LATTICE POLYNOMIALS AND POLYTOPES

by

Jacob M. Farinholt
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Mathematics

Committee:

____________________ Dr. James Lawrence, Dissertation Director

____________________ Dr. Valeriu Soltan, Committee Member

____________________ Dr. Geir Agnarsson, Committee Member

____________________ Dr. David Marchette, Committee Member

____________________ Dr. David Walnut, Department Chair

____________________ Dr. Donna M. Fox, Associate Dean, Office of Student Affairs & Special Programs, College of Science

____________________ Dr. Ali Andalibi, Interim Dean, The College of Science

Date: ____________________ Summer Semester 2020
George Mason University
Fairfax, VA

Lattice Polynomials and Polytopes

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Jacob M. Farinholt
Master of Science
George Mason University, 2012
Bachelor of Science
University of Mary Washington, 2009

Director: Dr. James Lawrence, Professor
Department of Mathematical Sciences

Summer Semester 2020
George Mason University
Fairfax, VA

Dedication

I dedicate this dissertation to my wife Sarah, and my two daughters, Rachel and Olivia.

Acknowledgments

There are so many I would like to thank for their assistance, guidance, patience, and support throughout the (very) long process of obtaining my PhD. While I will do my best to thank many of you here, I am certain that some names will inevitably be omitted. For those of you I have not listed, please accept my sincerest apologies, and be assured that this omission is only representative of my ever-forgetful nature, and not in any way a sign of ingratitude.

I would like to thank my parents for being living examples of fortitude and perseverance. I would like to thank all of my many professors over the years who have helped to cultivate the love of mathematics I now have. Special thanks to my dissertation advisor, Dr. Lawrence, and my committee, Drs. Agnarsson, Soltan, and Marchette, for the patience and guidance you provided. I would like to thank my many friends and coworkers who have been forced to listen to me as I struggled through this thesis research, and helped me work through some of the problems.

I am especially grateful to my wife, Sarah, and two beautiful daughters, Rachel and Olivia. Sarah, you are the epitome of love and patience, and have been by my side through it all. You have been my lighthouse to guide me through many of life's storms, my personal motivator, my teammate, and my better half. Rachel and Olivia, you have been my personal cheerleaders from the beginning! I will always cherish and proudly wear all of the many math bowties, math socks, math t-shirts, math pajamas, and math jackets I've received from you, as well as the magnetic shape toys, "to help you with your polytopes, Daddy." Thank you for reminding me daily of the importance of curiosity and the joy of discovery.

I am grateful for the financial support for this degree provided by both the Naval Surface Warfare Center, Dahlgren Division and Booz Allen Hamilton.

Table of Contents

List of Figures

Abstract

LATTICE POLYNOMIALS AND POLYTOPES
Jacob M. Farinholt, PhD
George Mason University, 2020
Dissertation Director: Dr. James Lawrence

A polyhedron, expressed in its canonical form as the intersection of half-spaces, is implicitly defined with respect to a particular "component-wise" partial order. This partial order on $\mathbb{R}^n$ is a distributive lattice. While a polyhedron with this partial order may not be a sublattice of $\mathbb{R}^n$, it may still nevertheless retain some of its lattice structure. This thesis characterizes and classifies polyhedra in $\mathbb{R}^n$ according to how much lattice structure is retained. This is done by investigating their closure under convex clones of lattice polynomials. In addition, we investigate the join irreducibles of join semilattice polytopes, and show that they necessarily form faces of the polytope. We then characterize various attributes of these "join irreducible faces."

Chapter 1: Introduction

1.1 Background

Linear programming refers to the process of optimizing a linear function over $\mathbb{R}^n$ given a set of linear constraints. The canonical form of a linear program is given by the following:

$$\begin{aligned} \text{Minimize} \quad & f(\mathbf{x}) \\ \text{Subject to} \quad & A\mathbf{x} \leq \mathbf{b}, \end{aligned}$$

where $\mathbf{x} = (x_1, x_2, \ldots, x_n)^t$ is a vector in $\mathbb{R}^n$, $f(\mathbf{x}) = \sum_{i=1}^{n} c_i x_i$ for some vector $\mathbf{c} = (c_1, c_2, \ldots, c_n)^t \in \mathbb{R}^n$, A is an $m \times n$ real matrix, and $\mathbf{b}$ is a vector in $\mathbb{R}^m$. Many different problems can be characterized as linear programs, including mission planning, resource optimization, network flow, and so forth; and hence the application area is wide.

The collection of constraints in a linear program define a convex polyhedron in $\mathbb{R}^n$. The component-wise partial order on $\mathbb{R}^n$ imparts a natural lattice structure on this space. While previous efforts have characterized the collection of polyhedra that completely preserve this lattice, this dissertation provides initial progress on generalizing these results by attempting to determine just how much structure from the lattice is preserved by a given polyhedron.

On the other hand, *lattice programming*, an area of operations research pioneered by Arthur Veinott, concerns itself with characterizing how an optimal solution to a program changes as global parameters change whenever the problem domain is a lattice [17]. This often provides valuable qualitative information without significant computational overhead. More precisely, lattice programming aims to determine when an optimal solution $\mathbf{s} = s_t^0 \in \mathbb{R}^n$ of the program $\min_{\mathbf{s} \in L_t} f(\mathbf{s}, \mathbf{t})$ is increasing in the parameter $\mathbf{t} \in \mathbb{R}^m$, where L_t is a lattice, and $f(\mathbf{s}, \mathbf{t})$ is a *subattitive* real function, that is, it satisfies

$$f(\mathbf{a} \vee \mathbf{b}) + f(\mathbf{a} \wedge \mathbf{b}) \leq f(\mathbf{a}) + f(\mathbf{b}), \tag{1.1}$$

for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n \times \mathbb{R}^m$.

Linear programming overlaps with lattice programming only when the polytope that defines the linear constraints is also a sublattice of $\mathbb{R}^n$. As we shall see later (and as was proved originally by Veinott [16]), this only occurs under rather strict conditions, and hence the role of lattice programming within linear programming is rather limited.

Nevertheless, it is of interest to the author to consider how much of the lattice structure is preserved by a polytope in $\mathbb{R}^n$. While the global constraints of a lattice program are (not surprisingly) assumed to be a lattice, in certain applications, the full structure of the lattice may not be utilized. If we relax the condition that the constraints form a lattice, then we may find that certain aspects of lattice programming may have a wider range of applicability within the framework of linear programming.

To that end, this thesis is dedicated largely to the study of polytopes, their closure under various lattice polynomials, and characterizations therein. Before introducing any new results, we provide necessary background material in the remaining sections of this chapter. Chapter 2 investigates the clone lattice over $\mathbb{R}^n$ and motivates the study of convex geometry to characterize certain properties of this lattice. Chapter 3 completely characterizes the set of all convex clones of lattice polynomials on any finitely-generated free distributive lattice. Using the motivation from Chapter 2 and the results from Chapter 3, Chapter 4 then shows exactly when a convex lattice polynomial clone over $\mathbb{R}^n$ can be represented by a convex polyhedron, or more precisely, when a polyhedron is closed under a lattice polynomial clone. In Chapter 5 we step away from clone theory, and restrict ourselves to the study of join semilattice polytopes and their join irreducibles. Finally, in Chapter 6 we will discuss conclusions and possible avenues of future research.

1.2 Polytopes and Posets

Recall that a partially ordered set, or *poset*, is a set A, along with a binary relation (called the partial order) $\leq$ that is reflexive, antisymmetric, and transitive. That is to say:

(i) For all $x \in A$, $x \leq x$,

(ii) For all $x, y \in A$, if $x \leq y$ and $y \leq x$, then $x = y$,

(iii) For all $x, y, z \in A$, if $x \leq y$ and $y \leq z$, then $x \leq z$.

Any totally ordered subset of a poset is called a *chain*. A poset $(A, \leq)$ satisfying the property that for any pair $x, y \in A$ there exists a greatest lower bound (called the *meet* and denoted $x \wedge y$) is called a *meet semilattice*. The meet semilattice is called *complete* if every arbitrary subset (not just pairs) of elements of A has a greatest lower bound in A. Likewise, one satisfying the property that for any pair $x, y \in A$ there exists a least upper bound (called the *join* and denoted $x \vee y$) is called a *join semilattice*. The join semilattice is *complete* if every arbitrary subset of A has a least upper bound in A. A poset that is both a meet and join semilattice is called a *lattice*. A lattice is *complete* if it is both a complete join semilattice and a complete meet semilattice. In particular, every finite lattice is trivially complete. As another example, we have the following known result, which is a common exercise in introductory courses on this topic, the proof of which we include for completeness:

Proposition 1.2.1. *Every complete meet semilattice with a greatest element is a complete lattice.*

Proof. Let S be a complete semilattice with greatest element, and for each $x \in S$, let $\uparrow x := \{y \in S \mid x \leq y\}$. Now let B be any subset of S, and let B^u denote the upper bounds of B, that is, $B^u := \{x \in S \mid x \geq b \text{ for all } b \in B\} = \cap_{b \in B} \uparrow b$, which is necessarily nonempty since S contains a greatest element. But then

$$\bigvee B = \bigwedge B^u = \bigwedge \left(\bigcap_{b \in B} \uparrow b \right). \tag{1.2}$$

Hence S is closed under arbitrary joins, making S a complete lattice. $\square$

Alternatively, a lattice may be equivalently defined algebraically as a set A with two binary operations, $\wedge$ and $\vee$, satisfying the following three identities for all $a, b, c \in A$:

commutativity:

$$\begin{aligned} a \vee b &= b \vee a \\ a \wedge b &= b \wedge a, \end{aligned} \tag{1.3}$$

associativity:

$$\begin{aligned} a \wedge (b \wedge c) &= (a \wedge b) \wedge c \\ a \vee (b \vee c) &= (a \vee b) \vee c, \end{aligned} \tag{1.4}$$

and *absorption*:

$$a \wedge (a \vee b) = a \vee (a \wedge b) = a. \tag{1.5}$$

Another property that lattices satisfy, which can be derived from repeated applications of the absorption identity is *idempotence*:

$$\begin{aligned} a \wedge a &= a \\ a \vee a &= a. \end{aligned} \tag{1.6}$$

To see that idempotence is derived from the absorption identity, note that the absorption identity holds for all $b \in A$, and hence we may let $b = a \vee c$ for some $c \in A$. Then by the absorption identity we have:

$$a = a \vee (a \wedge (a \vee c)) = a \vee a. \tag{1.7}$$

The identity $a = a \wedge a$ can be obtained similarly.

Proposition 1.2.2. *These binary operators give rise to a partial order from the two equivalent relations:*

$$a \leq b \;\Leftrightarrow\; a \wedge b = a \;\Leftrightarrow\; a \vee b = b. \tag{1.8}$$

Proof. We will explicitly show this for the meet operator. The argument for the join operator follows a nearly identical argument. To see that this partial order is reflexive, observe that

$$a = a \wedge a \Leftrightarrow a \leq a. \tag{1.9}$$

To see that it is antisymmetric, observe that $x \leq y \Leftrightarrow x = x \wedge y$ and $y \leq x \Leftrightarrow y = x \wedge y$. And hence

$$x = x \wedge y = y. \tag{1.10}$$

To see that it is transitive, observe that $x \leq y \Leftrightarrow x = x \wedge y$ and $y \leq z \Leftrightarrow y = y \wedge z$, and hence

$$\begin{aligned} x \leq y \Leftrightarrow x &= x \wedge y \\ &= x \wedge (y \wedge z) \Leftrightarrow y \leq z \\ &= (x \wedge y) \wedge z \\ &= x \wedge z \Leftrightarrow x \leq z. \end{aligned} \tag{1.11}$$

□

Algebraically, a semilattice may alternatively be defined as a set A with a single binary operator $*$ that is commutative, associative, and idempotent.

Let $\mathcal{L}$ be a lattice, and let $a, b \in \mathcal{L}$ be distinct elements. The element b is said to *cover* a if $a \leq b$ and if $a \leq c \leq b$ for some $c \in \mathcal{L}$, then either $c = a$ or $c = b$. A lattice is *bounded* if it contains a greatest element 1 and least element 0. The *atoms* of a bounded lattice $\mathcal{L}$ are the elements that cover 0, and the *coatoms* are the elements covered by 1.

For any lattice $\mathcal{L} = (L, \wedge, \vee)$, there always exists a unique *dual* lattice $\mathcal{L}^* = (L, \vee, \wedge)$ obtained by reversing the partial order. That is $a \leq b$ in $\mathcal{L}$ if and only if $b \leq a$ in $\mathcal{L}^*$. The meet operator in $\mathcal{L}$ becomes the join operator in $\mathcal{L}^*$, and conversely. Likewise, if $\mathcal{L}$ contains a 0 and 1 element, then the atoms of $\mathcal{L}$ are the coatoms of $\mathcal{L}^*$ and the coatoms of $\mathcal{L}$ are the atoms of $\mathcal{L}^*$.

Let $(A, \leq)$ be any poset and let $B \subseteq A$. Then clearly $(B, \leq)$ is a poset, with partial order inherited from A. However, if $(A, \leq)$ is a (semi-)lattice and $B \subseteq A$, it is generally not the case that $(B, \leq)$ is also a (semi-)lattice. We note that, for simplicity, if the partial order is clear, we will often denote a poset $(A, \leq)$ simply by the set A.

Recall that a topological space is a set T and a collection τ of open subsets of T (the topology) satisfying:

- The empty set and T are both in τ,
- τ is closed under finite intersection, and
- τ is closed under arbitrary union.

Observe that the union is idempotent, commutative, and transitive, and hence the topology is a complete join semilattice, with join given by union. Moreover, the topology contains a least element, ϕ, and hence by the dual of Proposition 1.2.1, it follows that the topology is necessarily a complete lattice.

Another example, which will play a prominent role throughout the remainder of this dissertation, is the following. Consider the set $\mathbb{R}^n$ with the component-wise partial order. That is, for $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$ in $\mathbb{R}^n$, we have $\mathbf{x} \leq \mathbf{y}$ if and only if (iff) $x_i \leq y_i$ for all $1 \leq i \leq n$. It is straightforward to see that there exist well-defined meet and join given by $\mathbf{x} \wedge \mathbf{y} = (z_1, \ldots, z_n)$ and $\mathbf{x} \vee \mathbf{y} = (v_1, \ldots, v_n)$ where $z_i = \min(x_i, y_i)$ and

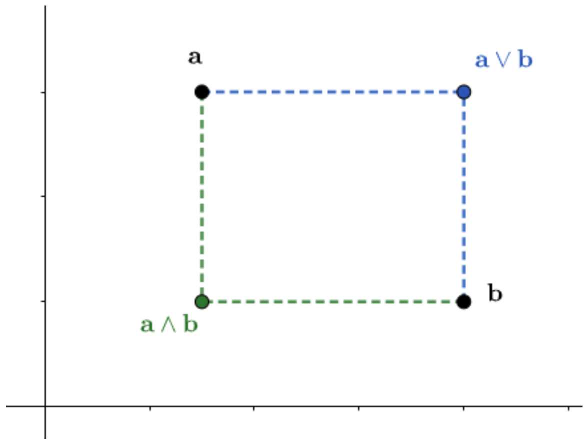

Figure 1.1: An example of a meet and join of two elements in $\mathbb{R}^2$ under the component-wise partial order.

$v_i = \max(x_i, y_i)$ for all $1 \leq i \leq n$ (see Figure 1.1). We see, then, that $\mathbb{R}^n$ with this partial order is, in fact, a lattice, which we will denote $\mathbb{L}^n$. With a little effort, one can show that this lattice is distributive, that is, it satisfies one (and hence both) of the following two (equivalent) properties: For all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathcal{L}$,

$$\mathbf{x} \wedge (\mathbf{y} \vee \mathbf{z}) = (\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{z}) \tag{1.12}$$

$$\mathbf{x} \vee (\mathbf{y} \wedge \mathbf{z}) = (\mathbf{x} \vee \mathbf{y}) \wedge (\mathbf{x} \vee \mathbf{z}). \tag{1.13}$$

However, $\mathbb{L}^n$ is not complete. Consider the line l in $\mathbb{L}^2$ given by the set $x = y$. It is a totally ordered, but unbounded subset of $\mathbb{L}^n$. So $\mathbb{L}^n$ contains neither $\bigwedge l$ nor $\bigvee l$.

A *polyhedron* in $\mathbb{R}^n$ is defined as the intersection of finitely many closed half-spaces in $\mathbb{R}^n$. If in addition, this intersection is bounded, then the polyhedron is called a *polytope*. Equivalently (albeit quite surprisingly), a polytope may also be defined as the convex closure of finitely many points in $\mathbb{R}^n$.

For any vector $\mathbf{a} \in \mathbb{R}^n$ and real number b, the closed half-space $H_{\mathbf{a},b}$ is given by the set of all $\mathbf{x} \in \mathbb{R}^n$ such that $\langle \mathbf{a}, \mathbf{x} \rangle \leq b$. But then if a polyhedron $\mathcal{P}$ in $\mathbb{R}^n$ is the intersection of closed half-spaces, $\mathcal{P} = H_{\mathbf{a}_1,b_1} \cap H_{\mathbf{a}_2,b_2} \cap \cdots \cap H_{\mathbf{a}_k,b_k}$, then we may write

$$\mathcal{P} = \mathcal{P}_{A,\mathbf{b}} = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}, \tag{1.14}$$

where A is a $k \times n$ matrix satisfying $A^t = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_k]$, and $\mathbf{b} = (b_1, b_2, \ldots, b_k)^t \in \mathbb{R}^k$. That is to say, the columns of A^t are the vectors $\mathbf{a}_i$ that define the half-spaces, and the entries in $\mathbf{b}$ are the associated b_i. We note in particular that the partial order $\leq$ in Equation (1.14) is the component-wise partial order described above. Hence, one may be led to believe that there *might* be some relationship between polyhedra and this particular partial order that may be leveraged to characterize and classify polyhedra, or conversely, we may be able to utilize polyhedra to better understand and characterize new order-theoretic attributes of this partial order on $\mathbb{R}^n$.

With a goal of doing precisely this, we are interested in analyzing polyhedra within the component-wise partial order $(\mathbb{R}^n, \leq)$ on $\mathbb{R}^n$ defined above. More precisely, let $\mathcal{P}$ be any polyhedron in $\mathbb{R}^n$. Then since $\mathcal{P} \subseteq \mathbb{R}^n$, it follows that $(\mathcal{P}, \leq)$ is a well-defined partially ordered set. We now have an entirely new avenue from which to study polyhedra. What properties of the poset $(\mathbb{R}^n, \leq)$ carry over to $(\mathcal{P}, \leq)$? What are necessary and sufficient conditions on $\mathcal{P}$ such that $(\mathcal{P}, \leq)$ also forms a lattice?

Let us give an example. Define a cone $\mathcal{C}$ in $(\mathbb{R}^3, \leq)$ by the collection of vectors $\mathbf{x} = (x_1, x_2, x_3)$ satisfying $x_1, x_2, x_3 \geq 0$ and

$$\begin{aligned} x_1 + x_2 &\geq x_3 \\ x_1 + x_3 &\geq x_2 \\ x_2 + x_3 &\geq x_1. \end{aligned} \tag{1.15}$$

(Recall that a cone $C \subseteq \mathbb{R}^n$ is a convex set satisfying the condition that, for any $\mathbf{x}, \mathbf{y} \in C$, $\alpha\mathbf{x} + \beta\mathbf{y} \in C$ for all $\alpha, \beta \geq 0$.) It is a straightforward exercise to see that $\mathcal{C}$ is closed under component-wise join: If $\mathbf{x}, \mathbf{y} \in \mathcal{C}$, then so is $\mathbf{x} \vee \mathbf{y}$. However, $\mathcal{C}$ is not closed under component-wise meet. For example $(1,6,7)$ and $(6,1,7)$ are both in $\mathcal{C}$, but $(1,6,7) \wedge (6,1,7) = (1,1,7) \notin \mathcal{C}$. Thus, $\mathcal{C}$ is a subsemilattice of $\mathbb{L}^3$ (viewed as a join semilattice), but not a sublattice of $\mathbb{L}^3$. One is naturally led to ask what other closure properties remain on $\mathcal{C}$? To answer this question, we turn to universal algebra.

1.3 Algebras on Polyhedra

In what follows, we will use the term "algebra" in the very general sense given by Birkhoff [2] which served as a prequel to the area now known as universal algebra. Namely, we define an algebra $\mathcal{A} = [A, F]$ to be a pair, where A is some nonempty set and F a collection of finitary operations (maps) onto A. More precisely, let $\mathcal{I} \subseteq \mathbb{N}_+$ be some (possibly infinite) subset of the positive integers, and for each $k \in \mathcal{I}$ let $F^{(k)}$ be some collection of maps $f^{(k)} : A^k \to A$. Let $F = \bigcup_{k \in \mathcal{I}} F^{(k)}$. Then $[A, F]$ defines an algebra. We say that maps of the form $f^{(k)}$ have "arity" k, or are "k-ary" maps on A.

