A tropical approach to Ising problems

by

Jose A. Ortiz

B.S., Portland State University, 2017

M.S., Portland State University, 2019

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Mathematics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Abstract

Tropical geometry offers a piecewise-linear framework that translates algebraic and combinatorial structures into polyhedral geometry. In this dissertation, we explore applications of tropical and polyhedral methods to the analysis of spin systems, particularly the Ising model on graphs.

Each spin configuration of a graph with n vertices corresponds to a vertex of the n-cube, and the energy associated to these configurations defines a Laurent polynomial whose Newton polytope encodes all interaction and external field parameters. We show that the tropicalization of this polynomial captures the loci in parameter space where multiple spin configurations minimize the system's energy—so-called degeneracy loci. These loci are described by tropical hypersurfaces, whose combinatorial types are determined by faces of the secondary polytope of the n-cube.

Through this connection, vertex-state interactions naturally parameterize regular subdivisions of the cube, and ground-state degeneracies are encoded by the dual secondary fan. We further construct new polytope J_Γ, whose vertices reflect both spin states and interaction parities, and provide a facet classification for graphs built from trees and cycles. This polyhedral perspective reveals a natural moduli space for studying phase transitions, optimization, and combinatorial symmetries in discrete physical systems.

A tropical approach to Ising problems

by

Jose A. Ortiz

B.S., Portland State University, 2017

M.S., Portland State University, 2019

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Mathematics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Approved by:

Major Professor
Gabe Kerr

Copyright

Abstract

Tropical geometry offers a piecewise-linear framework that translates algebraic and combinatorial structures into polyhedral geometry. In this dissertation, we explore applications of tropical and polyhedral methods to the analysis of spin systems, particularly the Ising model on graphs.

Each spin configuration of a graph with n vertices corresponds to a vertex of the n-cube, and the energy associated to these configurations defines a Laurent polynomial whose Newton polytope encodes all interaction and external field parameters. We show that the tropicalization of this polynomial captures the loci in parameter space where multiple spin configurations minimize the system's energy—so-called degeneracy loci. These loci are described by tropical hypersurfaces, whose combinatorial types are determined by faces of the secondary polytope of the n-cube.

Through this connection, vertex-state interactions naturally parameterize regular subdivisions of the cube, and ground-state degeneracies are encoded by the dual secondary fan. We further construct new polytope J_Γ, whose vertices reflect both spin states and interaction parities, and provide a facet classification for graphs built from trees and cycles. This polyhedral perspective reveals a natural moduli space for studying phase transitions, optimization, and combinatorial symmetries in discrete physical systems.

Table of Contents

List of Figures

Chapter 1

Background

1.1 Tropical Geometry

1.1.1 Introduction

Tropical geometry is a combinatorial shadow of classical geometry, wherein non-linear structures are replaced with piecewise-linear structures. Tropical geometry connects the fields of algebraic geometry, polyhedral geometry, and linear programming in interesting and novel ways. But before delving into applications of tropical geometry, we will need to introduce some background material. There are two main conventions that are used in tropical geometry: min-plus and max-plus. We use either convention where useful, but will generally default to min-plus for the introductory section and max-plus for later sections.

1.1.2 Valuations

Let k be a field. A **valuation** on k is a function $\mathrm{val} : k \to \mathbb{R} \cup \{\infty\}$ such that

1. $\mathrm{val}(a) = \infty \Leftrightarrow a = 0$
2. $\mathrm{val}(a \cdot b) = \mathrm{val}(a) + \mathrm{val}(b)$
3. $\mathrm{val}(a + b) \geq \min(\mathrm{val}(a), \mathrm{val}(b))$

4. If $\mathrm{val}(a) \neq \mathrm{val}(b)$, then $\mathrm{val}(a+b) = \min(\mathrm{val}(a), \mathrm{val}(b))$.

The image of a valuation forms an additive group called the **value group** of val, denoted Γ_{val}; throughout this work, we will assume that our valuations have a value group as a subgroup of the extended real numbers $\mathbb{R} \cup \{\infty\}$ (or $\mathbb{R} \cup \{-\infty\}$ in the case of $-\mathrm{val}$). Valuations can be thought of as a generalization of logarithms and were historically used in three settings: the trivial valuation, the p-adic valuation, and the Puiseux valuation.

