-3.10 implies the following theorem.

Theorem 10-6.1. *For any finite Coxeter group W and any orientation c of the diagram of W, the c-Cambrian lattice* (*the quotient of the weak order on W modulo the c-Cambrian congruence $\boldsymbol{\theta}_c$*) *is semidistributive, congruence uniform, and polygonal.*

While Theorem 10-6.1 asserts several strong properties of Cambrian lattices, it gives little insight into the nature of Cambrian lattices. The goal of the remainder of this section is to provide more insight, first by giving several examples, and then by quoting several theorems on the combinatorics and geometry of Cambrian congruences, lattices, and fans. References for the quoted results are found in the Notes at the end of the chapter.

Example 10-6.2. If W is of type A_2 in the notation of Figure 10-2.1 (that is, if $S = \{s_1, s_2\}$ and $m(s_1, s_2) = 3$), then there are exactly two orientations of the diagram of W. Figure 10-6.1 shows the Cambrian congruences, Cambrian lattices, and Cambrian fans for W of type A_2.

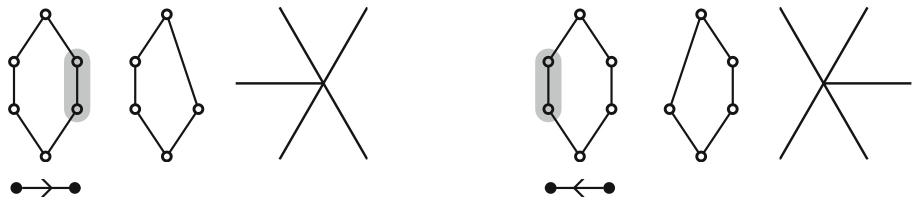

Figure 10-6.1: The two Cambrian congruences in type A_2 and their corresponding Cambrian lattices and Cambrian fans

Example 10-6.3. Suppose W is of type A_3 in the notation of Figure 10-2.1. The weak order on W appears in Figure 10-3.1, and is shown again in Figure 10-5.2, written in terms of permutations. The Coxeter arrangement for W is shown in Figures 10-2.2 and 10-5.1. There are four Cambrian congruences on the weak order on W, shown in Figure 10-6.2. (Compare Example 9-6.6.) Figure 10-6.3 shows the corresponding Cambrian lattices, and Figure 10-6.4 shows the corresponding Cambrian fans. In the Cambrian fan pictures, the Cambrian cones are shown by the thicker black lines, and the decomposition of Cambrian cones into regions is shown by the thinner gray lines. Figures 10-5.1 and 10-5.2 will be helpful in verifying these examples. Exercises 10.33 and 10.34 ask the reader to find the Cambrian congruences, lattices and fans of type B_3 and of type H_3.

Two of the Cambrian lattices in Figure 10-6.3 are isomorphic. The same two are also isomorphic to their duals. The other two Cambrian lattices are dual to each other. (This duality is not immediately obvious in the pictures, because the pictures are drawn to highlight the relationship to Figure 10-6.2 rather than to highlight the duality.) These isomorphisms and anti-isomorphisms are examples of a general result.

$\diamondsuit$ **Theorem 10-6.4.** *An isomorphism of oriented Coxeter diagrams induces an isomorphism of Cambrian lattices. An isomorphism of Cambrian lattices restricts to an isomorphism of oriented Coxeter diagrams. The same is true for anti-isomorphisms of Cambrian lattices and of oriented Coxeter diagrams.*

Since the identity map on the diagram is an anti-isomorphism from the orientation c to the opposite orientation c^{-1}, Theorem 10-6.4 implies that the c-Cambrian lattice is anti-isomorphic to the c-Cambrian lattice. The

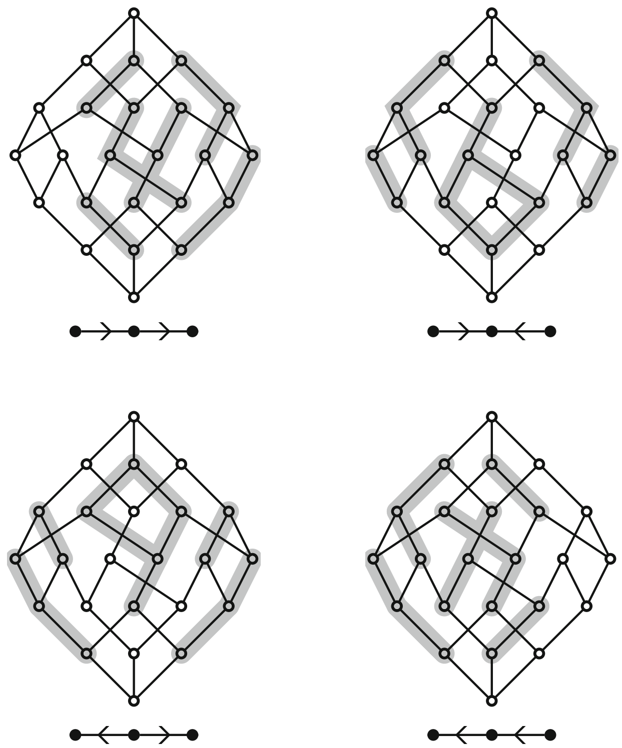

Figure 10-6.2: The four Cambrian congruences in type A_3

anti-isomorphism is related to the anti-automorphism $w \mapsto ww_0$ of the weak order discussed in Exercise 10.15. Specifically:

$\diamondsuit$ **Theorem 10-6.5.** *Given a Coxeter group W and an orientation c of the diagram for W, the c-Cambrian congruence $\boldsymbol{\theta}_c$ and the c^{-1}-Cambrian congruence $\boldsymbol{\theta}_{c^{-1}}$ are related by the map $w \mapsto ww_0$.*

That is, if v and w are in W, then $v \equiv w \pmod{\boldsymbol{\theta}_c}$ if and only if $vw_0 \equiv ww_0 \pmod{\boldsymbol{\theta}_{c^{-1}}}$. Since $w \mapsto ww_0$ corresponds to the antipodal map on regions (Exercise 10.15), Theorem 10-6.5 immediately implies the following result.

$\diamondsuit$ **Theorem 10-6.6.** *Given a Coxeter group W and an orientation c of the diagram for W, the c-Cambrian fan and the c^{-1}-Cambrian fan are related by the antipodal map.*

To see Theorem 10-6.6 in Figure 10-6.4, it it helpful to remember that each circle in the figures is the projection of a great circle. Given two projected circles in the plane, their two points of intersection are the projections of two antipodal points in the sphere.

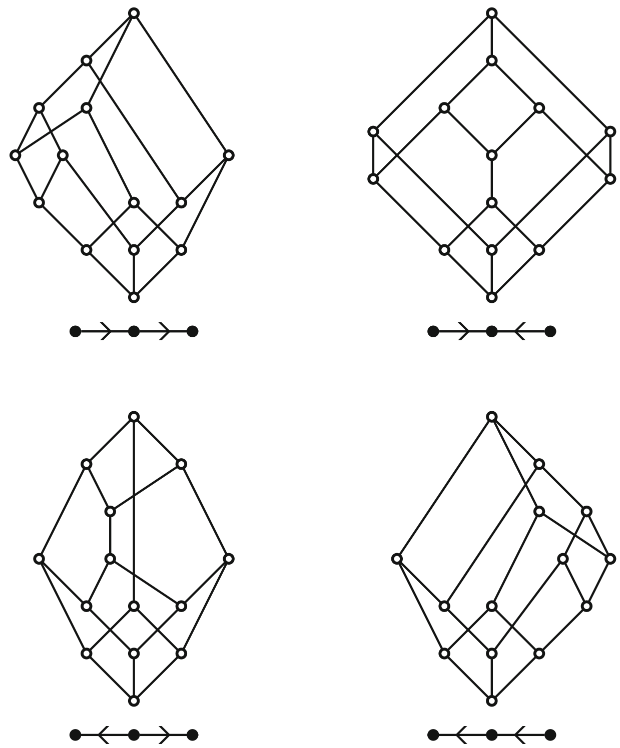

Figure 10-6.3: The four Cambrian lattices in type A_3

An examination of Figure 10-6.3 reveals that the Hasse diagrams of the four Cambrian lattices of type A_3 all define the same underlying graph. The patient reader who carries out Exercises 10.33 and 10.34 will see that the same is true for type B_3 and and for type H_3. In general, the underlying graph of the Hasse diagram of a Cambrian lattice depends only on the unoriented diagram (that is, on the Coxeter group). Looking at examples in rank 3, one also gets the impression of a 3-dimensional solid traced out by the Hasse diagrams of Cambrian lattices. Indeed, there is a polytope called a *generalized associahedron* for each Coxeter group W such that the following result holds.

$\diamondsuit$ **Theorem 10-6.7.** *Suppose W is a finite Coxeter group and suppose c is any orientation of the Coxeter diagram for W. The Hasse diagram of the c-Cambrian lattice is isomorphic, as a graph, to the graph of vertices and edges of the* (*simple*) *generalized associahedron.*

The generalized associahedron is defined in terms of a root system associated to W (in the sense of Section 10-3.4). For more details, see the Notes to this section. Theorem 10-6.7 is a consequence of Theorem 9-8.9 and the following stronger theorem.

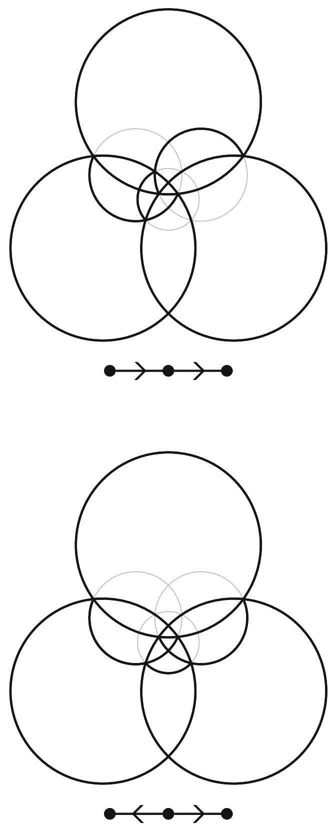

Figure 10-6.4: The four Cambrian fans in type A_3

$\diamondsuit$ **Theorem 10-6.8.** *Suppose W is a finite Coxeter group. The c-Cambrian fan for any orientation c of the Coxeter diagram for W is isomorphic to the normal fan of the (simple) generalized associahedron.*

In fact, there is a polytope, combinatorially isomorphic to the generalized associahedron, whose normal fan *is* the c-Cambrian fan. See the Notes for details and references.

10-6.2 Cambrian lattices of type A

We now characterize Cambrian lattices of type A in terms of permutations, continuing the notation of Section 10-5. In particular, W is a Coxeter group of type A_{n-1} and the defining generators S are $s_1, \ldots, s_{n-1}$ with $m(s_{i-1}, s_i) = 3$ for all $i = 2, \ldots, n-1$. An oriented Coxeter diagram for W has a directed edge $s_{i-1} \to s_i$ or $s_{i-1} \leftarrow s_i$ for each $i = 2, \ldots, n-1$. We encode this choice as a *barring* of the elements $2, \ldots, n-1$. If $s_{i-1} \to s_i$, then i is *lower-barred* and we write $\underline{i}$. If $s_{i-1} \leftarrow s_i$, then i is *upper-barred* and we write $\overline{i}$. Fixing some barring, we say that a permutation π *avoids the pattern* $31\underline{2}$ if there exists no subsequence kij of the one-line notation for π such that $i < j < k$ (in the usual numerical order on integers) and such that j is lower-barred. Similarly, π *avoids the pattern* $\overline{2}31$ if there exists no subsequence jki of the one-line notation for π such that $i < j < k$ and such that j is upper-barred.

Theorem 10-6.9. *Given a Coxeter group W of type A_n and an orientation c of the diagram of W, encoded as a barring, the c-Cambrian lattice is isomorphic to the subposet of the weak order on W induced by permutations avoiding the patterns $31\underline{2}$ and $\overline{2}31$.*

It should be emphasized that the avoidance condition in Theorem 10-6.9 depends on the choice of orientation c because the barring depends on c.

Proof. The c-Cambrian congruence is the smallest congruence contracting the join-irreducible element $s_i s_{i-1}$ whenever $s_{i-1} \to s_i$ and contracting the join-irreducible element $s_{i-1}s_i$ whenever $s_i \to s_{i-1}$. These join-irreducible elements are $\tau(i+1, i-1, R_i)$, where $R_i = \{i\}$ if i is lower-barred and $R_i = \varnothing$ if i is upper-barred. Theorem 10-5.11 completes the proof. $\square$

The avoidance conditions in Theorem 10-6.9 are conditions on the inversion set of a permutation. For example, a permutation $\pi \in \mathfrak{S}_n$ avoids $31\underline{2}$ if there exist no integers $1 \le i < j < k \le n$ with j lower-barred such that both (k,i) and (k,j) are inversions of π but (j,i) is not an inversion of π. These conditions are a special case of general conditions called *c-alignment* characterizing Cambrian lattices of all finite types in terms of the geometry of roots and certain local “orientations” of the root system. See the Notes for references.

On the other hand, a special case of the avoidance conditions in type A yields the Tamari lattices. A permutation π *avoids the pattern* 312 if there exists no subsequence cab of the one-line notation for π such that $a < b < c$. Similarly, π *avoids the pattern* 231 if there exists no subsequence bca of the one-line notation for π such that $a < b < c$. The Tamari lattice can be realized as the subposet (and in fact sublattice) of the weak order on permutations induced by the permutations that avoid 312, or alternately as the subposet induced by 231-avoiding permutations. The following is an immediate corollary of Theorem 10-6.9.

Corollary 10-6.10. *Let c be the orientation of the type-A_{n-1} Coxeter diagram with unique source on one end of the diagram and unique sink on the other end. The c-Cambrian lattice is isomorphic to the Tamari lattice* (*denoted* $\mathsf{A}(n)$ *in Section 7-2.2*).

Since there are two orientations, opposite to each other, that satisfy the hypotheses of Corollary 10-6.10, Theorem 10-6.4 and Corollary 10-6.10 combine to prove the well known self-duality of the Tamari lattice (cf. Proposition 7-4.14).

By analogy to Corollary 10-6.10, one can define a Cambrian lattice using the linear orientation of the B_n diagram. There are two anti-isomorphic linear orientations of the B_n diagram, so there are two reasonable (dual) candidates for the name "type-B_n Tamari lattice." These were defined as Cambrian lattices in [368, Section 7], where they were also realized in terms of centrally-symmetric triangulations of polygons. Independently, and at about the same time, they were defined by Thomas in [433] in terms of centrally-symmetric triangulations and in terms of bracket vectors, analogous to the bracket vectors that realize the usual Tamari lattice. (See Definition 7-4.9.) One of the Tamari lattices of type B_3 is shown in Figure 10-6.5. The other is dual.

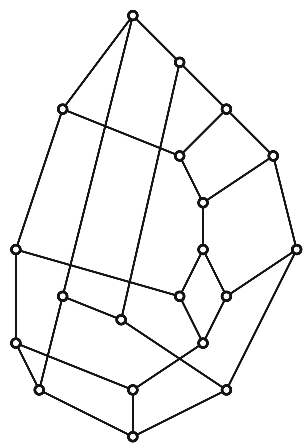

Figure 10-6.5: A Tamari lattice of type B_3

10-6.3 Sortable elements

By Proposition 9-5.5, the Cambrian lattice is the subposet of the weak order induced by the bottom elements of Cambrian congruence classes. These elements turn out to have a pleasant and useful combinatorial description in terms of the combinatorics of reduced words. These elements are called *sortable elements* -or *Coxeter-sortable element*. Most of what can be proved in general about Cambrian congruences/lattices/fans is proved using sortable elements.

In this section, we define sortable elements and quote some difficult theorems that link them to Cambrian lattices. We also discuss some of the tools and methods that contribute to the study of sortable elements. The first step is to encode orientations of the diagram of W by certain elements of W called Coxeter elements.

Definition 10-6.11. Suppose W is a Coxeter group with defining generators $S = \{s_1, \ldots, s_n\}$ and $|S| = n$. There are $n!$ ways to write a product $c = s_1 \cdots s_n$. Every choice is a reduced word for an element of W called a *Coxeter element*. There are typically fewer than $n!$ Coxeter elements of W, because different total orders on S can give rise to the same Coxeter element. Specifically, we can interpret every acyclic orientation of the diagram of W as a partial order on S by taking reflexive-transitive closure. Every linear extension $s_1, \ldots, s_n$ of this partial order is a reduced word for a Coxeter element, and conversely, every total order $s_1, \ldots, s_n$ arises in this way. (To see why, note that the total order $s_1, \ldots, s_n$ is a linear extension of the acyclic orientation of the diagram induced by the total order $s_1, \ldots, s_n$.) Exercise 10.36 is to verify that two total orders on S define the same Coxeter element if and only if they induce the same orientation on the diagram. In this case, the two total orders are related by a sequence of commutations of adjacent commuting entries (braid moves replacing $s_i s_j$ by $s_j s_i$ with $m(s_i, s_j) = 2$). Thus Coxeter elements in W are in bijection with acyclic orientations of the Coxeter diagram of W. In what follows, some constructions will depend on a Coxeter element c while other constructions will depend on a chosen reduced word for c.

To interpret an acyclic directed graph as a partial order, there is a choice to be made. As one might expect when various authors make the same choice independently, conventions vary in the literature. We follow [369, Section 1] in taking the convention that an arrow $r \to s$ appears in the diagram if and only if r precedes s in every reduced word for the corresponding Coxeter element c.

When W is a finite Coxeter group, we see from Theorem 10-2.19 that the diagram of W is a forest (a union of trees). In this case, all orientations of the diagram are acyclic, so Coxeter elements are in bijection with orientations of the diagram. This justifies our reuse of the letter c to denote both a Coxeter element and an orientation of the Coxeter diagram. As we proceed, we will identify an orientation of the diagram with the corresponding Coxeter element. Thus Cambrian congruences $\boldsymbol{\theta}_c$ on the weak order on W are indexed by Coxeter elements c of W.

Definition 10-6.12. Suppose W is a Coxeter group and suppose $s_1 \cdots s_n$ is a reduced word for a Coxeter element c. Consider the half-infinite word

$$(s_1 \cdots s_n)^\infty = s_1 \cdots s_n | s_1 \cdots s_n | s_1 \cdots s_n | s_1 \cdots s_n | \cdots$$

consisting of infinitely many copies of the word $s_1 \cdots s_n$. The symbols "|" are "dividers" which serve to mark the positions of the copies of $s_1 \cdots s_n$. Now suppose $w \in W$. Since every element of S appears infinitely many times in $(s_1 \cdots s_n)^\infty$, each reduced word for w appears infinitely many times as a subword of $(s_1 \cdots s_n)^\infty$. For each such appearance of a reduced word for w, one can record the sequence of positions (from left to right) of the letters chosen from $(s_1 \cdots s_n)^\infty$ to form the subword. For example, if $n = 2$, $m(s_1, s_2) = 3$, and w has a reduced word $s_1 s_2 s_1$, then this reduced word can appear in

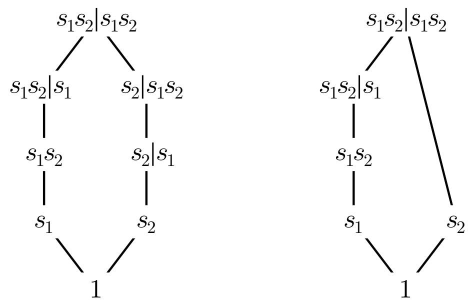

Figure 10-6.6: s_1s_2-sorting words and the s_1s_2-Cambrian lattice in type B_2

$(s_1 \cdots s_n)^\infty = s_1s_2s_1s_2\cdots$ in many sequences of positions, including $(1,2,3)$, $(1,2,5)$, $(1,4,9)$, $(5,16,19)$, etc. Out of all reduced words $a_1 \cdots a_k$ for w and all appearances of $a_1 \cdots a_k$ as a subword of $(s_1 \cdots s_n)^\infty$, there is one subword which occupies the lexicographically smallest sequence of positions. This subword is the $s_1 \cdots s_n$*-sorting word* for w. Since the $s_1 \cdots s_n$-sorting word for w occupies a particular position in $(s_1 \cdots s_n)^\infty$, it determines a sequence of subsets of S: the set U_1 of letters of $a_1 \cdots a_k$ appearing before the first divider, the set U_2 of letters of $a_1 \cdots a_k$ appearing between the first and second dividers, etc. The element w is called $s_1 \cdots s_n$*-sortable* if this sequence of sets is weakly nested $U_1 \supseteq U_2 \supseteq U_3 \supseteq \cdots$. Exercise 10.38 verifies that this notion depends only on c, not on the reduced word chosen for c. Thus we define w to be c-sortable if and only if it is $s_1 \cdots s_n$-sortable for some (equivalently every) reduced word $s_1 \cdots s_n$ for c. Since the main point is usually the choice of c rather than a reduced word for c, we will say "a c-sorting word for w" as shorthand for "the $s_1 \cdots s_n$-sorting word for w where $s_1 \cdots s_n$ is some reduced word for c."

Example 10-6.13. Let W be the Coxeter group (of type B_2) with defining generators $S = \{s_1, s_2\}$ and $m(s_1, s_2) = 4$. Let $c = s_1s_2$. The s_1s_2-sorting words for elements of W are

$$\text{the empty word},\ s_1,\ s_1s_2,\ s_1s_2|s_1,\ s_1s_2|s_1s_2,\ s_2,\ s_2|s_1,\ s_2|s_1s_2.$$

The first six of these represent c-sortable elements, but $s_2|s_1$ and $s_2|s_1s_2$ do not, because $\{s_2\} \not\supseteq \{s_1\}$) and $\{s_2\} \not\supseteq \{s_1, s_2\}$. Figure 10-6.6 shows the weak order on W with elements represented by their s_1s_2-sorting words, and also shows the restriction of the weak order to c-sortable elements. The latter coincides with the c-Cambrian lattice. Our next result shows that this coincidence is not accidental.

$\diamondsuit$ **Theorem 10-6.14.** *Suppose W is a finite Coxeter group and suppose c is a Coxeter element of W. Then an element $w \in W$ is the bottom element of its $\boldsymbol{\theta}_c$-class if and only if it is c-sortable.*

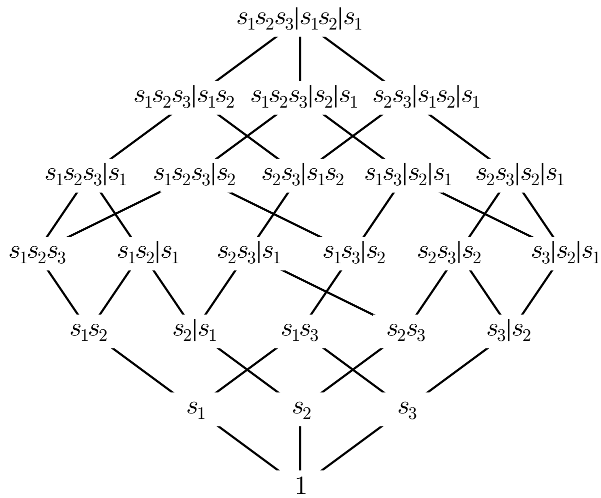

Figure 10-6.7: $s_1s_2s_3$-sorting words for elements of $\mathfrak{S}_4$

We will prove the type-A case of Theorem 10-6.14 as Theorem 10-6.25. By Proposition 9-5.5, we have the following corollary to Theorem 10-6.14.

Corollary 10-6.15. *Suppose W is a finite Coxeter group and suppose c is a Coxeter element of W. Then the c-Cambrian lattice is the subposet of the weak order on W induced by the c-sortable elements of W.*

In fact, more is true. Write $\pi_\downarrow^c$ as an abbreviation for $\pi_\downarrow^{\boldsymbol{\theta}_c}$ and recall from Proposition 9-5.8 that the image of W under $\pi_\downarrow^c$ need not induce a sublattice of the weak order. Cambrian congruences are special in this regard.