For example, any lattice $\mathcal{L}$ is an algebra with two commutative and associative binary operations: $f(x, y) = x \wedge y$ and $g(x, y) = x \vee y$ that satisfy the absorption relation $f(x, g(x, y)) = g(x, f(x, y)) = x$. Similarly, a semilattice is an algebra with only one commutative and associative binary operation. A complemented lattice is a bounded lattice with an additional unary operation $c(x) = \dot{x}$ satisfying the relation $f(c(x), x) = 0$ and $g(c(x), x) = 1$ for all $x \in \mathcal{L}$. A group with 0 is an algebra with an associative binary operation $f(x, y) = x * y$ and a unary operation $g(x) = -x$ satisfying the relation $f(x, g(x)) = f(g(x), x) = 0$. Consequently this notion of an algebra is an abstraction of many general algebraic concepts.

Likewise, the notion of a *homomorphism* between algebras naturally generalizes the standard group, ring, lattice, etc. definitions of a homomorphism. Namely, if $\mathcal{A} = [A, F_A]$ and $\mathcal{B} = [B, F_B]$ are two algebras, then a map $h : \mathcal{A} \to \mathcal{B}$ is a homomorphism if for each

$f_A \in F_A$ and corresponding $f_B \in F_B$ of, say, arity k, we have:

$$h(f_A(a_1, a_2, \ldots, a_k)) = f_B(h(a_1), h(a_2), \ldots, h(a_k)). \tag{1.16}$$

Given an algebra $\mathcal{A} = [A, F]$, a subalgebra $\mathcal{T} \subseteq \mathcal{A}$ is a subset $T \subseteq A$ that is closed with respect to the operations of F, or F-closed. In the example in the previous section, we saw that the cone $\mathcal{C}$ is not a subalgebra of the algebra $[\mathbb{R}^3, \{\wedge, \vee\}]$ obtained from the lattice $\mathcal{L}^{(3)}$. However, $\mathcal{C}$ is a subalgebra of $[\mathbb{R}^3, \vee]$. Thus, more generally, to say that a polyhedron $\mathcal{P}$ satisfies some closure properties on a partially ordered set $(\mathbb{R}^n, \leq)$ is equivalent to saying that there is some algebra on $\mathbb{R}^n$ defined with respect to the partial order that admits a subalgebra on $\mathcal{P}$. It is interesting to consider whether the algebraic properties of a polyhedron either determine or are determined by its combinatorial properties.

For the sake of comparison, consider the following example. Let $\mathcal{O}$ be the poset cone in $\mathbb{L}^3$ given by the collection of $\mathbf{x} = (x_1, x_2, x_3) \in \mathbb{R}^3$ having all nonnegative entries:

$$\begin{aligned} x_1 &\geq 0 \\ x_2 &\geq 0 \\ x_3 &\geq 0 \end{aligned} \tag{1.17}$$

This cone corresponds to the nonnegative orthant in $\mathbb{R}^3$,

$$\mathcal{O} = \text{Cone}\left((1,0,0), (0,1,0), (0,0,1)\right). \tag{1.18}$$

It is not too difficult to verify that $\mathcal{O}$ is closed under *both* meet and join, and hence is a sublattice of $\mathbb{L}^3$. Compare this with the cone $\mathcal{C}$ defined in the previous section, which is only a subsemilattice of $\mathbb{L}^3$. However, we may equivalently write $\mathcal{C}$ as

$$\mathcal{C} = \text{Cone}\left((1,1,0), (0,1,1), (1,0,1)\right). \tag{1.19}$$

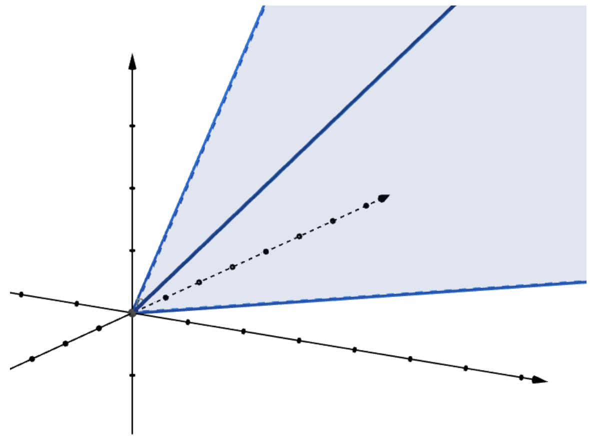

Figure 1.2: The cone $\mathcal{C}$ from Eq. (1.20).

In many ways, $\mathcal{C}$ and $\mathcal{O}$ are similar, even though their algebras are different. For example, they are both cones generated by three linearly independent vectors in $\mathbb{R}^3$, and their face posets are isomorphic. However, by writing both as systems of linear inequalities, we see that there are some distinguishing features. To make the representation of $\mathcal{C}$ more consistent with $\mathcal{O}$, we may rewrite it as

$$\begin{aligned} x_1 + x_2 - x_3 &\geq 0 \\ x_1 - x_2 + x_3 &\geq 0 \\ -x_1 + x_2 + x_3 &\geq 0. \end{aligned} \tag{1.20}$$

The left-hand side of each inequality defining $\mathcal{O}$ has only one variable, whereas the left-hand side of each inequality defining $\mathcal{C}$ in Eq. (1.20) has three variables. All of the variables defining $\mathcal{O}$ have positive scalars, whereas in the above representation of $\mathcal{C}$ each row has a variable with a negative scalar. It is perhaps the linear relations defining the polyhedra that determine the properties of their algebra. Indeed, in what follows, we will show that it is in fact only the signs of these scalars that determine when a given poset polyhedron is a sublattice of $\mathbb{L}^n$.

There is yet another direction that can be pursued in studying algebras on polyhedra. Thus far, we have considered the cases in which we have imposed an algebra on $\mathbb{R}^n$ and then studied and compared the algebraic properties that are inherited by various polyhedra in $\mathbb{R}^n$. Alternatively, given a poset polyhedron, one may consider classifying it according to its universal algebra. That is to say, can we leverage the (combinatorial, geometric) properties of a given poset polyhedron $\mathcal{P}$ to infer a characterization/classification of all algebras on $\mathcal{P}$? While we do not hope to provide such an analysis in its entirety in this thesis, it is our hope that some of the initial groundwork towards such an analysis can be provided herein. To perform such an analysis, we must leverage an area of universal algebra known as *clone theory*. Before we do this, however, we provide some historical context.

1.4 Post's Logic System

In 1920, a paper by Emil Post [11] developed an algebraic framework for the study of Boolean logic. This initial result was greatly expanded by him in another paper in 1941 [13]. In that paper, it was shown that logical propositions on a system of variables could be uniquely identified by their corresponding truth tables. As such, these propositions do not depend on their particular representation as compositions of basic propositional functions, with two propositions (that is, Boolean functions) on n variables being equivalent if the corresponding truth tables for these functions were identical.

For example, it is a straightforward exercise to verify that the two Boolean functions $f(x, y)$ and $g(x, y)$ on two variables x, y are equivalent, where $f(x, y) = (y \wedge \neg\, x) \vee (x \wedge y \wedge \neg\, x) \vee (\neg\, y \wedge \neg\, x)$ and $g(x, y) = \neg\, x$ (here, the notation "$\neg\, x$" means the complement of x in a Boolean lattice). The right-hand side of both of the above equations is called a *proposition* or *formula* in logic terms. Post's defining observation that motivated the entire paper was the fact that, because the two propositions generate the exact same truth tables, they correspond to the same function on two variables (though, obviously, in the above example the two functions on two variables end up actually only depending on the first variable - a property not obvious from the formulation of $f(x, y)$ but trivial to see once its equivalence with $g(x, y)$ is established).

His work had a profound influence on both logic and universal algebra, as it ultimately demonstrated that the underlying structure of propositional functions coincided with the free Boolean algebra Ω generated by a countably infinite set $X = \{x_1, x_2, \dots\}$ of variables. While Post's results were written for logicians and suffer from a lack of standard nomenclature in lattice theory and universal algebra at the time they were written, other authors have since made the mathematical formalism more explicit [1]. Separately, Post developed a more general m-valued logic system [12] (that is, a logic system whose propositions may have an integral number $m \geq 2$ of different truth values as opposed to the standard 2-valued "true," "false" Boolean system), which begged the question of whether the algebraic structure in the Boolean case could be applied more generally. The underlying algebra of these m-valued systems was developed by Rosenbloom [15] and further refined by Epstein [6].

Of particular note in [13] was that, in addition to mapping logic propositions to functions on a Boolean algebra, he further classified these functions into what he called distinct "classes," and demonstrated that these classes formed a lattice ordered by inclusion.

1.5 Universal Algebras and Clones

Generalizing Post's original results to more general algebras has become an area of considerable investigation in the field of universal algebra, and sometime in the 1960s this area of research became known as "clone theory." More specifically, we have the following definition [8]:

Definition 1.5.1. *For a set A and integer $n \geq 0$, let $O_A^{(n)}$ denote the set of n-ary operations on the set A, and set $O_A := \bigcup_{n>0} O_A^{(n)}$. A subset $C \subseteq O_A$ is called a* clone *if it contains all of the projection mappings $\pi_i^k : A^k \to A : (x_1, x_2, \dots, x_k) \mapsto x_i$, and is closed with respect to functional composition (also sometimes referred to as the "superposition" of operations):*

For $f \in O_A^{(n)} \bigcap C$ *and* $f_1, f_2, \ldots, f_n \in O_A^{(k)} \bigcap C$*, the* k*-ary operation* $f(f_1, f_2, \ldots, f_n)$ *defined by setting*

$$f(f_1, f_2, \ldots, f_n)(x_1, x_2, \ldots, x_k) := f(f_1(x_1, \ldots, x_k), \ldots f_n(x_1, \ldots, x_k)), \tag{1.21}$$

is also in C*.*

The base set A in the above definition is usually referred to as the *signature set.* Clones provide a way to study the behavior of algebras independent from their signatures. Almost all research in the study of clones are with respect to algebras having a finite signature set $|A| < \infty$. Some well-known examples of clones are:

- For any signature set A, the set O_A is a clone, as is the set J_A of all projections on A. These are called the full clone and trivial clone, respectively.

- Given an algebra $\mathbf{A} = (A, F)$, where F is a family of finitary operations on A, the set of finitary homomorphisms $\bigcup_{n>0} \mathrm{Hom}(\mathbf{A}^n, \mathbf{A})$ is a clone on A, called the *centralizer clone* of $\mathbf{A}$.

- A function f on A is said to be idempotent if $f(x, x, \ldots, x) = x$ for all $x \in A$. The collection of all idempotent functions is a clone on A.

- Given a partially ordered set $(A, \leq)$, all operations on A monotone in each variable with respect to that partial order form a clone, called the clone of that partial order.

It is known that, on any given domain of signatures A, the set of all clones of A forms a complete lattice with respect to inclusion, denoted $\mathcal{L}_A$. This lattice is clone theory's main object of study. For $|A| = 1$ the lattice has only one element. For $|A| = 2$ we have the case characterized by Post. The lattice is countably infinite, has 8 atoms and 5 coatoms, and is infinite only because of the existence of 8 infinite chains. Interestingly, however, moving to $|A| \geq 3$, very little is known about the corresponding lattice, as it is no longer countable, and there are no nontrivial lattice identities satisfied by $\mathcal{L}_A$ [5]. Even less still is known about the cases when A has infinite cardinality. A relatively recent review article on clones on infinite sets can be found in [7].

Chapter 2: The Clone Lattice, Galois Connections, and Polyhedra

Recall from Section 1.1 that one of the primary goals of this dissertation is to characterize how much of the lattice structure from $\mathbb{R}^n$ with the component-wise partial order may be preserved by a given polyhedron. In order to do this, we will leverage clone theory, but we would like to make the connection between clones on $\mathbb{R}^n$ and polyhedra a bit more explicit. This chapter will do this by leveraging a tool used heavily in clone theory - the Galois connection Inv - Pol between sets of finitary functions and sets of relations. With this structure in place, we can directly relate polyhedra to clones that preserve them.

2.1 Closure Operators and Closure Systems

Before we get into the main result of this section, we will take a slight detour to discuss closure operators and closure systems in more detail. These are well-known concepts, dating back to at least 1936 [3], but we review them here for completeness. In what follows, we will use the notation $\mathcal{P}(S)$ to denote the power set of S, that is, the collection of all subsets of S.

Definition 2.1.1. *A* closure system *on a set S is a family $\mathcal{F} \subseteq \mathcal{P}(S)$ of subsets of S closed under arbitrary intersection and containing S.*

Definition 2.1.2. *A* closure operator *h on a set S is a map $h : \mathcal{P}(S) \to \mathcal{P}(S)$ that satisfies, for all $A, B \in \mathcal{P}(S)$:*

(1) $A \subseteq h(A)$ *(extensive)*

(2) $A \subseteq B \Rightarrow h(A) \subseteq h(B)$ *(isotone)*

(3) $h(h(A)) = h(A)$ *(idempotent)*

It is well known that closure operators and closure systems are cryptomorphic. More precisely, for each closure operator h on a set S, there exists a unique corresponding closure system $\mathcal{F}_h$ on S given by the fixed points of h:

$$\mathcal{F}_h = \{A \in \mathcal{P}(S) \mid h(A) = A\}. \tag{2.1}$$

Conversely, for each closure system $\mathcal{F}$ on S, we may define a corresponding closure operator $h_{\mathcal{F}}$ on S satisfying, for any subset $K \subseteq S$,

$$h_{\mathcal{F}}(K) = \bigcap\{A \in \mathcal{F} \mid K \subseteq A\}. \tag{2.2}$$

It is straightforward to see that the elements of a closure system, ordered under inclusion, form a complete meet-semilattice, with partial order given by set inclusion. By definition, a closure system on a set S contains a greatest element, namely, S, and hence by Proposition 1.2.1, a closure system is also a complete lattice.

2.2 Galois Connections

A powerful tool for studying the clone lattice is Galois connections. Most notably, Galois connections were used to characterize the set of all coatoms of the clone lattice when the signature set A was finite [14]. These are well-known concepts, and we will briefly review them here. A brief introduction on this material can be found in [4], and the topic is also covered in [9].

Definition 2.2.1. *Let $\mathcal{X}$ and $\mathcal{Y}$ be nonempty sets. A pair of operators $\alpha : \mathcal{P}(\mathcal{X}) \to \mathcal{P}(\mathcal{Y})$ and $\beta : \mathcal{P}(\mathcal{Y}) \to \mathcal{P}(\mathcal{X})$ is called a* Galois connection *between $\mathcal{X}$ and $\mathcal{Y}$ if for all $X, X_1, X_2 \subseteq \mathcal{X}$ and $Y, Y_1, Y_2 \subseteq \mathcal{Y}$, the following hold:*

(1) $X_1 \subseteq X_2 \Rightarrow \alpha(X_1) \supseteq \alpha(X_2)$,

(2) $Y_1 \subseteq Y_2 \Rightarrow \beta(Y_1) \supseteq \beta(Y_2)$, *and*

(3) $X \subseteq \beta\alpha(X)$ *and* $Y \subseteq \alpha\beta(Y)$.

It is relatively easy to see that Galois connections give rise to closure operators.

Proposition 2.2.2. *Let* $\alpha - \beta$ *be a Galois connection between sets* $\mathcal{X}$ *and* $\mathcal{Y}$*. Then* $\beta\alpha : \mathcal{P}(\mathcal{X}) \to \mathcal{P}(\mathcal{X})$ *and* $\alpha\beta : \mathcal{P}(\mathcal{Y}) \to \mathcal{P}(\mathcal{Y})$ *are closure operators on* $\mathcal{X}$ *and* $\mathcal{Y}$*, respectively.*

Proof. That $\beta\alpha$ and $\alpha\beta$ satisfy the extensive and isotone properties of closure operators follows immediately from definition. We need only show idempotence. To do this, we first show that for all $Y \subseteq \mathcal{Y}$, we have $\beta\alpha\beta(Y) = \beta(Y)$. To see this, note that by property (3) of the definition of Galois connections, we have that $Y \subseteq \alpha\beta(Y)$. But then from property (2) of Galois connections it follows that $\beta\alpha\beta(Y) \subseteq \beta(Y)$. But by applying property (3) to $\beta(Y)$ we get that $\beta(Y) \subseteq \beta\alpha\beta(Y)$, giving us the desired equivalence. Now for any $X \subseteq \mathcal{X}$, let $Y = \alpha(X)$. Then since $\beta\alpha\beta(Y) = \beta(Y)$, it follows that $\beta\alpha\beta\alpha(X) = \beta\alpha(X)$, and hence $\beta\alpha$ satisfies idempotence, as we wanted to show. Similarly $\alpha\beta$ satisfies idempotence, and hence both $\beta\alpha$ and $\alpha\beta$ are closure operators on $\mathcal{X}$ and $\mathcal{Y}$, respectively. $\square$

Suppose $\alpha - \beta$ is a Galois connection between $\mathcal{X}$ and $\mathcal{Y}$ as above. Let $L_{\mathcal{X}}^{\beta\alpha}$ and $L_{\mathcal{Y}}^{\alpha\beta}$ be the closure systems corresponding to the closure operators $\beta\alpha$ and $\alpha\beta$, respectively. As we have already shown, a closure system, ordered under inclusion, is a complete lattice. A well-known result about Galois connections is that these lattices are connected. More precisely:

Theorem 2.2.3 ([4])**.** *The lattices* $L_{\mathcal{X}}^{\beta\alpha}$ *and* $L_{\mathcal{Y}}^{\alpha\beta}$ *are dually isomorphic. The dual isomorphisms are* $\alpha : L_{\mathcal{X}}^{\beta\alpha} \to L_{\mathcal{Y}}^{\alpha\beta}$ *and* $\beta : L_{\mathcal{Y}}^{\alpha\beta} \to L_{\mathcal{X}}^{\beta\alpha}$.

We now introduce the notion of a *relation*.

Definition 2.2.4. *Let* A *be any set. A* k*-ary* relation φ *on* A *is a subset* $\varphi \subseteq A^k$*. The set of all* k*-ary relations on* A *is denoted* $\mathbf{R}_A^{(k)}$*. The collection of all finitary relations on* A *is* *given by*

$$\mathbf{R}_A := \bigcup_{k\in\mathbb{N}^+} \mathbf{R}_A^{(k)}. \tag{2.3}$$

Let $f \in O_A$ be an n-ary function on A. Then for any $m \in \mathbb{N}^+$, we may extend f to a function $f : (A^m)^n \to A^m$ that acts component-wise. Namely for each m-tuple $\mathbf{a}$ in A^m we denote by $\mathbf{a}(i)$ the i-th component of $\mathbf{a}$. Then for $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n \in A^m$ we define

$$f\begin{pmatrix} \mathbf{a}_1(1) & \mathbf{a}_2(1) & \cdots & \mathbf{a}_n(1) \\ \mathbf{a}_1(2) & \mathbf{a}_2(2) & \cdots & \mathbf{a}_n(2) \\ \vdots & \vdots & & \vdots \\ \mathbf{a}_1(m) & \mathbf{a}_2(m) & \cdots & \mathbf{a}_n(m) \end{pmatrix} := \begin{pmatrix} f(\mathbf{a}_1(1), \ldots, \mathbf{a}_n(1)) \\ f(\mathbf{a}_1(2), \ldots, \mathbf{a}_n(2)) \\ \vdots \\ f(\mathbf{a}_1(m), \ldots, \mathbf{a}_n(m)) \end{pmatrix}. \tag{2.4}$$

Definition 2.2.5. *Let $f \in O_A$ be an n-ary map on a set A and let $\varphi \in \mathbf{R}_A$ be a k-ary relation on A. Then f is said to* preserve *φ, or equivalently that φ is an* invariant relation *or that f is a* polymorphism *of φ if for all $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n \in \varphi$ we have*

$$f(\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n) \in \varphi. \tag{2.5}$$

We are now ready to define the two particular Galois connections of interest to us: Pol - Inv.