Example 1.1. The trivial valuation $\mathrm{val}_0 : k \to \{0, \infty\}$ is defined as

$$\mathrm{val}_0(x) = \begin{cases} 0 & x \in k^* \\ \infty & x = 0. \end{cases}$$

In this case, $\Gamma_{\mathrm{val}_0} = \{0, \infty\}$. $\square$

Example 1.2. The p-adic valuation $\mathrm{val}_p : \mathbb{Q} \to \mathbb{Z}$ is given by

$$\mathrm{val}_p(x) = \begin{cases} \max\{n \in \mathbb{Z} \mid p^n \text{ divides } x\} & x \in \mathbb{Q}^* \\ \infty & x = 0 \end{cases}$$

So for example, if $p = 2$, then

$$\mathrm{val}_2(12) = \mathrm{val}_2(4) + \mathrm{val}_2(3) = 2 + 0 = 2$$
$$\mathrm{val}_2(16/4) = \mathrm{val}(16) + \mathrm{val}(4^{-1}) = 4 - 2 = 2$$
$$\mathrm{val}_2(11) = \mathrm{val}(3+8) = \min(\mathrm{val}(3),\ \mathrm{val}(8)) = 0$$

So the p-adic valuation measures how divisible a rational number is by p. The value group for the p-adic valuation is $\Gamma_{\mathrm{val}_p} = \mathbb{Z} \cup \{\infty\}$. $\square$

Example 1.3. The Puiseaux series valuation $\mathrm{val} : \mathbb{C}\{\{t\}\} \to \mathbb{R}$. Here $\mathbb{C}\{\{t\}\}$ denotes the field of Puiseaux series which contains all formal power series of the form $\sum_{n=n_0}^{\infty} a_n t^{m_n}$ where

$a_n \in \mathbb{C}^*$ and $m_n \in \mathbb{Q}$ $\forall n \geq n_0 \in \mathbb{Z}$. The valuation on $\mathbb{C}\{\{t\}\}$ is defined as

$$\operatorname{val}(x(t)) = \begin{cases} \min\{m_n \mid n \geq n_0\} & x(t) \in \mathbb{C}\{\{t\}\}^* \\ \infty & x(t) = 0 \end{cases}$$

So for example,

$$\operatorname{val}\left(2t^{-1} + \sum_{n \geq 0} t^n\right) = -1$$
$$\operatorname{val}\left(\sum_{n \geq 3} t^n\right) = 3$$

In other words, the Puiseaux valuation measures the degree of vanishing at $t = 0$. The value group for this valuation is $\mathbb{Q} \cup \{\infty\}$. $\square$

1.1.3 Tropicalization

To motivate the study of tropical semirings, it may be helpful to consider the following application of valuations.

With k still a field with valuation val, the ring of **Laurent polynomials** $k[x^{\pm}] = k[x_1^{\pm 1}, x_2^{\pm 1}, \ldots, x_n^{\pm 1}]$ contains polynomials of the form

$$f(x) = \sum_{u \in S \subset \mathbb{Z}^n} a_u x^u$$

where x^u denotes the monomial $x_1^{u_1} \cdot x_2^{u_2} \cdots x_n^{u_n}$ for any given n-tuple $u = (u_1, u_2, \ldots, u_n)$ from a fixed and finite set of points $S \subset \mathbb{Z}^n$, and every coefficient $a_u \in k^*$ is nonzero. A Laurent polynomial carves out a curve in the **algebraic torus** $T^n = (k^*)^n$ through its **zero locus**, or, **hypersurface** $\mathcal{V}(f) = \{x \in k^n \mid f(x) = 0\}$, so a natural question to ask is how a valuation behaves on a zero locus. The following observation was due to Kapranov.

Suppose k is an algebraically closed field with nontrivial valuation $\operatorname{val} : k \to \mathbb{R}$. If

$x \in \mathcal{V}(f) \subset (k^*)^n$ for any nonzero Laurent polynomial $f \in k[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, then $f(x) = 0$ and the first three properties of valuations gives the following

$$\begin{aligned}\infty = \mathrm{val}(f(x)) = \mathrm{val}\left(\sum_{u \in S \subset \mathbb{Z}^n} a_u x^u\right) \\ &> \min_{u \in S}(\mathrm{val}(a_u x^u)) \\ &= \min_{u \in S}(\mathrm{val}(a_u) + \mathrm{val}(x^u)) \\ &= \min_{u \in S}(\mathrm{val}(a_u) + u \cdot \mathrm{val}(x)).\end{aligned}$$

Next, note that the contrapositive of the fourth property of valuations says that if $\mathrm{val}(a + b) \neq \min(\mathrm{val}(a), \mathrm{val}(b))$, then $\mathrm{val}(a) = \mathrm{val}(b)$. So, the inequality above states that at least two of the terms $\mathrm{val}(a_u) + u \cdot \mathrm{val}(x)$ must achieve the minimum at $w = \mathrm{val}(x) \in \mathbb{R}^n$. Therefore, a valuation can be seen as measuring the failure of linearity in the function $\min_{u \in S}(\mathrm{val}(a_u) + u \cdot w)$. This motivates the following definitions.