$\diamondsuit$ **Theorem 10-6.16.** *Suppose W is a finite Coxeter group and suppose c is a Coxeter element of W. The c-Cambrian lattice, realized as the subposet induced by c-sortable elements, is a sublattice of W.*

Thus the Cambrian lattice (the subposet consisting of c-sortable elements) is a retract of the weak order. The retraction map is $\pi_\downarrow^c$. (Compare Proposition 7-6.9.) Exercise 10.40 is to prove the type-A case of Theorem 10-6.16 using Theorem 10-6.25.

Example 10-6.17. Figure 10-6.7 shows the $s_1s_2s_3$-sorting words for elements of $\mathfrak{S}_4$. The elements of $\mathfrak{S}_4$ are arranged as in the weak order (Figure 10-5.2). Figure 10-6.8 shows the restriction of the weak order to c-sortable elements for $c = s_1s_2s_3$. By Corollary 10-6.15, this is the c-Cambrian lattice. Compare the top-left picture in Figure 10-6.3.

The Cambrian congruences are also special in terms of the shards they remove, or equivalently the join-irreducible elements they contract.

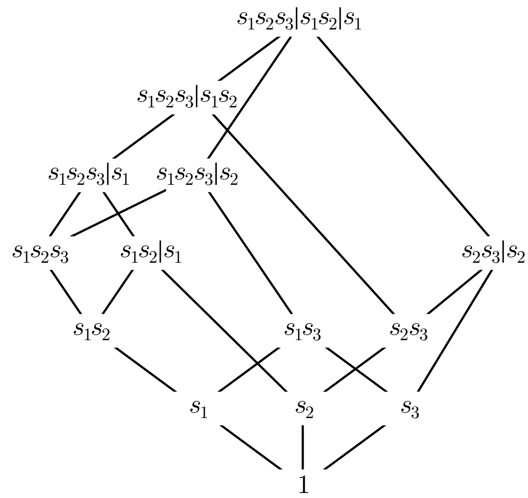

Figure 10-6.8: The c-Cambrian lattice for $c = s_1s_2s_3$ in $\mathfrak{S}_4$

◇ **Theorem 10-6.18.** *Let W be a finite Coxeter group and let c be a Coxeter element of W. For each reflection t, there is exactly one c-sortable join-irreducible element having t as its cover reflection. The Cambrian congruence $\boldsymbol{\theta}_c$ leaves exactly one unremoved shard in each hyperplane of the Coxeter arrangement.*

The two assertions of the theorem are equivalent by Corollary 10-6.15 and Proposition 9-7.8.

A lattice is *extremal* (in the sense of Markowsky) if the following three quantities are equal: the length of the longest chain in the lattice; the number of join-irreducible elements of the lattice; and the number of meet-irreducible elements of the lattice.

◇ **Theorem 10-6.19.** *Every Cambrian lattice is extremal.*

Exercise 10.41 asks the reader to prove Theorem 10-6.19 from the other diamond theorems. One more fact will be needed to complete the exercise. Recall that the maximal element of the weak order on a finite Coxeter group is called w_0, the longest element of W. Exercise 10.42 is to prove the following theorem from some other other diamond theorems.

◇ **Theorem 10-6.20.** *For any Coxeter element c in a finite Coxeter group, the longest element w_0 is c-sortable.*

10-6.4 Induction on length and rank

The most basic and most powerful tool for dealing with sortable elements is an inductive argument made possible by two simple lemmas that we prove below. Most of the results quoted here on sortable elements and Cambrian lattices ultimately rely on this inductive argument. We will give an example of the inductive argument in the proof of Theorem 10-6.25.

Definition 10-6.21. Given a Coxeter element c, a generator $s \in S$ is *initial in c* if there exists a reduced word $s_1 \cdots s_n$ for c with $s_1 = s$. In this case $scs = s_2 \cdots s_n s_1$ is a Coxeter element of W. A generator $s \in S$ is *final in c* if there exists a reduced word $s_1 \cdots s_n$ for c with $s_n = s$. In this case again, scs is a Coxeter element of W.

Definition 10-6.22. Let W be a Coxeter group with defining generators S. Given $I \subseteq S$, the subgroup of W generated by I is called a (*standard*) *parabolic subgroup* and written W_I. Exercise 10.44 is to verify that W_I is a Coxeter group with defining generators I. Most important is the case where I is $S \setminus \{s\}$ for some $s \in S$. We write $\langle s \rangle$ for $S \setminus \{s\}$ and thus $W_{\langle s \rangle}$ for $W_{S \setminus \{s\}}$. In the special case where s is initial in c, the element sc is a Coxeter element for $W_{\langle s \rangle}$.

Lemma 10-6.23. *Let s be initial in c and suppose $w \not\geq s$. Then w is c-sortable if and only if it is an sc-sortable element of $W_{\langle s \rangle}$.*

Proof. Suppose $s_1 \cdots s_n$ is a reduced word for c with $s_1 = s$. The hypothesis that $w \not\geq s$ says that no reduced word for w has s_1 as its first letter. In particular, the first letter of the $s_1 \cdots s_n$-sorting word for w is not s_1. That means that $s_i \notin U_1$, in the notation of Definition 10-6.12. Therefore if w is c-sortable, s_1 does not appear in the $s_1 \cdots s_n$-sorting word for w, so w is in $W_{\langle s \rangle}$. We have $sc = s_2 \cdots s_n$, and the $s_2 \cdots s_n$-sorting word for w is exactly the same sequence of letters as the $s_1 \cdots s_n$-sorting word for w. The corresponding sequences of subsets also agree, and we conclude that w is sc-sortable.

Conversely, suppose w is an sc-sortable element of $W_{\langle s \rangle}$. Since w is in $W_{\langle s \rangle}$, there is some word for w as a product of generators in $\langle s \rangle$. Theorem 10-4.1 implies that every reduced word for w contains only generators in $\langle s \rangle$. (Compare Exercise 10.23.) In particular, the $s_1 \cdots s_n$-sorting word for w contains only generators in $\langle s \rangle$. Considering $(s_2 \cdots s_n)^\infty$ as a subword of $(s_1 \cdots s_n)^\infty$ in the natural way, we see that the $s_1 \cdots s_n$-sorting word for w (the lexicographically leftmost reduced word for w in $(s_1 \cdots s_n)^\infty$) coincides with the $s_1 \cdots s_n$-sorting word for w (the lexicographically leftmost reduced word for w in $(s_2 \cdots s_n)^\infty$). Since w is sc-sortable, we conclude that w is c-sortable. $\square$

Lemma 10-6.24. *Let s be initial in c and suppose $w \geq s$. Then w is c-sortable if and only if sw is scs-sortable.*

Proof. Again write $c = s_1 \cdots s_n$ with $s_1 = s$. The Coxeter element $s_2 \cdots s_n s_1$ is scs. Since $w \geq s$, the operation of attaching the letter s to the beginning of a word establishes a bijection between the set of reduced words for sw

and the set of reduced words for w starting with s. We conclude that the $s_2 \cdots s_n s_1$-sorting word for sw is obtained by deleting the first letter (s) from the $s_1 \cdots s_n$-sorting word for w. Knowing that the $s_1 \cdots s_n$-sorting word for w starts with $s = s_1$, the criterion for w to be c-sortable is exactly the criterion for sw to be scs-sortable. □

Lemmas 10-6.23 and 10-6.24 make possible inductive arguments on the length of an element of W and on the rank of W. Induction on length comes in the case where Lemma 10-6.24 applies and we pass from considering w to considering sw. (The condition $w \geq s$ means that w has a reduced word starting with s, which is equivalent to the condition that $\ell(sw) < \ell(w)$.) Induction on rank comes in the case where Lemma 10-6.23 applies and we pass from considering W to considering $W_{\langle s \rangle}$. The induction is not circular: When we appeal to the inductive hypothesis by shortening the length of an element, we do so within W, so the rank is unchanged. The induction involves passing to different Coxeter elements of W and passing to parabolic subgroups of W. As a base case for the induction, we can take the case of the identity element in a trivial Coxeter group W with $S = \varnothing$, although it may be simpler in practice to take the base case to be the identity element in any W. We see that the sets of c-sortable elements, for all c, are uniquely defined by taking the identity element to be c-sortable for any c and by Lemmas 10-6.23 and 10-6.24.

10-6.5 Sortable elements of type A

As an example of how to apply Lemmas 10-6.23 and 10-6.24 to obtain properties of sortable elements and Cambrian congruences/lattice/fans, we prove Theorem 10-6.14 for Coxeter groups of type A. For convenience, we argue in type A_{n-1} so that we can consider permutations in $\mathfrak{S}_n$. To prove Theorem 10-6.9, we showed that a permutation $\pi \in \mathfrak{S}_n$ is the bottom element of its $\boldsymbol{\theta}_c$-class if and only if it avoids both $31\underline{2}$ and $\overline{2}31$. Thus the type-A case of Theorem 10-6.14 is an immediate corollary of the following theorem.

Theorem 10-6.25. *Suppose c is a Coxeter element of $\mathfrak{S}_n$, encoded as a barring. A permutation in $\mathfrak{S}_n$ is c-sortable if and only if it avoids the patterns $31\underline{2}$ and $\overline{2}31$.*

Proof. The proof will use several aspects of the combinatorics of the symmetric group proved in Section 10-5 and a few additional exercises. Exercise 10.45 shows that $s_i \leq \pi$ in the weak order if and only if the value i appears after the value $i+1$ in the sequence $\pi_1 \cdots \pi_n$. Exercise 10.46 shows that $\pi = \pi_1 \cdots \pi_n$ is in $W_{\langle s_i \rangle}$ if and only if $\{\pi_1, \ldots, \pi_i\} = \{1, \ldots, i\}$.

We argue by induction on length and rank as described above. As a base case, the identity permutation $12 \cdots n$ satisfies the avoidance conditions and is c-sortable. Let $\pi = \pi_1 \cdots \pi_n$ be a permutation in $\mathfrak{S}_n$.

Suppose s_i is initial in c. Then the edge(s) in the diagram incident to s_i are oriented away from s_i in the orientation c. The barring for c has $\underline{i+1}$ (if $i<n$) and $\overline{i}$ (if $i>1$).

If $s_i \le \pi$ then the value i appears after the value $i+1$ in $\pi_1 \cdots \pi_n$. Consider the permutation $s_i\pi$, obtained from π by swapping the values i and $i+1$ in the sequence $\pi_1 \cdots \pi_n$. The orientation s_ics_i is the same as c except that the orientation of edges incident to s_i is reversed. The barring for s_ics_i agrees with the barring for c except that the barring for s_ics_i is $\overline{i+1}$ (if $i<n$) and $\underline{i}$ (if $i>1$). We claim that π contains a pattern $31\underline{2}$ or $\overline{2}31$ in the barring for c if and only if $s_i\pi$ contains a pattern $31\underline{2}$ or $\overline{2}31$ in the barring for s_ics_i. Suppose def is a subsequence of π and let jkl be the subsequence of $s_i\pi$ occupying the same positions. If $\{i,i+1\}\cap\{d,e,f\}=\varnothing$, then def and jkl are the same subsequence and have the same barring in π as in $s_i\pi$. If $\{i,i+1\}\cap\{d,e,f\}=\{i\}$, then jkl is obtained from def by replacing i with $i+1$. The barring of i in π is the same as the barring of $i+1$ in $s_i\pi$, so jkl is a $31\underline{2}$- or $\overline{2}31$-pattern if and only if def is. If $\{i,i+1\}\cap\{d,e,f\}=\{i+1\}$ then we argue similarly. If $\{i,i+1\}\cap\{d,e,f\}=\{i,i+1\}$, then def cannot form a $31\underline{2}$- or $\overline{2}31$-pattern since $\underline{i+1}$ appears before $\overline{i}$ in π. In this case, $\{i,i+1\}\cap\{j,k,l\}=\{i,i+1\}$ as well, and jkl cannot form a $31\underline{2}$- or $\overline{2}31$-pattern since $\underline{i}$ appears before $\overline{i+1}$ in $s_i\pi$. We have proved the claim.

By induction on length, $s_i\pi$ is s_ics_i-sortable if and only if it satisfies to avoidance conditions with the barring associated to s_ics_i. We have shown that this avoidance condition on $s_i\pi$ is equivalent to the analogous avoidance condition on π relative to c, so by Lemma 10-6.24, we conclude that π is c-sortable if and only if it satisfies the avoidance condition.

On the other hand, if $s_i \not\le \pi$ then the value i appears before the value $i+1$ in $\pi_1 \cdots \pi_n$. The parabolic subgroup $W_{\langle s_i\rangle}$ is isomorphic to $\mathfrak{S}_i \times \mathfrak{S}_{n-i}$. We can harmlessly realize $\mathfrak{S}_{n-i}$ as the group of permutations of $\{i+1,\ldots,n\}$ and adjust all definitions accordingly. The Coxeter element s_ic of $W_{\langle s_i\rangle}$ can be written c_1c_2, where c_1 is a Coxeter element of $\mathfrak{S}_i$ and c_2 is a Coxeter element of $\mathfrak{S}_{n-i}$. The barring defined by c_1 and by c_2 agrees with the barring defined by c except that we may ignore the barring on i and on $i+1$. We claim that π avoids the patterns $31\underline{2}$ and $\overline{2}31$ if and only if $\{\pi_1,\ldots,\pi_i\}=\{1,\ldots,i\}$ and the subsequences $\pi_1\cdots\pi_i$ and $\pi_{i+1}\cdots\pi_n$ both avoid $31\underline{2}$ and $\overline{2}31$.

If $\{\pi_1,\ldots,\pi_i\}=\{1,\ldots,i\}$, then any occurrence of $31\underline{2}$ or $\overline{2}31$ in π must occur within the subsequence $\pi_1\cdots\pi_i$ or the subsequence $\pi_{i+1}\cdots\pi_n$, so we have proved the "if" direction. Conversely, suppose π avoids $31\underline{2}$ and $\overline{2}31$. Then any subsequence of π avoids these patterns. Suppose some element $b\in\{i+1,\ldots,n\}$ precedes some element $a\in\{1,\ldots,i\}$ in π. If a and b both precede $\underline{i+1}$, then $ba\underline{(i+1)}$ is a $31\underline{2}$-pattern in π. If they both follow $\overline{i}$, then $\overline{i}ba$ is a $\overline{2}31$-pattern in π. Since i precedes $i+1$ in π, the only remaining possibility is that the four elements are distinct and form a subsequence $b\overline{i}\underline{(i+1)}a$. In this case, $b\overline{i}\underline{(i+1)}$ is a $31\underline{2}$-pattern in π. We conclude by these contradictions that $\{\pi_1, \ldots, \pi_i\} = \{1, \ldots, i\}$, and we have proved the claim.

Suppose now that π is c-sortable. Then Lemma 10-6.23 says that π is in $W_{\langle s_i \rangle}$ and is $s_i c$-sortable. Exercise 10.48 says that $\pi_1 \cdots \pi_i$ is c_1-sortable and $\pi_{i+1} \cdots \pi_n$ is c_2-sortable. By induction on rank, $\pi_1 \cdots \pi_i$ and $\pi_{i+1} \cdots \pi_n$ both avoid $31\underline{2}$ and $\overline{2}31$. The claim implies that π avoids $31\underline{2}$ and $\overline{2}31$. Conversely, suppose π avoids $31\underline{2}$ and $\overline{2}31$. The claim says that π is in $W_{\langle s_i \rangle}$ and that $\pi_1 \cdots \pi_i$ and $\pi_{i+1} \cdots \pi_n$ both avoid $31\underline{2}$ and $\overline{2}31$. By induction on rank, $\pi_1 \cdots \pi_i$ is c_1-sortable and $\pi_{i+1} \cdots \pi_n$ is c_2-sortable, so Exercise 10.48 says that π is $s_i c$-sortable. Finally, Lemma 10-6.23 says that π is c-sortable. □

10-6.6 Sortable elements and the Cambrian fan

Another benefit of realizing Cambrian lattices in terms of sortable elements is that each sortable element has a sorting word that contains a lot of readily available and meaningful combinatorial information. For example, sorting words contain information that allows a direct combinatorial construction of the Cambrian fan. Specifically, we show in this section how to read off, from a c-sorting word for a c-sortable element v, vectors that define the $\boldsymbol{\theta}_c$-cone for the $\boldsymbol{\theta}_c$-class of v.

Let $a_1 \cdots a_k$ be a c-sorting word for v and let $r \in S$. Recall that $a_1 \cdots a_k$ is the lexicographically leftmost subword of $(s_1 \cdots s_n)^\infty$ which is a reduced word for v, where $s_1 \cdots s_n$ is some reduced word for c. Among all instances of r in $(s_1 \cdots s_n)^\infty$, there is a leftmost instance that is not in the subword $a_1 \cdots a_k$. There is some i such that this leftmost instance of r in $(s_1 \cdots s_n)^\infty$ occurs between the location of a_i and the location of a_{i+1}. (If the leftmost instance of r occurs before the position of a_1, then we set $i = 0$. If the leftmost instance of r occurs after the position of a_k, then we set $i = k$.) We say that v *skips r in position i*. Define a set of vectors (in fact roots) by

$$C_c(v) = \{a_1 \cdots a_i \alpha_r \mid r \in S,\ v \text{ skips } r \text{ in position } i\},$$

where α_r is the simple root associated to r. The set $C_c(v)$ can also be defined recursively by induction on length and rank, thereby making possible inductive proofs using Lemmas 10-6.23 and 10-6.24. See Exercise 10.49.

The set $C_c(v)$ is linearly independent for each c-sortable element v (Exercise 10.50), so we can define a simplicial cone

$$\mathrm{Cone}_c(v) = \bigcap_{\beta \in C_c(v)} \{x \in \mathbb{R}^n \mid \langle x, \beta \rangle \geq 0\},$$

where $\langle \cdot\,, \cdot \rangle$ is the Euclidean inner product. Recall that $\boldsymbol{\theta}_c$ is the c-Cambrian congruence and that each $\boldsymbol{\theta}_c$-cone is the union of cones wB over some $\boldsymbol{\theta}_c$-class. (Here $\mathcal{A}$ is an associated Coxeter arrangement and B a base region as in Section 10-3.) In particular, by Theorem 10-6.14, each $\boldsymbol{\theta}_c$-cone contains exactly one cone vB such that v is c-sortable.

v	$C_c^{s_1}(v)$	$C_c^{s_2}(v)$	$C_c(v)$
1	α_1	α_2	$\{\alpha_1, \alpha_2\}$
s_1	$s_1\alpha_1$	$s_1\alpha_2$	$\{-\alpha_1, 2\alpha_1 + \alpha_2\}$
s_1s_2	$s_1s_2\alpha_1$	$s_1s_2\alpha_2$	$\{\alpha_1 + \alpha_2, -2\alpha_1 - \alpha_2\}$
$s_1s_2s_1$	$s_1s_2s_1\alpha_1$	$s_1s_2s_1\alpha_2$	$\{-\alpha_1 - \alpha_2, \alpha_2\}$
$s_1s_2s_1s_2$	$s_1s_2s_1s_2\alpha_1$	$s_1s_2s_1s_2\alpha_2$	$\{-\alpha_1, -\alpha_2\}$
s_2	α_1	$s_2\alpha_2$	$\{\alpha_1, -\alpha_2\}$

Figure 10-6.9: The map C_c for W of type B_2 and $c = s_1s_2$.

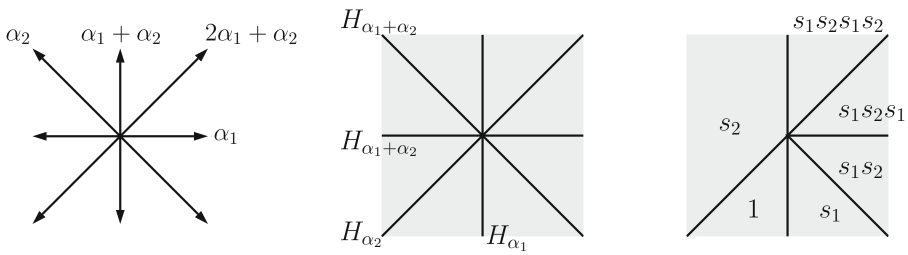

Figure 10-6.10: The root system, Coxeter arrangement, and c-Cambrian fan for W of type B_2 and $c = s_1s_2$ (Example 10-6.27)

$\diamondsuit$ **Theorem 10-6.26.** *Suppose W is a finite Coxeter group and c is a Coxeter element. If v is c-sortable, then the $\boldsymbol{\theta}_c$-cone associated to v is* $\mathrm{Cone}_c(v)$. *That is, for $w \in W$, we have $v \equiv w \pmod{\boldsymbol{\theta}_c}$ if and only if $wB \in \mathrm{Cone}_c(v)$.*

Example 10-6.27. Figure 10-6.9 shows the map C_c on c-sortable elements for W and c as in Example 10-6.13. The root system, Coxeter arrangement and c-Cambrian fan are shown in Figure 10-6.10. Each maximal cone $C_c(v)$ of the c-Cambrian fan is labeled with v in the figure Exercise 10.51 is to carry out the same computations for $W = \mathfrak{S}_4$ and $c = s_1s_2s_3$ as in Example 10-6.17.

10-6.7 Coxeter-Catalan combinatorics

The sorting words for sortable elements also contain information that allows us to connect sortable elements bijectively to other (*a priori* unrelated) objects. In particular, sortable elements and Cambrian lattices enter into the realm of *Coxeter-Catalan combinatorics*,the study of families of objects counted by the *Coxeter-Catalan numbers*.

Definition 10-6.28. As verified in Exercise 10.52, all Coxeter elements in a finite Coxeter group W are in the same conjugacy class. Thus the order h of any Coxeter element c is a well-defined invariant of W, called the *Coxeter number* of W. The eigenvalues of c, as a linear transformation, are of the form $e^{\frac{2\pi ki}{h}}$ for integers k between 1 and $h-1$. The n values of k appearing are called the *exponents* of W and written $e_1, \ldots, e_n$. (These might not all be distinct.)

Example 10-6.29. For W of type A_n, every Coxeter element is an $(n+1)$-cycle in $\mathfrak{S}_{n+1}$. Thus the Coxeter number is $n+1$. The exponents are $1, 2, \ldots, n$. In type B_n, the Coxeter number is $2n$ and the exponents are $1, 3, \ldots, 2n-1$. In type D_n, the Coxeter number is $2n-2$ and the exponents are $1, 3, \ldots, 2n-3$ and $n-1$. See the Notes for references to more information about exponents.

Definition 10-6.30. Suppose W is a finite irreducible Coxeter group. The *Coxeter-Catalan number* for W, also known as the *W-Catalan number* is

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

Example 10-6.31. The Coxeter-Catalan numbers for finite irreducible Coxeter groups are shown here. In particular, the type-A Catalan number is the usual Catalan number.

A_n	B_n	D_n	E_6	E_7	E_8	F_4	H_3	H_4	$I_2(m)$
$\frac{1}{n+2}\binom{2n+2}{n+1}$	$\binom{2n}{n}$	$\frac{3n-2}{n}\binom{2n-2}{n-1}$	833	4160	25080	105	32	280	$m+2$

The W-Catalan number counts, among other things, combinatorial clusters (vertices of the generalized associahedron), c-noncrossing partitions, antichains in the root poset (a certain partial order on the positive roots), and c-sortable elements. Initially, these counts were established independently and the coincidence between the counts in different contexts was unexplained. Coxeter-Catalan combinatorics is a name for the (still unfinished) program to try to explain all of the numerical congruences and furthermore to look for deeper relationships between the underlying structures in different contexts.