Definition 2.2.6. *For any $F \subseteq O_A$ and $R \subseteq \mathbf{R}_A$ we define:*

$$\text{Inv } F := \{\varphi \in \mathbf{R}_A \mid \forall\ f \in F,\ f \textit{ preserves } \varphi\} \tag{2.6}$$

$$\text{Pol } R := \{f \in O_A \mid \forall\ \varphi \in R,\ f \textit{ preserves } \varphi\}. \tag{2.7}$$

As we mentioned at the beginning of this chapter, we are interested in using closure operators as a means to better characterize properties of clones. We now have the structure to do this. The universal algebra O_A on a set A and the set $\mathbf{R}_A$ of all relations on A will play the role of $\mathcal{X}$ and $\mathcal{Y}$ in the Galois connection Pol - Inv. Observe that, for any set of relations $R \subseteq \mathbf{R}_A$ on a set A, the set Pol R is a clone. Consequently, for any subset $F \subset O_A$, its closure Pol Inv F is a clone. Unfortunately, however, it is known [8] that, except when $|A| < \infty$, not every clone can be expressed in such a way. Nevertheless, because every set of relations determines a clone, it may be illuminating to characterize classes of clones by sets of relations that generate them under Pol.

Indeed, our goal is not to characterize all clones. Rather, in the next chapter we will be restricting ourselves to convex clones of lattice polynomials. Thankfully, restricted sets $\mathcal{M} \subseteq \mathcal{X}$ may inherit much of the Galois $\alpha - \beta$ closure:

Theorem 2.2.7 ([4])**.** *If $\mathcal{M} \subseteq \mathcal{X}$, then the operator pair $\alpha - \beta_{\mathcal{M}}$ given by*

$$\alpha : \mathcal{P}(\mathcal{M}) \to \mathcal{P}(\mathcal{Y}), \qquad X \mapsto \alpha(X) \tag{2.8}$$

$$\beta_{\mathcal{M}} : \mathcal{P}(\mathcal{Y}) \to \mathcal{P}(\mathcal{M}), \qquad Y \mapsto \mathcal{M} \cap \beta(Y), \tag{2.9}$$

forms a Galois connection between $\mathcal{M}$ and $\mathcal{Y}$. A subset $M \subseteq \mathcal{M}$ is Galois-closed under $\beta_{\mathcal{M}}\alpha$ if and only if $M = \mathcal{M} \cap X$ for some $X \in L_{\mathcal{X}}^{\beta\alpha}$. Furthermore $L_{\mathcal{Y}}^{\alpha\beta_{\mathcal{M}}} \subseteq L_{\mathcal{Y}}^{\alpha\beta}$ and $\alpha\beta\alpha\beta_{\mathcal{M}} = \alpha\beta_{\mathcal{M}}$.

2.3 Clones on $\mathbb{R}^n$ and Polyhedra

In what follows, our signature set, or base set A on which we are defining finitary maps $f \in O_A$ and relations $\varphi \in \mathbf{R}_A$, will be $\mathbb{L}^n = (\mathbb{R}^n, \leq)$ (recall that this is the space $\mathbb{R}^n$ with the component-wise partial order). Thus, for simplicity, we will drop the subscript A, using the notation $O := O_A$ and $\mathbf{R} := \mathbf{R}_A$. Recall that a polyhedron $\mathcal{P}_{A,\mathbf{b}}$ in $\mathbb{R}^n$ is defined with respect to an $m \times n$ matrix A and vector $\mathbf{b} \in \mathbb{R}^m$:

$$\mathcal{P}_{A,\mathbf{b}} := \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}. \tag{2.10}$$

Consequently every polyhedron is a relation on $\mathbb{L}^n$. That is, $\mathcal{P}_{A,\mathbf{b}} \in \mathbf{R}$, and hence we may consider clones on $\mathbb{L}^n$ of the form Pol $\mathcal{P}_{A,\mathbf{b}}$. This is the collection of all finitary maps on $\mathbb{L}^n$ that preserve $\mathcal{P}_{A,\mathbf{b}}$. We see now why clones provide the necessary key to characterizing the extent to which a given polyhedron preserves the lattice structure on $\mathbb{L}^n$.

Consider the following example. Let $\mathbf{d} = (1,1)^t \in \mathbb{R}^2$. Then the closed half-space $H_{\mathbf{d},0}$ is the collection of all $\mathbf{x} \in \mathbb{R}^2$ satisfying $\mathbf{x}_1 + \mathbf{x}_2 \leq 0$. Observe that for any pair $\mathbf{a}, \mathbf{b} \in H_{\mathbf{d},0}$, their component-wise meet $\mathbf{a} \wedge \mathbf{b}$ is also in $H_{\mathbf{d},0}$ (See Fig. 2.1).

Now let $f : \mathbb{R}^2 \to \mathbb{R}$ be the min function on $\mathbb{R}$; that is, $f(a,b) = \min(a,b)$. Note that f is a binary map on $\mathbb{R}$, and $H_{\mathbf{d},0}$ defines a relation on $\mathbb{R}^2$. We may extend f as in (2.4) to a map from $(\mathbb{R}^2)^2$ to $\mathbb{R}^2$. Let $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2)^t$ and $\mathbf{b} = (\mathbf{b}_1, \mathbf{b}_2)^t$ be two elements of $H_{\mathbf{d},0}$. Then we may consider:

$$f\begin{pmatrix} \mathbf{a}_1 & \mathbf{b}_1 \\ \mathbf{a}_2 & \mathbf{b}_2 \end{pmatrix} = \begin{pmatrix} f(\mathbf{a}_1, \mathbf{b}_1) \\ f(\mathbf{a}_2, \mathbf{b}_2) \end{pmatrix} = \begin{pmatrix} \min(\mathbf{a}_1, \mathbf{b}_1) \\ \min(\mathbf{a}_2, \mathbf{b}_2) \end{pmatrix} = \mathbf{a} \wedge \mathbf{b}. \tag{2.11}$$

As we can see, the binary min function on pairs of elements in $\mathbb{R}$ acts as the component-wise meet operator over relations in $\mathbb{R}^2$. As we have already discussed, $H_{\mathbf{d},0}$ is preserved by component-wise meet, and hence $f \in$ Pol $H_{\mathbf{d},0}$. Moreover, every relation in Inv Pol $H_{\mathbf{d},0}$ is necessarily preserved by component-wise meet. In other words, a necessary condition for a polyhedron $\mathcal{P}_{A,\mathbf{b}}$ to be in Inv Pol $H_{\mathbf{d},0}$ is that it be preserved by component-wise meet.

We see, then, that the Inv-Pol connection between finitary functions and relations on $\mathbb{R}$ is a powerful tool that allows us to both characterize the collection of functions that preserve a polyhedron, and classify polyhedra according to the functions that preserve them.

The fact that functions on $\mathbb{R}$ operate component-wise on relations in $\mathbb{R}^n$ immediately gives us the following useful result.

Lemma 2.3.1. *Let $f : \mathbb{R}^k \to \mathbb{R}$ be any k-ary function on $\mathbb{R}$. If f preserves relation R_m in $\mathbb{R}^m$ and relation R_n in $\mathbb{R}^n$, then f preserves the relation $R_m \times R_n$ in $\mathbb{R}^{m+n}$.*

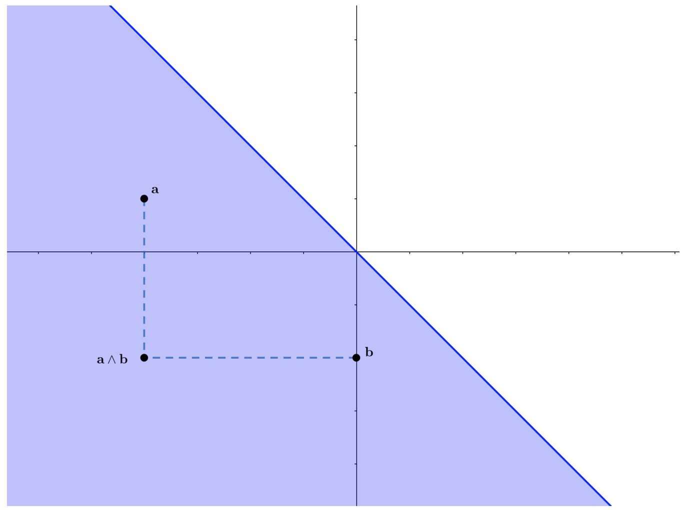

Figure 2.1: The closed half-space $\mathbf{x}_1 + \mathbf{x}_2 \leq 0$. The map that defines the component-wise meet operation preserves this half-space.

Let us look at another example. Consider the polytope given in Figure 2.2. It is easily seen to be preserved by neither component-wise meet nor component-wise join. However, as we will see, it is preserved by the ternary "median" function

$$m(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = (\mathbf{x}_1 \wedge \mathbf{x}_2) \vee (\mathbf{x}_1 \wedge \mathbf{x}_3) \vee (\mathbf{x}_2 \wedge \mathbf{x}_3), \tag{2.12}$$

where the meet and join are the component-wise meet and join operators. In fact, we may completely characterize all polyhedra that are preserved by the median function m.

In order to do this, we will first state the following result, a much stronger version of which will be stated and proved in Chapter 4 (Corollary 4.1.5) and so we will omit the proof here.

Lemma 2.3.2. *Suppose a polyhedron $\mathcal{P}$ is preserved by the median function m. Then every affine shift $\mathbf{k} + \mathcal{P}$ is also preserved by m.*

With the above Lemma, we may now show the following.

Lemma 2.3.3. *Suppose $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^2$. Then $m(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \operatorname{conv}(\mathbf{x}, \mathbf{y}, \mathbf{z})$.*

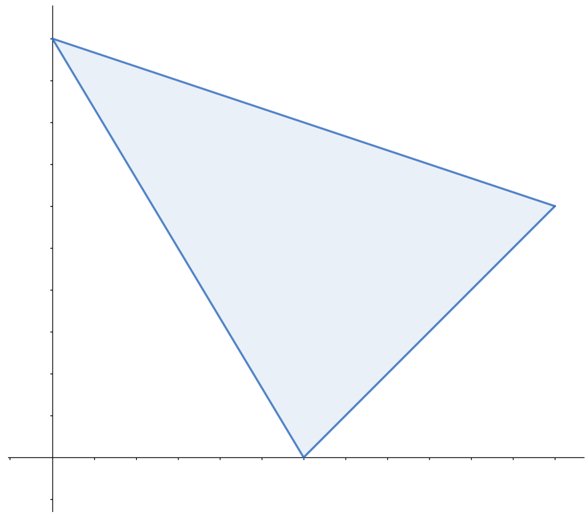

Figure 2.2: A polytope in $\mathbb{R}^2$ that is preserved by neither component-wise meet nor component-wise join.

Proof. By Lemma 2.3.2, we may assume without loss of generality that $\mathbf{x} = (0,0)^t$. Let us suppose, without any loss of generality, that $x_1 \leq y_1 \leq z_1$. Then we may write $y_1 = (1-k_1)x_1 + k_1 z_1$ for some $k_1 \in [0,1]$. Now, if $x_2 \leq y_2 \leq z_2$ or $z_2 \leq y_2 \leq x_2$, then $m_3(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathbf{y}$ and hence is in $\mathrm{conv}(\mathbf{x},\mathbf{y},\mathbf{z})$. Supposing this is not the case, then, we have four possible remaining orderings on x_2, y_2, and z_2. We will explicitly work through one of the cases here, but the remaining three cases follow an almost identical argument. Let us suppose that $x_2 \leq z_2 < y_2$, so that $m(\mathbf{x},\mathbf{y},\mathbf{z}) = (y_1, z_2)^t$. Note that, if $z_1 = 0$, then $x_1 = y_1 = z_1$ and hence, $m(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathbf{z} \in \mathrm{conv}(\mathbf{x},\mathbf{y},\mathbf{z})$, so let us assume that $z_1 > 0$. Note also that we may write $z_2 = (1-k_2)x_2 + k_2 y_2$ for some $k_2 \in [0,1)$. We want to show that there exist $a_1, a_2, a_3 \geq 0$, $\sum_i a_i = 1$, such that $a_1\mathbf{x} + a_2\mathbf{y} + a_3\mathbf{z} = m(\mathbf{x},\mathbf{y},\mathbf{z})$. Thus, we need to find such an a_1, a_2, a_3 that satisfy

$$a_1 x_1 + a_2 y_1 + a_3 z_1 = y_1 \tag{2.13}$$

$$a_1 x_2 + a_2 y_2 + a_3 z_2 = z_2. \tag{2.14}$$

Since by assumption we have $x_1 = x_2 = 0$ the above simplifies to

$$a_3 z_1 = (1 - a_2) y_1 \tag{2.15}$$

$$a_2 y_2 = (1 - a_3) z_2. \tag{2.16}$$

Recalling that $y_1 = (1 - k_1)x_1 + k_1 z_1$ for some $k_1 \in [0, 1]$ and $z_2 = (1 - k_2)x_2 + k_2 y_2$ for some $k_2 \in [0, 1)$, we may reduce the above further:

$$a_3 z_1 = (1 - a_2) k_1 z_1 \tag{2.17}$$

$$a_2 y_2 = (1 - a_3) k_2 y_2. \tag{2.18}$$

Since $z_1, y_2 \neq 0$ by assumption, we reduce the above further:

$$a_3 = (1 - a_2) k_1 \tag{2.19}$$

$$a_2 = (1 - a_3) k_2. \tag{2.20}$$

Solving, we find that $a_3 = 1 - \frac{1-k_1}{1-k_1 k_2}$ and $a_2 = \frac{(1-k_1)k_2}{1-k_1 k_2}$. It is easily seen that for all $k_1 \in [0, 1]$ and $k_2 \in [0, 1)$, we have $a_2, a_3 \geq 0$ and $a_2 + a_3 \leq 1$. Letting $a_1 = 1 - (a_2 + a_3)$, we obtain the desired result. The remaining cases to consider are $(z_2 \leq x_2 < y_2)$, $(y_2 < x_2 \leq z_2)$, and $(y_2 < z_2 \leq x_2)$. Following the same approach as above for these cases, we may conclude that, in all cases, $m(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathrm{conv}(\mathbf{x}, \mathbf{y}, \mathbf{z})$, as we wanted to show. $\square$

We immediately have:

Corollary 2.3.4. *Every convex polyhedron in $\mathbb{R}^2$ is preserved by m.*

Leveraging Lemma 2.3.1, we may now begin characterizing polyhedra in $\mathbb{R}^n$ for $n > 2$ that are preserved by m.

Proposition 2.3.5. *Let $\mathcal{P}$ be a convex polyhedron in $\mathbb{R}^2$ and let $\mathcal{Q} = \mathcal{P} \times \mathbb{R}^{n-2}$ in $\mathbb{R}^n$. Then $\mathcal{Q}$ is preserved by m.*

Proof. This result follows immediately from Corollary 2.3.4 and Lemma 2.3.1. $\square$

We may now prove the following:

Theorem 2.3.6. *Let m be the median function from Eq.* (2.12)*, and let $\mathcal{P}$ be any polyhedron. Then $\mathcal{P} \in \operatorname{Inv} m$ if and only if it is the intersection of the inverse images of its projections onto the two-dimensional coordinate planes.*

Proof. First let us suppose that a convex polyhedron $\mathcal{P}$ is the intersection of the inverse images of its projection to the two-dimensional coordinate planes. Then $\mathcal{P}$ is a finite intersection of polyhedra of the form given in Proposition 2.3.5, and hence is closed under m. Conversely, suppose $\mathcal{P}$ is preserved by m. We want to show that $\mathcal{P}$ is the intersection of the inverse images of its projection to the two-dimensional coordinate planes. By Corollary 2.3.4, the result holds trivially when $\mathcal{P}$ resides in $\mathbb{R}^2$. We prove the more general case when $\mathcal{P}$ resides in $\mathbb{R}^n$ for any $n \geq 2$ inductively. That is, suppose the inductive hypothesis holds for $n = k$ for some $k > 2$. Now consider the case of $n = k+1$. Let $\mathcal{D}$ denote the intersection of the inverse images of the projection of $\mathcal{P}$ to the two-dimensional coordinate planes. We want to show that $\mathcal{P} = \mathcal{D}$. Clearly, $\mathcal{P} \subseteq \mathcal{D}$. Let $\mathbf{x} \in \mathcal{D}$. Let $\Pi_{\hat{i}}$ denote the coordinate projector that excludes the i-th coordinate (not to be confused with the projectors from clone theory). Observe that $\Pi_{\hat{i}}[\mathcal{P}]$ is a convex polyhedron in $\mathbb{R}^k$, and since $\mathcal{P}$ is preserved by m, so is $\Pi_{\hat{i}}[\mathcal{P}]$. By the inductive hypothesis, it follows that $\Pi_{\hat{i}}(\mathbf{x}) \in \Pi_{\hat{i}}[\mathcal{P}]$. Then there exists some vector $\mathbf{p}_i \in \mathcal{P}$ such that for all $l \in [n] \backslash \{i\}$, $\mathbf{p}_{i_l} = \mathbf{x}_l$. We may identify vectors $\mathbf{p}_j$ and $\mathbf{p}_k$ in $\mathcal{P}$ defined similarly, where i, j, k are all distinct. But $m(\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k) = \mathbf{x}$. Since $\mathcal{P}$ is preserved by m by assumption, it follows that $\mathbf{x} \in \mathcal{P}$, and hence $\mathcal{D} \subseteq \mathcal{P}$. Thus, $\mathcal{P} = \mathcal{D}$, concluding the proof. $\square$

As we can see, the median function is a ternary function that is constructed from compositions of component-wise meet and join operators. Moreover, there are many polyhedra that are preserved by this function, while being preserved by neither component-wise meet nor join. Consequently, it is possible for polyhedra to preserve some of the structure from the component-wise lattice $\mathbb{L}^n = (\mathbb{R}^n, \leq)$ without necessarily being a sublattice of $\mathbb{L}^n$. In the next chapter, we will make precise what is meant by "lattice structure." In particular, we will introduce the concept of *lattice polynomials*, and consider the collection of convex clones generated by lattice polynomials.

Chapter 3: Convex Clones of Lattice Polynomials on a Free Distributive Lattice

Given a lattice L and variables $x_1, x_2, \ldots, x_n \in L$, recall that a *lattice polynomial* in the variables $x_1, x_2, \ldots, x_n$ is defined as follows:

(1) $x_1, x_2, \ldots, x_n$ are lattice polynomials,

(2) if p and q are lattice polynomials in $x_1, x_2, \ldots, x_n$, then so are $p \vee q$ and $p \wedge q$, and

(3) every lattice polynomial is formed by finitely many applications of (1) and (2).

Now let $FD(n)$ be the free distributive lattice on n generators $g_1, g_2, \ldots, g_n$. This is the lattice of equivalence classes of lattice polynomials over n generators, where two lattice polynomials are equivalent if they define the same function $L^n \to L$ on every distributive lattice L. Hence, for any distributive lattice L, it follows that the set $FD(n)$ forms a clone over L, which we will denote by $\mathcal{LP}$.

Recall that an upper semi-ideal, or up-set, of a lattice L is any subset $S \subseteq L$ such that if $p \in S$ and $q \geq p$, then $q \in S$. The lower semi-ideals, or down-sets, are defined dually. A *composition ideal* on $FD(n)$ is a semi-ideal that is also a clone. More precisely, an upper composition ideal is a set $J \subseteq FD(n)$ which

(1) contains the generators g_i,

(2) is an upper semi-ideal, and

(3) is closed under functional composition.

Let $S \subseteq FD(n)$, and let $\langle S \rangle^{\uparrow}$ denote the smallest (upper) composition ideal containing S. Observe that $\langle \emptyset \rangle^{\uparrow} = \{p \in FD(n) \mid \exists\ i \in [n] \text{ such that } p \geq g_i\}$ is the smallest upper composition ideal. For $k = 1, \ldots, n-1$ define

$$h_k := \bigvee_{1 \leq i < j \leq k+1} (g_i \wedge g_j). \tag{3.1}$$

For $k = 1, \ldots, n-1$, let $J_k^{\uparrow} = \langle h_k \rangle^{\uparrow}$, and define $J_n^{\uparrow} = \langle \emptyset \rangle^{\uparrow}$. In [10], it was shown that every upper composition ideal of $FD(n)$ is of the form $J_k^{\uparrow}$ for $k = 1, \ldots, n$.