For a Laurent polynomial $f(x) = \sum_{u \in S \subset \mathbb{Z}^n} a_u x^u$, its **tropicalization** $\mathrm{trop}(f)$ is the piecewise-linear concave function

$$\mathrm{trop}(f)(w) = \min_{u \in S}(\mathrm{val}(a_u) + u \cdot w) \quad \text{for } w \in \mathbb{R}^n.$$

When instead given a hypersurface $\mathcal{V}(f) \subset (k^*)^n$, its **tropicalization** $\mathrm{trop}(\mathcal{V}(f))$ is the set

$$\mathrm{trop}(\mathcal{V}(f)) = \{w \in \mathbb{R}^n \mid \text{the minimum occurs at least twice in } \mathrm{trop}(f)(w)\}$$

which is the set of points where $\mathrm{trop}(f)$ is not linear. The linear functions in $\mathrm{trop}(f)$ that achieve the minimum at any given point $w \in \mathbb{R}^n$ lift to an **initial form** of f at w

$$\mathrm{in}_w(f) = \sum_{\substack{u \in \mathbb{N}^n \text{ s.t.} \\ \mathrm{trop}(f)(w) = \\ \mathrm{trop}(a_u x^u)(u)}} \overline{a_u t^{-\mathrm{val}(a_u)}} x^u \in (R/\mathfrak{m})[x_1, \ldots, x_n]$$

where $R = \{x \in k \mid \mathrm{val}(x) \geq 0\}$ and $\mathfrak{m} = \{x \in k \mid \mathrm{val}(x) > 0\}$. That is, the initial form of f at a given point $w \in \mathbb{R}^n$ is a polynomial with valuation-zero terms.

Theorem 1.1 (Kapranov's Theorem[10]). *Let k be an algebraically closed field with nontrivial valuation $\mathrm{val} : k \to \mathbb{R}$. For any fixed nonzero Laurent polynomial $f(x) = \sum_{u \in S} a_u x^u \in k[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ where $S \subset \mathbb{Z}^n$ is finite and $a_u \neq 0$ for all $u \in S$. Then, the following sets are equivalent:*

1. *the tropical hypersurface $\mathrm{trop}(\mathcal{V}(f)) \subset \mathbb{R}^n$*
2. *the set $cl(\{\mathrm{val}(x) \mid x \in \mathcal{V}(f)\})$, where val is applied coordinatewise*
3. *the set $\{w \in \mathbb{R}^n \mid in_w(f)$ is not a monomial$\}$*

Furthermore, if f is irreducible and $w \in \mathrm{trop}(\mathcal{V}(f)) \cap \Gamma_{val}^n$, then the fiber $\{x \in \mathcal{V}(f) \mid \mathrm{val}(x) = w\}$ is Zariski dense in $\mathcal{V}(f) \subset (k^*)^n$.

Example 1.4. To see the connections between the three sets described in Kapranov's Theorem, consider the Laurent polynomial $f \in k[x_1^{\pm 1}, x_2^{\pm 1}]$ defined as

$$\begin{aligned} f(x) &= x_2 - x_1 - t \\ &= 1 \cdot x^{(0,1)} + (-1) \cdot x^{(1,0)} + (-t) \cdot x^{(0,0)} \end{aligned}$$

with coefficients in $k = \mathbb{C}\{\{t\}\}$. Here we have $S = \{u_1, u_2, u_3\} = \{(0,1), (1,0), (0,0)\}$ and $a_{u_1} = 1$, $a_{u_2} = -1$, $a_{u_3} = -t$. The zero-locus of f in the algebraic torus is $\mathcal{V}(f) = \{(x_1, x_1 + t) \mid x_1 \in k^*\} \subset (k^*)^2$.

For the first viewpoint, consider

$$\begin{aligned} \mathrm{trop}(f)(w) &= \min_{u \in S} (\mathrm{val}(a_u) + u \cdot w) \\ &= \min (0 + w_2,\, 0 + w_1,\, 1) \end{aligned}$$

If the minimum is achieved by at least two of the linear functions in $\mathrm{trop}(f)$ at any given $w \in \mathbb{R}$, then there are three cases: $w_2 = w_1 \leq 1$, $w_2 = 1 \leq w_1$, and $w_1 = 1 \leq w_2$.