We will not give precise definitions of any of these Coxeter-Catalan objects here, but references are given in the Notes. Instead, our goal is to pique the reader's curiosity by showing how c-noncrossing partitions and combinatorial clusters each arise in the context of sortable elements. We will quote two theorems, each of which separately implies the following theorem.

♢ **Theorem 10-6.32.** *Suppose W is a finite irreducible Coxeter group. For any Coxeter element c, the number of c-sortable elements is* $\mathrm{Cat}(W)$.

The c-noncrossing partitions are certain elements of W which generalize the classical noncrossing partitions first defined in [287]. The definition of c-noncrossing partitions also produces a lattice structure on c-noncrossing partitions called the c-noncrossing partition lattice. We define a map nc_c from c-sortable elements to elements of W by sending v to the product of the cover reflections of v. (See Definition 10-3.8.) The order of multiplication

v	Product of cover reflections	$\mathrm{nc}_c(v)$
1		1
s_1	s_1	s_1
s_1s_2	$s_1s_2s_1$	$s_1s_2s_1$
$s_1s_2\vert s_1$	$s_1s_2s_1s_2s_1$	$s_2s_1s_2$
$s_1s_2\vert s_1s_2$	$s_1 \cdot s_1s_2s_1s_2s_1s_2s_1$	s_1s_2
s_2	s_2	s_2

Figure 10-6.11: The map nc_c

is determined as follows. Fix a c-sorting word $a_1 \cdots a_k$ for v. Each cover reflection t is in particular an inversion of v, and thus by Proposition 10-3.4 equals $t_i = a_1 \cdots a_i \cdots a_1$ for a unique i between 1 and k. If the cover reflections of v are $t_{i_1}, \ldots, t_{i_j}$ with $i_1 < \cdots < i_j$, then $\mathrm{nc}_c = t_{i_1} \cdots t_{i_j}$.

$\diamondsuit$ **Theorem 10-6.33.** *The map* nc_c *is a bijection from c-sortable elements to c-noncrossing partitions.*

Theorem 10-6.33 is closely related to the following theorem. Recall from Remark 9-8.15 that each quotient of the weak order defines a join-subsemilattice of the shard intersection order on elements of W.

$\diamondsuit$ **Theorem 10-6.34.** *The c-sortable elements of W constitute a sublattice of the shard intersection order on W. The map* nc_c *is an isomorphism from this sublattice to the c-noncrossing partition lattice.*

Example 10-6.35. Figure 10-6.11 shows the map nc_c on c-sortable elements for W of type B_2 and $c = s_1s_2$ as in Example 10-6.13. Exercise 10.53 is to carry out the same computations for $W = \mathfrak{S}_4$ and $c = s_1s_2s_3$ as in Example 10-6.17. The exercise also relates the output of the calculations to the classical noncrossing partitions.

Finally, recall that Theorem 10-6.7 implies a bijection between the elements of the c-Cambrian lattice (the c-sortable elements) and the vertices of the generalized associahedron. These vertices are in bijection with *combinatorial clusters*. Each combinatorial cluster consists of n linearly independent roots, and thus the nonnegative linear span of the cluster is a full-dimensional cone. Together, these cones define a complete fan called the c-cluster fan. Given a c-sortable element w, the corresponding combinatorial cluster is easily read off from a c-sorting word for w as follows.

Suppose w is c-sortable and let $a_1 \cdots a_k$ be a c-sorting word for w. We define a root $\mathrm{cl}_c^s(w)$ for each $s \in S$. If s occurs as a letter in $a_1 \cdots a_k$, then take i to be the position of the last (rightmost) occurrence of s in $a_1 \cdots a_k$ and define $\mathrm{cl}_c^s(w)$ to be $a_1 \cdots a_{i-1}\alpha_s$. Otherwise, define $\mathrm{cl}_c^s(w)$ to be $-\alpha_s$. Here as before, α_s is the simple root associated to s. Define $\mathrm{cl}_c(w) = \{\mathrm{cl}_c^s(w) \mid s \in S\}$.

v	$\mathrm{cl}_c^{s_1}(v)$	$\mathrm{cl}_c^{s_2}(v)$	$\mathrm{cl}_c(v)$
1	$-\alpha_1$	$-\alpha_2$	$\{-\alpha_1, -\alpha_2\}$
s_1	α_1	$-\alpha_2$	$\{\alpha_1, -\alpha_2\}$
s_1s_2	α_1	$s_1\alpha_2 = 2\alpha_1 + \alpha_2$	$\{\alpha_1, 2\alpha_1 + \alpha_2\}$
$s_1s_2s_1$	$s_1s_2\alpha_1 = \alpha_1 + \alpha_2$	$s_1\alpha_2 = 2\alpha_1 + \alpha_2$	$\{\alpha_1 + \alpha_2, 2\alpha_1 + \alpha_2\}$
$s_1s_2s_1s_2$	$s_1s_2\alpha_1 = \alpha_1 + \alpha_2$	$s_1s_2s_1\alpha_2 = \alpha_2$	$\{\alpha_1 + \alpha_2, \alpha_2\}$
s_2	$-\alpha_1$	α_2	$\{-\alpha_1, \alpha_2\}$

Figure 10-6.12: The map cl_c

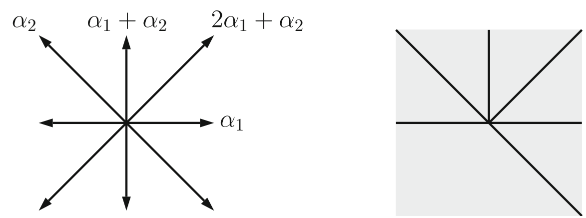

Figure 10-6.13: The root system and c-cluster fan for Example 10-6.37

The following theorem adds detail to Theorem 10-6.7.

$\diamondsuit$ **Theorem 10-6.36.** *The map* cl_c *is a bijection from c-sortable elements to combinatorial clusters (with respect to c). The bijection extends to a graph isomorphism from the unoriented Hasse diagram of the c-Cambrian lattice to the vertex-edge graph of the generalized associahedron. The bijection extends further to an isomorphism from the c-Cambrian fan to the c-cluster fan.*

Example 10-6.37. Figure 10-6.12 shows the map cl_c on c-sortable elements for W and c as in Example 10-6.13. The simple root associated to s_i is α_i, and the root system is shown as the left picture in Figure 10-6.13. The c-cluster fan appears as the right picture in Figure 10-6.13. Exercise 10.54 is to carry out the same computations for $W = \mathfrak{S}_4$ and c as in Example 10-6.17. The exercise also illustrates the isomorphism of fans.

10-7. Some other lattice quotients of the weak order

The contents of this chapter suggest that lattice congruences on the weak order "know" a lot of combinatorics and geometry related to Coxeter groups. We conclude the chapter with some other interesting congruences and quotients of the weak order. The goal is to spark interest, so we give few details and no proofs. References to additional information are found in the Notes.

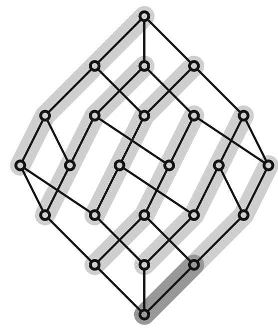

Figure 10-7.1: A congruence on the weak order of type A_3 whose quotient is the weak order of type A_2

Example 10-7.1 (**Parabolic congruences**). Suppose W is a finite Coxeter group with defining generators S, and let I be a subset of S. Then the parabolic subgroup W_I (see Definition 10-6.22) is a Coxeter group in its own right with defining generators I (Exercise 10.44). The weak order on W_I is a lower interval in the weak order on W (Exercise 10.55), and thus in particular a sublattice. In fact, the weak order on W_I is a retract of the weak order on W. Specifically, given $w \in W$, there exists a unique element w_I of W_I such that $\mathrm{inv}(w_I) = \mathrm{inv}(w) \cap W_I$. The map $w \mapsto w_I$ is a surjective lattice homomorphism from W to W_I with $(w_I)_I = w_I$. Therefore the weak order on W_I is the quotient of the weak order on I modulo some congruence $\boldsymbol{\alpha}$, called a *parabolic congruence*. The congruence $\boldsymbol{\alpha}$ has a very simple description: it is $\bigvee_{s\in(S\setminus I)} \mathrm{con}(1, s)$, the smallest congruence contracting each $s \in S \setminus I$. Figure 10-7.1 shows the parabolic congruence on the weak order in type A_3 whose quotient is the weak order in type A_2. This is the smallest congruence contracting the edge shaded darker in the picture.

Example 10-7.2 (**Diagram homomorphisms**). If W and W' are finite Coxeter systems with the same defining generators S and with $m'(s,t) \le m(s,t)$ for each pair $s, t \in S$, then we say that W *dominates* W'. That is, the Coxeter diagram of W' is obtained from the Coxeter diagram of W by decreasing the labels on edges or by erasing edges. Surprisingly, in this case, the weak order on W' is the quotient of the weak order on W modulo a congruence that admits a simple description. We call a surjective lattice homomorphism from W to W' a *diagram homomorphism*.

Recall from Section 10-6.1 that at the bottom of the weak order, for each $r, s \in S$ there is a polygon $[1, r \vee s]$ consisting of $2m(r,s)$ elements. For example, for W of type B_3, these intervals are an octagon, a hexagon, and a square, and for W' of type A_3 they are two hexagons and a square. One can turn the octagonal interval into a hexagonal interval by contracting two of its

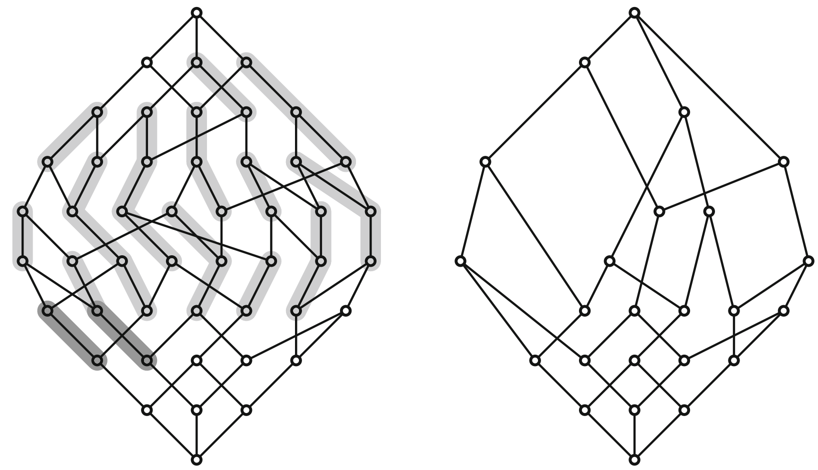

Figure 10-7.2: A congruence on the weak order of type B_3 and its quotient, the weak order of type A_3

side edges, one on each side. There are four choices of how to do this. For three of the four choices, the quotient of W modulo the smallest congruence contracting the chosen edges is isomorphic to the weak order on W'. (For the other choice, to obtain W' we must choose one additional edge to contract elsewhere in W.) The left picture of Figure 10-7.2 shows a quotient on the weak order on W of type B_3. This is the smallest congruence contracting the two edges shaded darker in the picture. The right picture of the figure shows the quotient modulo the congruence from the left picture. Comparison with Figure 10-3.1 shows that this quotient is isomorphic to the weak order on W' of type A_3.

The lattice homomorphisms between weak orders on different finite Coxeter groups descend to lattice homomorphisms between different Cambrian lattices.

The examples above show that lattice theory pierces to the heart of the structure of Coxeter groups. The next few examples show other combinatorial structures arising surprisingly from lattice congruences on the weak order on permutations. In each case, there is a partial order (a lattice quotient of the weak order) on the objects described. This partial order is shown in a figure, but we leave the combinatorial description of the partial order to the primary sources cited in the Notes to this chapter.

Example 10-7.3 (**Sashes**)**.** A *sash* is a tiling of a $1 \times n$ rectangle by white 1×1 squares, black 1×1 squares, and/or white 1×2 rectangles. Sashes are counted by the Pell numbers. There is a quotient of the weak order on $\mathfrak{S}_{n+1}$ whose elements are labeled by sashes in a $1 \times n$ rectangle and whose cover relations correspond to simple modifications of sashes. An example is shown

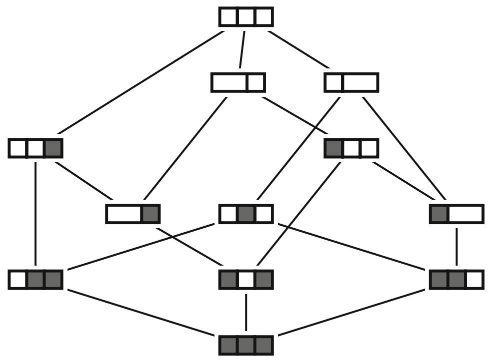

Figure 10-7.3: A lattice of sashes

in Figure 10-7.3.

Example 10-7.4 (**Diagonal rectangulations**)**.** Start with a square with horizontal and vertical sides. A *diagonal rectangulation* is a decomposition of the square into rectangles such that the interior of each rectangle intersects the top-left-to-bottom-right diagonal of the square. (We consider diagonal rectangulations up to combinatorial equivalence.) The diagonal rectangulations are counted by the Baxter numbers. There is a lattice quotient of the weak order on $\mathfrak{S}_n$ whose elements are labeled by diagonal rectangulations with n rectangles and whose cover relations are "pivots" that replace vertical segments by horizontal segments. This lattice appears in Figure 10-7.4 for $n = 4$.

Example 10-7.5 (**Generic rectangulations**)**.** A *generic rectangulation* is a decomposition of the square into rectangles such that no four rectangles have a common corner. Again, we consider these up to combinatorial equivalence. Every diagonal rectangulation is a generic rectangulation, but not *vice versa*. The generic rectangulations also occur as a lattice quotient of the weak order on permutations in $\mathfrak{S}_n$. This partial order is shown in Figure 10-7.5 for $n = 4$. The partial order shown coincides with the weak order on $\mathfrak{S}_4$, but for larger n, the lattice of generic rectangulations is the quotient of the weak order modulo a nontrivial congruence.

The previous two examples concern lattice congruences on the weak order on permutations. The next two examples concern more general finite Coxeter groups.

Example 10-7.6 (**Coxeter-biCatalan combinatorics**)**.** Given a finite Coxeter group W and a Coxeter element c, the *c-biCambrian congruence* $\boldsymbol{\beta}_c$ is defined to be the meet, in $\mathrm{Con}(W)$, of the c-Cambrian congruence $\boldsymbol{\theta}_c$ and the c^{-1}-Cambrian congruence $\boldsymbol{\theta}_{c^{-1}}$. That is, two elements of W are congruent

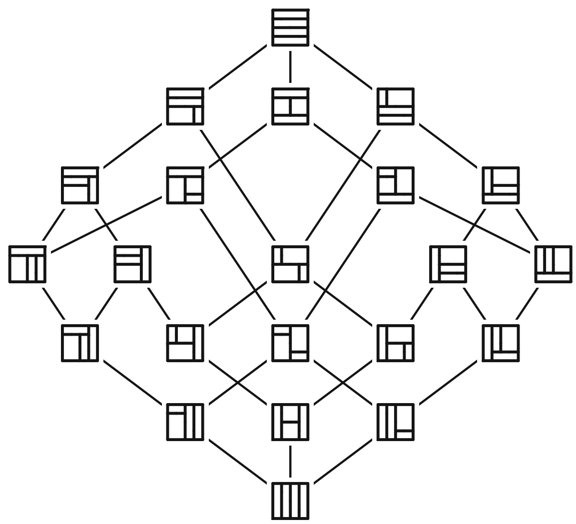

Figure 10-7.4: A lattice of diagonal rectangulations

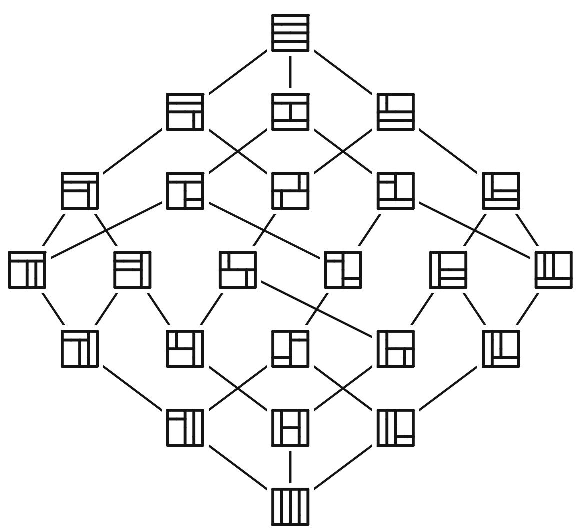

Figure 10-7.5: A lattice of generic rectangulations

modulo $\boldsymbol{\beta}_c$ if and only if they are congruent modulo $\boldsymbol{\theta}_c$ and congruent modulo $\boldsymbol{\theta}_{c^{-1}}$. The number of c-biCambrian congruence classes depends on the choice of c.

The idea of the biCambrian congruence is lattice-theoretically simple, but one might not expect it to have any combinatorial significance. Surprisingly, it does. A motivating case of the c-biCambrian congruence is the case where W is of type A_n and c orients the diagram as a directed path. In that case, the c-biCambrian congruence is the congruence considered in Example 10-7.4, so c-biCambrian congruence classes are counted by the Baxter number. Since not all Coxeter diagrams are paths, it is hard to generalize this example to arbitrary Coxeter groups.

However, there is a different, more uniform way to choose a Coxeter element for a given finite Coxeter group. Since, by Theorem 10-2.19, the Coxeter diagram of a finite Coxeter group is a forest, it is in particular a bipartite graph. Writing $S = S_+ \cup S_-$ for a bipartition of the diagram, we construct a *bipartite Coxeter element* by orienting each edge of the diagram from S_- to S_+. The bipartite biCambrian congruence offers a combinatorial surprise: Congruence classes in the bipartite biCambrian congruence are in bijection with antichains in the *doubled root poset*. (This is the union of the root poset with a dual copy of the root poset, identified on the simple roots.) In fact, the standard Coxeter-Catalan objects (e.g. noncrossing partitions, clusters) each have "biCatalan" analogs, and thus there is an entire theory of *Coxeter-biCatalan combinatorics* that parallels Coxeter-Catalan combinatorics.

In type A, the bipartite Cambrian congruence classes are in bijection with a subset of the noncrossing arc diagrams called *alternating arc diagrams*. These consist of points on a vertical line and noncrossing arcs connecting the points, with each arc alternating between left and right as it passes the intervening points. The bipartite biCambrian lattice of type A_3 (the quotient of the weak order of type A_3 modulo the bipartite biCambrian congruence) is shown in Figure 10-7.6 with congruence classes represented by alternating arc diagrams.

We conclude with an example linking lattice congruences of the weak order to representation theory.

Example 10-7.7 (**Algebraic congruences**). A finite Coxeter group is *simply laced* if $m(s,t) \in \{2,3\}$ for distinct $s,t \in S$. Associated to each simply laced finite Coxeter group W is a *preprojective algebra* Π. This is a quotient of the algebra of paths in the Coxeter diagram with product given by concatenation. The weak order on W is isomorphic to the inclusion order on the *torsion classes* of Π. If Π/I is an quotient (in the usual algebraic sense) of Π, then the inclusion order on torsion classes of Π/I is naturally isomorphic to a quotient of the weak order on W. Not all lattice quotients of W arise in this way. The quotients that arise are called *algebraic quotients*. One naturally wonders which lattice quotients of W are algebraic. Some examples of algebraic quotients include the Cambrian lattices. In type A, the bipartite biCambrian lattices are

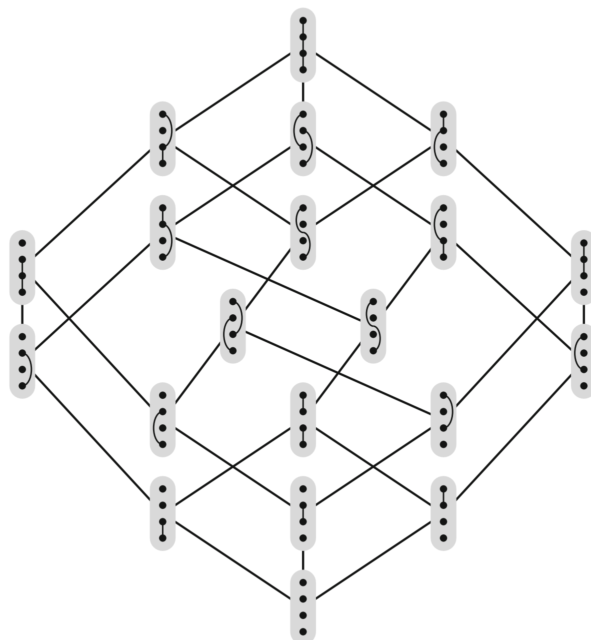

Figure 10-7.6: The bipartite biCambrian lattice of type A_3

also algebraic quotients, but at the time of this writing it is unknown whether the same holds outside of type A. The quotient in Example 10-7.4 associated with diagonal rectangulations is not algebraic. In type A, algebraic quotients are completely characterized: A quotient is algebraic if and only if the graph underlying its Hasse diagram is regular. Furthermore, this happens if and only if the corresponding congruence is the smallest congruence contracting some set of double join-irreducible elements. (A *double join-irreducible element* is a join-irreducible element j such that j_* is either join-irreducible or equals 0.)

10-8. Exercises

Finite reflection groups

10.1. Suppose t is a reflection acting on $\mathbb{R}^n$ with reflecting hyperplane H_t and suppose w is an orthogonal transformation on $\mathbb{R}^n$. Show that wtw^{-1} is a reflection with reflecting hyperplane wH_t.

10.2. Recall from the proof of Proposition 10-2.3 that a line of $\mathcal{A}$ is a 1-dimensional linear subspace of $\mathbb{R}^n$ that is the intersection of some collection of hyperplanes in $\mathcal{A}$. Suppose $\mathcal{A}$ is an essential hyperplane arrangement and suppose $H \in \mathcal{A}$. Show that there exists a line in $\mathcal{A}$ that is not contained in H.

10.3. Prove Proposition 10-2.6.

10.4. Prove Proposition 10-2.11.

10.5. Let W be a Coxeter group with defining generators S, and suppose r and s are distinct elements of S. Prove that r and s commute if and only if $m(r, s) = 2$. (*One direction of this proof needs Proposition 10-2.17, which was proved here only for W finite. See the Notes for a reference to a proof of Proposition 10-2.17 in general.*)

10.6. Let W be a Coxeter group with defining generators S. Suppose S is the disjoint union $S_1 \cup S_2$ with $m(s, t) = 2$ for all $s \in S_1$ and $t \in S_2$. In other words, S_1 and S_2 are the vertices of two (not necessarily connected) components of the Coxeter diagram of W. Let W_1 be the subgroup of W generated by S_1 and let W_2 be the subgroup of W generated by S_2. Show that W_1 and W_2 are Coxeter groups and that W is isomorphic to the direct product $W_1 \times W_2$.