It follows from duality that the lower composition ideals are given by $J_k^{\downarrow}$, where

$$J_n^{\downarrow} = \langle \emptyset \rangle^{\downarrow} := \{p \in FD(n) \mid \exists\, i \in [n] \text{ such that } p \leq g_i\}, \tag{3.2}$$

and for $k = 1, \ldots, n-1$, $J_k = \langle f_k \rangle^{\downarrow}$, where

$$f_k := \bigwedge_{1 \leq i < j \leq k+1} (g_i \vee g_j). \tag{3.3}$$

Observe that for $k = 1$, h_1 is the meet operator and f_1 is the join operator, and hence $h_1 < f_1$. Then the upper composition ideal $J_1^{\uparrow}$ includes all lattice polynomials larger than meet, as well as their closure under composition. We see then, that $J_1^{\uparrow}$ is the clone of all lattice polynomials on $FD(n)$. Likewise, $J_1^{\downarrow}$ is also the clone of lattice polynomials on $FD(n)$, so they define the same convex clone. Observe also that $J_n^{\downarrow} \cap J_n^{\uparrow}$ is simply the set of generators, or equivalently, the projectors onto the generators.

In the case $k = 2$, by distributing join over meet we see that

$$h_2 = (g_1 \wedge g_2) \vee (g_1 \wedge g_3) \vee (g_2 \wedge g_3) \tag{3.4}$$

$$= (g_1 \vee g_2) \wedge (g_1 \vee g_3) \wedge (g_2 \vee g_3) \tag{3.5}$$

$$= f_2. \tag{3.6}$$

The lattice polynomial $h_2 = f_2$ is commonly referred to as the *median*. Now, we observe that $h_{k-1} \leq h_k$ for all $1 < k \leq n-1$. This is because we may always express h_k as

$$h_k = h_{k-1} \bigvee_{i\in[k]} (g_i \wedge g_{k+1}). \tag{3.7}$$

By duality, it follows that $f_k \leq f_{k-1}$ for all $1 < k \leq n-1$. Thus, we have shown:

Lemma 3.0.1. *The lattice polynomials h_k and f_k defined in* (3.1) *and* (3.3)*, respectively, satisfy the order relation*

$$f_{n-1} \leq f_{n-2} \leq \cdots \leq f_2 = h_2 \leq \cdots \leq h_{n-2} \leq h_{n-1}. \tag{3.8}$$

We also have the following result that applies to all lattice polynomials.

Lemma 3.0.2. *Lattice polynomials in $FD(n)$ are monotone under composition.*

Proof. Let $p, q \in FD(n)$ and suppose $p \leq q$. Then for any $f_1, f_2, \ldots, f_k \in FD(n)$, it follows that $p(f_1, f_2, \ldots, f_k) \leq q(f_1, f_2, \ldots, f_k)$. Furthermore, if $f_i \leq r_i$ for all $i \in [k]$, then $q(f_1, f_2, \ldots, f_k) \leq q(r_1, r_2, \ldots, r_k)$. □

Definition 3.0.3. *Let L be any poset and let $S \subseteq L$ be any subset of L. We say that S is* convex *in L if for all $s_1, s_2 \in S$ with $s_1 \leq s_2$, the set $[s_1, s_2]_L$ is also in S, where*

$$[s_1, s_2]_L := \{f \in L \mid s_1 \leq f \leq s_2\}. \tag{3.9}$$

A clone of lattice polynomials in $FD(n)$ will be called *convex* if it is convex as a subset of $FD(n)$.

Proposition 3.0.4. *Let C be any convex clone of lattice polynomials in $FD(n)$. Then $C = \langle C\rangle^{\uparrow} \cap \langle C\rangle^{\downarrow}$.*

Proof. Clearly $C \subseteq \langle C\rangle^{\uparrow} \cap \langle C\rangle^{\downarrow}$. Let $f \in (\langle C\rangle^{\uparrow} \cap \langle C\rangle^{\downarrow})$. Lemma 3.0.2 implies that every element in $\langle C\rangle^{\uparrow}\backslash C$ is greater than or equal to some element in C, and every element in

$\langle C\rangle^{\downarrow}\backslash C$ is less than or equal to some element in C. Consequently, either $f \in C$ or there exists a $c_1, c_2 \in C$ such that $c_1 \leq f \leq c_2$. But since C is convex in $FD(n)$ by assumption, the latter case then implies that $f \in C$, and hence $\langle C\rangle^{\uparrow} \cap \langle C\rangle^{\downarrow} \subseteq C$. Thus $C = \langle C\rangle^{\uparrow} \cap \langle C\rangle^{\downarrow}$, as we wanted to show. $\square$

The next result follows trivially from Proposition 3.0.4 and the containment property $J_n^{\uparrow} \subseteq \cdots \subseteq J_1^{\uparrow}$ and $J_n^{\downarrow} \subseteq \cdots \subseteq J_1^{\downarrow}$.

Corollary 3.0.5. *All convex clones of lattice polynomials in $FD(n)$ can be expressed as $J_i^{\uparrow} \cap J_k^{\downarrow}$ for $i, k \in [n]$.*

We note trivially that, since $J_1^{\downarrow} = J_1^{\uparrow}$, it follows from the containment property that $J_i^{\downarrow} = J_i^{\downarrow} \cap J_1^{\uparrow}$, and $J_k^{\uparrow} = J_1^{\downarrow} \cap J_k^{\uparrow}$ for all $i, k \in [n]$.

In Figure 3.1 we draw the entire lattice of convex clones of the lattice polynomials on the free distributive lattice $FD(n)$. Let C be any collection of lattice polynomials. Then we use the notation $\langle C\rangle_c$ to denote the convex lattice polynomial clone generated by C. We have the following result:

Corollary 3.0.6. *Let p be a lattice polynomial that is minimally in $J_i^{\downarrow} \cap J_k^{\uparrow}$ (that is, $p \in J_i^{\downarrow} \cap J_k^{\uparrow}$, and there is no convex lattice polynomial clone $C \subset J_i^{\downarrow} \cap J_k^{\uparrow}$ containing p). Then $J_i^{\downarrow} \cap J_k^{\uparrow} = \langle\{p\}\rangle_c$.*

Proof. On the one hand, clearly $\langle\{p\}\rangle_c \subseteq J_i^{\downarrow} \cap J_k^{\uparrow}$. On the other hand, since all convex lattice polynomial clones are of the form $J_i^{\downarrow} \cap J_k^{\uparrow}$ by Corollary 3.0.5, it follows that $\langle\{p\}\rangle_c = J_r^{\downarrow} \cap J_s^{\uparrow}$ for some $r \leq i$ and $s \leq k$. By our minimality assumption, the result follows. $\square$

While little is known about the full clone lattice over nearly any signature set with more than two elements, what we have shown is that when the signature set is a distributive lattice, then it is possible to completely characterize all *convex* subclones of the clone of lattice polynomials over this set. In the next chapter, we investigate combinatorial and geometric methods of characterizing these clones within the framework of lattice polynomials in $\mathbb{R}^n$.

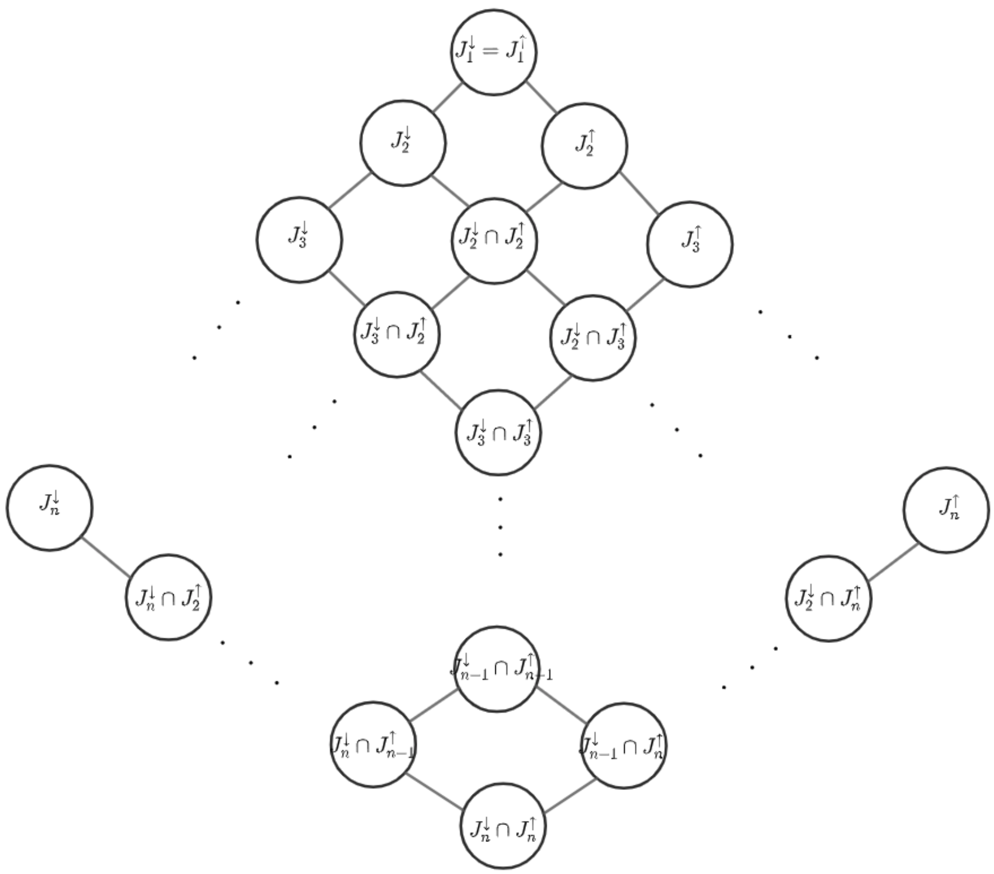

Figure 3.1: The lattice of convex clones of lattice polynomials on the free distributive lattice $FD(n)$.

Chapter 4: Convex Lattice Polynomial Clones on $\mathbb{R}^n$

In the previous chapter, we showed that all convex lattice polynomial clones on a distributive lattice are the intersection of finitely generated composition ideals. In this chapter, we will alternatively characterize each convex lattice polynomial clone C over $\mathbb{R}$ by instead characterizing the set of all polyhedra in Inv C. By a lattice polynomial over $\mathbb{R}$, we are specifically referring to finitary maps $p : \mathbb{R} \times \mathbb{R} \times \cdots \times \mathbb{R} \to \mathbb{R}$ that can be expressed as a finite composition of pairwise min and max operations. Note that, given a relation $S \subseteq \mathbb{R}^n$, then a k-ary lattice polynomial on $\mathbb{R}$ is extended to a k-ary lattice polynomial on $\mathbb{R}^n$ as in Eq. (2.4). In so doing, f is a lattice polynomial over $\mathbb{R}^n$ with the component-wise partial order. Let $\mathcal{LP}$ be the set of all such lattice polynomials over $\mathbb{R}$. Then a set P of polyhedra is said to *determine the clone* $C \subseteq \mathcal{LP}$ if $\mathrm{Pol}_{\mathcal{LP}}\ P = C$, where $\mathrm{Pol}_{\mathcal{LP}}$ is the restriction of the set of all finitary maps on $\mathbb{R}$ to the set $\mathcal{LP}$ of lattice polynomials.

Recall from Chapter 2 that any polyhedron $\mathcal{P}$ in $\mathbb{R}^n$ is a relation on $\mathbb{L}^n$. In keeping with the same terminology from Chapter 2, we will say that a k-ary map $f : (\mathbb{L}^n)^k \to \mathbb{L}^n$ *preserves* $\mathcal{P}$, or that $\mathcal{P}$ is *preserved by* f, if for any $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k \in \mathcal{P}$, $f(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k) \in \mathcal{P}$.

4.1 Problem Space Reduction

In order to find representative polyhedra, we first provide several results on the properties of lattice polynomials that we will leverage to reduce our problem space significantly. We provide the following definition.

Definition 4.1.1. *Let A and B be two convex sets in $\mathbb{R}^n$. We say that $A \cap B$ is a* nontrivial *intersection if* $\dim(A \cap B) = \min(\dim(A),\ \dim(B))$.

Proposition 4.1.2. *Suppose $f : (\mathbb{R}^n)^k \to \mathbb{R}^n$ satisfies*

$$f[h(\mathbf{x}_1), h(\mathbf{x}_2), \ldots, h(\mathbf{x}_k)] = h[f(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k)], \tag{4.1}$$

for every positive homothety h. Let $\mathcal{P}$ be any polyhedron in $\mathbb{R}^n$. Then $\mathcal{P}$ is preserved by f if and only if each bounding half-space of $\mathcal{P}$ is preserved by f.

Proof. The first direction is trivial. If each half-space is preserved by f, then so is their finite intersection. Conversely, let $H_{\mathbf{a},b} = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{a}, \mathbf{x} \rangle \leq b\}$ be any half-space in $\mathbb{R}^n$. We will show that if there is any set of vectors $\mathbf{x}_1, \ldots, \mathbf{x}_k \in H_{\mathbf{a},b}$ such that $f(\mathbf{x}_1, \ldots, \mathbf{x}_k) \notin H_{\mathbf{a},b}$ and B_ϵ is any ϵ-ball having nontrivial intersection with $H_{\mathbf{a},b}$ at its boundary, then $B_\epsilon \cap H_{\mathbf{a},b}$ contains such a set.

Let $\mathbf{x}_1, \ldots, \mathbf{x}_k$ be defined as above. Let $\mathbf{q} \in H_{\mathbf{a},b}$ be on the boundary of the half-space, and define the half-space $S := \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{a}, \mathbf{x} \rangle \leq 0\}$, so that $H_{\mathbf{a},b} = \mathbf{q} + S$. Then we may expand each $\mathbf{x}_i = \mathbf{q} + \mathbf{s}_i$ for some $\mathbf{s}_i \in S$. By our assumption on f, it follows that $f(\mathbf{x}_1, \ldots, \mathbf{x}_k) = \mathbf{q} + f(\mathbf{s}_1, \ldots, \mathbf{s}_k)$, and hence $f(\mathbf{x}_1, \ldots, \mathbf{x}_k) \notin H_{\mathbf{a},b} \Leftrightarrow \langle \mathbf{a}, f(\mathbf{s}_1, \ldots, \mathbf{s}_k) \rangle > 0$.

Now, choose an appropriate $\mathbf{q}^*$ on the boundary of $H_{\mathbf{a},b}$ and $\delta > 0$ such that each $\mathbf{q}^* + \delta \mathbf{s}_i \in B_\epsilon \cap H_{\mathbf{a},b}$. Then $f[(\mathbf{q}^* + \delta \mathbf{s}_1), \ldots, (\mathbf{q}^* + \delta \mathbf{s}_k)] = \mathbf{q}^* + \delta f(\mathbf{s}_1, \ldots, \mathbf{s}_k)$, and hence $\langle \mathbf{a}, f[(\mathbf{q}^* + \delta \mathbf{s}_1), \ldots, (\mathbf{q}^* + \delta \mathbf{s}_k)] \rangle = b + \delta \langle \mathbf{a}, f(\mathbf{s}_1, \ldots, \mathbf{s}_k) \rangle > b$. Hence $f[(\mathbf{q}^* + \delta \mathbf{s}_1), \ldots, (\mathbf{q}^* + \delta \mathbf{s}_k)] \notin H_{\mathbf{a},b}$.

It follows that if $H_{\mathbf{a},b}$ is not preserved by f, then the nontrivial intersection of $H_{\mathbf{a},b}$ with any other half-space necessarily contains a collection of vectors that violate the preservation. It follows from contraposition that if a polyhedron is preserved by f, then so are each of the half-spaces defining it. □

We now show that lattice polynomials all share the property of the function f given in the above proposition.

Lemma 4.1.3. *Let f be any t-ary lattice polynomial and h any positive homothety. Then*

$$f[h(\mathbf{x}_1), h(\mathbf{x}_2), \ldots, h(\mathbf{x}_t)] = h[f(\mathbf{x}_1, \ldots, \mathbf{x}_t)]. \tag{4.2}$$

Proof. Observe that for all $\mathbf{k}, \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\alpha > 0$,

$$\mathbf{k} + (\mathbf{x} \wedge \mathbf{y}) = \{k_i + \min(x_i, y_i)\}_i \tag{4.3}$$

$$= \{\min((k_i + x_i), (k_i + y_i))\}_i \tag{4.4}$$

$$= (\mathbf{k} + \mathbf{x}) \wedge (\mathbf{k} + \mathbf{y}). \tag{4.5}$$

Likewise,

$$\mathbf{k} + (\mathbf{x} \vee \mathbf{y}) = (\mathbf{k} + \mathbf{x}) \vee (\mathbf{k} + \mathbf{y}). \tag{4.6}$$

In a similar manner, we observe that

$$\alpha(\mathbf{x} \wedge \mathbf{y}) = \{\alpha \min(x_i, y_i)\}_i \tag{4.7}$$

$$= \{\min(\alpha x_i, \alpha y_i)\}_i \tag{4.8}$$

$$= (\alpha \mathbf{x}) \wedge (\alpha \mathbf{y}). \tag{4.9}$$

Likewise,

$$\alpha(\mathbf{x} \vee \mathbf{y}) = (\alpha \mathbf{x}) \vee (\alpha \mathbf{y}). \tag{4.10}$$

Since any lattice polynomial is a finite composition of meets and joins, the above properties extend to any lattice polynomial. □

It follows from Proposition 4.1.2 and Lemma 4.1.3 that we may infer the lattice polynomial polymorphisms of an arbitrary polyhedron, that is, the lattice polynomials that preserve the polyheldron, by studying instead the lattice polynomial polymorphisms of its bounding half-spaces. In fact, because lattice polynomials act independently on each coordinate, we may generalize Lemma 4.1.3 further in a way that permits a further reduction to considering only lattice polynomial polymorphisms of half-spaces of the form $S_{\mathbf{a}} := \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{a}, \mathbf{x} \rangle \leq 0\}$, where $\mathbf{a} \in \{-1, 0, 1\}^n$. We provide the following generalization of scalar multiplication.

Definition 4.1.4. *Let* $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. *The* Schur product *of* $\mathbf{x}$ *and* $\mathbf{y}$ *is given by* $\mathbf{xy} := \{\mathbf{x}_i \mathbf{y}_i\}_{i=1}^n \in \mathbb{R}^n$. *That is to say, the Schur product of two vectors is the component-wise product of the two.*

Following a nearly identical proof to that of Lemma 4.1.3, we have the following corollary.

Corollary 4.1.5. *Let f be any t-ary lattice polynomial. Then*

$$f[(\mathbf{k} + \beta \mathbf{x}_1), \ldots, (\mathbf{k} + \beta \mathbf{x}_t)] = \mathbf{k} + \beta f(\mathbf{x}_1, \ldots, \mathbf{x}_t), \tag{4.11}$$

for all $\mathbf{k}, \mathbf{x}_1, \ldots, \mathbf{x}_t \in \mathbb{R}^n$ *and* $\beta \in (\mathbb{R}^n)^+$.

Let $\beta \in (\mathbb{R}^n)^+$ (that is, β has all strictly positive entries). Let $\beta^{-1} := \{\beta_i^{-1}\}_{i=1}^n$. Observe that, for any $\mathbf{a}, \mathbf{x} \in \mathbb{R}^n$,

$$\langle \mathbf{a}, \mathbf{x} \rangle = \langle \mathbf{a}\beta, \beta^{-1}\mathbf{x} \rangle. \tag{4.12}$$

Now, given any $\mathbf{a} \in \mathbb{R}^n$, define the vector $\mathbf{q}_a \in (\mathbb{R}^n)^+$ as

$$\mathbf{q}_{ai} := \begin{cases} |\mathbf{a}_i|^{-1}, & \text{if } \mathbf{a}_i \neq 0, \\ 1, & \text{otherwise.} \end{cases} \tag{4.13}$$

It then follows from Corollary 4.1.5 that

$$\langle \mathbf{a}, f(\mathbf{x}_1, \ldots, \mathbf{x}_t) \rangle \leq 0 \tag{4.14}$$

$$\Leftrightarrow \langle \mathbf{a}\mathbf{q}_a, \mathbf{q}_a^{-1} f(\mathbf{x}_1, \ldots, \mathbf{x}_t) \rangle \leq 0 \tag{4.15}$$

$$\Leftrightarrow \langle \mathbf{a}\mathbf{q}_a, f(\mathbf{q}_a^{-1}\mathbf{x}_1, \ldots, \mathbf{q}_a^{-1}\mathbf{x}_t) \rangle \leq 0, \tag{4.16}$$

for all $\mathbf{x}_1, \ldots, \mathbf{x}_t \in \mathbb{R}^n$ and any t-ary lattice polynomial f.

Observe that $\mathbf{aq}_a \in \{-1, 0, 1\}^n$. It follows that the lattice polynomial polymorphisms of a half-space $S_{\mathbf{a}}$ are determined by the sign vector associated with $\mathbf{a}$. Consequently, we may reduce ourselves further to considering only lattice polynomial polymorphisms of half-spaces of the form $S_{\mathbf{a}} = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{a}, \mathbf{x} \rangle \leq 0\}$, where $\mathbf{a} \in \{-1, 0, 1\}^n$.