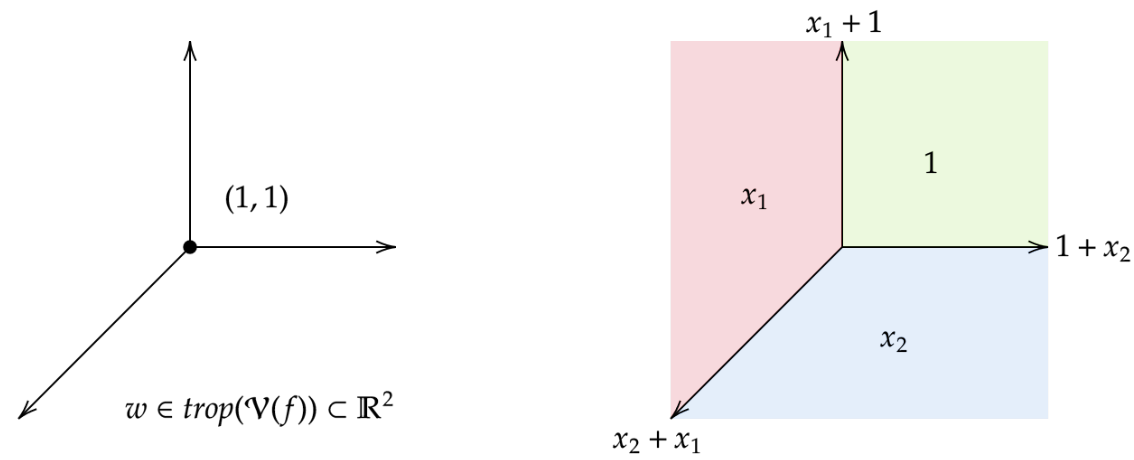

Figure 1.1: *A tropical hypersurface and a complex of initial forms of f.*

This results in the three half-open rays intersecting at $(1, 1) \in \mathbb{R}^2$: $\{(w_1, w_1) \mid w_1 \leq 1\}$, $\{(w_1, 1) \mid w_1 \geq 1\}$, and $\{(1, w_2) \mid w_2 \geq 1\}$, respectively.

For the second viewpoint, consider applying val to $\mathcal{V}(f) = \{(x_1, x_1 + t) \mid x_1 \in k^*\}$ pointwise; so we will consider points of the form $(\text{val}(x_1), \text{val}(x_1 + t))$. For any given nonzero Puiseaux series $x_1 \in k^*$, there are three possibilities: $\text{val}(x_1) < 1$, $\text{val}(x_1) > 1$, or $\text{val}(x_1) = 1$.

In case one, the lowest degree term of x_1 is less than one, so adding t does not affect that term; thus $\text{val}(x + t) < 1$. In case two, we have the opposite case; the term of lowest degree in x_1 is greater than one, so the lowest-degree term in $x_1 + t$ is t; so $\text{val}(x + t) = 1$. Lastly, in case three, the degree of vanishing of x_1 is one, so adding t may or may not cancel the term of degree one. If adding t cancels terms in x_1, then the term with the lowest degree will have a strictly higher degree. But if adding t does not cancel terms in x_1, then the lowest-degree term still has degree one. Considering both of these subcases, we get $\text{val}(x_1 + t) \geq 1$. (Note that, in this example, taking the closure of the union of the three rays adds no additional points.)

For the third point of view, consider the complement $\mathbb{R}^2 \setminus \text{trop}(\mathcal{V}(f))$ composed of three open regions: $R_1 = \{(w_1, w_2) \mid w_1, w_2 > 1\}$, $R_2 = \{(w_1, w_2) \mid w_2 < w_1 \text{ and } w_2 < 1\}$, and $R_3 = \{(w_1, w_2) \mid w_1 > w_2 \text{ and } w_1 < 1\}$. Over each of these three regions, we have the following monomial initial forms of f

$$\begin{aligned}
\mathrm{in}_w(f) &= 1, \text{ for } w \in R_1, \\
\mathrm{in}_w(f) &= x_2, \text{ for } w \in R_2, \\
\mathrm{in}_w(f) &= x_1, \text{ for } w \in R_3.
\end{aligned}$$