10.7. Let W be a Coxeter group with defining generators S. Suppose S is the disjoint union $S_1 \cup S_2$, but don't make any *a priori* assumptions on the function $m(\cdot\,,\,\cdot)$. Define W_1, and W_2 as in Exercise 10.6. Suppose there is an isomorphism $W \cong W_1 \times W_2$ such that $s \mapsto (s, 1)$ for each $s \in S_1$ and $s \mapsto (1, s)$ for each $s \in S_2$. Prove that $m(s, t) = 2$ for all $s \in S_1$ and $t \in S_2$. (*This exercise requires Proposition 10-2.17. See the parenthetical comment to Exercise 10.5.*)

10.8. Find the smallest (in terms of the number of hyperplanes) simplicial arrangement that is not combinatorially isomorphic to some Coxeter arrangement. (*Use Theorem 10-2.21. Look in rank* 3 *and add a hyperplane to some Coxeter arrangement listed in Figure 10-2.1. You might use Theorem 10-2.19 and Exercise 10.6 to rule out smaller arrangements. See also Example 10-2.20.*)

The weak order and the poset of regions

10.9. Prove Proposition 10-3.11.

10.10. Prove Proposition 10-3.12. (*Apply Proposition 10-3.11 to the non-reduced word $a_1 \cdots a_k s$ and argue that one of the letters deleted must be s.*)

10.11. Prove Proposition 10-3.13.

10.12. Prove Proposition 10-3.14.

10.13. Prove Proposition 10-3.15.

10.14. Suppose W is a finite Coxeter group. Recall that the maximal element of the weak order is called w_0. Show that w_0 is an involution. (That is, $w_0 w_0$ is the identity in W.)

10.15. Taking W and w_0 as in Exercise 10.14, show that the map $w \mapsto ww_0$ is an antiautomorphism of the weak order. (*Show that a reduced word* *for w_0 corresponds to a path from B to $-B$. Act on the regions in this path by w to conclude that $ww_0B = -wB$ and apply Exercise 9.4.*)

10.16. Taking W and w_0 as in Exercise 10.14, show that the map $w \mapsto w_0w$ is also an antiautomorphism of the weak order on W. (*Use Exercise 10.1 and Proposition 10-3.5 to show that* $\mathrm{inv}(w_0w)$ *is the complement* $T \setminus \{w_0tw_0 \mid t \in \mathrm{inv}(w)\}$*, where T is the set of reflections in W.*)

10.17. Taking W and w_0 as in Exercise 10.14, show that the map $w \mapsto w_0ww_0$ is an automorphism of the weak order on W. For which rank-two Coxeter groups $I_2(m)$ is this map the identity?

10.18. Taking W and w_0 as in Exercise 10.14, show that the map $s \mapsto w_0sw_0$ is a permutation of S and induces an automorphism of the Coxeter diagram of W. That is, $w_0sw_0 \in S$ for all $s \in S$ and $m(w_0sw_0, w_0tw_0) = m(s,t)$ for all $s,t \in S$. (*Use Exercise 10.17.*)

10.19. Prove Proposition 10-3.22.

10.20. Interpret the equivalence of (i) and (iii) in Theorem 10-3.24 as a statement about inversion sets of permutations. Use Theorem 10-3.25(ii) to describe the join of permutations in terms of their inversion sets. (Use the combinatorial definition $\mathrm{inv}(\pi) = \{(\pi_i, \pi_j) \mid i < j, \pi_i > \pi_j\}$ for inversions.)

The Word Problem for finite Coxeter groups

10.21. Prove Lemma 10-4.2. (*It may be useful to look at the proof of Theorem 10-2.9.*)

10.22. Use Theorem 10-4.1 to construct the weak orders on all of the finite Coxeter groups of rank 3. This is an infinite problem, so instead of doing every Coxeter group of type $I_2(m) \times A_1$, use $B_2 \times A_1$ as a representative as in Examples 10-2.20 and 10-3.2. Do not rely on the pictures in Figures 10-3.1 and 10-3.2, but rather use your calculations to confirm the correctness of the pictures independent of Theorem 10-3.1. (*Construct the restriction of the weak order to elements of length $\leq k$, and let k increase until you have determined the entire weak order. Constructing the weak order for type B_3 in this way is an exercise for a patient person, and doing type H_3 this way requires considerably more patience.*)

10.23. Let W be a finite Coxeter group with defining generators S. Given an element $w \in W$ and a reduced word $a_1 \cdots a_k$ for w, the *support* of w is the set $\{a_1, \ldots, a_k\} \subseteq S$ of generators appearing in the word. Prove that the support of w is well-defined (i.e., does not depend on the reduced word chosen).

Coxeter groups of type A

10.24. Prove Proposition 10-5.1.

10.25. Prove Proposition 10-5.3.

10.26. Prove Proposition 10-5.4.

10.27. Use Theorem 10-3.9 to prove Theorem 10-5.6.

10.28. Find the canonical join representations of all elements of $\mathfrak{S}_4$ in two ways: Using Theorem 10-5.6 and using Figure 10-5.2.

10.29. Consider the Coxeter arrangement $\mathcal{A}$ described in Section 10-5 for a Coxeter group of type A_n. What are the rank-two subarrangements of $\mathcal{A}$?

10.30. Show that the cutting relation on hyperplanes in the Coxeter arrangement $\mathcal{A}$ for a Coxeter group of type A_{n-1} is as follows: A hyperplanes cuts $(\mathbf{e}_k - \mathbf{e}_i)^\perp$ if and only if it is $(\mathbf{e}_j - \mathbf{e}_i)^\perp$ or $(\mathbf{e}_k - \mathbf{e}_j)^\perp$ for some j with $1 \le i < j < k \le n$. (*Use Exercise 10.29.*)

10.31. Prove Proposition 10-5.8.

Cambrian lattices and sortable elements

10.32. Let W be a Coxeter group with defining generators S. Suppose r and s are distinct elements of S. Show that the join of r and s in the weak order is the element with two distinct reduced words $rsrs\cdots$ and $srsr\cdots$, each of length $m(r,s)$.

10.33. Use Figures 10-2.3 and 10-3.2 to find the Cambrian congruences, lattices and fans of type B_3, just as Example 10-6.3 does for type A_3. (See Figures 10-6.2, 10-6.3, and 10-6.4.)

10.34. Use Figures 10-2.4 and 10-3.2 to find the Cambrian congruences, lattices and fans of type H_3. *See the comment on Exercise 10.33.*

10.35. By Corollary 7-6.15, the Cambrian lattices of type A_{n-1} are exactly the quotients of the weak order on $\mathfrak{S}_n$ by the minimal meet-irreducible congruences on the weak order. Show that this statement is not true for general finite Coxeter groups by identifying the minimal meet-irreducible congruences on the weak order on a Coxeter group of type B_2.

10.36. Let W be a Coxeter group with defining generators S. Show that two total orders on S define the same Coxeter element of W if and only if they induce the same orientation on the diagram of W. (*Use Exercise 10.5 and Theorem 10-4.1.*)

10.37. Suppose $s_1 \cdots s_n$ and $s'_1 \cdots s'_n$ are words, each containing each element of S exactly once. Show that both are reduced. Show that the two define the same Coxeter element if and only if they are related by a sequence of moves, each of which swaps two letters that are adjacent in the word and commute in W.

10.38. Given an element w and two reduced words $s_1 \cdots s_n$ and $s'_1 \cdots s'_n$ for the same Coxeter element c, show that w is $s_1 \cdots s_n$-sortable if and only if it is $s'_1 \cdots s'_n$-sortable. Show also that if $s_1 \cdots s_n$ and $s'_1 \cdots s'_n$ are reduced words for different Coxeter elements, then there exists an element that is $s_1 \cdots s_n$-sortable but not $s'_1 \cdots s'_n$-sortable. (*Use Exercises 10.36 and 10.37.*)

10.39. Show that, in any finite lattice, the length of any maximal chain is less than or equal to the number of join-irreducible elements and also less than or equal to the number of meet-irreducible elements.

10.40. Prove Theorem 10-6.16 for Coxeter groups of type A_n by showing that the set of permutations avoiding the patterns $31\underline{2}$ and $\overline{2}31$ is a sublattice of the weak order on $\mathfrak{S}_n$. (*Use Proposition 9-5.8 and the dual of Exercise 10.20.*)

10.41. Prove Theorem 10-6.19. You may use diamond theorems. (*Use Theorem 10-6.20 to construct a chain of of length $|T|$ consisting of sortable elements. Then appeal to Theorem 10-6.18, Exercises 10.13 and 10.39 and Theorem 10-6.4.*)

10.42. Prove Theorem 10-6.20. You may use diamond theorems, but not Exercise 10.41. (*Use Theorem 10-6.4 and Corollary 10-6.15.*)

10.43. Suppose W is a finite Coxeter group and c is a Coxeter element with reduced word $s_1 \cdots s_n$. For each i from 0 to n, let x_i be the join $s_{n-i+1} \vee s_{n-i+2} \vee \cdots \vee s_n$. Show that

$$1 = x_0 \prec x_1 \prec x_2 \prec \cdots \prec x_n = w_0$$

is a maximal chain in the c-Cambrian lattice.

10.44. Suppose W is a Coxeter group and I is a subset of the defining generators S. Show that the parabolic subgroup W_I is a Coxeter group with defining generators I and with the quantities $m(r,s)$ inherited from the defining presentation of W. (*Use Theorem 10-4.1 to verify that every relation in W_I among the generators I is a consequence of the defining relations.*)

10.45. Show that $s_i \leq \pi$ in the weak order on $\mathfrak{S}_n$ if and only if the value i appears after the value $i+1$ in the sequence $\pi_1 \cdots \pi_n$.

10.46. For W of type A_{n-1}, show that $\pi = \pi_1 \cdots \pi_n$ is in $W_{\langle s_i \rangle}$ if and only if $\{\pi_1, \ldots, \pi_i\} = \{1, \ldots, i\}$.

10.47. Verify Theorem 10-6.25 for $W = \mathfrak{S}_4$ and $c = s_2 s_1 s_3$.

10.48. Suppose W is the direct product of two of its parabolic subgroups. That is, S is the disjoint union $S_1 \cup S_2$ and $W \cong W_{S_1} \times W_{S_2}$. (See Exercises 10.6 and 10.7.) Suppose c is a Coxeter element of W

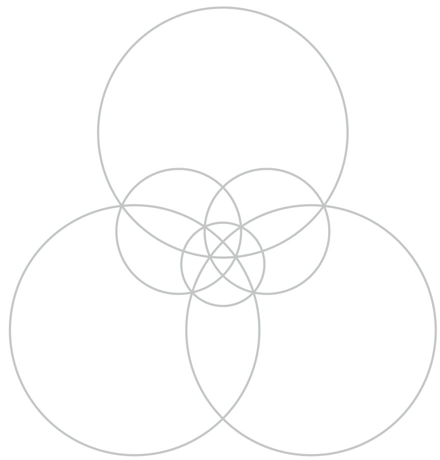

Figure 10-8.1: Sketching aid for Exercise 10.51

corresponding to $(c_1, c_2) \in W_{S_1} \times W_{S_2}$. Take $w \in W$ and let (w_1, w_2) be the corresponding element of $W_{S_1} \times W_{S_2}$. Show that w is c-sortable if and only if w_1 is c_1-sortable and w_2 is c_2-sortable.

10.49. For v a c-sortable element and s initial in c, show that

$$C_c(v) = \begin{cases} C_{sc}(v) \cup \{\alpha_s\} & \text{if } v \not\geq s \\ sC_{scs}(sv) & \text{if } v \geq s \end{cases}$$

10.50. Show that $C_c(v)$ is linearly independent for each c-sortable element v.

10.51. Compute $C_c(v)$ for each c-sortable element v with $W = \mathfrak{S}_4$ and $c = s_1 s_2 s_3$ as in Example 10-6.17. Also for each c-sortable element v, sketch the cone $\mathrm{Cone}_c(v)$ and label that cone with v. Verify Theorem 10-6.26 in this case. To assist the sketching, Figure 10-8.1 shows the Coxeter arrangement in light gray. The base region is

$$B = \mathrm{Cone}_c(1) = \bigcap_{i=1}^{3} \{x \in \mathbb{R}^n \mid \langle x, \alpha_i \rangle \geq 0\}$$

where α_i is the simple root $e_{i+1} - e_i$. The base region B projects to the triangle inside all circles in Figure 10-8.1. Besides the simple roots, the other roots are $e_3 - e_1$, $e_4 - e_2$, and $e_4 - e_1$. *Please see the comment to Exercise 10.33.*

10.52. Show that any two Coxeter elements of a finite Coxeter group W are conjugate in W. (*Show that the two can be related by a sequence of moves relating c to scs where s is initial in c. Interpret such*

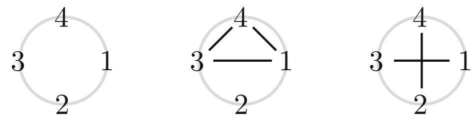

Figure 10-8.2: Illustrations for Exercise 10.53

moves in terms of orientations of the diagram of W and use the fact that each connected component of the diagram is a tree.)

10.53. Compute $\mathrm{nc}_c(v)$ for each c-sortable element v with $W = \mathfrak{S}_4$ and $c = s_1 s_2 s_3$ as in Example 10-6.17. (Compare Example 10-6.35 and Figure 10-6.11.) For each v, rewrite $\mathrm{nc}_c(v)$ as a permutation in disjoint cycle notation. Place the numbers 1, 2, 3, 4 in cyclic order clockwise on a circle as in the left picture of Figure 10-8.2. For each cycle in each $\mathrm{nc}_c(v)$, draw the cycle as a polygon on the circle. Thus, for example, the middle picture of Figure 10-8.2 shows how one would draw the permutation with cycle notation $(1\,3\,4)(2)$. Each disjoint cycle notation can be interpreted as a partition of the set $\{1, 2, 3, 4\}$, so that for example $(1\,3\,4)(2)$ is the partition with two blocks $\{1, 3, 4\}$ and $\{2\}$. A partition is called *noncrossing* if when it is drawn on the cycle 1, 2, 3, 4, the blocks are non-intersecting. Every partition of $\{1, 2, 3, 4\}$ is noncrossing except the partition with blocks $\{1, 3\}$ and $\{2, 4\}$, illustrated in the right picture of Figure 10-8.2. Verify in this example that as v runs through all c-sortable elements, the cycle notations of the elements $\mathrm{nc}_c(v)$ produce each noncrossing partition of $\{1, 2, 3, 4\}$ exactly once. (It is possible for two different elements of W to give the same drawing on the circle. For example, $(1\,3\,4)(2)$ and $(1\,4\,3)(2)$. The elements $\mathrm{nc}_c(v)$ are exactly the elements of W whose cycle notation defines a noncrossing partition *and* whose cycles read clockwise on the circle.)

Exercise 10.53 generalizes to $\mathfrak{S}_n$ for all n and to all Coxeter elements c. To properly generalize, one must write c in cycle notation. It will be an n-cycle, and this cycle is written clockwise around a circle. Other cases of Theorem 10-6.33 also yield nice combinatorial models. For example, for W of type B_n, the c-noncrossing partitions can be interpreted as the centrally-symmetric noncrossing partitions of a cycle.

10.54. Compute $\mathrm{cl}_c(v)$ for each c-sortable v with $W = \mathfrak{S}_4$ and $c = s_1 s_2 s_3$ as in Example 10-6.17. (Compare Example 10-6.37 and Figure 10-6.12.) Sketch the fan whose maximal cones are spanned by the sets $\mathrm{cl}_c(v)$, and verify in this case that the result is isomorphic to the c-Cambrian fan. To assist the sketching, Figure 10-8.3 shows the projections of the rays of the cluster fan (labeled by the corresponding

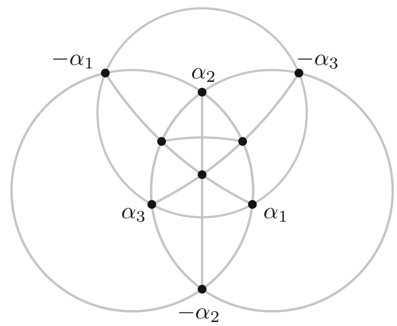

Figure 10-8.3: Sketching aid for Exercise 10.54

roots). The unlabeled roots are $\alpha_1 + \alpha_2$, $\alpha_2 + \alpha_3$, and $\alpha_1 + \alpha_2 + \alpha_3$. (These roots can be identified by coplanarity. For example, α_1, α_2, and $\alpha_1 + \alpha_2$ are coplanar, so their projections are on a common circle.) All arcs that might be drawn are shown in the figure in light gray. Complete the sketch by darkening those arcs that correspond to 2-dimensional faces of the c-cluster fan. *Please see the comment to Exercise 10.33.*

10.55. Suppose W is a Coxeter group and I is a subset of the defining generators S. Show that the subposet of the weak order on W induced by the elements of W_I is a down-set and is isomorphic to the weak order on W_I. Show that if W_I is finite, then this subposet is a lower interval in W. (*Use Exercise 10.23.*)

10-9. Notes

Coxeter groups and the weak order

Standard references on Coxeter groups include [84, 241]. For an approach emphasizing combinatorics, see [69]. For a treatment that expands on the combinatorial group theory aspects (and then moves deeply into geometric group theory), see [109]. Other books whose emphases are farther from the goals of this book include [33, 45, 224, 424].

Recursive unsolvability of the Word Problem for finitely presented groups is credited to Novikov [340] and Boone [75, 76, 77, 78, 79], while that result for the Finiteness problem is credited to Adyan [31, 32] and Rabin [363]. See Miller's survey [319] for a discussion of these results.

Finite reflection groups

Theorems 10-2.1 and 10-2.5 are well known. See Bourbaki [84, Proposition V.3.7] and [84, Theorem V.3.2]. The proof of Theorem 10-2.1 given here is mentioned in the paragraph before [70, Theorem 2.4]. Proposition 10-2.3 is due to Shannon [411], who in fact proved a sharp lower bound on the number of simplicial regions in an arrangement. The proof we give is essentially Shannon's, altered to avoid dealing with non-central arrangements. Propositions 10-2.7 and 10-2.8 are also standard. Theorems 10-2.9, 10-2.10, 10-2.18 and 10-2.19 go back to various papers of Coxeter in the mid 1930's, including [104, 105]. Proposition 10-2.17 appears, for example, as [69, Proposition 1.1.1] without the restriction to finite Coxeter groups. Theorem 10-2.21 has not, to our knowledge, appeared elsewhere.

The weak order and the poset of regions

Theorem 10-3.1 goes back to Edelman [139, Corollary 4.3]. The fact that the weak order is a lattice in general was pointed out without proof by Björner in [68]. The semidistributivity in Theorem 10-3.7 is due to Le Conte de Poly-Barbut [297, Lemme 9]. Congruence uniformity (boundedness) is due to Caspard, Le Conte de Poly-Barbut, and Morvan [91, Theorem 6]. See also [90]. The proof given here is similar to the proof in [364], and both owe much to the proof in [91], although the debt may be less apparent due to the generality of hyperplane arrangements and the geometric constructions related to shards. A slightly different proof of the acyclicity of the shard digraph for weak order is also in [364, Proposition 28]. The proof in [364] is a geometric argument using reflective symmetry, while the proof here substitutes a discrete metric (the depth) for a continuous metric on the sphere. The polygonality in Theorem 10-3.7 is due to Caspard, Le Conte de Poly-Barbut, and Morvan [91, Proposition 6]. Theorem 10-3.9 is due to Reading and Speyer [381, Theorem 8.1]. For Propositions 10-3.11 and 10-3.12, see Bourbaki [84, Exercise IV.1.1] and [84, Section IV.1.4] or [69, Section 1.4]. The biconvexity characterization of inversion sets in Theorem 10-3.24 is stated by Björner [68, Proposition 3]. The rank-two biconvexity characterization of inversion sets in Theorem 10-3.24 is due to Kostant [284, Proposition 5.10] for crystallographic root systems and was extended to noncrystallographic root systems in [73, Proposition 4.2] (See also [84, Exercise VI.1.16].) Various generalizations of Theorems 10-3.24 and 10-3.25 to infinite Coxeter groups have been made, including in [136, 227, 352].

The Word Problem for finite Coxeter groups

Theorem 10-4.1 is due to Tits [435]. Theorem 10-4.3 is well known, appearing for example in Humphreys as [241, Exercise 1.12.1]. For an in-depth discussion of computational issues in Coxeter groups, see Stembridge [425] and references therein.

Coxeter groups of type A

Early work on the weak order on permutations (also known as the permutation lattice or permutoèdre) is described in detailed notes at the end of [97] and [98]. Details on modeling Coxeter groups of type A by permutations are found throughout [69]. Theorem 10-5.6 follows from Theorem 10-3.9 and [367, Proposition 6.4] and appears as [376, Theorem 2.4]. Proposition 10-5.8 appeared, with different notation, in Bancroft [48], while Proposition 10-5.9 is a rephrasing of [366, Theorem 8.1]. Theorem 10-5.11 is essentially [376, Theorem 4.6].

Cambrian lattices and sortable elements

Cambrian congruences and Cambrian lattices have their earliest origins in the work of Björner and Wachs on the Tamari lattice [72, Section 9]. Their focus was not lattice-theoretic, but they assembled all the ingredients for a proof that the Tamari lattice is a lattice quotient of the weak order on permutations. (It is important to be aware of a clash of notation. It is standard in the Coxeter groups literature to speak of "quotients" of Coxeter groups in a combinatorial/group-theoretic sense. See [71]. These quotients are not lattice quotients of the weak order.)

Cambrian lattices were first defined in [368], and many of the results quoted about them here were conjectured there. Some of the conjectures were proved there for W of type A or B. Since the definition in [368] is the same as the definition here, Theorem 10-6.1 followed immediately from the definition, as here. Theorem 10-6.4 is [368, Theorem 3.5]. Theorem 10-6.5 is [370, Proposition 1.3]. Theorem 10-6.8 was proved by Reading and Speyer as [379, Theorem 1.1] and Theorem 10-6.7 follows. Theorem 10-6.9 follows from the explicit construction in [368, Section 6] of all Cambrian lattices of type A in terms of triangulations. This is a special case of the theorem obtained by directly concatenating [369, Theorem 4.1] with [370, Theorem 1.4], which characterizes the bottom elements of Cambrian lattices in terms of a certain "alignment" condition on their inversion sets. As mentioned earlier, the type-B Tamari lattice was constructed independently in [368] and by Thomas in [433].

Cambrian lattices admit an EL-labeling. This follows from results of Ingalls and Thomas by combining [243, Theorem 4.17] with [432, Proposition 3]. (The latter result is quoted from Liu [301].) EL-labelings were also given by Kallipoliti and Mühle in [264, Theorem 1.1] and by Pilaud and Stump in [350].

Sortable elements were defined in [369] and have been studied further by various researchers in [35, 36, 38, 229, 243, 264, 277, 361, 370, 380, 381]. Theorem 10-6.14 and Corollary 10-6.15 are [370, Theorem 1.1] and [370, Theorem 1.4]. Theorem 10-6.16 is [370, Theorem 1.2]. Theorem 10-6.18 is implicit as a special case of [369, Theorem 6.1]. See also [381, Theorem 8.9(iv)]. Something stronger than Theorem 10-6.19 was conjectured and partially proved by Thomas in [432] and then proved by Ingalls and Thomas in [243, Theorem 4.17]. The Tamari lattice case of Theorem 10-6.19 goes back to [304]. Theorem 10-6.20 is [369, Corollary 4.4]. Lemmas 10-6.23 and 10-6.24 are [369, Lemmas 2.4–2.5]. Theorem 10-6.25 can be obtained, *via* Theorem 10-6.14, from [368, Proposition 5.7] and [368, Theorem 6.2]. Alternately, it is a special case of [369, Theorem 4.1], proved separately as [369, Theorem 4.3]. A much more general statement, proved using the same strategy of induction on length and rank, is [381, Theorem 4.3]. Theorem 10-6.26 is the finite-type case of [381, Theorem 6.3].

Theorems 10-6.33 and 10-6.34 are [369, Theorem 6.1] and [372, Theorem 8.5]. The first statement of Theorem 10-6.36 is [369, Theorem 8.1], while the second statement follows from Theorem 10-6.8 and Corollary 10-6.15.

For more background on Coxeter-Catalan combinatorics, see [157] (especially [157, Lecture 5]) and the introductory chapters of [37]. Exercise 10.52 goes back at least to [60, Theorem 1.2], which is phrased in terms of oriented diagrams (quivers). For an exposition and computation of the Coxeter number and exponents for each finite Coxeter group, see for example [84, Chapter V.6] or [241, Sections 3.19–3.20]. These references also describe the beautiful and surprising connection between exponents and polynomial invariants of W.