4.2 Some Representatives

Let us now consider several half-spaces of the form $S_{\mathbf{a}} = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{a}, \mathbf{x} \rangle \leq 0\}$, where $\mathbf{a} \in \{-1, 0, 1\}^n$. We note that, since there are n^2 distinct convex clones in $FD(n)$, but 3^n sign vectors (that is, vectors in $\{-1, 0, 1\}^n$), not every sign vector corresponds to a unique convex clone.

Proposition 4.2.1. *The half-space $S_{-\mathbf{1}}$ in $\mathbb{R}^n$ given by $\mathbf{x} = (x_1, x_2, \ldots, x_n)^t \in \mathbb{R}^n$ such that $x_1 + x_2 + \cdots + x_n \geq 0$ is preserved by h_k if and only if $k \geq n$.*

Proof. Recall that

$$h_k := \bigvee_{1 \leq i < j \leq k+1} (\mathbf{x}_i \wedge \mathbf{x}_j) \tag{4.17}$$

is a polynomial defined on $k+1$ variables $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1} \in \mathbb{R}^n$. If $h_k(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1}) = \mathbf{y}$, then each y_j is given by the second largest value from the (possibly multi-) set $\{x_{1,j}, x_{2,j}, \ldots, x_{k+1,j}\}$, where the notation $x_{i,j}$ refers to the j-th entry in $\mathbf{x}_i$.

Let us first suppose that $k \geq n$. We now show that there necessarily exists an index $i \in [k+1]$ such that $\mathbf{x}_i \leq \mathbf{y}$. To see this, first note that for each $j \in [n]$, there are n vectors $\mathbf{x}_{i_1}, \mathbf{x}_{i_2}, \ldots, \mathbf{x}_{i_n}$ such that $x_{i_1,j},\ x_{i_2,j}, \ldots, x_{i_n,j} \leq y_j$. Let us suppose for contradiction that there exists no index $i \in [k+1]$ such that $\mathbf{x}_i \leq \mathbf{y}$. It follows that for each vector $\mathbf{x}_i$ there exists at least one j such that the entry $x_{i,j} > y_j$. But because each y_j is the second largest element in the set $\{x_{1,j}, x_{2,j}, \ldots, x_{k+1,j}\}$, it follows that for each j there exists exactly one vector $\mathbf{x}_i$ such that $x_{i,j} > y_j$. However, there are $k+1 > n$ vectors but j only runs from 1 through n, so this cannot occur.

Hence there necessarily exists a vector $\mathbf{x}_i$ from the set $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1}\}$ such that $\mathbf{x}_i \leq \mathbf{y}$. Now, since $\sum_{j=1}^{n} x_{i,j} \geq 0$ and each $y_j \geq x_{i,j}$, it follows that $\sum_{j=1}^{n} y_j \geq 0$ as well, so that $S_{-\mathbf{1}}$ is preserved by h_k whenever $k \geq n$.

Conversely, for any positive integer $k < n$, it is always possible to find a collection of $k+1$ vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1}$ in $\mathbb{R}^n$ such that there exists no index $i \in [k+1]$ satisfying $\mathbf{x}_i \leq h_k(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1})$. This is done by assuring that for each vector $\mathbf{x}_i$ there exists at least one j such that the entry $x_{i,j} > y_j$, where $\mathbf{y} = h_k(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{k+1})$. Consequently, $S_{-\mathbf{1}}$ is not closed over h_k for all $k < n$, concluding the proof. $\square$

The following is a straightforward generalization of the above proposition.

Corollary 4.2.2. *Let $\mathbf{a} \in \{-1, 0, 1\}^n$ contain no positive entries and m strictly negative entries. Then $S_{\mathbf{a}}$ is preserved by h_k if and only if $k \geq m$.*

By duality, we also have the following result.

Corollary 4.2.3. *Let $\mathbf{a} \in \{-1, 0, 1\}^n$ contain no negative entries and m strictly positive entries. Then $S_{\mathbf{a}}$ is preserved by f_k if and only if $k \geq m$.*

Corollaries 4.2.2 and 4.2.3 only address some of the polymorphisms of the half-space $S_{\mathbf{a}}$ when $\mathbf{a}$ contains only non-positive or non-negative entries. The following result shows us that, whenever $k \geq 2$, these results completely determine the invariance of $S_{\mathbf{a}}$ under h_k and f_k for all $\mathbf{a} \in \{-1, 0, 1\}^n$.

Proposition 4.2.4. *Suppose $\mathbf{a}$ has at least one positive entry and at least two negative entries. Then $S_{\mathbf{a}}$ is not preserved by h_k for all $k \geq 2$. Equivalently, if $\mathbf{a}$ has at least one negative entry and at least two positive entries, then $S_{\mathbf{a}}$ is not preserved by f_k for all $k \geq 2$.*

Proof. We will prove the first part of the proposition, as the second part follows from duality. Without loss of generality, let us suppose the first three entries of $\mathbf{a}$ are given by -1, -1, and 1, respectively. Then define the vectors $\mathbf{x}_1 = (1, 0, 1, 0, \ldots, 0)^t$, $\mathbf{x}_2 = (2, 0, 2, 0, \ldots, 0)^t$, and $\mathbf{x}_3 = (0, 2, 2, 0, \ldots, 0)^t$. And if $k > 2$, we let $\mathbf{x}_4 = \cdots = \mathbf{x}_{k+1} = (0, 0, \ldots, 0)^t$. We see that $\mathbf{x}_1, \ldots, \mathbf{x}_{k+1} \in S_{\mathbf{a}}$. Calculating, we find that $h_k(\mathbf{x}_1, \ldots, \mathbf{x}_{k+1}) = (1, 0, 2, 0, \ldots, 0)^t$, which is not in $S_{\mathbf{a}}$. □

Thus far we have provided conditions under which a half-space is preserved by either h_k or f_k for each k. If a half-space is preserved by a set of lattice polynomials, then it is also preserved by their functional composition. Thus, if a half-space is preserved by, say h_k for some k, then it is also preserved by the clone $\langle h_k \rangle$ generated by h_k. Note that this is slightly different than saying it is preserved by the *convex* clone generated by h_k and it is certainly different than saying it is preserved by $J_k^{\uparrow} = \langle h_k \rangle^{\uparrow}$, the upper composition ideal generated by h_k, which is the closure under composition of $\langle h_k \rangle$ with all lattice polynomials p greater than or equal to some lattice polynomial in $\langle h_k \rangle$. However, we are able to resolve this with the following result.

Lemma 4.2.5. *Let* $\mathbf{a}$ *be any sign vector with only nonpositive entries. If* p *is any lattice polynomial polymorphism of the half-space* $S_{\mathbf{a}}$*, then* $S_{\mathbf{a}}$ *is preserved by* $\langle p \rangle^{\uparrow}$*, the upper composition ideal generated by* p*.*

Dually, let $\mathbf{a}$ *be any sign vector with only nonnegative entries. If* p *is any lattice polynomial polymorphism of the half-space* $S_{\mathbf{a}}$*, then* $S_{\mathbf{a}}$ *is preserved by* $\langle p \rangle^{\downarrow}$*, the lower composition ideal generated by* p*.*

Proof. This follows from the fact that for any half-space $S_{\mathbf{a}}$ for which $\mathbf{a}$ contains only nonpositive entries, if $\mathbf{x} \in S_{\mathbf{a}}$ and $\mathbf{y} \geq \mathbf{x}$, then $\mathbf{y} \in S_{\mathbf{a}}$. Since $p \leq q$, we have that $p(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k) \leq q(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k)$ for all $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k \in S_{\mathbf{a}}$. But since $p(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k) \in S_{\mathbf{a}}$ by assumption, it follows that $q(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k) \in S_{\mathbf{a}}$.

As already discussed, since $S_{\mathbf{a}}$ is preserved by p, it is necessarily preserved by the clone $\langle p \rangle$. Now let q be a lattice polynomial in $\langle p \rangle^{\uparrow}$, the upper composition ideal generated by p. Then by monotonicity of lattice polynomials, it follows that there is a lattice polynomial $r \in \langle p \rangle$ such that $r \leq q$. Since $S_{\mathbf{a}}$ is preserved by r, it follows that it is preserved by q. □

We now have enough to prove our culminating result of this chapter, which characterizes many classes of polyhedra preserved by convex clones of lattice polynomials.

Theorem 4.2.6. *Let $\mathcal{P}_{A,\mathbf{b}}$ be a polyhedron in $\mathbb{R}^n$. Then we have the following.*

(1) For all $k \in [n]$, $\mathcal{P}_{A,\mathbf{b}} \in \mathrm{Inv}\ J_k^{\uparrow}$ if and only if for each row $\mathbf{a}$ of A, either

(a) $\mathbf{a}$ has one positive entry and at most one negative entry, or

(b) $\mathbf{a}$ has no positive entries and at most k negative entries.

Dually, $\mathcal{P}_{A,\mathbf{b}} \in \mathrm{Inv}\ J_k^{\downarrow}$ if and only if for each row $\mathbf{a}$ of A, either

(a') $\mathbf{a}$ has one negative entry and at most one positive entry, or

(b') $\mathbf{a}$ has no negative entries and at most k positive entries.

(2) Suppose $\mathcal{P}_{A,\mathbf{b}}$ can be expressed as $\mathcal{P}_{A,\mathbf{b}} = \mathcal{Q} \cap \mathcal{R}$ for some $\mathcal{Q} \in \mathrm{Inv}\ J_i^{\downarrow}$ and some $\mathcal{R} \in \mathrm{Inv}\ J_j^{\uparrow}$. Then $\mathcal{P}_{A,\mathbf{b}} \in \mathrm{Inv}\ J_r^{\downarrow} \cap J_s^{\uparrow}$ for all $r \geq i$ and $s \geq j$.

(3) $\mathcal{P}_{A,\mathbf{b}} \in \mathrm{Inv}\ J_2^{\downarrow} \cap J_2^{\uparrow}$ if and only if each row $\mathbf{a}$ of A has at most two nonzero entries.

Proof. Not that for part (1) we only need to prove the first part (up until "Dually"), as the second part follows from duality. Thus, to prove part (1), first note that a polyhedron $\mathcal{P}$ is preserved by $J_1^{\uparrow} = J_1^{\downarrow}$ if and only if $\mathcal{P}$ is a sublattice of $\mathbb{R}^n$, the characterization of which was proved by Veinott [16]. Namely, $\mathcal{P}_{A,\mathbf{b}}$ is preserved by $J_1^{\uparrow} = J_1^{\downarrow}$ if and only if each row $\mathbf{a}$ of A has at most one positive entry and at most one negative entry. Of course, by the containment properties $J_{k+1}^{\uparrow} \subseteq J_k^{\uparrow}$ and $J_{k+1}^{\downarrow} \subseteq J_k^{\downarrow}$, if any bounding half space of $\mathcal{P}_{A,\mathbf{b}}$ is preserved by $J_1^{\uparrow} = J_1^{\downarrow}$, it is necessarily preserved by all other convex lattice polynomial clones. Hence, if case (a) of part (1) holds for a given row $\mathbf{a}_i$ of A, then the bounding half-space $H_{\mathbf{a}_i,b_i}$ is preserved by all convex lattice polynomial clones. If case (b) holds in part (1), then the result follows from Corollary 4.2.2. Conversely, if $\mathcal{P}_{A,b}$ is preserved by $J_k^{\uparrow}$, then in particular it is preserved by h_k. By Corollary 4.2.2, it follows that if a row $\mathbf{a}$ in A has no positive entries, then it has at most k negative entries, and it follows from Proposition 4.2.4 that if a row $\mathbf{a}$ in A has one positive entry then it can have at most one negative entry and if it has more than one negative entry, then it can have no positive entries. Consequently, either (a) or (b) are satisfied, concluding the proof for part (1).

For part (2), observe that $J_r^{\downarrow} \cap J_s^{\uparrow} \subseteq J_i^{\downarrow} \cap J_j^{\uparrow}$ for all $r \geq i$ and $s \geq j$, and hence $\text{Inv } J_i^{\downarrow} \cap J_j^{\uparrow} \subseteq \text{Inv } J_r^{\downarrow} \cap J_s^{\uparrow}$ for all $r \geq i$ and $s \geq j$. Hence it suffices to show that $\mathcal{P}_{A,\mathbf{b}} \in J_i^{\downarrow} \cap J_j^{\uparrow}$, as the rest follows from containment. Now, let p be any k-ary lattice polynomial in $J_i^{\downarrow} \cap J_j^{\uparrow}$ and $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k \in \mathcal{P}_{A,\mathbf{b}}$. Then $p(\mathbf{x}_2, \ldots, \mathbf{x}_k) \in \mathcal{Q}$ since $p \in J_i^{\downarrow}$ and $\mathcal{Q}$ is preserved by $J_i^{\downarrow}$. Likewise $p(\mathbf{x}_2, \ldots, \mathbf{x}_k) \in \mathcal{R}$ since $p \in J_j^{\uparrow}$ and $\mathcal{R}$ is preserved by $J_j^{\uparrow}$. Hence $p(\mathbf{x}_2, \ldots, \mathbf{x}_k) \in \mathcal{Q} \cap \mathcal{R}$. Thus, $\mathcal{P}_{A,\mathbf{b}} \in \text{Inv } J_i^{\downarrow} \cap J_j^{\uparrow}$.

To prove part (3), suppose first that $\mathcal{P}_{A,\mathbf{b}}$ is preserved by $J_2^{\downarrow} \cap J_2^{\uparrow}$. Then in particular it is preserved by $h_2 = f_2$. Let $\mathbf{a}$ be any row of A and consider the bounding half-space $S_{\mathbf{a}}$. Note that Corollary 4.2.2 implies that if a row $\mathbf{a}$ of A has no positive entries, then the bounding half-space $S_{\mathbf{a}}$ is preserved by $h_2 = f_2$ if and only if it has at most 2 negative entries. Corollary 4.2.3 implies that if a row $\mathbf{a}$ of A has no negative entries, then $S_{\mathbf{a}}$ is preserved by $f_2 = h_2$ if and only if it has at most 2 positive entries. From Proposition 4.2.4 it follows that if $\mathbf{a}$ has exactly one negative entry, then it can have at most one positive entry, and if $\mathbf{a}$ has exactly one positive entry, then it can have at most one negative entry. Hence, if $S_{\mathbf{a}}$ is preserved by $J_2^{\downarrow} \cap J_2^{\uparrow}$, it follows that $\mathbf{a}$ must have at most two nonzero entries.

Conversely, if $\mathbf{a}$ has at most two nonzero entries, then if both are positive or both negative, it follows from Corollaries 4.2.2 and 4.2.3 that $S_{\mathbf{a}}$ is preserved by $f_2 = h_2$, and hence preserved by the clone $\langle h_2 \rangle$. Moreover, from Lemma 4.2.5, we see that if there are lattice polynomials $p, q \in \langle h_2 \rangle$ and some lattice polynomial s such that $p \leq s \leq q$, then $S_{\mathbf{a}}$ is also preserved by s (if both entries are negative, this follows from the fact that $s \geq p$ and if both entries are positive, this follows from the fact that $s \leq q$) and hence $S_{\mathbf{a}}$ is preserved by the convex lattice polynomial clone generated by $f_2 = h_2$, which is given by $J_2^{\downarrow} \cap J_2^{\uparrow}$. If there are two nonzero entries of opposite sign, or only one nonzero entry, then $S_{\mathbf{a}}$ is preserved by $J_1^{\uparrow} = J_1^{\downarrow}$ and hence preserved by all lattice polynomial clones, thus concluding our proof of part (3). $\square$

As an interesting aside, we note that a consequence from the proof of part (3) of the above theorem is that a polyhedron is preserved by h_2 if and only if it is preserved by the *convex* lattice polynomial clone generated by h_2. In turn, this implies that the characterization of polyhedra in part (3) is in fact equivalent to the class of polyhedra described in Theorem 2.3.6.

Unfortunately, we do not currently have enough tools to determine whether the converse of part (2) in Theorem 4.2.6 holds. This ultimately points to some of the limitations of using Galois Pol-Inv connections over infinite sets. As mentioned in Section 2.2, for every set R of relations, the set Pol R is a clone, but not every clone may be expressed as Pol R for some set of relations R. What this means here is that the most we can say is that $J_i^{\downarrow} \cap J_j^{\uparrow} \subseteq \mathrm{Pol}_{\mathcal{LP}}\ \mathrm{Inv}\ J_i^{\downarrow} \cap J_j^{\uparrow}$, and hence we cannot leverage the dual isomorphism of Theorem 2.2.3. In particular, the lattice of lattice polynomial clones on $\mathbb{R}$ may not even be a sublattice of the lattice of Galois closed sets of lattice polynomial clones on $\mathbb{R}$.

We note that part (3) of Theorem 4.2.6 is a particular exception to the above discussion. That is, part (3) is equivalent to saying the converse of part (2) holds in the case of $J_2^{\downarrow} \cap J_2^{\uparrow}$. In general, however, this is likely not the case. There is no reason to believe that, simply because a half-space $S_{\mathbf{a}}$ is preserved by the convex lattice polynomial clone $\langle C\rangle_c = \langle C\rangle^{\downarrow} \cap \langle C\rangle^{\uparrow}$, it is therefore necessarily preserved by either $\langle C\rangle^{\downarrow}$ or $\langle C\rangle^{\uparrow}$. Indeed, suppose $p \in J_i^{\downarrow} \cap J_j^{\uparrow}$ is chosen as in Corollary 3.0.6 so that $J_i^{\downarrow} \cap J_j^{\uparrow} = \langle\{p\}\rangle_c = \langle\{p\}\rangle^{\downarrow} \cap \langle\{p\}\rangle^{\uparrow}$. We note that, in particular, contrary to the $i = j = 2$ case, for $i, j > 2$, a half-space $S_{\mathbf{a}}$ cannot be preserved by both f_i and h_j. While being preserved by either h_j or f_i suffices, we cannot yet rule out the possibility that there is an $S_{\mathbf{a}}$ that is preserved by $\langle\{p\}\rangle_c$ but is preserved by neither f_i nor h_j.

While Theorem 4.2.6 is not a complete characterization of all polyhedra preserved by convex lattice polynomial clones, this nevertheless significantly generalizes the results of Veinott [16], and we hope this may serve as a first step in generalizing lattice programming and broadening the applicability of this area to linear programming.

Chapter 5: Join Irreducibles of Polytopes

In this chapter, we wish to characterize polytopes in $\mathbb{R}^n$ by their join irreducible elements. We will once again restrict ourselves to considering the component-wise join, as in previous chapters. Unless specified otherwise, when a polytope or polyhedron is referred to as a "join semilattice," it is to be assumed that the join operator is the component-wise join. That is, it is to be assumed that the polytope or polyhedron is a subsemilattice of $\mathbb{L}^n = (\mathbb{R}^n, \leq)$.

Before characterizing the join irreducibles, it is helpful to recall their definition:

Definition 5.0.1. *Let L be a join semilattice. An element $\mathbf{x} \in L$ is* join irreducible *if whenever $\mathbf{x} = \mathbf{y} \vee \mathbf{z}$ in L, then either $\mathbf{x} = \mathbf{y}$ or $\mathbf{x} = \mathbf{z}$. In a meet semilattice,* meet irreducibles *are defined dually.*

We note that the term "join irreducible" is not synonymous with minimal. An element $\mathbf{x}$ is *minimal* in L if there exist no other $\mathbf{y} \in L$ such that $\mathbf{y} < \mathbf{x}$. Clearly, every minimal element is join irreducible, but the converse is not true in general.

We note that Theorem 14 of [16] provides an irreducible representation of polyhedral sublattices of $\mathbb{R}^n$. In that article, the representation was called *irreducible* because the elements of the representation (the tangent half-spaces defining the polyhedral sublattice) were both meet and join irreducible *in the lattice of closed convex sublattices of* $\mathbb{R}^n$. In this chapter, we are interested instead in characterizing the elements that are join irreducible *in a join semilattice polytope* $\mathcal{P}$. Thus, while the objects of study in [16] and this Chapter are similar (polyhedral sublattices vs. polyhedral join semilattices), the irreducible elements characterized in [16] and those characterized in this Chapter differ significantly.

In what follows, we will establish the following notation. Let S be any subset of $\mathbb{R}^n$. Denote by $P_{i_1,i_2,\ldots,i_k}(S)$ the coordinate projection of S onto its $i_1, i_2, \ldots, i_k$ coordinates, and let $P^{-1}_{i_1,i_2,\ldots,i_k}(P_{i_1,i_2,\ldots,i_k}(S))$ denote the inverse image of the coordinate projection. For simplicity, if $k = n-1$, then we may choose to express $P_{i_1,i_2,\ldots,i_k}(S) \equiv P_{\widehat{j}}(S)$, where $\{j\} = [n]\backslash\{i_1, i_2, \ldots, i_k\}$.