If we instead look at the overlap of the closed regions, we get non-monomial initial forms:

$$\begin{aligned}
\mathrm{in}_w(f) &= 1 + x_2, \text{ for } w \in \mathrm{cl}(R_1) \cap \mathrm{cl}(R_2), \\
\mathrm{in}_w(f) &= x_2 + x_1, \text{ for } w \in \mathrm{cl}(R_2) \cap \mathrm{cl}(R_3), \\
\mathrm{in}_w(f) &= x_1 + 1, \text{ for } w \in \mathrm{cl}(R_3) \cap \mathrm{cl}(R_1) \\
\mathrm{in}_w(f) &= 1 + x_2 + x_1, \text{ for } w \in \mathrm{cl}(R_1) \cap \mathrm{cl}(R_2) \cap \mathrm{cl}(R_3)
\end{aligned}$$

We can see that the three different viewpoints coincide in 1.1. Although interesting, in practice one computes tropical hypersurfaces more efficiently through the use of duality, which will be introduced later. For now, notice that a tropical hypersurface can include more information than just the set $\mathrm{trop}(\mathcal{V}(f))$; the regions R_1, R_2, R_3 together with $\mathrm{trop}(\mathcal{V}(f))$ form a **tropical complex** Σ_{trop}. $\square$

For hypersurfaces, we have $\mathcal{V}(\mathrm{trop}(f)) = \mathrm{trop}(\mathcal{V}(f))$, but in the more general setting, for any given generating set $I = (f_1, f_2, \ldots, f_k) \subset k[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ and zero-locus $X = \mathcal{V}(I) \subset (k^*)^n$, we have

$$\mathrm{trop}(X) = \mathrm{trop}(\mathcal{V}(I)) = \bigcap_{f \in I} \mathrm{trop}(\mathcal{V}(f)) \subseteq \bigcap_{i=1}^{k} \mathrm{trop}(\mathcal{V}(f_i)).$$

Where the inclusion above comes from the fact that a **tropical variety**–a set of the form $\mathrm{trop}(X)$, where X is a subvariety of an algebraic torus–is not necessarily a finite intersection

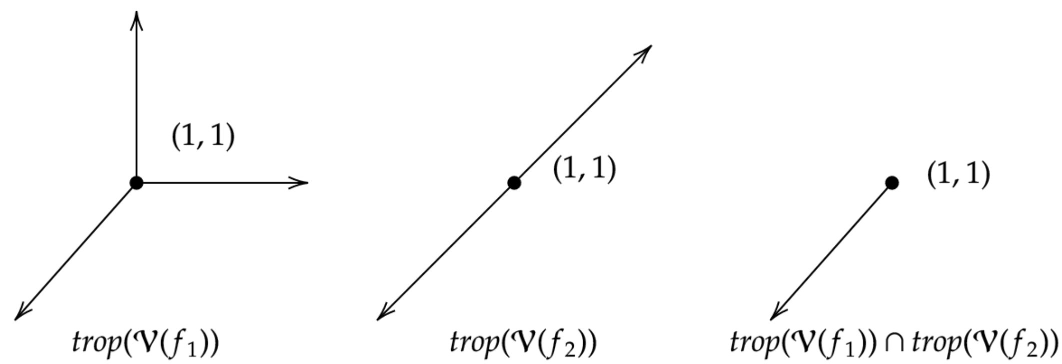

Figure 1.2: *Two tropical hypersurfaces generally intersect as a tropical prevariety.*

of tropical hypersurfaces; the latter is called a **tropical prevariety**. The following example illustrates how the tropicalization map does not distribute across intersections.

Example 1.5. Consider the Laurent polynomials $f_1(x) = x_2 - x_1 - t$ and $f_2(x) = x_2 - 2x_1$. We have $\mathcal{V}(f_1) = \{(x_1, x_1 + t) \mid x_1 \in k^*\}$ as in 1.4, and $\mathcal{V}(f_2) = \{(x_1, 2x_1) \mid x_1 \in k^*\}$. Then, with $I = (f_1, f_2)$ and $X = \mathcal{V}(I) = \mathcal{V}(f_1) \cap \mathcal{V}(f_2) = \{(t, 2t)\}$, we have $\mathrm{trop}(X) = \mathrm{cl}(\mathrm{val}(X)) = \{(1, 1)\}$. However, we have $\mathrm{trop}(\mathcal{V}(f_1))$ as in 1.4 and $\mathrm{trop}(\mathcal{V}(f_2)) = \{(w_1, w_1) \mid w_1 \in \mathbb{R}\}$.