The classical noncrossing partitions were first defined by Kreweras [287]. For the definition of c-noncrossing partitions, see [37] or (in different terminology) [62, 85]. Exercise 10.53 refers to centrally-symmetric noncrossing partitions. For details, see [384]. For type-D models, see [41], and for a general discussion of planar models for c-noncrossing partitions, see [371].

For a definition of combinatorial clusters, see [159, 306, 369]. The generalized associahedra were defined by Fomin and Zelevinsky in [159] and shown in [158] to be the combinatorial structure underlying *cluster algebras of finite type*. The original definition defines only a simplicial sphere (dual to the simple generalized associahedron), but Chapoton, Fomin, and Zelevinsky proved polytopality in [99]. Another family of polytopal realizations of the generalized associahedra, using c-sortable elements for any c, is given by Hohlweg, Lange, and Thomas in [228, 229]. That realization has a normal fan that actually *coincides* with the c-Cambrian fan for each c. The Cambrian fan also coincides exactly with a fan arising in the theory of cluster algebras, called the $\mathbf{g}$-vector fan. (This was conjectured and proved in a special case by Reading and Speyer in [379, Section 10]. It was later proved by Yang and Zelevinsky in [459] and then by Reading and Speyer in [382].) There is a compatible but different approach to generalized associahedra in terms of the *brick polytopes* of Pilaud and Santos [349]. See Pilaud and Stump [351] and references therein.

The definition of sortable elements does not require W to be finite. The article [381] explores the combinatorics of sortable elements in infinite Coxeter groups; it also tidies up some of the proofs of several key lemmas of [369] by providing proofs that do not involve a type-by-type analysis using the classification of finite Coxeter groups.

The bijection between Coxeter elements and acyclic orientations of the Coxeter diagram was pointed out by Shi in [412]. Exercise 10.39 is [304, Lemma 1].

Some other lattice quotients of the weak order

Parabolic congruences were studied (in the generality of simplicial arrangements) in [366, Section 6]. The fact that $w \mapsto w_I$ is a lattice homomorphism was pointed out in by Jedlička [247] and in [366]. Diagram homomorphisms are the subject of [377] and appear in the context of matrix mutation in [378]. Sashes and rectangulations also index the bases of certain combinatorial Hopf algebras, as part of a general lattice-theoretic construction [367]. Sashes were studied by Law in her thesis [291, 292], while rectangulations are the subject of [293] and [375]. Diagonal rectangulations are in bijection with the *Baxter permutations* of Chung, Graham, Hoggatt, and Kleiman [100]. Rectangulations and related constructions have also been studied in other contexts, for example in [1, 2, 39, 101, 132, 152, 187, 234, 460]. Coxeter-biCatalan combinatorics is considered in [52]. The connection between preprojective modules and the weak order was made by Mizuno [320]. Algebraic lattice congruences are the topic of [245].

10-10. Open problems

Problem 10.1. Characterize the lattices that appear as weak orders on finite Coxeter groups. This is a special case of Problem 9.2 in Section 9-11. See that problem for details. The equational theory of the weak order on the symmetric group is proved to be recursively solvable in [398].

Problem 10.2. Find the order dimension of the weak order on a finite Coxeter group. This is a special case of Problem 9.3 in Section 9-11. The problem is solved for some finite Coxeter groups in [154, 365].

Problem 10.3. Find a formula for the number of lattice congruences of the weak order on the symmetric group $\mathfrak{S}_n$. We expect this problem to be hard, as it involves counting order ideals in a poset (described in Proposition 10-5.9), and such problems are often hard. (Compare Dedekind's Problem of counting the elements of the free distributive lattice, or equivalently counting antichains in a finite Boolean algebra. See for example [278].) A more realistic problem is to give an asymptotic estimate. Exact values, up to $n = 7$, are found in [366, Section 9], and from these values, one is led to guess that the number of congruences on $\mathfrak{S}_{n+1}$ is roughly the square of the number of congruences on $\mathfrak{S}_n$.

Problem 10.4. Find a formula or asymptotic estimate for the number of congruences of the weak order on a finite Coxeter group of type B_n. Exact values, up to $n = 4$, are found in [366, Section 9]. For these small values, the number of congruences on B_{n+1} is very roughly the cube of the number of congruences on $\mathfrak{S}_n$.

10-11. Acknowledgments for Chapters 9 and 10

The author thanks Friedrich Wehrung for catching many errors, for making many excellent suggestions, and for contributing to the mathematical content of the chapters. The author also thanks Emily Barnard, Christophe Hohlweg, Joseph Kung, Thomas McConville, Vic Reiner, Hugh Thomas, and Thomas Zaslavsky for valuable comments.

Bibliography

[1] Eyal Ackerman, Gill Barequet, and Ron Y. Pinter, *A bijection between permutations and floorplans, and its applications*, Discrete Appl. Math. **154** (2006), no. 12, 1674–1684. MR 2233287 (2007a:05001)

[2] ______, *On the number of rectangulations of a planar point set*, J. Combin. Theory Ser. A **113** (2006), no. 6, 1072–1091. MR 2244135 (2007f:05001)

[3] Michael E. Adams and Wiesław Dziobiak, *Q-universal quasivarieties of algebras*, Proc. Amer. Math. Soc. **120** (1994), no. 4, 1053–1059. MR 1172942 (94f:08019)

[4] Kira V. Adaricheva, *Semidistributive and coalgebraic lattices of subsemilattices*, Algebra i Logika **27** (1988), no. 6, 625–640, 736. MR 1038102 (91b:06010)

[5] ______, *The structure of finite lattices of subsemilattices*, Algebra i Logika **30** (1991), no. 4, 385–404, 507. MR 1202734 (94b:06011)

[6] ______, *Two embedding theorems for lower bounded lattices*, Algebra Universalis **36** (1996), no. 4, 425–430. MR 1419356 (97k:06015)

[7] ______, *Join-semidistributive lattices of relatively convex sets*, Contributions to general algebra. 14, Heyn, Klagenfurt, 2004, improved version at arxiv:math/0403104, pp. 1–14. MR 2058597

[8] ______, *Lattices of algebraic subsets*, Algebra Universalis **52** (2004), no. 2-3, 167–183. MR 2161648 (2006d:06009)

[9] ______, *On the prevariety of perfect lattices*, Algebra Universalis **65** (2011), no. 1, 21–39. MR 2784235 (2012a:06006)

[10] ______, *Optimum basis of finite convex geometry*, arXiv:1205.3236, 2012.

G. Grätzer, F. Wehrung (eds.), *Lattice Theory: Special Topics and Applications*,
DOI 10.1007/978-3-319-44236-5

[11] ———, *Supersolvable semidistributive lattices*, preprint, 2012.

[12] ———, *Algebraic convex geometries revisited*, arXiv:1406.3721, 2014.

[13] ———, *Representing finite convex geometries by relatively convex sets*, European J. Combin. **37** (2014), 68–78. MR 3138592

[14] Kira V. Adaricheva and Viktor A. Gorbunov, *Equational closure operator and forbidden semidistributive lattices*, Sibirsk. Mat. Zh. **30** (1989), no. 6, 7–25. MR 1043428 (91g:06006)

[15] ———, *On lower bounded lattices*, Algebra Universalis **46** (2001), no. 1-2, 203–213, The Viktor Aleksandrovich Gorbunov memorial issue. MR 1835794 (2002c:06004)

[16] Kira V. Adaricheva, Viktor A. Gorbunov, and Wiesław Dziobiak, *Algebraic point lattices of quasivarieties*, Algebra i Logika **36** (1997), no. 4, 363–386, 478. MR 1601515 (99c:08013)

[17] ———, *Congruence properties of lattices of quasivarieties*, Algebra i Logika **36** (1997), no. 6, 605–620, 721. MR 1657297 (2000c:08005)

[18] Kira V. Adaricheva, Viktor A. Gorbunov, and V. I. Tumanov, *Join-semidistributive lattices and convex geometries*, Adv. Math. **173** (2003), no. 1, 1–49. MR 1954454 (2003m:06004)

[19] Kira V. Adaricheva, Miklós Maróti, Ralph N. McKenzie, James B. Nation, and Eric R. Zenk, *The Jónsson-Kiefer property*, Studia Logica **83** (2006), no. 1-3, 111–131. MR 2250107 (2007f:06007)

[20] Kira V. Adaricheva and James B. Nation, *Equaclosure operators on join semidistributive lattices*, working document.

[21] ———, *Largest extension of a finite convex geometry*, Algebra Universalis **52** (2004), no. 2-3, 185–195. MR 2161649 (2006b:06007)

[22] ———, *Reflections on lower bounded lattices*, Algebra Universalis **53** (2005), no. 2-3, 307–330. MR 2148303 (2006b:06010)

[23] ———, *Lattices of quasi-equational theories as congruence lattices of semilattices with operators, Part I*, Internat. J. Algebra Comput. **22** (2012), no. 7, 1250065, 27. MR 2999371

[24] ———, *Lattices of quasi-equational theories as congruence lattices of semilattices with operators, Part II*, Internat. J. Algebra Comput. **22** (2012), no. 7, 1250066, 19. MR 2999372

[25] ———, *On implicational bases of closure systems with unique critical sets*, Discrete Appl. Math. **162** (2014), 51–69. MR 3128509

[26] ———, *A class of infinite convex geometries*, arXiv:1501.4174, 2015.

[27] Kira V. Adaricheva, James B. Nation, and Robert Rand, *Ordered direct implicational basis of a finite closure system*, Discrete Appl. Math. **161** (2013), no. 6, 707–723. MR 3027961

[28] Kira V. Adaricheva and Maurice Pouzet, *On scattered convex geometries*, Proceedings of TACL-2011, 64–68; also arXiv:1505.03023, 2015.

[29] Kira V. Adaricheva and Friedrich Wehrung, *Embedding finite lattices into finite biatomic lattices*, Order **20** (2003), no. 1, 31–48. MR 1993408 (2004g:06004)

[30] Kira V. Adaricheva and Marcel Wild, *Realization of abstract convex geometries by point configurations*, European J. Combin. **31** (2010), no. 1, 379–400. MR 2552617 (2011h:52001)

[31] Sergeĭ I. Adyan, *Finitely presented groups and algorithms*, Dokl. Akad. Nauk SSSR (N.S.) **117** (1957), 9–12. MR 0095873 (20 #2371)

[32] ———, *Unsolvability of some algorithmic problems in the theory of groups.*, Trudy Moskov. Mat. Obšč. **6** (1957), 231–298. MR 0095872 (20 #2370)

[33] Marcelo Aguiar and Swapneel Mahajan, *Coxeter groups and Hopf algebras*, Fields Institute Monographs, vol. 23, American Mathematical Society, Providence, RI, 2006, With a foreword by Nantel Bergeron. MR 2225808 (2008d:20072)

[34] M. A. Aĭzerman and A. V. Malishevski, *Structural properties in the theory of choice of variants*, Proceedings of the 7th of All Union Seminar on Control Problems, Book 2, 1977, pp. 93–97.

[35] Claire Amiot, *A derived equivalence between cluster equivalent algebras*, J. Algebra **351** (2012), 107–129. MR 2862200

[36] Claire Amiot, Osamu Iyama, Idun Reiten, and Gordana Todorov, *Preprojective algebras and c-sortable words*, Proc. Lond. Math. Soc. (3) **104** (2012), no. 3, 513–539. MR 2900235

[37] Drew Armstrong, *Generalized noncrossing partitions and combinatorics of Coxeter groups*, Mem. Amer. Math. Soc. **202** (2009), no. 949, x+159. MR 2561274 (2011c:05001)

[38] ———, *The sorting order on a Coxeter group*, J. Combin. Theory Ser. A **116** (2009), no. 8, 1285–1305. MR 2568800 (2010m:20057)

[39] Andrei Asinowski, Gill Barequet, Mireille Bousquet-Mélou, Toufik Mansour, and Ron Y. Pinter, *Orders induced by segments in floorplans and (2-14-3, 3-41-2)-avoiding permutations*, Electron. J. Combin. **20** (2013), no. 2, Paper 35, 43. MR 3084577

[40] Christos A. Athanasiadis, Paul H. Edelman, and Victor Reiner, *Monotone paths on polytopes*, Math. Z. **235** (2000), no. 2, 315–334. MR 1795510 (2002a:52011)

[41] Christos A. Athanasiadis and Victor Reiner, *Noncrossing partitions for the group D_n*, SIAM J. Discrete Math. **18** (2004), no. 2, 397–417 (electronic). MR 2112514 (2006b:06004)

[42] Christos A. Athanasiadis and Francisco Santos, *Monotone paths on zonotopes and oriented matroids*, Canad. J. Math. **53** (2001), no. 6, 1121–1140. MR 1863845 (2003a:52031)

[43] Giorgio Ausiello, Alessandro D'Atri, and Domenico Saccà, *Minimal representation of directed hypergraphs*, SIAM J. Comput. **15** (1986), no. 2, 418–431. MR 837593 (87g:05162)

[44] Sherwin P. Avann, *Application of the join-irreducible excess function to semi-modular lattices*, Math. Ann. **142** (1960/1961), 345–354. MR 0122741 (23 #A81)

[45] Patrick Bahls, *The isomorphism problem in Coxeter groups*, Imperial College Press, London, 2005. MR 2162146 (2006e:20069)

[46] Kirby A. Baker, *Equational classes of modular lattices*, Pacific J. Math. **28** (1969), 9–15. MR 0244118 (39 #5435)

[47] ______, *Some non-finitely-based varieties of lattices*, Universal algebra (Esztergom, 1977), Colloq. Math. Soc. János Bolyai, vol. 29, North-Holland, Amsterdam-New York, 1982, pp. 53–59. MR 660847 (83g:06009)

[48] Erin Elizabeth Bancroft, *Shard Intersections and Cambrian Congruence Classes in Type A*, Ph.D. thesis, North Carolina State University, Raleigh, NC, 2011, 63 pages.

[49] ______, *Shard intersections and Cambrian congruence classes in type A*, arXiv:1103.1910, 2011.

[50] Hans-Jürgen Bandelt and Henry Martyn Mulder, *Distance-hereditary graphs*, J. Combin. Theory Ser. B **41** (1986), no. 2, 182–208. MR 859310 (87k:05142)

[51] Marc Barbut and Bernard Monjardet, *Ordre et Classification: Algèbre et Combinatoire. Tomes I et II*, Librairie Hachette, Paris, 1970, Méthodes Mathématiques des Sciences de l'Homme, Collection Hachette Université. MR 0419312 (54 #7334)

[52] Emily Barnard and Nathan Reading, *Coxeter-biCatalan combinatorics*, preprint, 2016.

[53] Mary K. Bennett, *Lattices of convex sets*, Trans. Amer. Math. Soc. **234** (1977), no. 1, 279–288. MR 0463051 (57 #3014)

[54] ———, *Convexity closure operators*, Algebra Universalis **10** (1980), no. 3, 345–354. MR 564121 (81c:54002)

[55] Mary K. Bennett and Garrett Birkhoff, *Two families of Newman lattices*, Algebra Universalis **32** (1994), no. 1, 115–144. MR 1287019 (95m:06022)

[56] François Bergeron, *Combinatorics of r-Dyck paths, r-Parking functions, and the r-Tamari lattices*, arXiv:1202.6269, 2012.

[57] François Bergeron, Gilbert Labelle, and Pierre Leroux, *Combinatorial Species and Tree-Like Structures*, Encyclopedia of Mathematics and its Applications, vol. 67, Cambridge University Press, Cambridge, 1998, Translated from the 1994 French original by Margaret Readdy, with a foreword by Gian-Carlo Rota. MR 1629341 (2000a:05008)

[58] Clifford Bergman, *Amalgamation classes of some distributive varieties*, Algebra Universalis **20** (1985), no. 2, 143–166. MR 806610 (87g:08016)

[59] ———, *Nonaxiomatizability of the amalgamation class of modular lattice varieties*, Order **6** (1989), no. 1, 49–58. MR 1020456 (90j:06011)

[60] Joseph N. Bernšteĭn, Izrail′ M. Gel′fand, and Sergei I. Gel′fand, *Schubert cells, and the cohomology of the spaces G/P*, Uspehi Mat. Nauk **28** (1973), no. 3 (171), 3–26. MR 0429933 (55 #2941)

[61] Karell Bertet and Bernard Monjardet, *The multiple facets of the canonical direct unit implicational basis*, Theoret. Comput. Sci. **411** (2010), no. 22-24, 2155–2166. MR 2662512 (2011g:68061)

[62] David Bessis, *The dual braid monoid*, Ann. Sci. École Norm. Sup. (4) **36** (2003), no. 5, 647–683. MR 2032983 (2004m:20071)

[63] Garrett Birkhoff, *On the combination of subalgebras*, Math. Proc. Cambridge Philos. Soc. **29** (1933), no. 4, 441–464.

[64] ———, *On the structure of abstract algebras*, Proc. Camb. Philos. Soc. **31** (1935), 433–454 (English).

[65] ———, *Lattice Theory*, third ed., American Mathematical Society Colloquium Publications, vol. 25, American Mathematical Society, Providence, R. I., 1979. MR 598630 (82a:06001)

[66] Garrett Birkhoff and Mary K. Bennett, *The convexity lattice of a poset*, Order **2** (1985), no. 3, 223–242. MR 824696 (87d:06004)

[67] Anders Björner, *Orderings of Coxeter groups*, Combinatorics and Algebra (Boulder, Colo., 1983), Contemp. Math., vol. 34, Amer. Math. Soc., Providence, RI, 1984, pp. 175–195. MR 777701 (86i:05024)

[68] ———, *Orderings of Coxeter groups*, Combinatorics and Algebra (Boulder, Colo., 1983), Contemp. Math., vol. 34, Amer. Math. Soc., Providence, RI, 1984, pp. 175–195. MR 777701 (86i:05024)

[69] Anders Björner and Francesco Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, vol. 231, Springer, New York, 2005. MR 2133266 (2006d:05001)

[70] Anders Björner, Paul H. Edelman, and Günter M. Ziegler, *Hyperplane arrangements with a lattice of regions*, Discrete Comput. Geom. **5** (1990), no. 3, 263–288. MR 1036875 (90k:51036)

[71] Anders Björner and Michelle L. Wachs, *Generalized quotients in Coxeter groups*, Trans. Amer. Math. Soc. **308** (1988), no. 1, 1–37. MR 946427 (89c:05012)

[72] ———, *Shellable nonpure complexes and posets. II*, Trans. Amer. Math. Soc. **349** (1997), no. 10, 3945–3975. MR 1401765 (98b:06008)

[73] Dieter Blessenohl, Christophe Hohlweg, and Manfred Schocker, *A symmetry of the descent algebra of a finite Coxeter group*, Adv. Math. **193** (2005), no. 2, 416–437. MR 2137290 (2005m:20089)

[74] Peter A. Bloniarz, Harry B. Hunt, III, and Daniel J. Rosenkrantz, *On the computational complexity of algebra on lattices*, SIAM J. Comput. **16** (1987), no. 1, 129–148. MR 873256 (88c:68054)

[75] William W. Boone, *Certain simple, unsolvable problems of group theory. I*, Nederl. Akad. Wetensch. Proc. Ser. A. **57** (1954), 231–237 = Indag. Math. 16, 231–237 (1954). MR 0066372 (16,564d)

[76] ———, *Certain simple, unsolvable problems of group theory. II*, Nederl. Akad. Wetensch. Proc. Ser. A. **57** (1954), 492–497 = Indag. Math. 16, 492–497 (1954). MR 0066373 (16,564e)

[77] ———, *Certain simple, unsolvable problems of group theory. III*, Nederl. Akad. Wetensch. Proc. Ser. A. **58** (1955), 252–256 = Indag. Math. 17, 252–256 (1955). MR 0066374 (16,564f)

[78] ———, *Certain simple, unsolvable problems of group theory. IV*, Nederl. Akad. Wetensch. Proc. Ser. A. 58 = Indag. Math. **17** (1955), 571–577. MR 0098775 (20 #5230)

[79] ———, *Certain simple, unsolvable problems of group theory. V, VI*, Nederl. Akad. Wetensch. Proc. Ser. A. 60 = Indag. Math. **19** (1957), 22–27, 227–232. MR 0098776 (20 #5231)

[80] Endre Boros, Ondřej Čepek, Alexander Kogan, and Petr Kučera, *A subclass of Horn CNFs optimally compressible in polynomial time*, Ann. Math. Artif. Intell. **57** (2009), no. 3-4, 249–291. MR 2671573 (2012f:06016)

[81] ———, *Exclusive and essential sets of implicates of Boolean functions*, Discrete Appl. Math. **158** (2010), no. 2, 81–96. MR 2563275 (2010j:94088)

[82] François Boulier, Florent Hivert, Daniel Krob, and Jean-Christophe Novelli, *Pseudo-permutations. II. Geometry and representation theory*, Discrete Models: Combinatorics, Computation, and Geometry (Paris, 2001), Discrete Math. Theor. Comput. Sci. Proc., AA, Maison Inform. Math. Discrèt. (MIMD), Paris, 2001, pp. 123–132 (electronic). MR 1888767 (2003i:05002)

[83] Nicolas Bourbaki, *Éléments de Mathématique. Fasc. XXXIV. Groupes et Algèbres de Lie. Chapitre IV: Groupes de Coxeter et Systèmes de Tits. Chapitre V: Groupes Engendrés par des Réflexions. Chapitre VI: Systèmes de Racines*, Actualités Scientifiques et Industrielles, No. 1337, Hermann, Paris, 1968. MR 0240238 (39 #1590)

[84] ———, *Lie Groups and Lie Algebras. Chapters 4–6*, Elements of Mathematics (Berlin), Springer-Verlag, Berlin, 2002, Translated from the 1968 French original by Andrew Pressley. MR 1890629 (2003a:17001)

[85] Thomas Brady and Colum Watt, *Non-crossing partition lattices in finite real reflection groups*, Trans. Amer. Math. Soc. **360** (2008), no. 4, 1983–2005. MR 2366971 (2008k:20088)

[86] Peter Bruyns and Henry Rose, *Varieties with cofinal sets: examples and amalgamation*, Proc. Amer. Math. Soc. **111** (1991), no. 3, 833–840. MR 1039528 (91k:08007)

[87] J. Richard Büchi, *On a decision method in restricted second order arithmetic*, Logic, Methodology and Philosophy of Science (Proc. 1960 Internat. Congr.), Stanford Univ. Press, Stanford, Calif., 1962, pp. 1–11. MR 0183636 (32 #1116)

[88] A. I. Budkin and Viktor A. Gorbunov, *On the theory of quasivarieties of algebraic systems*, Algebra i Logika **14** (1975), no. 2, 123–142, 240. MR 0396373 (53 #240)

[89] Stanley Burris, *Polynomial time uniform word problems*, Math. Logic Quart. **41** (1995), no. 2, 173–182. MR 1330852 (96j:08004)

[90] Nathalie Caspard, *The lattice of permutations is bounded*, Internat. J. Algebra Comput. **10** (2000), no. 4, 481–489. MR 1776052 (2001d:06008)

[91] Nathalie Caspard, Claude Le Conte de Poly-Barbut, and Michel Morvan, *Cayley lattices of finite Coxeter groups are bounded*, Adv. in Appl. Math. **33** (2004), no. 1, 71–94. MR 2064358 (2005b:06006)

[92] Nathalie Caspard, Bruno Leclerc, and Bernard Monjardet, *Finite Ordered Sets*, Encyclopedia of Mathematics and its Applications, vol. 144, Cambridge University Press, Cambridge, 2012, Concepts, results and uses. MR 2919901

[93] Nathalie Caspard and Bernard Monjardet, *The lattices of closure systems, closure operators, and implicational systems on a finite set: a survey*, Discrete Appl. Math. **127** (2003), no. 2, 241–269, The 1998 Conference on Ordinal and Symbolic Data Analysis (OSDA '98) (Amherst, MA). MR 1984087 (2004b:06006)

[94] ______, *Some lattices of closure systems on a finite set*, Discrete Math. Theor. Comput. Sci. **6** (2004), no. 2, 163–190 (electronic). MR 2041845 (2004m:06012)

[95] Ivan Chajda and Václav Snášel, *Congruences in ordered sets*, Math. Bohem. **123** (1998), no. 1, 95–100. MR 1618652 (98k:06002)

[96] Célestin Chameni-Nembua, *Permutoèdre et choix social*, Ph.D. thesis, Université Paris 5, 1989, Thèse de Mathématiques Appliquées aux Sciences Sociales.