With this notation, we may now characterize necessary and sufficient conditions for an element of a join semilattice polytope to be join irreducible. An important observation for the remainder of this chapter is that, if $\mathcal{P}$ is any polytope in $\mathbb{R}^n$ and $s \in \mathcal{P}$, then $P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{s\})) \cap \mathcal{P}$ is a chain, that is, a totally ordered subset, and since $\mathcal{P}$ is bounded, this chain has a least element, $\min\{P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{s\})) \cap \mathcal{P}\}$.

5.1 A Characterization of Join Irreducibles

We begin this section with the following result, once again noting that the join operator of interest is the component-wise join.

Proposition 5.1.1. *Let $\mathcal{P}$ be a polytope in $\mathbb{R}^n$ that is also a join semilattice. If an element $\mathbf{k}$ in $\mathcal{P}$ is join irreducible, then it is the least element in $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap \mathcal{P}$ for some $i \in [n]$.*

Proof. Let $\mathbf{x}$ be the least element in $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\}))\cap\mathcal{P}$ and $\mathbf{y}$ the least element in $P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{\mathbf{k}\}))\cap \mathcal{P}$ for some $i, j \in [n]$. Since $\mathbf{x} \leq \mathbf{k}$ and $\mathbf{y} \leq \mathbf{k}$, and $\mathbf{x}$ and $\mathbf{y}$ each agree with $\mathbf{k}$ on $n-1$ coordinates, it is straightforward to see that $\mathbf{k} = \mathbf{x} \vee \mathbf{y}$. But by assumption, $\mathbf{k}$ is join irreducible, and hence either $\mathbf{x} = \mathbf{k}$ or $\mathbf{y} = \mathbf{k}$. □

An almost immediate consequence of the above proposition is the following proposition, which completely characterizes the join irreducibles in a polytope.

Proposition 5.1.2. *Let $\mathcal{P}$ be a polytope in $\mathbb{R}^n$ that is also a join semilattice. An element $\mathbf{k}$ in $\mathcal{P}$ is join irreducible in $\mathcal{P}$ if and only if there is at most one coordinate $i \in [n]$ for which $\mathbf{k} \neq \min\{P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap \mathcal{P}\}$.*

Proof. If $\mathbf{k}$ is join irreducible in $\mathcal{P}$ then the result follows immediately from Proposition 5.1.1. Conversely, suppose there exists at most one coordinate $i \in [n]$ for which $\mathbf{k} \neq$

$\min\{P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\}))\cap\mathcal{P}\}$. Let us suppose for contradiction that $\mathbf{k}$ is not join irreducible. Then there exists $\mathbf{x}, \mathbf{y} \in \mathcal{P}\backslash\{\mathbf{k}\}$ such that $\mathbf{x}\vee\mathbf{y} = \mathbf{k}$. Note that, since $\mathbf{k} = \min\{P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{\mathbf{k}\}))\cap\mathcal{P}\}$ for all $j \in [n]\backslash\{i\}$ it is necessarily the case that $\mathbf{x}_i < \mathbf{k}_i$. Otherwise, by minimality, we would necessarily have $\mathbf{x} = \min\{P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{\mathbf{x}\}))\cap\mathcal{P}\} = \min\{P_{\widehat{j}}^{-1}(P_{\widehat{j}}(\{\mathbf{k}\}))\cap\mathcal{P}\} = \mathbf{k}$ for all $j \in [n]$, and hence $\mathbf{x} = \mathbf{k}$. But if $\mathbf{x}_i < \mathbf{k}_i$, then we must have $\mathbf{y}_i = \mathbf{k}_i$, and hence again by minimality, it follows that $\mathbf{y} = \mathbf{k}$, giving us our desired contradiction. □

An immediate corollary to Proposition 5.1.2 is the following result when we restrict ourselves to the case of $\mathbb{R}^2$.

Corollary 5.1.3. *Let $\mathcal{P}$ be a polytope in $\mathbb{R}^2$ that is also a join semilattice. Then for any $\mathbf{k} \in \mathcal{P}$ and $i \in \{1, 2\}$, the element $\min\{P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\}))\cap\mathcal{P}\}$ is join irreducible in $\mathcal{P}$. Moreover, every element $\mathbf{k} \in \mathcal{P}$ that is not itself join irreducible, may be expressed as the join of exactly two join irreducibles, namely*

$$\mathbf{k} = \min\{P_{\widehat{1}}^{-1}(P_{\widehat{1}}(\{\mathbf{k}\}))\cap\mathcal{P}\} \vee \min\{P_{\widehat{2}}^{-1}(P_{\widehat{2}}(\{\mathbf{k}\}))\cap\mathcal{P}\}. \tag{5.1}$$

We mentioned at the beginning of this chapter the difference between join irreducible and minimal elements. We may characterize this difference another way with the following result.

Lemma 5.1.4. *Let $\mathcal{P}$ be any polytope in $\mathbb{L}^n = (\mathbb{R}^n, \leq)$. An element $\mathbf{k} \in \mathcal{P}$ is minimal in $\mathcal{P}$ if and only if $\mathbf{k} = \min\{P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\}))\cap\mathcal{P}\}$ for every coordinate i.*

Proof. Suppose $\mathbf{k} = \min\{P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\}))\cap\mathcal{P}\}$ for every coordinate i. Then if $\mathbf{x} \in \mathcal{P}$ is less than $\mathbf{k}$, then each coordinate $\mathbf{x}_i \leq \mathbf{k}_i$ with at least one coordinate j satisfying $\mathbf{x}_j < \mathbf{k}_j$, a contradiction. Conversely, if $\mathbf{k}$ is minimal, then there exists no $\mathbf{x} < \mathbf{k}$ in $\mathcal{P}$, and hence every coordinate $\mathbf{k}_i$ of $\mathbf{k}$ is minimal. The result follows. □

We note that, in the above lemma, we did not need to assume that $\mathcal{P}$ was a semilattice in $\mathbb{R}^n$. That it was a poset sufficed.

Combining Proposition 5.1.2 and Lemma 5.1.4, we may generalize Corollary 5.1.3. Before doing so, we provide the following notation which will be utilized in the next result. For an element $\mathbf{k} \in \mathbb{R}^n$, we let $\mathcal{H}_{\mathbf{k}_i}$ denote the hyperplane $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_i = \mathbf{k}_i\}$. That is, it is the hyperplane that fixes the i-th coordinate as $\mathbf{k}_i$.

Proposition 5.1.5. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$. Then any element $\mathbf{x} \in \mathcal{P}$ may be expressed as the join of at most n join irreducibles in $\mathcal{P}$.*

Proof. For each coordinate i, consider the polytope $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_i}$. If $\mathbf{x}$ is minimal in $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_i}$, then there is at most one coordinate, namely i, for which $\mathbf{x} \neq \min\{P_{\hat{i}}^{-1}(P_{\hat{i}}(\{\mathbf{x}\})) \cap \mathcal{P}\}$. Then by Proposition 5.1.2, it follows that $\mathbf{x}$ is join irreducible, and we are done. Suppose, then, that $\mathbf{x}$ is not minimal in $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_i}$. Since $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_i}$ is a polytope, and hence bounded, it follows that there necessarily exists at least one minimal element $\mathbf{k} \in \mathcal{P} \cap \mathcal{H}_{\mathbf{x}_i}$. Since $\mathbf{x}$ is not minimal by assumption, we may choose $\mathbf{k}$ such that $\mathbf{k} \leq \mathbf{x}$. Then there is at most one coordinate, namely i, for which $\mathbf{k} \neq \min\{P_{\hat{i}}^{-1}(P_{\hat{i}}(\{\mathbf{k}\})) \cap \mathcal{P}\}$, and hence by Proposition 5.1.2, it follows that $\mathbf{k}$ is join irreducible. For each coordinate i, enumerate each such join irreducible element found in this way by $\mathbf{k_i}$. Observe that, for each i, we have $\mathbf{k_i}_i = \mathbf{x}_i$ and $\mathbf{k_i}_j \leq \mathbf{x}_j$ for all $j \neq i$. It follows that $\mathbf{k_1} \vee \mathbf{k_2} \vee \cdots \vee \mathbf{k_n} = \mathbf{x}$. Hence, $\mathbf{x}$ may be expressed as the join of n (not necessarily distinct) join irreducibles. $\square$

For illustrative purposes, consider the polytope $\mathcal{P}$ from Figure 5.1, given by

$$\mathcal{P} = \operatorname{conv}\left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}. \tag{5.2}$$

It is easy to see that $\mathcal{P}$ is a join semilattice. The join irreducibles are all of the

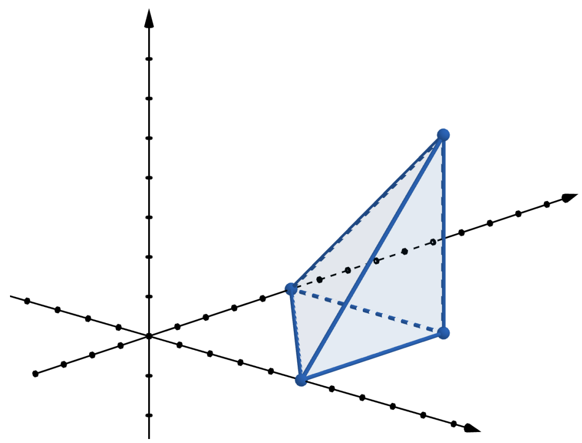

Figure 5.1: The polytope $\mathcal{P}_{A,\mathbf{b}}$ from Eq. (5.2).

points on the plane $\operatorname{conv}\big((1,0,0)^t,(0,1,0)^t,(1,1,1)^t\big)$. Perhaps somewhat counterintuitively, the point $(1,1,1)^t$ is join irreducible even though it is also the greatest element. The point $\mathbf{x} = (1,1,0)^t$ is not join irreducible, so following the method from the proof of Proposition 5.1.5, we begin by finding a minimal element in the polytope $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_1} = \operatorname{conv}\big((1,0,0)^t,(1,1,0)^t,(1,1,1)^t\big)$. There is exactly one: $(1,0,0)^t$. Similarly, we find a minimal element in the polytope $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_2} = \operatorname{conv}\big((0,1,0)^t,(1,1,0)^t,(1,1,1)^t\big)$. Once again, there is exactly one: $(0,1,0)^t$. Both $(1,0,0)^t$ and $(0,1,0)^t$ are join irreducible, and already we see that $(1,0,0)^t \vee (0,1,0)^t = (1,1,0)^t$, so we do not need to consider the last case. Nevertheless, for the sake of this example, we may consider finding a minimal element in the polytope $\mathcal{P} \cap \mathcal{H}_{\mathbf{x}_3} = \operatorname{conv}\big((1,0,0)^t,(0,1,0)^t,(1,1,0)^t\big)$. In this case all of the points along the edge $\operatorname{conv}\big((1,0,0)^t,(0,1,0)^t\big)$ are minimal and less than $\mathbf{x} = (1,1,0)^t$, so we may pick any one of them as our join irreducible. Again, only two join irreducibles were necessary. In fact, it does not take too much effort to see that any point in this polytope that is not itself a join irreducible may be decomposed as the join of exactly two join irreducibles.

Note that, because Propositions 5.1.1 and 5.1.2 apply to polytopes, which are bounded,

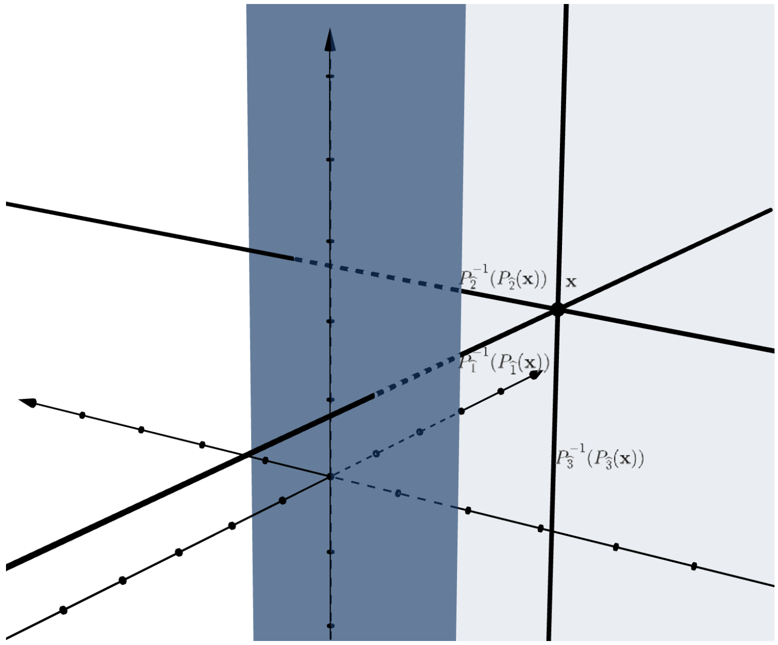

Figure 5.2: A point $\mathbf{x}$ in the half-space $H_{\mathbf{a},b}$ for $\mathbf{a} = (1, -1, 0)^t$ and $b = 0$. Since the first coordinate of $\mathbf{a}$ is positive, $P_{\hat{1}}^{-1}(P_{\hat{1}}(\mathbf{x})) \cap H_{\mathbf{a},b}$ has a maximum but no minimum; since the second coordinate is negative, $P_{\hat{2}}^{-1}(P_{\hat{2}}(\mathbf{x})) \cap H_{\mathbf{a},b}$ has a minimum but no maximum; and since the third coordinate is zero, $P_{\hat{3}}^{-1}(P_{\hat{3}}(\mathbf{x})) \cap H_{\mathbf{a},b}$ never intersects the boundary of $H_{\mathbf{a},b}$ and so it has neither a maximum nor a minimum.

their proofs implicitly assumed that $\min\{P_{\hat{i}}^{-1}(P_{\hat{i}}(\mathbf{x})) \cap \mathcal{P}\}$ existed. This was only for convenience, and the same result can easily be restated to apply to half-spaces as well, as long as we acknowledge that in this case, $P_{\hat{i}}^{-1}(P_{\hat{i}}(\mathbf{x})) \cap H$ may be totally or partially unbounded when H is a half-space and $\mathbf{x} \in H$, and hence $\min\{P_{\hat{i}}^{-1}(P_{\hat{i}}(\mathbf{x})) \cap H\}$ may not always exist in this case. More precisely, we have the following:

Corollary 5.1.6. *Let H be a half-space in $\mathbb{R}^n$ that is also a join semilattice. An element $\mathbf{k}$ in H is join irreducible in H if and only if there is at most one coordinate $i \in [n]$ for which $\mathbf{k} \neq \min\{P_{\hat{i}}^{-1}(P_{\hat{i}}(\{\mathbf{k}\})) \cap H\}$.*

An immediate consequence of the above corollary is that some join semilattice polyhedra do not have join irreducibles. The half space $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_1 \geq 0\}$ is one such example. In turn, given an arbitrary join semilattice polyhedron, we cannot assume that every element may be decomposed as the join of join irreducibles. Hence, Proposition 5.1.5 cannot be generalized to join semilattice polyhedra.

Nevertheless, it is now straightforward to see the requirements for a half-space to contain join-irreducibles. Let $H_{\mathbf{a},b} := \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{x}, \mathbf{a} \rangle \leq b\}$. Then for any $\mathbf{k} \in H_{\mathbf{a},b}$ and $i \in [n]$, we consider $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap H_{\mathbf{a},b}$. There are exactly three possible cases, depending on the sign of $\mathbf{a}_i$. Namely, when $\mathbf{a}_i > 0$, then $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap H_{\mathbf{a},b}$ has a maximum but no minimum; when $\mathbf{a}_i < 0$, then $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap H_{\mathbf{a},b}$ has a minimum but no maximum; and when $\mathbf{a}_i = 0$, then $P_{\widehat{i}}^{-1}(P_{\widehat{i}}(\{\mathbf{k}\})) \cap H_{\mathbf{a},b}$ has neither. See Figure 5.2 for an example. Thus, we may equivalently reformulate the above Corollary:

Corollary 5.1.7. *The half-space $H_{\mathbf{a},b}$ contains join irreducibles if and only if $\mathbf{a}$ has at most one nonnegative entry. Moreover, if $H_{\mathbf{a},b}$ contains join irreducibles, they are exactly the boundary of $H_{\mathbf{a},b}$.*

By a similar argument, we may say the following:

Corollary 5.1.8. *The half-space $H_{\mathbf{a},b}$ contains minimal elements if and only if $\mathbf{a}$ has all negative entries. Moreover, if $H_{\mathbf{a},b}$ contains minimal elements, they are exactly the boundary of $H_{\mathbf{a},b}$.*

By the above discussion, we see that it is not the values themselves, but the signs of the entries in the vector $\mathbf{a}$ defining the half-space that determine the existence of join irreducibles and minimal elements. More precisely, to each half-space $H_{\mathbf{a},b}$ we may associate another half-space $S_{\mathrm{sgn}(\mathbf{a})} := \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{x}, \mathrm{sgn}(\mathbf{a}) \rangle \leq 0\}$, where $\mathrm{sgn}(\mathbf{a}) \in \{-1, 0, 1\}^n$ is the sign vector of $\mathbf{a}$. Then the next result follows immediately.

Corollary 5.1.9. *The half-space $H_{\mathbf{a},b}$ in $\mathbb{R}^n$ contains join irreducibles if and only the half-space $S_{\mathrm{sgn}(\mathbf{a})}$ contains join irreducibles, or equivalently if and only if $\mathrm{sgn}(\mathbf{a})$ contains at least $n-1$ negative entries. Similarly, the half-space $H_{\mathbf{a},b}$ in $\mathbb{R}^n$ contains minimal elements if* *and only the half-space $S_{\mathrm{sgn}(\mathbf{a})}$ contains minimal elements, or equivalently, if* $\mathrm{sgn}(\mathbf{a}) = -\mathbf{1}$. *The join irreducibles of both $H_{\mathbf{a},b}$ and $S_{\mathrm{sgn}(\mathbf{a})}$ coincide with their boundaries when they exist.*

Of course, if $\mathbf{k}$ is join irreducible in H for some half-space H, then it is also join irreducible in $H_1 \cap H_2 \cap \cdots \cap H_m$, where we set $H_1 = H$. So now we may consider decomposing a polytope into its join irreducibles. Recall that a half-space $H_{\mathbf{a},b}$ is preserved by componentwise join if and only if $\mathbf{a}$ has at most one strictly positive entry. However, $H_{\mathbf{a},b}$ contains join irreducibles if and only if $\mathbf{a}$ contains at most one nonnegative entry. Therefore is it certainly possible to have a join semilattice polytope that is an intersection of hyperplanes, none of which satisfy the conditions of the above Corollary. Take for example the unit n-cube in $\mathbb{R}^n$. It is a join semilattice (indeed, it is a sublattice of $\mathbb{R}^n$), but none of its bounding half-spaces satisfy the conditions of Corollary 5.1.7. Does that mean that the n-cube contains no join-irreducibles? Quite to the contrary, one easily sees that the join irreducibles are precisely the intersection of the n-cube with its n coordinate axes. This provides the intuition for our characterization of the join irreducibles of an arbitrary join semilattice polytope in $\mathbb{R}^n$.

Definition 5.1.10. *A nonempty face F of a join semilattice polytope $\mathcal{P}$ is called a* join irreducible face *if all of its elements are join irreducible in $\mathcal{P}$, and if F' is another face in $\mathcal{P}$ all of whose elements are join irreducible in $\mathcal{P}$ and $F \subseteq F'$, then $F = F'$.*

The second part of the above definition imparts a maximality requirement on the face. Of course, if F is a face, all of whose elements are join irreducible in $\mathcal{P}$, then any subface of F will also only contain join irreducibles, and hence it is of particular interest to characterize the maximal such sets. We will set up our next result in the following way. Let $\mathcal{P} = \mathcal{P}_{A,\mathbf{b}}$ be a polytope in $\mathbb{R}^n$ with bounding half-spaces $H_{\mathbf{a_1},b_1}, \ldots, H_{\mathbf{a_m},b_m}$. For each half-space $H_{\mathbf{a_i},b_i}$, denote by $\mathcal{H}_i$ the corresponding boundary hyperplane. Any face F other than the interior of a full-dimensional polytope (that is, any proper face) lies tangent to a maximal subset of these hyperplanes: $F \subseteq \mathcal{H}_{i_1} \cap \cdots \cap \mathcal{H}_{i_j}$. Thus, for any proper face F, we may associate a matrix $\mathcal{A}$, the rows of which correspond to the vectors $\mathbf{a_1}, \ldots, \mathbf{a_j}$. We will denote this face by $F_{\mathcal{A}}$. Then we have the following result.