If we instead compute the tropicalizations first, then take intersections, we get a strictly larger set

$$\mathrm{trop}(X) = \bigcap_{f \in I} \mathrm{trop}(\mathcal{V}(f)) = \{(1,1)\} \subsetneq \bigcap_{i=1}^{2} \mathrm{trop}(\mathcal{V}(f_i)) = \{(w_1, w_1) \mid w_1 \leq 1\}.$$

□

1.1.4 Tropical Semiring

1.2 Polyhedral Geometry

1.2.1 Introduction

In order to get a deeper appreciation of tropical geometry, we need to introduce some concepts from polyhedral geometry. Polyhedral geometry has many connections to seemingly unrelated fields, such as convex optimization and tropical geometry. In this section, we mainly follow the conventions set by Ziegler's text [15].

1.2.2 Polytopes

A set of points $X \subseteq \mathbb{R}^n$ is **convex** if it is closed under all line segments: if $\lambda x + (1-\lambda)y \in X$ for all $x, y \in X$ and $\lambda \in [0,1]$. If we have an arbitrary set X, we may consider the smallest convex set containing X, this is called the **convex hull** of X and is denoted $\operatorname{conv}(X)$. Note that this description of $\operatorname{conv}(X)$ also coincides with the intersection of all convex sets containing X:

$$\operatorname{conv}(X) = \bigcap_{\substack{X \subseteq U \\ U \text{ convex}}} U.$$

In the more common case where X is a finite set of points, say $X = \{x_1, \ldots, x_r\}$, the convex hull can be written as

$$P = \operatorname{conv}(X) = \left\{ \sum_{i=1}^{r} \lambda_i x_i \;\middle|\; \lambda_i \in [0,1] \,\forall i = 1, \ldots, r, \text{ and } \sum_{i=1}^{r} \lambda_i = 1 \right\}.$$

In this case, we say that $P = \operatorname{conv}(X)$ is a **polytope**. Each sum of the form $\sum_{i=1}^{r} \lambda_i x_i$ with $\lambda_i \in [0,1]$ and $\sum_{i=1}^{r} \lambda_i = 1$ is called a **convex combination**. Similarly to convex hulls, a **cone** over a set X is the **positive hull** of X; ie, the set of all positive combinations of

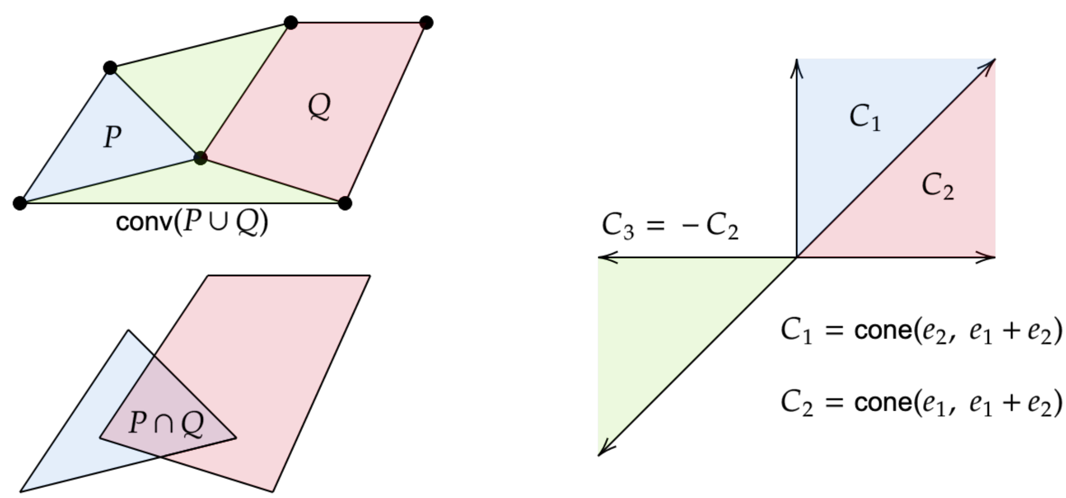

Figure 1.3: (L) *The union of convex sets may not be convex, but the convex hull of the union is. The intersection of convex sets is convex.* (R) *Three cones in* $\mathbb{R}^2$.

elements in X. When X is finite, the cone over X can be written as

$$C = \mathrm{cone}(X) = \left\{ \sum_{i=1}^{r} \lambda_i x_i \;\middle|\; \lambda_i > 0 \,\forall i = 1, \ldots, r \right\}.$$

In this case, we say that $C = \mathrm{cone}(X)$ is a **polyhedral cone**.