[97] Célestin Chameni-Nembua and Bernard Monjardet, *Les treillis pseudocomplémentés finis*, European J. Combin. **13** (1992), no. 2, 89–107. MR 1158803 (93d:06011)

[98] ______, *Finite pseudocomplemented lattices and "permutoèdre"*, Discrete Math. **111** (1993), no. 1-3, 105–112, Graph theory and combinatorics (Marseille-Luminy, 1990). MR 1210087 (94b:06019)

[99] Frédéric Chapoton, Sergey Fomin, and Andrei Zelevinsky, *Polytopal realizations of generalized associahedra*, Canad. Math. Bull. **45** (2002), no. 4, 537–566, Dedicated to Robert V. Moody. MR 1941227 (2003j:52014)

[100] Fan Chung, Ronald L. Graham, Verner E. Hoggatt Jr., and Mark Kleiman, *The number of Baxter permutations*, J. Combin. Theory Ser. A **24** (1978), no. 3, 382–394. MR 491652 (82b:05011)

[101] Jim Conant and Tim Michaels, *On the Number of Tilings of a Square by Rectangles*, Ann. Comb. **18** (2014), no. 1, 21–34. MR 3167602

[102] Raul Cordovil, *A combinatorial perspective on the non-Radon partitions*, J. Combin. Theory Ser. A **38** (1985), no. 1, 38–47. MR 773553 (87e:52003a)

[103] Raul Cordovil and M. L. Moreira, *A homotopy theorem on oriented matroids*, Discrete Math. **111** (1993), no. 1-3, 131–136, Graph theory and combinatorics (Marseille-Luminy, 1990). MR 1210090 (94d:52016)

[104] H. S. M. Coxeter, *Discrete groups generated by reflections*, Ann. of Math. (2) **35** (1934), no. 3, 588–621. MR 1503182

[105] ______, *The complete enumeration of finite groups of the form* $R_i^2 = (R_iR_j)^{k_{ij}} = 1$, J. London Math. Soc. **s1-10** (1935), no. 3, 21–25.

[106] Gábor Czédli and Ralph Freese, *On congruence distributivity and modularity*, Algebra Universalis **17** (1983), no. 2, 216–219. MR 726275 (85e:08011)

[107] Brian A. Davey, Werner Poguntke, and Ivan Rival, *A characterization of semi-distributivity*, Algebra Universalis **5** (1975), 72–75. MR 0382103 (52 #2991)

[108] Brian A. Davey and Hilary A. Priestley, *Introduction to Lattices and Order*, second ed., Cambridge University Press, New York, 2002. MR 1902334 (2003e:06001)

[109] Michael W. Davis, *The Geometry and Topology of Coxeter Groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008. MR 2360474 (2008k:20091)

[110] Alan Day, *A simple solution to the word problem for lattices*, Canad. Math. Bull. **13** (1970), 253–254. MR 0268092

[111] ______, *Splitting lattices and congruence modularity*, Contributions to universal algebra (Colloq., József Attila Univ., Szeged, 1975), North-Holland, Amsterdam, 1977, pp. 57–71. Colloq. Math. Soc. János Bolyai, Vol. 17. MR 0552770 (58 #27694)

[112] ______, *Splitting lattices generate all lattices*, Algebra Universalis **7** (1977), no. 2, 163–169. MR 0434897 (55 #7861)

[113] ______, *Characterizations of finite lattices that are bounded-homomorphic images of sublattices of free lattices*, Canad. J. Math. **31** (1979), no. 1, 69–78. MR 518707 (81h:06004)

[114] ______, *Doubling constructions in lattice theory*, Canad. J. Math. **44** (1992), no. 2, 252–269. MR 1162342 (93f:06008)

[115] ______, *The lattice theory of functional dependencies and normal decompositions*, Internat. J. Algebra Comput. **2** (1992), no. 4, 409–431. MR 1189671 (93i:68051)

[116] ______, *Congruence normality: the characterization of the doubling class of convex sets*, Algebra Universalis **31** (1994), no. 3, 397–406. MR 1265350 (95a:06007)

[117] Alan Day and Ralph Freese, *A characterization of identities implying congruence modularity. I*, Canad. J. Math. **32** (1980), no. 5, 1140–1167. MR 596102 (82b:08009)

[118] Alan Day and Jaroslav Ježek, *The amalgamation property for varieties of lattices*, Trans. Amer. Math. Soc. **286** (1984), no. 1, 251–256. MR 756038 (86d:06012)

[119] Alan Day and James B. Nation, *Congruence normal covers of finitely generated lattice varieties*, Canad. Math. Bull. **35** (1992), no. 3, 311–320. MR 1184008 (93h:06006)

[120] Alan Day, James B. Nation, and Steve Tschantz, *Doubling convex sets in lattices and a generalized semidistributivity condition*, Order **6** (1989), no. 2, 175–180. MR 1031653 (91f:06007)

[121] Richard A. Dean, *Component subsets of the free lattice on n generators*, Proc. Amer. Math. Soc. **7** (1956), 220–226. MR 0078957 (18,6a)

[122] ______, *Free lattices generated by partially ordered sets and preserving bounds*, Canad. J. Math. **16** (1964), 136–148. MR 0157916 (28 #1144)

[123] Pierre Deligne, *Les immeubles des groupes de tresses généralisés*, Invent. Math. **17** (1972), 273–302. MR 0422673 (54 #10659)

[124] William DeMeo, *Congruence lattices of finite algebras*, Ph.D. thesis, University of Hawai'i, 2012.

[125] Brenda L. Dietrich, *Matroids and antimatroids—a survey*, Discrete Math. **78** (1989), no. 3, 223–237. MR 1026344 (90h:05037)

[126] Robert P. Dilworth, *Lattices with unique irreducible decompositions*, Ann. of Math. (2) **41** (1940), 771–777. MR 0002844 (2,120e)

[127] ______, *The structure of relatively complemented lattices*, Ann. of Math. (2) **51** (1950), 348–359. MR 0033795 (11,489b)

[128] Robert P. Dilworth and Peter Crawley, *Decomposition theory for lattices without chain conditions*, Trans. Amer. Math. Soc. **96** (1960), 1–22. MR 0118690 (22 #9461)

[129] ______, *Algebraic Theory of Lattices*, Prentice-Hall, New Jersey, 1973.

[130] Gerhard Dorfer, *Lattice-extensions by means of convex sublattices*, Contributions to General Algebra, 9 (Linz, 1994), Hölder-Pichler-Tempsky, Vienna, 1995, pp. 127–132. MR 1484432 (99c:06008)

[131] William F. Dowling and Jean H. Gallier, *Linear-time algorithms for testing the satisfiability of propositional Horn formulae*, J. Logic Programming **1** (1984), no. 3, 267–284. MR 770156 (86g:68144)

[132] Serge Dulucq and Olivier Guibert, *Stack words, standard tableaux and Baxter permutations*, Proceedings of the 6th Conference on Formal Power Series and Algebraic Combinatorics (New Brunswick, NJ, 1994), vol. 157, 1996, pp. 91–106. MR 1417289 (98a:05009)

[133] Vincent Duquenne, *The core of finite lattices*, Discrete Math. **88** (1991), no. 2-3, 133–147, Combinatorics of ordered sets (Oberwolfach, 1988). MR 1108010 (92g:06011)

[134] Vincent Duquenne and Ameziane Cherfouh, *On permutation lattices*, Math. Social Sci. **27** (1994), no. 1, 73–89. MR 1267689 (95g:20004)

[135] Ben Dushnik and Edwin W. Miller, *Partially ordered sets*, Amer. J. Math. **63** (1941), 600–610. MR 0004862 (3,73a)

[136] Matthew Dyer, *On the weak order of Coxeter groups*, arXiv:1108.5557, 2011.

[137] Wiesław Dziobiak, *On atoms in the lattice of quasivarieties*, Algebra Universalis **24** (1987), no. 1-2, 32–35. MR 921527 (89a:08008)

[138] Paul H. Edelman, *The Bruhat order of the symmetric group is lexicographically shellable*, Proc. Amer. Math. Soc. **82** (1981), no. 3, 355–358. MR 612718 (82j:06004)

[139] ______, *A partial order on the regions of $\mathbf{R}^n$ dissected by hyperplanes*, Trans. Amer. Math. Soc. **283** (1984), no. 2, 617–631. MR 737888 (85k:51030)

[140] Paul H. Edelman and Curtis Greene, *Combinatorial correspondences for Young tableaux, balanced tableaux, and maximal chains in the weak Bruhat order of S_n*, Combinatorics and Algebra (Boulder, Colo., 1983), Contemp. Math., vol. 34, Amer. Math. Soc., Providence, RI, 1984, pp. 155–162. MR 777699 (86j:05014)

[141] Paul H. Edelman and Robert E. Jamison, *The theory of convex geometries*, Geom. Dedicata **19** (1985), no. 3, 247–270. MR 815204 (87f:52002)

[142] Paul H. Edelman and David G. Larman, *On characterizing collections arising from n-gons in the plane*, Geom. Dedicata **33** (1990), no. 1, 83–89. MR 1042627 (91b:52003)

[143] Paul H. Edelman and Victor Reiner, *h-shellings and h-complexes*, Adv. Math. **106** (1994), no. 1, 36–64. MR 1275865 (95h:52017)

[144] Paul H. Edelman and James W. Walker, *The homotopy type of hyperplane posets*, Proc. Amer. Math. Soc. **94** (1985), no. 2, 221–225. MR 784167 (86j:52004)

[145] H. G. Eggleston, *Convexity*, Cambridge Tracts in Mathematics and Mathematical Physics, No. 47, Cambridge University Press, New York, 1958. MR 0124813 (23 #A2123)

[146] Kimmo Eriksson, *Polygon posets and the weak order of Coxeter groups*, J. Algebraic Combin. **4** (1995), no. 3, 233–252. MR 1331744 (96c:05186)

[147] Yu. L. Ershov, *Lattice properties preserved under free products*, Algebra Log. **39** (2000), no. 1, 66–73, 121. MR 1782317 (2001h:06010)

[148] Trevor Evans, *The word problem for abstract algebras*, J. London Math. Soc. **26** (1951), 64–71. MR 0038958 (12,475c)

[149] ———, *Embeddability and the word problem*, J. London Math. Soc. **28** (1953), 76–80. MR 0053915 (14,839d)

[150] Ulrich Faigle, *Exchange properties of combinatorial closure spaces*, Discrete Appl. Math. **15** (1986), no. 2-3, 249–260, Applications of combinatorial methods in mathematical programming (Gainesville, Fla., 1985). MR 865005 (87k:06011)

[151] Siemion Fajtlowicz and Jürgen Schmidt, *Bézout families, join-congruences, and meet-irreducible ideals*, Lattice theory (Proc. Colloq., Szeged, 1974), North-Holland, Amsterdam, 1976, pp. 51–76. Colloq. Math. Soc. János Bolyai, Vol. 14. MR 0444533 (56 #2883)

[152] Stefan Felsner, Éric Fusy, Marc Noy, and David Orden, *Bijections for Baxter families and related objects*, J. Combin. Theory Ser. A **118** (2011), no. 3, 993–1020. MR 2763051 (2012f:05016)

[153] Luca Ferrari and Renzo Pinzani, *Lattices of lattice paths*, J. Statist. Plann. Inference **135** (2005), no. 1, 77–92. MR 2202340 (2006j:06002)

[154] Sigrid Flath, *The order dimension of multinomial lattices*, Order **10** (1993), no. 3, 201–219. MR 1267187 (95f:06002)

[155] Dominique Foata and Guo-Niu Han, *Signed words and permutations. I. A fundamental transformation*, Proc. Amer. Math. Soc. **135** (2007), no. 1, 31–40. MR 2280171 (2007j:05004)

[156] Dominique Foata and Doron Zeilberger, *Graphical major indices*, J. Comput. Appl. Math. **68** (1996), no. 1-2, 79–101. MR 1418752 (98b:05004)

[157] Sergey Fomin and Nathan Reading, *Root systems and generalized associahedra*, Geometric combinatorics, IAS/Park City Math. Ser., vol. 13, Amer. Math. Soc., Providence, RI, 2007, pp. 63–131. MR 2383126

[158] Sergey Fomin and Andrei Zelevinsky, *Cluster algebras. II. Finite type classification*, Invent. Math. **154** (2003), no. 1, 63–121. MR 2004457 (2004m:17011)

[159] ———, *Y-systems and generalized associahedra*, Ann. of Math. (2) **158** (2003), no. 3, 977–1018. MR 2031858 (2004m:17010)

[160] Ralph Freese, *The structure of modular lattices of width four with applications to varieties of lattices*, Mem. Amer. Math. Soc. **9** (1977), no. 181, vii+91. MR 0441803 (56 #197)

[161] ———, *The variety of modular lattices is not generated by its finite members*, Trans. Amer. Math. Soc. **255** (1979), 277–300. MR 542881 (81g:06003)

[162] ———, *Free modular lattices*, Trans. Amer. Math. Soc. **261** (1980), no. 1, 81–91. MR 576864 (81k:06010)

[163] ———, *Finitely presented lattices: canonical forms and the covering relation*, Trans. Amer. Math. Soc. **312** (1989), no. 2, 841–860. MR 949899 (90b:06011)

[164] ———, *Finitely presented lattices: continuity and semidistributivity*, Lattices, semigroups, and universal algebra (Lisbon, 1988), Plenum, New York, 1990, pp. 67–70. MR 1085067 (91j:06010)

[165] ———, *Finitely based modular congruence varieties are distributive*, Algebra Universalis **32** (1994), no. 1, 104–114. MR 1287018 (95h:06006)

[166] ———, *Alan Day's early work: congruence identities*, Algebra Universalis **34** (1995), no. 1, 4–23. MR 1344951 (97a:01050)

[167] ———, *Equations implying congruence n-permutability and semidistributivity*, Algebra Universalis **70** (2013), no. 4, 347–357. MR 3127978

[168] Ralph Freese, Christian Herrmann, and András P. Huhn, *On some identities valid in modular congruence varieties*, Algebra Universalis **12** (1981), no. 3, 322–334. MR 624298 (82m:08007)

[169] Ralph Freese, Jennifer Hyndman, and James B. Nation, *On the lattice of subquasivarieties of a locally finite quasivariety of finite type*, preprint, 2015.

[170] Ralph Freese, Jaroslav Ježek, and James B. Nation, *Free Lattices*, Mathematical Surveys and Monographs, vol. 42, American Mathematical Society, Providence, RI, 1995. MR 1319815 (96c:06013)

[171] ———, *Lattices with large minimal extensions*, Algebra Universalis **45** (2001), no. 2-3, 221–309, Conference on Lattices and Universal Algebra (Szeged, 1998). MR 1810550 (2002e:06008)

[172] Ralph Freese and Bjarni Jónsson, *Congruence modularity implies the Arguesian identity*, Algebra Universalis **6** (1976), no. 2, 225–228. MR 0472644 (57 #12340)

[173] Ralph Freese, Keith A. Kearnes, and James B. Nation, *Congruence lattices of congruence semidistributive algebras*, Lattice theory and its applications (Darmstadt, 1991), Res. Exp. Math., vol. 23, Heldermann, Lemgo, 1995, pp. 63–78. MR 1366865 (97a:08015)

[174] Ralph Freese, George F. McNulty, and James B. Nation, *Inherently nonfinitely based lattices*, Ann. Pure Appl. Logic **115** (2002), no. 1-3, 175–193. MR 1897025 (2003c:06005)

[175] Ralph Freese and James B. Nation, *Congruence lattices of semilattices*, Pacific J. Math. **49** (1973), 51–58. MR 0332590 (48 #10916)

[176] ———, *Projective lattices*, Pacific J. Math. **75** (1978), no. 1, 93–106. MR 500031 (80c:06012)

[177] ———, *Covers in free lattices*, Trans. Amer. Math. Soc. **288** (1985), no. 1, 1–42. MR 773044 (86f:06016)

[178] Louis Frey, *Parties distributives du treillis des permutations*, Ordres Totaux Finis, Gauthier-Villars, 1971, Mathématiques et Sciences de l'Homme XII, Paris, pp. 115–126.

[179] Haya Friedman and Dov Tamari, *Problèmes d'associativité: Une structure de treillis finis induite par une loi demi-associative*, J. Combinatorial Theory **2** (1967), 215–242. MR 0238984 (39 #344)

[180] Fred Galvin and Bjarni Jónsson, *Distributive sublattices of a free lattice*, Canadian J. Math. **13** (1961), 265–272.

[181] Bernhard Ganter, *Algorithmen zur formalen Begriffsanalyse*, Beiträge zur Begriffsanalyse (Darmstadt, 1986), Bibliographisches Inst., Mannheim, 1987, report 831 TH Darmstadt, FB Mathematik, 1984, pp. 241–254. MR 949457

[182] Bernhard Ganter and Rudolf Wille, *Formal concept analysis*, Springer-Verlag, Berlin, 1999, Mathematical foundations, Translated from the 1996 German original by Cornelia Franzke. MR 1707295 (2000i:06002b)

[183] Winfried Geyer, *Intervallverdopplung und verwandte Konstruktion bei Verbänden*, Tech. report, Fachbereich Mathematik der Technische Hochschule Darmstadt, 1992, Dissertation.

[184] ———, *The generalized doubling construction and formal concept analysis*, Algebra Universalis **32** (1994), no. 3, 341–367. MR 1294253 (95m:06020)

[185] ———, *On Tamari lattices*, Discrete Math. **133** (1994), no. 1-3, 99–122. MR 1298967 (95m:06021)

[186] Gerhard Gierz, Karl H. Hofmann, Klaus Keimel, Jimmy D. Lawson, Michael Mislove, and Dana S. Scott, *Continuous Lattices and Domains*, Encyclopedia of Mathematics and its Applications, vol. 93, Cambridge University Press, Cambridge, 2003. MR 1975381 (2004h:06001)

[187] Samuele Giraudo, *Algebraic and combinatorial structures on pairs of twin binary trees*, J. Algebra **360** (2012), 115–157. MR 2914637

[188] Jacob E. Goodman and Richard Pollack, *Multidimensional sorting*, SIAM J. Comput. **12** (1983), no. 3, 484–507. MR 707408 (85c:68082)

[189] Viktor A. Gorbunov, *Canonical decompositions in complete lattices*, Algebra i Logika **17** (1978), no. 5, 495–511, 622. MR 555259 (81e:06005)

[190] ———, *The cardinality of subdirectly irreducible systems in quasivarieties*, Algebra i Logika **25** (1986), no. 1, 3–50, 116. MR 892837 (89b:08010)

[191] ———, *The structure of the lattices of quasivarieties*, Algebra Universalis **32** (1994), no. 4, 493–530. MR 1300484 (95j:08009)

[192] ———, *Algebraic Theory of Quasivarieties*, Siberian School of Algebra and Logic, Consultants Bureau, New York, 1998, Translated from the Russian. MR 1654844 (2001a:08004)

[193] Viktor A. Gorbunov and V. I. Tumanov, *A class of lattices of quasivarieties*, Algebra i Logika **19** (1980), no. 1, 59–80, 132–133. MR 604658 (82f:08011)

[194] ———, *On the structure of lattices of quasivarieties*, Dokl. Akad. Nauk SSSR **254** (1980), no. 2, 272–275. MR 587152 (81m:08016)

[195] ———, *Construction of lattices of quasivarieties*, Mathematical logic and the theory of algorithms, Trudy Inst. Mat., vol. 2, "Nauka" Sibirsk. Otdel., Novosibirsk, 1982, pp. 12–44. MR 720198 (85e:08013)

[196] George Grätzer, *Equational classes of lattices*, Duke Math. J. **33** (1966), 613–622. MR 0199128 (33 #7278)

[197] ———, *Universal algebra*, Trends in Lattice Theory (Sympos., U.S. Naval Academy, Annapolis, Md., 1966), Van Nostrand Reinhold, New York, 1970, pp. 173–210. MR 0281674

[198] ———, *Lattice Theory. First Concepts and Distributive Lattices*, W. H. Freeman and Co., San Francisco, Calif., 1971. MR 0321817 (48 #184)

[199] ———, *A property of transferable lattices*, Proc. Amer. Math. Soc. **43** (1974), 269–271. MR 0335378

[200] ———, *Universal Algebra*, second ed., Springer-Verlag, New York-Heidelberg, 1979. MR 538623 (80g:08001)

[201] ———, *General Lattice Theory*, second ed., Birkhäuser Verlag, Basel, 1998, New appendices by the author with B. A. Davey, R. Freese, B. Ganter, M. Greferath, P. Jipsen, H. A. Priestley, H. Rose, E. T. Schmidt, S. E. Schmidt, F. Wehrung and R. Wille. MR 1670580 (2000b:06001)

[202] ———, *A technical lemma for congruences of finite lattices*, arXiv:1308.4607, 2013.

[203] ———, *Congruences and prime-perspectivities in finite lattices*, arXiv:1312.2537, 2013.

[204] George Grätzer, András P. Huhn, and Harry Lakser, *On the structure of finitely presented lattices*, Canad. J. Math. **33** (1981), 404–411.

[205] George Grätzer, Bjarni Jónsson, and Harry Lakser, *The amalgamation property in equational classes of modular lattices*, Pacific J. Math. **45** (1973), 507–524. MR 0366768 (51 #3014)

[206] George Grätzer and Harry Lakser, *The structure of pseudocomplemented distributive lattices. II. Congruence extension and amalgamation*, Trans. Amer. Math. Soc. **156** (1971), 343–358. MR 0274359 (43 #124)

[207] George Grätzer and Robert W. Quackenbush, *The variety generated by planar modular lattices*, Algebra Universalis **63** (2010), no. 2-3, 187–201. MR 2728135 (2011i:06016)

[208] George Grätzer and Friedrich Wehrung, *A new lattice construction: the box product*, J. Algebra **221** (1999), no. 1, 315–344. MR 1722915 (2000i:06015)

[209] George Grätzer and Friedrich Wehrung (eds.), *Lattice Theory: Special Topics and Applications. Vol. 1*, Birkhäuser/Springer, Cham, 2014. MR 3307662

[210] Oliver A. Gross, *Preferential arrangements*, Amer. Math. Monthly **69** (1962), 4–8. MR 0130837 (24 #A691)

[211] Peter M. Gruber, *The space of convex bodies*, Handbook of convex geometry, Vol. A, B, North-Holland, Amsterdam, 1993, pp. 301–318. MR 1242983 (95c:52005)

[212] Jean-Louis Guigues and Vincent Duquenne, *Familles minimales d'implications informatives résultant d'un tableau de données binaires*, Math. Sci. Humaines (1986), no. 95, 5–18, 83. MR 868423 (87k:08009)

[213] Georges Th. Guilbaud and Pierre Rosenstiehl, *Analyse algébrique d'un scrutin*, Math. Sci. Hum. **4** (1963), 9–33.