Theorem 5.1.11. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$. Then the join irreducible faces of $\mathcal{P}$ are precisely the proper faces $F_{\mathcal{A}}$ with associated matrix $\mathcal{A}$ satisfying*

(a) at most one column of $\mathcal{A}$ contains only nonnegative entries, and

(b) if any row of $\mathcal{A}$ is removed, then condition (a) will not be satisfied.

Proof. Condition (b) excludes subfaces of join irreducible faces, so that the maximality requirement of Definition 5.1.10 is satisfied.

For Condition (a), we may assume without loss of generality that the rows of $\mathcal{A}$, denoted $\mathbf{a_1}, \mathbf{a_2}, \ldots, \mathbf{a_j}$, are all in $\{-1, 0, 1\}^n$, and that the face $F_{\mathcal{A}}$ is the collection of points $\mathbf{x} \in \mathbb{R}^n$ satisfying $\mathcal{A}\mathbf{x} = \mathbf{0}$. Now, suppose Condition (a) is satisfied. Without loss of generality, suppose the first $n-1$ columns each contain a -1, and let $k = \max_i\{\mathbf{a_i}_n\}$. Then if $\mathbf{x}$ is in $F_{\mathcal{A}}$, so that $\mathcal{A}\mathbf{x} = \mathbf{0}$, it follows that $\langle \mathbf{k}, \mathbf{x} \rangle = 0$, where $\mathbf{k} = (-1, -1, \ldots, k)^t$. But then $\mathbf{x}$ is an element of the boundary of the half-space $H_{\mathbf{k},0}$. Since $\mathbf{k}$ satisfies the conditions of Corollary 5.1.7, it follows that $\mathbf{x}$ is a join irridiucible of $H_{\mathbf{k},0}$, and hence a join irreducible of $H_{\mathbf{k},0} \cap H_{\mathbf{a_1},0} \cap \cdots \cap H_{\mathbf{a_j},0} = H_{\mathbf{a_1},0} \cap \cdots \cap H_{\mathbf{a_j},0}$, so that $\mathbf{x}$ is a join irreducible element of $\mathcal{P}$.

Conversely, let us suppose $F_{\mathcal{A}}$ is a join irreducible face of $\mathcal{P}$, and assume for contradiction that at least two columns of $\mathcal{A}$ contain only nonnegative entries. Once again, without loss of generality, we may assume that $\mathbf{b} = \mathbf{0}$ and each $\mathbf{a_i} \in \{-1, 0, 1\}^n$, so that each $\mathcal{H}_{\mathbf{a_i}} = \{\mathbf{x} \in \mathbb{R}^n \mid \langle \mathbf{x}, \mathbf{a_i} \rangle = 0\}$. Let us assume even further without loss of generality that it is the first two columns of $\mathcal{A}$ that contain only nonnegative entries. Now there are exactly two cases we need to consider. The first case is that $(\min_i\{\mathbf{a_i}_1\}, \min_i\{\mathbf{a_i}_2\}) = (0, 0)$, and the second case is that $(\min_i\{\mathbf{a_i}_1\}, \min_i\{\mathbf{a_i}_2\}) = (0, 1)$. Note that the case $(\min_i\{\mathbf{a_i}_1\}, \min_i\{\mathbf{a_i}_2\}) = (1, 0)$ follows a nearly identical argument to the second case, and the case that $(\min_i\{\mathbf{a_i}_1\}, \min_i\{\mathbf{a_i}_2\}) = (1, 1)$ is excluded by the assumption that $\mathcal{P}$ is a join semilattice (and hence each $\mathbf{a_i}$ can have at most one strictly positive entry).

For the first case, observe that $\mathbf{0} \in F_{\mathcal{A}}$, and there are $\alpha_1, \alpha_2 < 0$ such that $(\alpha_1, 0, \ldots, 0)^t$, and $(0, \alpha_2, 0, \ldots, 0)^t$ are both elements of $F_{\mathcal{A}}$. However, $(\alpha_1, 0, \ldots, 0)^t \vee (0, \alpha_2, \ldots, 0)^t = \mathbf{0}$, and hence $\mathbf{0}$ is not join irreducible, a contradiction.

For the second case, observe once again that $\mathbf{0} \in F_{\mathcal{A}}$, and there is an $\alpha_1 < 0$ such that $(\alpha_1, 0, \ldots, 0)^t$ is an elements of $F_{\mathcal{A}}$. Also, there is necessarily an $\alpha_2 < 0$ such that $(0, \alpha_2, \ldots, 0)^t$ is an interior point in $\mathcal{P}$. But then $(\alpha_1, 0, \ldots, 0)^t \vee (0, \alpha_2, \ldots, 0)^t = \mathbf{0}$, and hence $\mathbf{0}$ is not join irreducible, a contradiction.

It follows that if $F_{\mathcal{A}}$ is a join irreducible face of $\mathcal{P}$, then Condition (a) must be satisfied, concluding the proof. $\square$

An immediate consequence of the above theorem is the following.

Corollary 5.1.12. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$ with $\dim(\mathcal{P}) \geq 2$. If $F_{\mathcal{A}}$ is a join irreducible face of $\mathcal{P}$, then $\dim(F_{\mathcal{A}}) \geq 1$.*

Proof. If, on the contrary, $\dim(F_{\mathcal{A}}) = 0$, so that $F_{\mathcal{A}}$ is a vertex, then in particular, $\mathcal{A}$ will be an $n \times n$ matrix. Condition (b) of Theorem 5.1.11 implies that removing any row will result in a matrix with at least two columns containing only nonnegative entries. It follows that each row must contain at least one negative entry in a position different from any other row. Since $\mathcal{A}$ has n rows, it follows that every column of $\mathcal{A}$ contains at least one negative entry, and furthermore, if one row is removed, then at most one column of $\mathcal{A}$ will contain only nonnegative entries, which violates condition (b), giving us our desired contradiction. $\square$

A nearly identical argument to that of Theorem 5.1.11 proves the following.

Corollary 5.1.13. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$. Then the faces of $\mathcal{P}$ that contain only minimal elements in $\mathcal{P}$ are precisely the proper faces $F_{\mathcal{A}}$ with associated matrix $\mathcal{A}$ such that every column of $\mathcal{A}$ contains at least one negative entry. The face is a maximal face of minimal elements if in addition, removing any row of $\mathcal{A}$ causes this condition to fail.*

In particular, from the above we see that Corollary 5.1.12 does not necessarily hold in the case of faces of minimal elements. Indeed, a minimal element may arise as the least element in a join irreducible face.

5.2 Join-Closed Join Irreducible Faces

Suppose $\mathcal{P}$ is a join semilattice polytope, and let $J \subseteq \mathcal{P}$ be the collection of join irreducibles. As we have shown in the previous section, the join irreducibles of $\mathcal{P}$ can be grouped into join irreducible faces. Let us supposed further that each one of these faces is closed under the join operation. That is to say, if $F \subseteq J$ is such a face and $f_1, f_2 \in F$, then $f_1 \vee f_2 \in F$. Note that, since f_1 and f_2 are both join irreducible in $\mathcal{P}$, it follows that $f_1 \vee f_2$ is either f_1 or f_2. In other words, such a face is necessarily totally ordered, and hence $\dim F \leq 1$. We know by Corollary 5.1.12 that it cannot be a vertex, and hence such a face is necessarily an edge. We will call such faces of $\mathcal{P}$ *join-closed join irreducible* faces. We note that the term "join irreducible" faces refers to the fact that all of the elements in these faces are join irreducible in $\mathcal{P}$. However, we will shortly see that these join-closed join irreducible faces are join irreducible in another sense as well.

Let $\mathbb{J}^n$ be the set of all join subsemilattices of $\mathbb{L}^n = (\mathbb{R}^n, \leq)$. Observe that if $F, G \in \mathbb{J}^n$ are both join subsemilattices of $\mathbb{L}^n$, then the set

$$F \vee_{\mathbb{J}} G := \{f \vee g \mid f \in F,\ g \in G\}, \tag{5.3}$$

is also a join subsemilattice of $\mathbb{L}^n$. Note that, because each $F \in \mathbb{J}^n$ is a join subsemilattice (and hence join-closed as a subset of $\mathbb{L}^n$), it follows that $F \vee_{\mathbb{J}} F = F$. Associativity and commutativity of the operator $\vee_{\mathbb{J}}$ defined in (5.3) are easily seen to hold, and hence $(\mathbb{J}^n, \vee_{\mathbb{J}})$ is a well-defined join semilattice. The partial order is given by $F \leq G \Leftrightarrow F \vee_{\mathbb{J}} G = G$.

For example, consider the simplest case of $(\mathbb{J}^1, \vee_{\mathbb{J}})$. Let A be any subset of $\mathbb{R}$, and suppose $a, b \in A$. Then $a \vee b = \max\{a, b\} \in A$, and hence every subset of $\mathbb{R}$ is a join subsemilattice of $\mathbb{L}^1$. So $\mathbb{J}^1 = \mathcal{P}(\mathbb{R})$, the set of all subsets of $\mathbb{R}$. However, the join operation $\vee_{\mathbb{J}}$ is very different from set union. For example, let $A = \{(0, 1) \cup \{2\}\}$ and let $B = [1, 2)$. Then $A \vee_{\mathbb{J}} B = [1, 2]$. More generally, if $A, B \subseteq \mathbb{R}$ are both nonempty and there exists an $x \in A$ such that $x < y$ for every $y \in B$, then $A \vee_{\mathbb{J}} B = (A \cup B) \setminus \{x \in A \mid x < y \text{ for all } y \in B\}$.

Note that the empty set $\emptyset$ is trivially a join subsemilattice of $\mathbb{L}^n$. Moreover, since the elements in $F \vee_{\mathbb{J}} G$ in Eq. (5.3) are pairwise joins of elements from F and G, then if either F or G is empty, then so must be $F \vee_{\mathbb{J}} G$. In other words, $\emptyset \vee_{\mathbb{J}} G = \emptyset$ for all $G \in \mathbb{J}^n$, and hence $\emptyset$ is the unique maximal element in the join semilattice $(\mathbb{J}^n, \vee_{\mathbb{J}})$.

Now if the set $\mathcal{J}$ of join irreducible faces of the join semilattice polytope $\mathcal{P}$ are also join-closed, then they are also elements of $\mathbb{J}^n$, as is the polytope $\mathcal{P}$. It is thus of interest to consider the structure of the subsemilattice $\langle \mathcal{J}, \vee_{\mathbb{J}} \rangle$ generated by $\mathcal{J}$ and its relation to $\mathcal{P}$. We first make the following observation.

Lemma 5.2.1. *Let $\mathcal{P} \subseteq \mathbb{R}^n$ be a join semilattice polytope. Let $\mathcal{J}$ be the set of join irreducible faces of $\mathcal{P}$ and suppose all of these faces are join-closed. Then $\mathcal{P}$ has a least element.*

Proof. Suppose for contradiction that there does not exist a least element in $\mathcal{P}$. Note that every minimal element in $\mathcal{P}$ is the least element in a join-closed join irreducible face. Hence there are finitely many minimal elements in $\mathcal{P}$. Let $\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_k}$ be the set of minimal elements in $\mathcal{P}$. Since we are assuming there is no least element, it follows that $k \geq 2$. Then let $\mathbf{z} = \frac{1}{k}(\mathbf{x_1} + \mathbf{x_2} + \cdots + \mathbf{x_k})$. By convexity, $\mathbf{z} \in \mathcal{P}$. However, none of the $\mathbf{x_i}$ are $\leq \mathbf{z}$ and hence $\mathbf{z}$ cannot be expressed as the join of join irreducibles, contradicting Proposition 5.1.5. ☐

Thus, the restriction that the set $\mathcal{J}$ of join irreducible faces of a join semilattice polytope $\mathcal{P}$ all be join-closed implies that $\mathcal{P}$ necessarily has a least element. However, we will now show that the set $\mathcal{J}$ viewed as poset has no least element.

Proposition 5.2.2. *Suppose $\mathcal{P}$ is a join semilattice polytope with $\dim(\mathcal{P}) \geq 2$. Suppose the set $\mathcal{J}$ of join irreducible faces of $\mathcal{P}$ are all join-closed. Then $\mathcal{J}$ contains no least element.*

Proof. Suppose for contradiction that $\mathcal{J}$ does contain a least element F^*. Since $\dim(\mathcal{P}) \geq 2$, it follows that $\mathcal{P}$ contains elements that are not join irreducible. Moreover, since every join irreducible face is greater than or equal to F^*, it follows that there exists an element $\mathbf{x}$ of

$\mathcal{P}$ that is not join irreducible such that the only join irreducible elements less than $\mathbf{x}$ are in F^*.

To see this, consider the least element $\mathbf{x}^*$ of F^*. It is strictly less than any join irreducible element $\mathbf{k}$ from any other join irreducible face $K \in \mathcal{J}\backslash\{F^*\}$. Then for each $K \in \mathcal{J}\backslash\{F^*\}$ there exists a sufficiently small ϵ-ball around $\mathbf{x}^*$ such that its intersection with $\mathcal{P}$ contains no elements greater than or equal to any element in K. We choose our $\mathbf{x}$ to be a non join irreducible element of the intersection of all of these ϵ-balls with $\mathcal{P}$.

By Proposition 5.1.5, $\mathbf{x}$ can be expressed as the join of join irreducibles, but since $\mathbf{x}$ is not itself join irreducible, it must be expressed as the join of at least two join irreducibles. Since the only join irreducibles less than or equal to $\mathbf{x}$ are in F^*, it follows that $\mathbf{x}$ is a join of join irreducibles from F^*. But since F^* is join-closed, it follows that $\mathbf{x} \in F^*$, a contradiction since $\mathbf{x}$ is not join irreducible. $\square$

Note that the subsemilattice $\langle \mathcal{J}, \vee_{\mathbb{J}} \rangle$ is finite (and hence complete). It follows from Proposition 5.2.2 that $\langle \mathcal{J}, \vee_{\mathbb{J}} \rangle$ has no least element. Let us add a zero element to $\langle \mathcal{J}, \vee_{\mathbb{J}} \rangle$ and call the resulting space $\mathcal{K}_{\mathcal{P}}$. Then by the dual of Proposition 1.2.1, it follows that $\mathcal{K}_{\mathcal{P}}$ is, in fact, a lattice. Note that, by construction, the join-closed join irreducible faces $\mathcal{J}$ of $\mathcal{P}$ are precisely the join irreducible elements of $\mathcal{K}_{\mathcal{P}}$ (providing somewhat of a double meaning to the "join irreducible" part of the term "join-closed join irreducible faces").

We stress that $\mathcal{K}_{\mathcal{P}}$ is well-defined if and only if each face in $\mathcal{J}$ is join-closed. Suppose we extend the definition (5.3) of the operator $\vee_{\mathbb{J}}$ to arbitrary subsets of $\mathbb{L}^n$ and not just join subsemilattices. In this larger space, the operator is still easily seen to be commutative and associative, but it is not, in general, idempotent. Hence this join operator on the larger space defines a semigroup, but not a semilattice. For this reason, we must restrict ourselves to considering join semilattice polytopes $\mathcal{P}$ such that their join irreducible faces $\mathcal{J}$ are all join-closed.

A simple example of how idempotence fails in the more general setting may be found by considering the triangle $\mathcal{P} = \operatorname{conv}\{(1,0)^t,\ (0,1)^t,\ (1,1)^t\}$. In this example, the only join irreducible face is $A = \operatorname{conv}\{(1,0)^t,\ (0,1)^t\}$, and $\mathcal{P} = A \vee_{\mathbb{J}} A$. In other words, each element of $\mathcal{P}$ is a join of at most two elements from the same join irreducible face.

For the remainder of this section, unless otherwise specified, we will assume that $\mathcal{P}$ is a join semilattice polytope, and all of its join irreducible faces are join-closed. Furthermore, with a slight abuse of notation, we may choose to omit the subscript $\mathbb{J}$ in the join operator $\vee_{\mathbb{J}}$. We have the following result.

Proposition 5.2.3. *$\mathcal{K}_{\mathcal{P}}$ is a distributive lattice.*

Proof. Each nonzero element of $\mathcal{K}_{\mathcal{P}}$ may be expressed as the join of elements in $\mathcal{J}$. Let $K_1, K_2, K_3 \in \mathcal{K}_{\mathcal{P}}$ and for each K_i let $\mathcal{F}_i \subseteq \mathcal{J}$ be chosen such that $K_i = \vee \mathcal{F}_i$ (we use the notation $\vee \mathcal{F}_i$ to denote the join of all of the elements in $\mathcal{F}_i$). For each $\mathcal{F}_i$ define

$$\widehat{\mathcal{F}}_i := \{F \in \mathcal{J} \cup \{0\} \mid \exists\, G \in \mathcal{F}_i \text{ such that } F \leq G\}. \tag{5.4}$$

In particular, 0 is an element of every $\widehat{\mathcal{F}}_i$. Clearly, $K_i = \vee \widehat{\mathcal{F}}_i$, and furthermore,

$$K_i \vee K_j = \bigvee \left(\widehat{\mathcal{F}}_i \cup \widehat{\mathcal{F}}_j \right). \tag{5.5}$$

Note that, for every element G in $\mathcal{K}_{\mathcal{P}}$ less than or equal to K_i, there exists a corresponding subset $\mathcal{G} \subseteq \widehat{\mathcal{F}}_i$ such that $G = \vee \mathcal{G}$. In turn, the meet $K_i \wedge K_j$ is given by the join of the elements that $\widehat{\mathcal{F}}_i$ and $\widehat{\mathcal{F}}_j$ have in common. More precisely,

$$K_i \wedge K_j = \bigvee \left(\widehat{\mathcal{F}}_i \cap \widehat{\mathcal{F}}_j \right). \tag{5.6}$$

It then follows that

$$
\begin{aligned}
K_1 \vee (K_2 \wedge K_3) &= \bigvee \left(\widehat{\mathcal{F}}_1 \cup (\widehat{\mathcal{F}}_2 \cap \widehat{\mathcal{F}}_3) \right) \\
&= \bigvee \left((\widehat{\mathcal{F}}_1 \cup \widehat{\mathcal{F}}_2) \cap (\widehat{\mathcal{F}}_1 \cup \widehat{\mathcal{F}}_3) \right) \qquad (5.7) \\
&= (K_1 \vee K_2) \wedge (K_1 \vee K_3).
\end{aligned}
$$

□

Proposition 5.2.4. *Let $\mathcal{P}$ be a join semilattice polytope, and let $\mathcal{J}$ denote the set of join irreducible faces of $\mathcal{P}$. Suppose each face in $\mathcal{J}$ is join-closed. Then*

$$
\mathcal{P} = \bigcup_{\mathcal{F} \subseteq \mathcal{J}} \left(\bigvee_{F \in \mathcal{F}} F \right). \qquad (5.8)
$$

Proof. Proposition 5.1.5 showed us that every element of $\mathcal{P}$ can be expressed as the join of join-irreducible elements. Theorem 5.1.11 showed us that the join-irreducible elements of $\mathcal{P}$ are faces of $\mathcal{P}$. If each join irreducible face is join-closed, then it follows that each element of $\mathcal{P}$ may be expressed as a join of join irreducible elements, each element of which is from a different join irreducible face. The result follows. □

Hence, when the join irreducible faces are all join-closed, we can consider "building up" the polytope $\mathcal{P}$ from the elements of $\mathcal{K}_{\mathcal{P}}$.

Another immediate consequence of Proposition 5.2.4 is the following refinement of Proposition 5.1.5 when each of the join irreducible faces is join-closed.

Corollary 5.2.5. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$ with join irreducibles, and let $\mathcal{J}$ denote the set of join irreducible faces of $\mathcal{P}$. Suppose each element of $\mathcal{J}$ is join-closed. Then each element of $\mathcal{P}$ can be expressed as the join of at most $\min\{n, |\mathcal{J}|\}$ join irreducibles.*

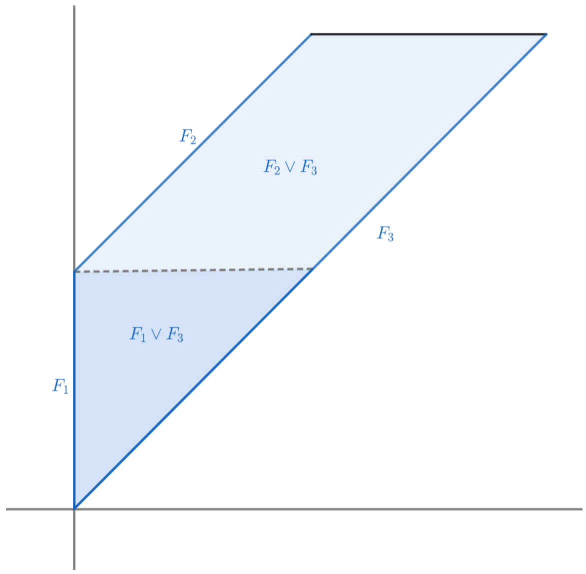

Figure 5.3: The polytope $\mathcal{P}_{A,\mathbf{b}}$ from Eq. (5.9) and its corresponding decomposition in terms of join irreducibles.