In this discussion, we mainly stick to the finite case for X; so "cones" should be assumed to be polyhedral cones unless otherwise specified. Furthermore, a polytope is always convex and bounded, while (polyhedral) cones are always convex and unbounded.

Every polytope P and polyhedral cone C is a collection of **faces** F of P, denoted $F \preceq P$. Each face is defined by a linear functional $w \in (\mathbb{R}^n)^\vee$ and we describe this relationship as the **face of P relative to w**:

$$\mathrm{face}_w(P) = \{x \in P \mid w(x - y) \leq 0 \,\forall y \in P\}$$

The faces of a cone are defined identically; with P replaced by C in the above definition.

There are two main ways to represent a polytope; either as its $\mathcal{V}$-representation (for *vertex*-representation) or its $\mathcal{H}$-representation (for *hyperplane*-representation). For the $\mathcal{V}$-representation, we view the polytope as the convex hull of its vertices. However, for the $\mathcal{H}$-representation, we view the polytope as the (bounded) intersection of its supporting closed half-spaces. These two representations coincide for any polytope [15]–so the word polytope is reserved for a bounded intersection of closed half-spaces, while **polyhedra** will refer to either a (polyhedral) cone or a (possibly unbounded) intersection of half-spaces.

If the linear functional w corresponds to an inward-pointing normal vector for the polyhedra P, then the faces of P relative to w have another representation

$$\mathrm{face}_w(P) = \{x \in P \mid w(x) = \min_{y \in P} w(y)\}$$

and the polyhedra can be represented as the region $P = \{x \in P \mid Ax \geq b\}$ where $A = [w_F]_{F \preceq P}$ is a matrix whose columns consist of inward-pointing normals and $b = [\min_{y\in P} w_F(y)]^T_{F\preceq P}$. In practice, the choice of whether to use the 'min' or 'max' notation for $\mathrm{face}_w(P)$ is used will only result in whether to include a negative sign; it will be clear from context which of the two choices we will make in each example. With this representation, the faces of P can be written as $F = \{x \in P \mid A'x = b'\}$ where $A'x \leq b'$ is some linear subsystem of $Ax \leq b$.

If instead w corresponds to an outward-pointing normal for P, then

$$\mathrm{face}_w(P) = \{x \in P \mid w(x) = \max_{y \in P} w(y)\}$$

and $P = \{x \in P \mid Ax \leq b\}$ where $A = [w_F]_{F\preceq P}$ is a matrix whose columns consist of outward-pointing normals and $b = [\max_{y\in P} w_F(y)]^T_{F\preceq P}$.

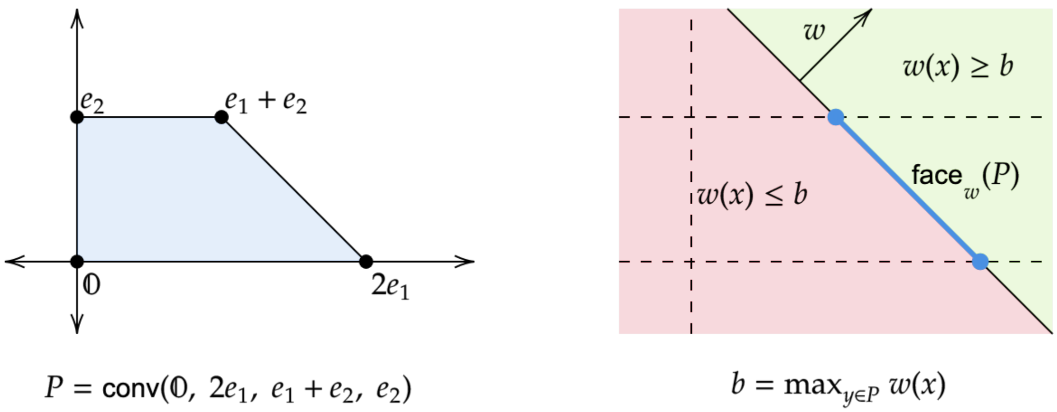

Figure 1.4: (L) *A convex hull P of vertices in $\mathbb{R}^2$.* (R) *A linear functional determines a face of P. (Shown in blue.)*