[214] Torsten Hahmann, *Model-Theoretic Analysis of Asher and Vieu's Mereotopology*, 2008, M.Sc. Thesis, University of Toronto.

[215] P. L. Hammer and Alexander Kogan, *Quasi-acyclic propositional Horn knowledge bases: optimal compression*, IEEE Transactions on Knowledge and Data Engineering **7** (1995), no. 5, 751–762.

[216] Guo-Niu Han, *Ordres bipartitionnaires et statistiques sur les mots*, Electron. J. Combin. **3** (1996), no. 2, Research Paper 3, approx. 5 pp. The Foata Festschrift. MR 1392488 (97c:05006)

[217] Peter Harremoës, *Functional dependences and Bayesian networks*, Proceedings WITMSE 2011, 2011, pp. 35–38.

[218] Jobst Heitzig and Jürgen Reinhold, *Counting finite lattices*, Algebra Universalis **48** (2002), no. 1, 43–53. MR 1930032

[219] Christian Herrmann, *Weak (projective) radius and finite equational bases for classes of lattices*, Algebra Universalis **3** (1973), 51–58. MR 0329983

[220] ______, *On the arithmetic of projective coordinate systems*, Trans. Amer. Math. Soc. **284** (1984), no. 2, 759–785. MR 743743 (85m:06019)

[221] Christian Herrmann and Anvar M. Nurakunov, *On locally finite modular lattice varieties of finite height*, Order **24** (2007), no. 1, 31–37. MR 2335851 (2008f:06017)

[222] Christian Herrmann, Douglas Pickering, and Michael Roddy, *A geometric description of modular lattices*, Algebra Universalis **31** (1994), no. 3, 365–396. MR 1265349 (95d:06005)

[223] Gábor Hetyei and Christian Krattenthaler, *The poset of bipartitions*, European J. Combin. **32** (2011), no. 8, 1253–1281. MR 2838014 (2012i:06005)

[224] Howard Hiller, *Geometry of Coxeter Groups*, Research Notes in Mathematics, vol. 54, Pitman (Advanced Publishing Program), Boston, Mass.-London, 1982. MR 649068 (83h:14045)

[225] David Hobby and Ralph N. McKenzie, *The Structure of Finite Algebras*, Contemporary Mathematics, vol. 76, American Mathematical Society, Providence, RI, 1988. MR 958685 (89m:08001)

[226] Hans-Jürgen Hoehnke, *Fully invariant algebraic closure systems of congruences and quasivarieties of algebras*, Lectures in universal algebra (Szeged, 1983), Colloq. Math. Soc. János Bolyai, vol. 43, North-Holland, Amsterdam, 1986, pp. 189–207. MR 860265 (87k:08012)

[227] Christophe Hohlweg and Jean-Philippe Labbé, *On inversion sets and the weak order in Coxeter groups*, arXiv:1502.06926, European J. Combin., to appear, 2015.

[228] Christophe Hohlweg and Carsten E. M. C. Lange, *Realizations of the associahedron and cyclohedron*, Discrete Comput. Geom. **37** (2007), no. 4, 517–543. MR 2321739 (2008g:52021)

[229] Christophe Hohlweg, Carsten E. M. C. Lange, and Hugh Thomas, *Permutahedra and generalized associahedra*, Adv. Math. **226** (2011), no. 1, 608–640. MR 2735770 (2012d:20085)

[230] Tristan Holmes, *Inflation of finite lattices along all-or-nothing sets*, Ph.D. thesis, University of Hawai'i, 2015, p. 54. MR 3427203

[231] Tristan Holmes and James B. Nation, *Inflation of finite lattices along all-or-nothing sets*, preprint.

[232] Dang Xuan Hong, *Covering relations among lattice varieties*, ProQuest LLC, Ann Arbor, MI, 1970, Thesis (Ph.D.)–Vanderbilt University. MR 2619629

[233] ______, *Covering relations among lattice varieties*, Pacific J. Math. **40** (1972), 575–603. MR 0306068 (46 #5195)

[234] Xianlong Hong, Gang Huang, Yici Cai, Jiangchun Gu, Sheqin Dong, Chung-Kuan Cheng, and Jun Gu, *Corner block list: An effective and efficient topological representation of non-slicing floorplan*, Proceedings of IEEE/ACM International Conference on Computer-Aided Design (ICCAD2000) (ed. ACM, IEEE), 5-9 Nov. 2000, 2000, pp. 8–12.

[235] Edward Howorka, *On metric properties of certain clique graphs*, J. Combin. Theory Ser. B **27** (1979), no. 1, 67–74. MR 539077 (80i:05032)

[236] Yuang Cheh Hsueh, *Lattices of width three generate a non-finitely based variety*, Algebra Universalis **17** (1983), no. 1, 132–134. MR 710006 (85b:06008)

[237] Samuel Huang and Dov 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. MR 0306064 (46 #5191)

[238] Danièle Huguet, *La structure du treillis des polyèdres de parenthésages*, Algebra Universalis **5** (1975), 82–87. MR 0434898 (55 #7862)

[239] András P. Huhn, *Schwach distributive Verbände. I*, Acta Sci. Math. (Szeged) **33** (1972), 297–305. MR 0337710 (49 #2479)

[240] ______, *On nonmodular n-distributive lattices: the decision problem for identities in finite n-distributive lattices*, Acta Sci. Math. (Szeged) **48** (1985), no. 1-4, 215–219. MR 810879 (87d:06027)

[241] James E. Humphreys, *Reflection groups and Coxeter groups*, Cambridge Studies in Advanced Mathematics, vol. 29, Cambridge University Press, Cambridge, 1990. MR 1066460 (92h:20002)

[242] Jennifer Hyndman, James B. Nation, and Joy Nishida, *Congruence lattices of semilattices with operators*, Studia Logica, to appear.

[243] Colin Ingalls and Hugh Thomas, *Noncrossing partitions and representations of quivers*, Compos. Math. **145** (2009), no. 6, 1533–1562. MR 2575093 (2010m:16021)

[244] Luisa Iturrioz, *A simple proof of a characterization of complete orthocomplemented lattices*, Bull. London Math. Soc. **14** (1982), no. 6, 542–544. MR 679931 (84a:06005)

[245] Osamu Iyama, Nathan Reading, Idun Reiten, and Hugh Thomas, *Algebraic lattice quotients of Weyl groups coming from preprojective algebras*, in preparation.

[246] Thomas Jech, *Set Theory*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, The third millennium edition, revised and expanded. MR 1940513 (2004g:03071)

[247] Přemysl Jedlička, *A combinatorial construction of the weak order of a Coxeter group*, Comm. Algebra **33** (2005), no. 5, 1447–1460. MR 2149070 (2006c:05143)

[248] Peter Jipsen and Nathan Lawless, *Generating all finite modular lattices of a given size*, Algebra Universalis **74** (2015), no. 3-4, 253–264. MR 3397437

[249] Peter Jipsen and Henry Rose, *Absolute retracts and amalgamation in certain congruence distributive varieties*, Canad. Math. Bull. **32** (1989), no. 3, 309–313. MR 1010069 (91b:08009)

[250] ———, *Varieties of Lattices*, Lecture Notes in Mathematics, vol. 1533, Springer-Verlag, Berlin, 1992, Out of print, available online at http://www1.chapman.edu/~jipsen/JipsenRoseVoL.html. MR 1223545 (94d:06022)

[251] Mark R. Johnson and Richard A. Dean, *An algebraic characterization of path independent choice functions*, Third International Meeting of the Society for Social Choice and Welfare, Maastricht, The Netherlands, 1996.

[252] ———, *Locally complete independent choice functions and their lattices*, Math. Social Sci. **42** (2001), no. 1, 53–87.

[253] ———, *Designer path independent choice functions*, Econom. Theory **26** (2005), no. 3, 729–740. MR 2214184 (2006j:91106)

[254] Bjarni Jónsson, *Sublattices of a free lattice*, Canad. J. Math. **13** (1961), 256–264. MR 0123493 (23 #A818)

[255] ———, *Algebras whose congruence lattices are distributive*, Math. Scand. **21** (1967), 110–121 (1968). MR 0237402 (38 #5689)

[256] ———, *Equational classes of lattices*, Math. Scand. **22** (1968), 187–196 (1969). MR 0246797 (40 #66)

[257] ———, *The variety covering the variety of all modular lattices*, Math. Scand. **41** (1977), no. 1, 5–14. MR 0463062 (57 #3025)

[258] ———, *Congruence varieties*, Algebra Universalis **10** (1980), no. 3, 355–394. MR 564122 (81e:08004)

[259] ______, *Amalgamation in small varieties of lattices*, J. Pure Appl. Algebra **68** (1990), no. 1-2, 195–208, Special issue in honor of B. Banaschewski. MR 1082790 (92b:06021)

[260] Bjarni Jónsson and James E. Kiefer, *Finite sublattices of a free lattice*, Canad. J. Math. **14** (1962), 487–497. MR 0137667 (25 #1117)

[261] Bjarni Jónsson and James B. Nation, *A report on sublattices of a free lattice*, Contributions to Universal Algebra (Colloq., József Attila Univ., Szeged, 1975), North-Holland, Amsterdam, 1977, pp. 223–257. Colloq. Math. Soc. János Bolyai, Vol. 17. MR 0472614 (57 #12310)

[262] Bjarni Jónsson and Ivan Rival, *Lattice varieties covering the smallest nonmodular variety*, Pacific J. Math. **82** (1979), no. 2, 463–478. MR 551703 (81j:06007)

[263] Gil Kalai, *A simple way to tell a simple polytope from its graph*, J. Combin. Theory Ser. A **49** (1988), no. 2, 381–383. MR 964396 (89m:52006)

[264] Myrto Kallipoliti and Henri Mühle, *On the topology of the Cambrian semilattices*, Electron. J. Combin. **20** (2013), no. 2, Paper 48, 21. MR 3084590

[265] Young Yug Kang, *Joins of finitely based lattice varieties*, ProQuest LLC, Ann Arbor, MI, 1987, Thesis (Ph.D.)–Vanderbilt University. MR 2635597

[266] Richard M. Karp, *Reducibility among combinatorial problems*, Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972), Plenum, New York, 1972, pp. 85–103. MR 0378476 (51 #14644)

[267] Kenji Kashiwabara and Masataka Nakamura, *The rooted circuits, the rooted cocircuits of convex geometries, the closure operators and monotone extensive operators*, slides of AAB II workshop, 2012.

[268] ______, *The prime stems of rooted circuits of closure spaces and minimum implicational bases*, Electron. J. Combin. **20** (2013), no. 1, Paper 22, 13. MR 3035032

[269] Kenji Kashiwabara, Masataka Nakamura, and Yoshio Okamoto, *The affine representation theorem for abstract convex geometries*, Comput. Geom. **30** (2005), no. 2, 129–144. MR 2107032 (2005m:52002)

[270] František Katrnoška, *On the representation of orthocomplemented posets*, Comment. Math. Univ. Carolin. **23** (1982), no. 3, 489–498. MR 677857 (84c:06003)

[271] David C. Kay and Gary Chartrand, *A characterization of certain ptolemaic graphs*, Canad. J. Math. **17** (1965), 342–346. MR 0175113 (30 #5298)

[272] Keith A. Kearnes and Emil W. Kiss, *The Shape of Congruence Lattices*, Mem. Amer. Math. Soc. **222** (2013), no. 1046, viii+169. MR 3076179

[273] Keith A. Kearnes and James B. Nation, *Axiomatizable and nonaxiomatizable congruence prevarieties*, Algebra Universalis **59** (2008), no. 3-4, 323–335. MR 2470584 (2009k:08004)

[274] Klaus Keimel and Jimmy D. Lawson, *Continuous and completely distributive lattices*, Lattice Theory: Selected Topics and Applications. Volume 1, Birkhäuser/Springer Basel AG, Basel, 2014, forthcoming, pp. 5–53.

[275] Thomas Kelly, *The myth of the skytale*, Cryptologia (1998), 244–260.

[276] Emil W. Kiss, László Márki, Péter Pröhle, and Walter Tholen, *Categorical algebraic properties. A compendium on amalgamation, congruence extension, epimorphisms, residual smallness, and injectivity*, Studia Sci. Math. Hungar. **18** (1982), no. 1, 79–140. MR 759319 (85k:18003)

[277] Mark Kleiner and Allen Pelley, *Admissible sequences, preprojective representations of quivers, and reduced words in the Weyl group of a Kac-Moody algebra*, Int. Math. Res. Not. IMRN (2007), no. 4, Art. ID rnm013, 28. MR 2338197 (2008f:16035)

[278] Daniel Kleitman, *On Dedekind's problem: The number of monotone Boolean functions*, Proc. Amer. Math. Soc. **21** (1969), 677–682. MR 0241334 (39 #2674)

[279] Donald E. Knuth, *The Art of Computer Programming. Volume 1: Fundamental Algorithms*, Second printing, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont, 1969. MR 0286317 (44 #3530)

[280] ———, *The Art of Computer Programming. Volume 3: Sorting and Searching. Second Edition*, Reading, Massachusetts: Addison-Wesley, 1998.

[281] M. Kolibiar, *Congruence relations and direct decompositions of ordered sets*, Acta Sci. Math. (Szeged) **51** (1987), no. 1-2, 129–135. MR 911564 (88m:06003)

[282] Bernhard Korte, László Lovász, and Rainer Schrader, *Greedoids*, Algorithms and Combinatorics, vol. 4, Springer-Verlag, Berlin, 1991. MR 1183735 (93f:90003)

[283] Gleb A. Koshevoy, *Choice functions and abstract convex geometries*, Math. Social Sci. **38** (1999), no. 1, 35–44. MR 1695061 (2001c:91026)

[284] Bertram Kostant, *Lie algebra cohomology and the generalized Borel-Weil theorem*, Ann. of Math. (2) **74** (1961), 329–387. MR 0142696 (26 #265)

[285] Alan Kostinsky, *Projective lattices and bounded homomorphisms*, Pacific J. Math. **40** (1972), 111–119. MR 0302519 (46 #1663)

[286] Christian Krattenthaler, *The enumeration of lattice paths with respect to their number of turns*, Advances in Combinatorial Methods and Applications to Probability and Statistics, Stat. Ind. Technol., Birkhäuser Boston, Boston, MA, 1997, pp. 29–58. MR 1456725 (98f:05006)

[287] G. Kreweras, *Sur les partitions non croisées d'un cycle*, Discrete Math. **1** (1972), no. 4, 333–350. MR 0309747 (46 #8852)

[288] Daniel Krob, Matthieu Latapy, Jean-Christophe Novelli, Ha Duong Phan, and Sylviane Schwer, *Pseudo-permutations I: First combinatorial and lattice properties*, FPSAC 2001, 2001, 13-th International Conference on Formal Power Series and Algebraic Combinatorics Arizona State University, USA May 20–26, 2001.

[289] Hans Kurzweil, *Endliche Gruppen mit vielen Untergruppen*, J. Reine Angew. Math. **356** (1985), 140–160. MR 779379 (86f:20024)

[290] Harry Lakser, *Exercises 30 and 31*, in *General Lattice Theory* [201], p. 68. MR 1670580 (2000b:06001)

[291] Shirley Elizabeth Law, *Combinatorial Realization of Certain Hopf Algebras of Pattern-Avoiding Permutations*, Ph.D. thesis, North Carolina State University, Raleigh, NC, 2013, p. 88.

[292] ———, *Combinatorial realization of the Hopf algebra of sashes*, arXiv:1407.4073, 2014.

[293] Shirley Elizabeth Law and Nathan Reading, *The Hopf algebra of diagonal rectangulations*, J. Combin. Theory Ser. A **119** (2012), no. 3, 788–824. MR 2871762

[294] Minh Ha Le and Ha Duong Phan, *Generalized pseudo-permutations*, FPSAC 2002, 2002, 14-th International Conference on Formal Power Series and Algebraic Combinatorics, The University of Melbourne, July 8–12, 2002.

[295] Claude Le Conte de Poly-Barbut, *Automorphismes du permutoèdre et votes de Condorcet*, Math. Inform. Sci. Humaines **111** (1990), 73–82. MR 1082275 (91j:90016)

[296] ———, *Le diagramme du treillis permutoèdre est intersection des diagrammes de deux produits directs d'ordres totaux*, Math. Inform. Sci. Humaines **112** (1990), 49–53, 61. MR 1096920 (92a:05004)

[297] ———, *Sur les treillis de Coxeter finis*, Math. Inform. Sci. Humaines **125** (1994), 41–57. MR 1281945 (95f:20064)

[298] Jeh Gwon Lee, *Almost distributive lattice varieties*, Algebra Universalis **21** (1985), no. 2-3, 280–304. MR 855745 (87i:06019)

[299] ———, *Joins of finitely based lattice varieties*, J. Korean Math. Soc. **22** (1985), no. 2, 125–133. MR 826438 (87e:06010)

[300] Leonid Libkin, *n-distributivity, dimension and Carathéodory's theorem*, Algebra Universalis **34** (1995), no. 1, 72–95. MR 1344955 (96h:06020)

[301] L. S.-C. Liu, *Left-modular elements and edge labellings*, Ph.D. thesis, Michigan State University, East Lansing, MI, 1999, p. 55.

[302] David Maier, *Minimum covers in the relational database model*, J. Assoc. Comput. Mach. **27** (1980), no. 4, 664–674. MR 594693 (82f:68025)

[303] Anatoliĭ I. Mal′cev, *Algebraic Systems*, Springer-Verlag, New York-Heidelberg, 1973, Posthumous edition, edited by D. Smirnov and M. Taĭclin, Translated from the Russian by B. D. Seckler and A. P. Doohovskoy, Die Grundlehren der mathematischen Wissenschaften, Band 192. MR 0349384 (50 #1878)

[304] George Markowsky, *Primes, irreducibles and extremal lattices*, Order **9** (1992), no. 3, 265–290. MR 1211380 (94a:06015)

[305] ———, *Permutation lattices revisited*, Math. Social Sci. **27** (1994), no. 1, 59–72. MR 1267688 (95g:20003)

[306] Robert Marsh, Markus Reineke, and Andrei Zelevinsky, *Generalized associahedra via quiver representations*, Trans. Amer. Math. Soc. **355** (2003), no. 10, 4171–4186. MR 1990581 (2004g:52014)

[307] René Mayet, *Une dualité pour les ensembles ordonnés orthocomplémentés*, C. R. Acad. Sci. Paris Sér. I Math. **294** (1982), no. 2, 63–65. MR 651787 (83b:06003)

[308] Thomas McConville, *Biclosed sets in real hyperplane arrangements*, arXiv:1411.1305, 2014.

[309] ———, *Crosscut-simplicial lattices*, arXiv:1409.6269, 2014.

[310] ———, *Lattice structure of grid-Tamari orders*, arXiv:1504.05213, 2015.

[311] William McCune, *Prover9 and Mace4 [computer software]*, 2005–2010.

[312] Ralph N. McKenzie, *Equational bases for lattice theories*, Math. Scand. **27** (1970), 24–38. MR 0274353 (43 #118)

[313] ______, *Equational bases and nonmodular lattice varieties*, Trans. Amer. Math. Soc. **174** (1972), 1–43. MR 0313141 (47 #1696)

[314] ______, *Finite forbidden lattices*, Universal algebra and lattice theory (Puebla, 1982), Lecture Notes in Math., vol. 1004, Springer, Berlin, 1983, pp. 176–205. MR 716183 (85b:06006)

[315] ______, *On minimal simple lattices*, Algebra Universalis **32** (1994), no. 1, 63–103. MR 1287017 (95i:06010)

[316] Ralph N. McKenzie, George F. McNulty, and Walter F. Taylor, *Algebras, Lattices, Varieties. Vol. I*, The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, 1987. MR 883644 (88e:08001)

[317] J. C. C. McKinsey, *The decision problem for some classes of sentences without quantifiers*, J. Symbolic Logic **8** (1943), 61–76. MR 0008991 (5,85a)

[318] Peter McNamara, *EL-labelings, supersolvability and 0-Hecke algebra actions on posets*, J. Combin. Theory Ser. A **101** (2003), no. 1, 69–89. MR 1953281 (2004c:05215)

[319] Charles F. Miller, III, *Decision problems for groups—survey and reflections*, Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989), Math. Sci. Res. Inst. Publ., vol. 23, Springer, New York, 1992, pp. 1–59. MR 1230627 (94i:20057)

[320] Yuya Mizuno, *Classifying τ-tilting modules over preprojective algebras of Dynkin type*, Math. Z. **277** (2014), no. 3-4, 665–690. MR 3229959

[321] Bernard Monjardet, *A use for frequently rediscovering a concept*, Order **1** (1985), no. 4, 415–417. MR 787552 (86e:06012)

[322] ______, *Arrowian characterizations of latticial federation consensus functions*, Math. Social Sci. **20** (1990), no. 1, 51–71. MR 1072291 (91j:90006)

[323] Bernard Monjardet and Vololonirina Raderanirina, *The duality between the anti-exchange closure operators and the path independent choice operators on a finite set*, Math. Social Sci. **41** (2001), no. 2, 131–150. MR 1806682 (2001j:91049)

[324] V. L. Murskiĭ, *The number of k-element algebras with a binary operation which do not have a finite basis of identities*, Problemy Kibernet. (1979), no. 35, 5–27, 208. MR 539884 (80i:08003)

[325] James B. Nation, *Varieties whose congruences satisfy certain lattice identities*, Algebra Universalis **4** (1974), 78–88. MR 0354501 (50 #6979)

[326] ———, *Some varieties of semidistributive lattices*, Universal algebra and lattice theory (Charleston, S.C., 1984), Lecture Notes in Math., vol. 1149, Springer, Berlin, 1985, pp. 198–223. MR 823017 (87d:06026)

[327] ———, *Lattice varieties covering $V(L_1)$*, Algebra Universalis **23** (1986), no. 2, 132–166. MR 896968 (88g:06009)

[328] ———, *An approach to lattice varieties of finite height*, Algebra Universalis **27** (1990), no. 4, 521–543. MR 1387900 (96m:06012)

[329] ———, *Addendum: "Congruence normality: the characterization of the doubling class of convex sets" [Algebra Universalis* **31** *(1994), no. 3, 397–406; MR1265350 (95a:06007)] by A. Day*, Algebra Universalis **31** (1994), no. 3, 407–410. MR 1265351 (95a:06008)

[330] ———, *Notes on Lattice Theory*, lecture notes, 1995.

[331] ———, *A counterexample to the finite height conjecture*, Order **13** (1996), no. 1, 1–9. MR 1402646 (97d:06008)

[332] ———, *Closure operators and lattice extensions*, Order **21** (2004), no. 1, 43–48 (2005). MR 2128033 (2005m:06009)

[333] ———, *Lattices of theories in languages without equality*, Notre Dame J. Form. Log. **54** (2013), no. 2, 167–175. MR 3028794

[334] James B. Nation and Joy Nishida, *A refinement of the equaclosure operator*, preprint, 2016.

[335] James B. Nation and James H. Schmerl, *The order dimension of relatively free lattices*, Order **7** (1990), no. 1, 97–99. MR 1086131 (92b:06023)

[336] Colin Naturman and Henry Rose, *Ultra-universal models*, Quaestiones Math. **15** (1992), no. 2, 189–195. MR 1185886 (93i:03049)

[337] Oscar Tivis Nelson, Jr., *Subdirect decompositions of lattices of width two*, Pacific J. Math. **24** (1968), 519–523. MR 0224515 (37 #114)

[338] R. Netter, *Eine Bemerkung zu Kongruenzverbänden*, preprint, 1986.