Consider for example the polytope $\mathcal{P}_{A,\mathbf{b}}$ in $\mathbb{R}^2$, where

$$A = \begin{bmatrix} 0 & 1 \\ -1 & 1 \\ -1 & 0 \\ 1 & -1 \end{bmatrix}, \qquad \mathbf{b} = \begin{pmatrix} 2 \\ 1 \\ 0 \\ 0 \end{pmatrix}. \tag{5.9}$$

By looking at A, we see that there are three join-irreducible faces; namely, those tangent to the lines $-x = 0$, $-x + y = 1$, and $x - y = 0$, which we label F_1, F_2, and F_3, respectively. Observe that $F_1 \leq F_2$, as $F_1 \vee F_2 = F_2$. A consequence of this is that $F_1 \vee F_2 \vee F_3 = F_2 \vee F_3 \neq \mathcal{P}$. Additionally, note that $F_1 \vee F_3$ is not convex, as $F_3 \subseteq F_1 \vee F_3$ (see Figures 5.3 and 5.4).

In the previous example, the polytope was a subset of $\mathbb{R}^2$, but it had three join irreducible faces. An interesting consequence of Corollary 5.2.5 is that since the number of join irreducible faces was larger than the dimension, there must be an order relation between

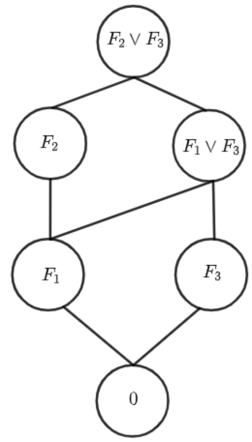

Figure 5.4: The lattice $\mathcal{K}_{\mathcal{P}}$, where $\mathcal{J}$ is the set of join irreducible faces of $\mathcal{P}_{A,\mathbf{b}}$ from Eq. (5.9).

at least two of the faces. Otherwise there would exist an element in $\mathcal{P}$ that required a minimum of three join irreducible elements to characterize it, violating Corollary 5.2.5.

In Theorem 5.1.11, we characterized the join irreducible faces. Building off of this, we may completely characterize the join-closed join irreducible faces.

Theorem 5.2.6. *Let $\mathcal{P}$ be a join semilattice polytope in $\mathbb{R}^n$. Then the join irreducible faces of $\mathcal{P}$ that are also join-closed are precisely the proper faces $F_{\mathcal{A}}$ of $\mathcal{P}$ with associated matrix $\mathcal{A}$ satisfying*

(a) At most one column of $\mathcal{A}$ contains only nonnegative entries,

(b) If any row of $\mathcal{A}$ is removed, then condition (a) will not be satisfied.

(c) Each row of $\mathcal{A}$ has at most two nonzero entries, and if a row contains exactly two nonzero entries, then they are of opposite signs.

Proof. From Theorem 5.1.11 we know that a proper face is join irreducible if and only if conditions (a) and (b) are satisfied. Now let $H_{\mathbf{a},b}$ denote a half-space with boundary hyperplane given by $\mathcal{H}$. Then $\mathcal{H} = H_{\mathbf{a},b} \cap H_{-\mathbf{a},-b}$, and hence $\mathcal{H}$ is join-closed if and only

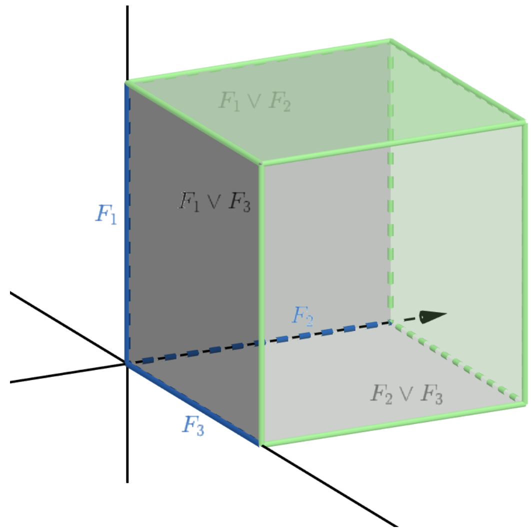

Figure 5.5: The unit 3-cube $\mathcal{D}$. The join irreducibles are the edges on the coordinate axes (colored in blue). Their pairwise joins are colored in grey. The join of all three is the entire cube.

if both $H_{\mathbf{a},b}$ and $H_{-\mathbf{a},-b}$ are join-closed. Condition (c) then follows immediately from this and Corollary 4.2.3 for the case $k = 1$. □

Recall that a half-space $H_{\mathbf{a},b}$ is join-closed if and only if its negative half-space $H_{-\mathbf{a},-b}$ is closed under meet, or "meet-closed." Consequently the hyperplane $\mathcal{H} = H_{\mathbf{a},b} \cap H_{-\mathbf{a},-b}$ is join-closed if and only if it is also meet-closed. Thus, an immediate consequence of part (c) of Theorem 5.2.6 is that the join-closed join irreducible faces of $\mathcal{P}$ are necessarily meet-closed as well.

For example, consider the unit 3-cube $\mathcal{D}$ in $\mathbb{R}^3$. Its join irreducible faces are the line segments $F_1 = \text{conv}\{(0,0,0)^t, (0,0,1)^t\}$, $F_2 = \text{conv}\{(0,0,0)^t, (0,1,0)^t\}$, and $F_3 = \text{conv}\{(0,0,0)^t, (1,0,0)^t\}$ (see Figure 5.5). Observe that $F_1 = F_{\mathcal{A}_1}$ has associated face matrix $\mathcal{A}_1$ given by

$$\mathcal{A}_1 = \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}. \tag{5.10}$$

Each row of $\mathcal{A}_1$ corresponds to a join-closed hyperplane (satisfying condition (c) of Theorem

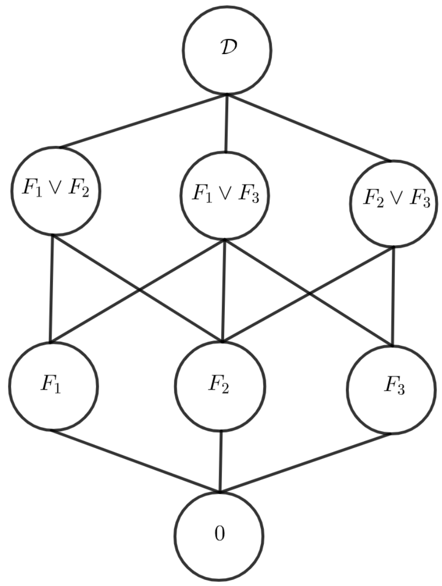

Figure 5.6: The lattice $\mathcal{K}_{\mathcal{D}}$, where $\mathcal{J}$ is the set of join irreducible faces of the three cube $\mathcal{D}$.

5.2.6). Additionally, only the last column contains only nonnegative entries, but if either row is removed there will be two columns with nonnegative entries, so Conditions (a) and (b) of Theorem 5.2.6 are also satisfied.

Our next result places a lower bound on the number of join irreducible faces of a polytope $\mathcal{P}$ when all of them are join-closed and $\mathcal{P}$ has maximal dimension.

Proposition 5.2.7. *Let $\mathcal{P} \subseteq \mathbb{R}^n$ be a join semilattice polytope with* $\dim(\mathcal{P}) = n$*. Let $\mathcal{J}$ be the set of join irreducible faces of $\mathcal{P}$ and suppose all of these faces are join-closed. Then $|\mathcal{J}| \geq n$.*

Proof. By Lemma 5.2.1, we know that $\mathcal{P}$ must have a least element, denoted $\mathbf{x}^*$. Since $\mathcal{P}$ is full-dimensional, it follows that $\mathbf{x}^*$ is the intersection of at least n facets, $\mathcal{H}_{\mathbf{a}_1}, \mathcal{H}_{\mathbf{a}_2}, \ldots, \mathcal{H}_{\mathbf{a}_d}$ for some $d \geq n$. Since $\mathbf{x}^*$ is clearly join irreducible, it is an element of a join irreducible face. Since all join irreducible faces of $\mathcal{P}$ are join-closed by assumption, it follows from Theorem 5.2.6 that $n-1$ of the facets defining $\mathbf{x}^*$ and having nonempty intersection, say $\mathcal{H}_{\mathbf{a}_1}, \mathcal{H}_{\mathbf{a}_2}, \ldots, \mathcal{H}_{\mathbf{a}_{n-1}}$, are join-closed, and their intersection defines a join-closed join irreducible face. In particular, because each of these $\mathcal{H}_{\mathbf{a}_i}$ is join-closed, it follows that each corresponding $\mathbf{a}_i$ has at most one negative entry. Consequently, the corresponding face matrix $\mathcal{A}$ defined by these $n-1$ facets has exactly one column which contains only nonnegative entries.

Moreover, $\mathbf{x}^*$ is a minimal element, and hence there exists another facet $\mathcal{H}_{\mathbf{a}_n}$ such that the face matrix $\mathcal{A}^*$ for $\mathbf{x}^*$ defined from these n facets satisfies the requirement of Corollary 5.1.13; namely, every column of $\mathcal{A}^*$ contains at least one negative entry.

Note that it immediate follows from the assumption that all join irreducible faces are join-closed (and hence have dimension 1) that $\mathbf{a}_n$ must have at most one negative entry, and it must correspond to the column in $\mathcal{A}$ with only nonnegative entries. To see that $\mathbf{a}_n$ cannot have more than one negative entry, note that, otherwise, it would be possible to construct a face matrix $\mathcal{A}'$ corresponding to the intersection of fewer than $n-1$ of these facets that satisfied the criteria of Theorem 5.1.11, and would thus be a join irriducible face of dimension > 1. Since all join-closed join irreducible faces have dimension 1, we see that this cannot happen, and thus, $\mathbf{a}_n$ must have at most one negative entry.

Now, it is straightforward to see that by removing any row from $\mathcal{A}^*$ gives us a matrix that satisfies the criteria of Theorem 5.2.6, and hence corresponds to a join-closed join irreducible face. There are n such combinations, and hence $\mathcal{P}$ has at least n join-closed join irreducible faces, as we wanted to show. $\square$

The following is a consequence of the above Proposition.

Proposition 5.2.8. *Suppose as in the above Proposition that $\mathcal{P}$ is full-dimensional in $\mathbb{R}^n$ with join irreducible set $\mathcal{J}$ such that all elements of $\mathcal{J}$ are join-closed. Then*

1. *There exist elements in $\mathcal{P}$ that must be expressed as the join of exactly n join irreducibles, and*

2. *If $|\mathcal{J}| = n$, then the largest element $\vee\mathcal{J}$ in the lattice $\mathcal{K}_{\mathcal{P}}$ is $\mathcal{P}$.*

Proof. The first result follows from the fact that, since there are necessarily n join irreducible faces all sharing the least element of $\mathcal{P}$, then these n join irreducible faces are all non-comparable.

The second result also follows from the fact that all of the join irreducible faces have the least element of $\mathcal{P}$ in common. In particular, not only do we recover all the elements of $\mathcal{P}$ that must be expressed as the join of exactly n join irreducibles, but whenever an element $\mathbf{z}$ of $\mathcal{P}$ can be expressed as an element of $k < n$ join irreducibles, then we may simply join $n-k$ copies of $\mathbf{x}^*$ to the decomposition, and then $\mathbf{z} \in \vee\mathcal{J}$. □

An immediate consequence of this is the following.

Corollary 5.2.9. *Suppose as in the above Proposition that $\mathcal{P}$ is full-dimensional in $\mathbb{R}^n$ with join irreducible set $\mathcal{J}$ such that all elements of $\mathcal{J}$ are join-closed. Then largest element $\vee\mathcal{J}$ in the lattice $\mathcal{K}_{\mathcal{P}}$ is $\mathcal{P}$ if and only if $|\mathcal{J}| = n$.*

Additionally, when $\mathcal{P}$ is full-dimensional in $\mathbb{R}^n$ with join irreducible set $\mathcal{J}$ containing only faces that are join-closed satisfying $|\mathcal{J}| = n$, then the join irreducible faces have no order relation between them, so that $\mathcal{K}_{\mathcal{P}} \simeq 2^{\mathcal{N}} \simeq 2^{[n]}$, where $\mathcal{N}$ is the join irreducible face set of the n-cube in $\mathbb{R}^n$, and $[n]$ denotes a set on n elements.

Chapter 6: Conclusion

Lattice programming studies how the optimal solution to a problem changes in a dynamic environment when the problem domain is a lattice. When exact solutions are not required, then the ability to do this is particularly valuable, as it can be far easier (with respect to computational resources) than recomputing the optimal solution at each time step.

Polyhedra are defined implicitly with respect to the component-wise partial order in the formalism:

$$\mathcal{P}_{A,\mathbf{b}} = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}. \tag{6.1}$$

Hence, when considering the overlap between lattice programming and linear programming, this partial order provides a natural framework from which to do so. Namely, we may consider polyhedra that are themselves lattices with respect to this partial order. Unfortunately, as was shown by Veinott [16], the overlap between lattice programming and linear programming within this framework is rather limited.

As a first step in attempting to generalize and extend lattice programming to be applicable to a wider range of linear optimization scenarios, we investigated how much more of an overlap there could be by relaxing the constraint that the domain be a lattice. When a polyhedron is a lattice, it is preserved by meet and join, and all compositions therein. But, as we have shown, it is possible for some polyhedra to be preserved by *some but not all* lattice polynomials. Hence it is possible for these polyhedra to retain some of the lattice structure in a manner that is made well-defined through the use of convex lattice polynomial clones. Thus the opportunity is opened to begin investigating the extent to which this relaxation of full lattice structure will permit a generalization or extension of lattice programming to a wider range of linear programming problems. This, however, will be reserved for future work.

In Veinott's same paper [16] that greatly motivated the work in this dissertation , another topic that was briefly addressed was that of irreducibles. In his work, Veinott was specifically pointing to the fact that polyhedral sublattices of $\mathbb{R}^n$ could be described in terms of a primitive set of irreducible elements in the lattice of convex sets, since the defining half planes are both join and meet irreducible in this lattice. The consideration by Veinott on irreducibles led this author to consider what the join irreducibles of a join semilattice polytope might look like.

Indeed, as we were able to show, the join irreducible elements always form faces of the corresponding polytope. Furthermore, if the join irreducible faces are all join-closed, then we may "build up" the entire polytope via joins of these join irreducible faces (which also happen to only be edges of the polytope). We discovered these properties in much the same way we went about characterizing the preservation of polyhedra under various convex lattice polynomial clones. Namely, we identified each of these faces with a particular matrix $\mathcal{A}$ and, once again, found that the characterization of join irreducible faces, and in particular the join-closed join irreducible faces, was determined by the associated sign vectors.

An interesting observation is that criteria (c) in Theorem 5.2.6 implies that $\mathcal{A}^t$ is the incidence matrix of a directed graph. This follows immediately from the fact that a non-interior face is preserved by join if and only if it is also preserved by meet. Hence, as Veinott proved in [16], the associated matrix is an incidence matrix for a generalized network flow graph. However, criteria (a) and (b) place restrictions on the associated graph. More precisely, they imply that the corresponding graph is one that has (a) at most one vertex that is either isolated or is a sink, and (b) if any edge is removed, then there is necessarily more than one vertex that is either isolated or a sink. In particular, criteria (b) implies that every vertex has at most a single edge leaving the vertex.

While the study of join irreducibles certainly revealed some interesting structure, these results are hardly the end. Over the course of this research, we have studied and characterized polyhedra preserved by a variety of different lattice polynomials, whereas in the study of join irreducibles, we needed to restrict ourselves to polyhedra that were preserved by join. An interesting question is whether we may generalize the notion of join irreducibility further. Namely, suppose a polytope $\mathcal{P}$ is preserved by a k-ary lattice polynomial p. Then we may call an $\mathbf{x}$ in $\mathcal{P}$ a p-irreducible element if whenever $\mathbf{x} = p(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k)$ for some $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k \in \mathcal{P}$, then $\mathbf{x} \in \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k\}$. When does a p-closed polyhedron $\mathcal{P}$ contain p-irreducibles? Can there be interior elements that are p-irreducible? Are there lattice polynomials p for which there exists no polyhedron preserved by p that contains p-irreducibles? We believe these are all interesting questions that are worth pursuing in future work.

Bibliography

[1] C. Benzaken. Post's closed systems and the weak chromatic number of hypergraphs. *Discrete Mathematics*, 23(2):77 – 84, 1978.

[2] G. Birkhoff. *Lattice Theory*, volume 25 of *American Mathematical Society colloquium publications*. American Mathematical Society, 1940.

[3] Garrett Birkhoff. On the combination of topologies. *Fundamenta Mathematicae*, 26(1):156–166, 1936.

[4] Ferdinand Börner. *Basics of Galois Connections*, pages 38–67. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[5] A. Bulatov. Identities in the lattices of closed classes. *Discrete Mathematics and Applications*, 3(6):601–610, 2009.

[6] George Epstein. The lattice theory of post algebras. *Transactions of the American Mathematical Society*, 95(2):300–317, 1960.

[7] Martin Goldstern and Michael Pinsker. A survey of clones on infinite sets. *Algebra universalis*, 59(3):365–403, Dec 2008.

[8] Sebastian Kerkhoff, Reinhard Pöschel, and Friedrich Martin Schneider. A short introduction to clones. *Electronic Notes in Theoretical Computer Science*, 303(Supplement C):107 – 120, 2014. Proceedings of the Workshop on Algebra, Coalgebra and Topology (WACT 2013).

[9] Dietlinde Lau. *Function Algebras on Finite Sets: A Basic Course on Many-Valued Logic and Clone Theory*. Springer, Berlin, Heidelberg, Berlin, Heidelberg, 2006.

[10] James F. Lawrence. Distributive lattices and hypergraph coloring. Technical Report NISTIR-5942, National Institute of Standards in Technology, February 1997.

[11] Emil L. Post. Determination of all closed systems of truth tables. *Bulletin of the American Mathematical Society*, 26:437, 1920.

[12] Emil L. Post. Introduction to a general theory of elementary propositions. *American Journal of Mathematics*, 43(3):163–185, 1921.

[13] Emil L. Post. *The Two-Valued Iterative Systems of Mathematical Logic. (AM-5)*. Princeton University Press, 1941.

[14] Ivo G. Rosenberg. Über die funktionale vollständigkeit in den mehrwertigen logiken. struktur der funktionen von mehreren veränderlichen auf endlichen mengen. *Rozpravy Česklovenské Akad. Věd Řada Mat. Přírod*, 80:93, 1970.

[15] Paul C. Rosenbloom. Post algebras. i. postulates and general theory. *American Journal of Mathematics*, 64(1):167–188, 1942.

[16] Arthur F. Veinott. Representation of general and polyhedral subsemilattices and sublattices of product spaces. *Linear Algebra and its Applications*, 114-115:681 – 704, 1989. Special Issue Dedicated to Alan J. Hoffman.

[17] Arthur F. Veinott. Lectures in supply-chain optimization. 2005.

Curriculum Vitae

Jake Farinholt received his B.S. in Mathematics from the University of Mary Washington in May, 2009. He began his career as a civilian scientist for the U.S. Navy at the Naval Surface Warfare Center, Dahlgren Division that same month. For the next ten years, Jake independently studied and conducted research in various aspects of quantum information science for the Navy, and obtained multiple competitive Navy research grants to do so. Towards the end of his career with the Navy, he led a team of approximately 20 research mathematicians, data scientists, and quantum information scientists. While doing this, he also pursued an M.S. in Mathematics from George Mason University, achieving this in August, 2012, with a Master's Thesis on quantum error correcting codes and finite projective geometries. He took a short break from graduate school after this to experience the newly discovered joys of fatherhood, but eventually returned to GMU to complete his PhD in Mathematics. In July of 2019, Jake left the Navy to become a Quantum Scientist for Booz Allen Hamilton, where he is currently employed. Jake received his PhD in Mathematics from GMU in the Summer of 2020.