[339] Norbert Newrly, *Lattices of equational theories are congruence lattices of monoids with one additional unary operation*, Algebra Universalis **30** (1993), no. 2, 217–220. MR 1223629 (94g:08010)

[340] Petr S. Novikov, *Ob algoritmičeskoĭ nerazrešimosti problemy toždestva slov v teorii grupp [On the algorithmic unsolvability of the word problem in group theory]* , Trudy Mat. Inst. im. Steklov. no. 44, Izdat. Akad. Nauk SSSR, Moscow, 1955. MR 0075197 (17,706b)

[341] Anvar M. Nurakunov, *Equational theories as congruences of enriched monoids*, Algebra Universalis **58** (2008), no. 3, 357–372. MR 2415287 (2009d:08011)

[342] Oystein Ore, *Combinations of closure relations*, Ann. of Math. (2) **44** (1943), 514–533. MR 0009596 (5,170d)

[343] Peter Ouwehand and Henry Rose, *Small congruence distributive varieties: retracts, injectives, equational compactness and amalgamation*, Period. Math. Hungar. **33** (1996), no. 3, 207–228. MR 1610655 (99i:08003)

[344] ——, *Lattice varieties with non-elementary amalgamation classes*, Algebra Universalis **41** (1999), no. 4, 317–336. MR 1710333 (2001a:06010)

[345] Jean-Marcel Pallo, *Enumerating, ranking and unranking binary trees*, Comput. J. **29** (1986), no. 2, 171–175. MR 841678

[346] Dona Papert, *Congruence relations in semi-lattices*, J. London Math. Soc. **39** (1964), 723–729. MR 0166123 (29 #3401)

[347] Peter Perkins, *Basic questions for general algebras*, Algebra Universalis **19** (1984), no. 1, 16–23. MR 748904 (85k:08010)

[348] T. Kyle Petersen, *On the shard intersection order of a Coxeter group*, SIAM J. Discrete Math. **27** (2013), no. 4, 1880–1912. MR 3123822

[349] Vincent Pilaud and Francisco Santos, *The brick polytope of a sorting network*, European J. Combin. **33** (2012), no. 4, 632–662. MR 2864447

[350] Vincent Pilaud and Christian Stump, *EL-labelings and canonical spanning trees for subword complexes*, Dis. Geom. Optim. **69** (2013), 213–248.

[351] ——, *Brick polytopes of spherical subword complexes and generalized associahedra*, Adv. Math. **276** (2015), 1–61. MR 3327085

[352] Annette Pilkington, *Convex geometries on root systems*, Comm. Algebra **34** (2006), no. 9, 3183–3202. MR 2252665 (2007c:52001)

[353] Craig R. Platt, *Finite transferable lattices are sharply transferable*, Proc. Amer. Math. Soc. **81** (1981), no. 3, 355–358. MR 597639 (82c:06012)

[354] Charles R. Plott, *Path independence, rationality, and social choice*, Econometrica **41** (1973), 1075–1091. MR 0441223 (55 #14086)

[355] S. V. Polin, *Identities in congruence lattices of universal algebras*, Mat. Zametki **22** (1977), no. 3, 443–451. MR 0491407 (58 #10663)

[356] Maurice Pouzet, Klaus Reuter, Ivan Rival, and Nejib Zaguia, *A generalized permutahedron*, Algebra Universalis **34** (1995), no. 4, 496–509. MR 1357480 (97a:06009)

[357] Walter Prenowitz, *Total lattices of convex sets and of linear spaces*, Ann. of Math. (2) **49** (1948), 659–688. MR 0025734 (10,57a)

[358] Pavel Pudlák and Jiří Tůma, *Yeast graphs and fermentation of algebraic lattices*, Lattice Theory (Proc. Colloq., Szeged, 1974), North-Holland, Amsterdam, 1976, pp. 301–341. Colloq. Math. Soc. János Bolyai, Vol. 14. MR 0441799 (56 #193)

[359] ______, *Every finite lattice can be embedded in a finite partition lattice*, Algebra Universalis **10** (1980), no. 1, 74–95. MR 552159 (81e:06013)

[360] George B. Purdy, *Some combinatorial problems in the plane*, Discrete Geometry and Convexity (New York, 1982), Ann. New York Acad. Sci., vol. 440, New York Acad. Sci., New York, 1985, pp. 65–68. MR 809192 (87b:52021)

[361] Y. Qiu, *C-sortable words as green mutation sequences*, arXiv:1205.0034, 2012.

[362] Robert W. Quackenbush, *Completeness theorems for universal and implicational logics of algebras via congruences*, Proc. Amer. Math. Soc. **103** (1988), no. 4, 1015–1021. MR 954975 (89f:03020)

[363] Michael O. Rabin, *Recursive unsolvability of group theoretic problems*, Ann. of Math. (2) **67** (1958), 172–194. MR 0110743 (22 #1611)

[364] Nathan Reading, *Lattice and order properties of the poset of regions in a hyperplane arrangement*, Algebra Universalis **50** (2003), no. 2, 179–205. MR 2037526 (2004m:06006)

[365] ______, *The order dimension of the poset of regions in a hyperplane arrangement*, J. Combin. Theory Ser. A **104** (2003), no. 2, 265–285. MR 2019275 (2004m:52057)

[366] ______, *Lattice congruences of the weak order*, Order **21** (2004), no. 4, 315–344 (2005). MR 2209128 (2007a:20036)

[367] ______, *Lattice congruences, fans and Hopf algebras*, J. Combin. Theory Ser. A **110** (2005), no. 2, 237–273. MR 2142177 (2006b:20054)

[368] ______, *Cambrian lattices*, Adv. Math. **205** (2006), no. 2, 313–353. MR 2258260 (2007g:05195)

[369] ———, *Clusters, Coxeter-sortable elements and noncrossing partitions*, Trans. Amer. Math. Soc. **359** (2007), no. 12, 5931–5958. MR 2336311 (2009d:20093)

[370] ———, *Sortable elements and Cambrian lattices*, Algebra Universalis **56** (2007), no. 3-4, 411–437. MR 2318219 (2008d:20073)

[371] ———, *Noncrossing partitions, clusters and the Coxeter plane*, Sém. Lothar. Combin. **63** (2010), Art. B63b, 32. MR 2734030 (2011j:20099)

[372] ———, *Noncrossing partitions and the shard intersection order*, J. Algebraic Combin. **33** (2011), no. 4, 483–530. MR 2781960 (2012a:06003)

[373] ———, *Coarsening polyhedral complexes*, Proc. Amer. Math. Soc. **140** (2012), no. 10, 3593–3605. MR 2929028

[374] ———, *From the Tamari lattice to Cambrian lattices and beyond*, Associahedra, Tamari Lattices and Related Structures: Tamari Memorial Festschrift (F. Müller-Hoissen, J. M. Pallo, and J. Stasheff, eds.), Progress in Mathematics - Birkhäuser, vol. 299, Springer Basel, 2012, pp. 293–322.

[375] ———, *Generic rectangulations*, European J. Combin. **33** (2012), no. 4, 610–623. MR 2864445

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

[377] ———, *Lattice homomorphisms between weak orders*, in preparation.

[378] ———, *Dominance, mutation-linear maps, and morphisms of cluster algebras*, in preparation.

[379] Nathan Reading and David E. Speyer, *Cambrian fans*, J. Eur. Math. Soc. (JEMS) **11** (2009), no. 2, 407–447. MR 2486939 (2011a:20102)

[380] ———, *Sortable elements for quivers with cycles*, Electron. J. Combin. **17** (2010), no. 1, Research Paper 90, 19. MR 2661393 (2011i:20061)

[381] ———, *Sortable elements in infinite Coxeter groups*, Trans. Amer. Math. Soc. **363** (2011), no. 2, 699–761. MR 2728584 (2011j:20100)

[382] ———, *Combinatorial frameworks for cluster algebras*, Int. Math. Res. Not. **2016** (2016), no. 1, 109–173.

[383] Nathan Reading and Debra J. Waugh, *The order dimension of Bruhat order on infinite Coxeter groups*, Electron. J. Combin. **11** (2004/06), no. 2, Research Paper 13, 26. MR 2120108 (2006a:06003)

[384] Victor Reiner, *Non-crossing partitions for classical reflection groups*, Discrete Math. **177** (1997), no. 1-3, 195–222. MR 1483446 (99f:06005)

[385] Vladimir B. Repnitskiĭ, *On finite lattices which are embeddable in subsemigroup lattices*, Semigroup Forum **46** (1993), no. 3, 388–397. MR 1206216 (94c:20101)

[386] H. Reppe, *Three generalizations of lattice distributivity: An fca approach*, Shaker Verlag, 2011, Ph.D Thesis, 178pp.

[387] Michael Richter and Luke G. Rogers, *Embedding convex geometries and a bound on convex dimension*, arXiv:1502.01941, 2015.

[388] Henry Rose, *Nonmodular lattice varieties*, Mem. Amer. Math. Soc. **47** (1984), no. 292, v+76. MR 727819 (85h:06018)

[389] W. Ruckelshausen, *Zur Hierarchie kleiner Verbandsvarietäten*, 1983, Thesis (Ph.D.)–Darmstadt.

[390] Anja Rusch and Rudolf Wille, *Knowledge spaces and formal concept analysis*, Data Analysis and Information Systems (H. H. Bock and W. Polasek, eds.), Studies in Classification, Data Analysis, and Knowledge Organization, Springer Verlag, Berlin, 1995, pp. 427–436.

[391] M. Salvetti, *Topology of the complement of real hyperplanes in* $\mathbf{C}^N$, Invent. Math. **88** (1987), no. 3, 603–618. MR 884802 (88k:32038)

[392] Luigi Santocanale, *A duality for finite lattices*, Preprint (2009). Available online at http://hal.archives-ouvertes.fr/hal-00432113/.

[393] ———, *On the join dependency relation in multinomial lattices*, Order **24** (2007), no. 3, 155–179. MR 2358079 (2008m:06008)

[394] Luigi Santocanale and Friedrich Wehrung, *Sublattices of associahedra and permutohedra*, Adv. in Appl. Math. **51** (2013), no. 3, 419–445. MR 3084507

[395] ———, *Varieties of lattices with geometric descriptions*, Order **30** (2013), no. 1, 13–38. MR 3018208

[396] ———, *The extended permutohedron on a transitive binary relation*, European J. Combin. **42** (2014), 179–206. MR 3240144

[397] ———, *Lattices of regular closed subsets of closure spaces*, Internat. J. Algebra Comput. **24** (2014), no. 7, 969–1030. MR 3286148

[398] ———, *The equational theory of the weak Bruhat order on finite symmetric groups*, 41 pages, available online at http://hal.archives-ouvertes.fr/hal-00986148, preprint.

[399] I. Richard Savage, *Contributions to the theory of rank order statistics: Applications of lattice theory*, Rev. Inst. Internat. Statist. **32** (1964), 52–64. MR 0183076 (32 #558)

[400] E. Tamás Schmidt, *Zur Charakterisierung der Kongruenzverbände der Verbände*, Mat. Časopis Sloven. Akad. Vied **18** (1968), 3–20. MR 0241335 (39 #2675)

[401] ______, *Kongruenzrelationen algebraischer Strukturen*, Mathematische Forschungsberichte, XXV, VEB Deutscher Verlag der Wissenschaften, Berlin, 1969. MR 0255474 (41 #136)

[402] Marina Schwidefsky, *The class of bounded lattices is not elementary*, Algebra and Logic, to appear, 2015.

[403] Milan Sekanina, *On a characterisation of the system of all regularly closed sets in general closure spaces*, Math. Nachr. **38** (1968), 61–66. MR 0234875 (38 #3189)

[404] Marina V. Semenova, *Lattices of suborders*, Sibirsk. Mat. Zh. **40** (1999), no. 3, 673–682, iv. MR 1709018 (2000h:06009)

[405] ______, *Lattices with unique irreducible decompositions*, Algebra Log. **39** (2000), no. 1, 93–103, 122. MR 1782320 (2001h:06008)

[406] ______, *On lattices that are embeddable into lattices of suborders*, Algebra Logika **44** (2005), no. 4, 483–511, 514. MR 2188936 (2006h:06008)

[407] ______, *On lattices embeddable into lattices of algebraic subsets*, Algebra Universalis **61** (2009), no. 3-4, 399–405. MR 2565861

[408] Marina V. Semenova and Friedrich Wehrung, *Sublattices of lattices of order-convex sets. II. Posets of finite length*, Internat. J. Algebra Comput. **13** (2003), no. 5, 543–564. MR 2027222 (2004j:06007)

[409] ______, *Sublattices of lattices of order-convex sets. I. The main representation theorem*, J. Algebra **277** (2004), no. 2, 825–860. MR 2067633 (2005e:06008)

[410] ______, *Sublattices of lattices of order-convex sets. III. The case of totally ordered sets*, Internat. J. Algebra Comput. **14** (2004), no. 3, 357–387. MR 2075158 (2005d:06008)

[411] Robert W. Shannon, *Simplicial cells in arrangements of hyperplanes*, Geom. Dedicata **8** (1979), no. 2, 179–187. MR 538524 (80g:52008)

[412] Jian-Yi Shi, *The enumeration of Coxeter elements*, J. Algebraic Combin. **6** (1997), no. 2, 161–171. MR 1436533 (98d:20048)

[413] V. M. Shiryaev, *Semilattices with semidistributive lattices of subsemilattices*, Vestnik Beloruss. Gos. Univ. Ser. I (1985), no. 1, 61–64, 80. MR 784209 (86g:06013)

[414] Bohuslav Sivák, *Representations of finite lattices by orders on finite sets*, Math. Slovaca **28** (1978), no. 2, 203–215. MR 526859 (80b:06007)

[415] Thoralf Skolem, *Logisch-Kombinatorische Untersuchungen über die Erfüllbarkeit oder Beweisbarkeit Mathematischer Sätze nebst einem Theoreme über dichte Mengen*, Skrifter. I. Mat.-naturv. klasse, 1920. no. 4. Utgit for Fridtjof Nansens Fond, Kristiania: in Kommission bei Jacob Dybwad, 1920 (Ink stamp: Printed in Norway), 1920.

[416] Václav Slavík, *Lattices with finite W-covers*, Algebra Universalis **36** (1996), no. 3, 286–315. MR 1408728 (97g:06008)

[417] Richard P. Stanley, *Supersolvable lattices*, Algebra Universalis **2** (1972), 197–217. MR 0309815 (46 #8920)

[418] ______, *Combinatorics and Commutative Algebra*, second ed., Progress in Mathematics, vol. 41, Birkhäuser Boston Inc., Boston, MA, 1996. MR 1453579 (98h:05001)

[419] Richard P. Stanley, *Enumerative Combinatorics. Vol. 2*, Cambridge Studies in Advanced Mathematics, vol. 62, Cambridge University Press, Cambridge, 1999, With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin. MR 1676282 (2000k:05026)

[420] ______, *Catalan addendum*, Regularly updated manuscript available online from Stanley's website on the monograph *Enumerative Combinatorics*, http://www-math.mit.edu/~rstan/ec/, July 2012.

[421] ______, *Enumerative Combinatorics. Volume 1*, second ed., Cambridge Studies in Advanced Mathematics, vol. 49, Cambridge University Press, Cambridge, 2012. MR 2868112

[422] James D. Stasheff, *Homotopy associativity of H-spaces. I*, Trans. Amer. Math. **108** (1963), 275–292. MR 0158400 (28 #1623)

[423] ______, *Homotopy associativity of H-spaces. II*, Trans. Amer. Math. **108** (1963), 293–312. MR 0158400 (28 #1623)

[424] Rafael Stekolshchik, *Notes on Coxeter Transformations and the McKay Correspondence*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2008. MR 2388772 (2009e:20089)

[425] John R. Stembridge, *Computational aspects of root systems, Coxeter groups, and Weyl characters*, Interaction of combinatorics and representation theory, MSJ Mem., vol. 11, Math. Soc. Japan, Tokyo, 2001, pp. 1–38. MR 1862148 (2002k:05244)

[426] ______, *A Maple Package for Root Systems and Finite Coxeter Groups*, February 28, 2004, Version 2.4; introduction at http://www.math.lsa.umich.edu/~jrs/software/coxeter.ps.

[427] Christian Stump, Hugh Thomas, and Nathan Williams, *Cataland: Why the Fuss?*, arXiv:1503.00710, 2015.

[428] Dov Tamari, *Monoïdes préordonnés et chaînes de Malcev*, Ph.D. thesis, Thèse, Université de Paris, 1951, (mimeographed), pp. iv+81. MR 0051833 (14,532b)

[429] ______, *The algebra of bracketings and their enumeration*, Nieuw Arch. Wisk. (3) **10** (1962), 131–146. MR 0146227 (26 #3749)

[430] Walter F. Taylor, *Residually small varieties*, Algebra Universalis **2** (1972), 33–53. MR 0314726 (47 #3278)

[431] ______, *Products of absolute retracts*, Algebra Universalis **3** (1973), 400–401. MR 0349538 (50 #2031)

[432] Hugh Thomas, *An analogue of distributivity for ungraded lattices*, Order **23** (2006), no. 2-3, 249–269. MR 2308910 (2008c:06012)

[433] ______, *Tamari lattices and noncrossing partitions in type B*, Discrete Math. **306** (2006), no. 21, 2711–2723. MR 2263728 (2008d:05014)

[434] Michael Tischendorf, *The representation problem for algebraic distributive lattices*, Ph.D. thesis, TH Darmstadt, 1992.

[435] Jacques Tits, *Le problème des mots dans les groupes de Coxeter*, Symposia Mathematica (INDAM, Rome, 1967/68), Vol. 1, Academic Press, London, 1969, pp. 175–185. MR 0254129 (40 #7339)

[436] William T. Trotter, *Combinatorics and Partially Ordered Sets*, Johns Hopkins Series in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1992, Dimension theory. MR 1169299 (94a:06001)

[437] V. I. Tumanov, *On quasivarieties of lattices*, Proc. of XVI-th All-Union Algebraic Conference, Leningrad, part 2, 1981, p. 135.

[438] ______, *Embedding theorems for join-semidistributive lattices*, Proc. 6-th All-Union Conference on Math. Logic, Tbilisi, 1982, p. 188.

[439] ———, *Finite distributive lattices of quasivarieties*, Algebra i Logika **22** (1983), no. 2, 168–181. MR 750708 (86a:08009)

[440] Alasdair Urquhart, *A topological representation theory for lattices*, Algebra Universalis **8** (1978), no. 1, 45–58. MR 0450150 (56 #8447)

[441] Carl G. Wagner, *Enumeration of generalized weak orders*, Arch. Math. (Basel) **39** (1982), no. 2, 147–152. MR 675654 (83k:06009)

[442] Nathalie Wahl, *Antimatroids of finite character*, J. Geom. **70** (2001), no. 1-2, 168–175. MR 1825557 (2002k:52014)

[443] Morgan Ward, *The closure operators of a lattice*, Ann. of Math. (2) **43** (1942), 191–196. MR 0006144 (3,261e)

[444] Bogdan Węglorz, *Equationally compact algebras. I*, Fund. Math. **59** (1966), 289–298. MR 0210575 (35 #1462)

[445] Friedrich Wehrung, *The dimension monoid of a lattice*, Algebra Universalis **40** (1998), no. 3, 247–411. MR 1668068 (2000i:06014)

[446] ———, *Sublattices of complete lattices with continuity conditions*, Algebra Universalis **53** (2005), no. 2-3, 149–173. MR 2148292 (2006j:06011)

[447] Friedrich Wehrung and Marina V. Semenova, *Sublattices of lattices of convex subsets of vector spaces*, Algebra Logika **43** (2004), no. 3, 261–290, 382. MR 2084037 (2005e:06029)

[448] Philip M. Whitman, *Free lattices*, Ann. of Math. (2) **42** (1941), 325–330. MR 0003614 (2,244f)

[449] ———, *Free lattices. II*, Ann. of Math. (2) **43** (1942), 104–115. MR 0006143 (3,261d)

[450] ———, *Lattices, equivalence relations, and subgroups*, Bull. Amer. Math. Soc. **52** (1946), 507–522. MR 0016750 (8,62b)

[451] Marcel Wild, *A theory of finite closure spaces based on implications*, Adv. Math. **108** (1994), no. 1, 118–139. MR 1293585 (95j:06004)

[452] ———, *Optimal implicational bases for finite modular lattices*, Quaest. Math. **23** (2000), no. 2, 153–161. MR 1795732 (2002e:06007)

[453] ———, *The joy of implications, aka pure horn formulas: Mainly a survey*, arXiv:1411.6432, Theoret. Comput. Sci., to appear, 2016.

[454] Rudolf Wille, *Primitive subsets of lattices*, Algebra Universalis **2** (1972), 95–98. MR 0311524 (47 #86)

[455] ———, *Subdirect decomposition of concept lattices*, Algebra Universalis **17** (1983), no. 3, 275–287. MR 729937 (85m:06012a)

[456] ———, *Tensorial decomposition of concept lattices*, Order **2** (1985), no. 1, 81–95. MR 794628 (87b:06016)

[457] C. Y. Wong, *Lattice varieties with weak distributivity*, 1989, Thesis (Ph.D.)–University of Hawaii.

[458] Takemi Yanagimoto and Masashi Okamoto, *Partial orderings of permutations and monotonicity of a rank correlation statistic*, Ann. Inst. Statist. Math. **21** (1969), 489–506. MR 0258209 (41 #2856)

[459] Shih-Wei Yang and Andrei Zelevinsky, *Cluster algebras of finite type via Coxeter elements and principal minors*, Transform. Groups **13** (2008), no. 3-4, 855–895. MR 2452619 (2009j:13029)

[460] Bo Yao, Hongyu Chen, Chung-Kuan Cheng, and Ronald L. Graham, *Floorplan representations: Complexity and connections*, ACM Trans. Des. Autom. Electron. Syst. **8** (2003), no. 1, 55–80.

[461] Mitsuru Yasuhara, *The amalgamation property, the universal-homogeneous models, and the generic models*, Math. Scand. **34** (1974), 5–36. MR 0371642 (51 #7860)

[462] Hiyori Yoshikawa, Hiroshi Hirai, and Kazuhisa Makino, *A representation of antimatroids by Horn rules and its application to educational systems*, arXiv:1508.05465, version 2, 2015.

[463] Günter M. Ziegler, *A new local criterion for the lattice property*, Algebra Universalis **31** (1994), no. 4, 608–610. MR 1274077 (94m:06003)

[464] ———, *Lectures on Polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995. MR 1311028 (96a:52011)

Corrections to STA1

- **Page 284 (footnote), line** -6: change "$\dot{x}_i \preceq \dot{y}_0 \vee \dot{y}_1$ and $\dot{x}_0 \wedge \dot{x}_1 \preceq \dot{y}_i$" to "$\dot{x}_i \preceq \dot{y}_0 \vee \dot{y}_1$ or $\dot{x}_0 \wedge \dot{x}_1 \preceq \dot{y}_i$".
- **Page 284 (footnote), line** -3: change "$\|\dot{x}_i \preceq \dot{y}_0 \vee \dot{y}_1\| \wedge \|\dot{x}_0 \wedge \dot{x}_1 \preceq \dot{y}_i\|$" to "$\|\dot{x}_i \preceq \dot{y}_0 \vee \dot{y}_1\| \vee \|\dot{x}_0 \wedge \dot{x}_1 \preceq \dot{y}_i\|$".
- **Page 290, line** 20 (line 11 of Exercise 7.16): change "$f_i\colon (A,\alpha) \to (B_1,\beta_1)$" to "$f_i\colon (A,\alpha) \to (B_i,\beta_i)$".
- **Page 291, line** 19 (line 3 of Exercise 7.20): change "$(\vee,0)$-semilattice" to "$(\vee,0)$-homomorphism".

Index

G. Grätzer, F. Wehrung (eds.), *Lattice Theory: Special Topics and Applications*,
DOI 10.1007/978-3-319-44236-5

Author index

G. Grätzer, F. Wehrung (eds.), *Lattice Theory: Special Topics and Applications*,
DOI 10.1007/978-3-319-44236-5