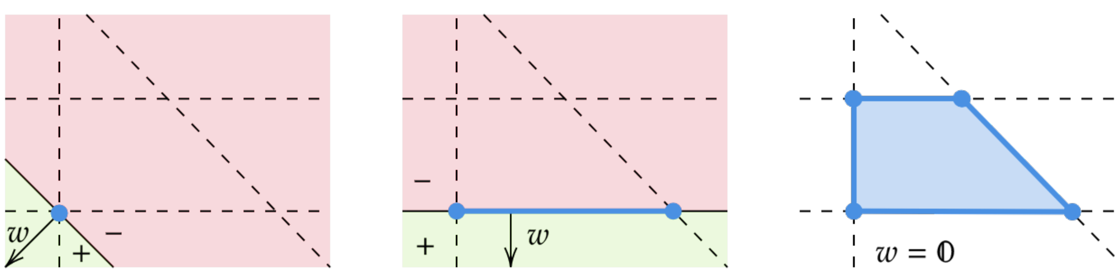

Figure 1.5: *Faces of a polytope P relative to different $w \in (\mathbb{R}^2)^\vee$.*

A **fan** Σ is a collection of cones such that

1. For any cone $C \in \Sigma$, if $F \preceq C$, then $F \in \Sigma$.

2. For all cones $C_1, C_2 \in \Sigma$, we have $C_1 \cap C_2 \preceq C_1$ and $C_1 \cap C_2 \preceq C_2$.

If we allow a fan Σ to contain more than just cones, we obtain a **polyhedral complex** Σ, which is a collection of polyhedra, such that

1. For any polyhedra $P \in \Sigma$, if $F \preceq P$, then $F \in \Sigma$.

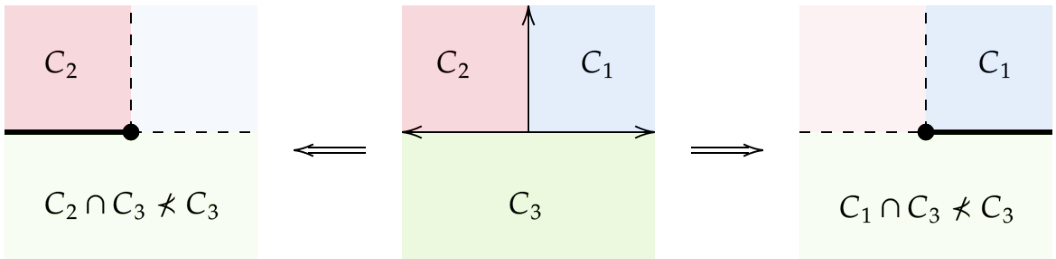

Figure 1.6: *A collection of 3 cones that do not form a fan.*

2. For all polyhedra $P_1, P_2 \in \Sigma$, we have either $P_1 \cap P_2 = \emptyset$ or $P_1 \cap P_2 \preceq P_1, P_2$.

If a face F is maximal with respect to the partial ordering $\preceq$ on Σ, then F is called a **facet**. A facet of a facet is called a **ridge**; in Figure 1.7, these are the black 1-dimensional line segments and rays. In Figure 1.7 the facets are the two-dimensional cells. The **support** of a complex Σ is the union of its elements; this is denoted by $\mathrm{supp}(\Sigma) := \bigcup_{P \in \Sigma} P$. Each of the three complexes in Figure 1.7 has a support equal to $\mathbb{R}^2$.

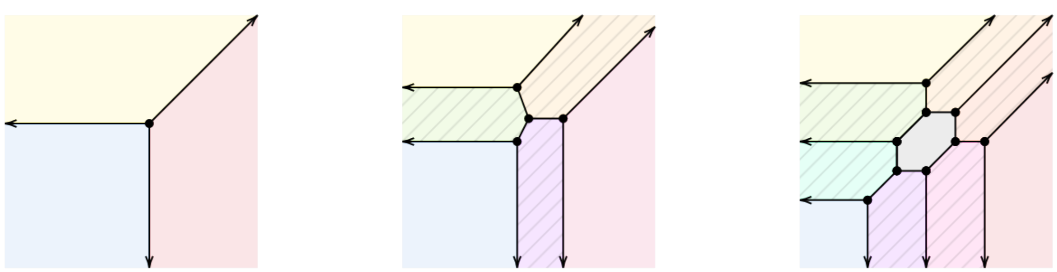

Figure 1.7: (L) *A polyhedral fan Σ_1 with 3 cones.* (M) *A polyhedral complex Σ_2 with 3 cones, 3 polyhedra (diagonal lines), and 0 polytopes.* (R) *A polyhedral complex Σ_3 with 3 cones, 6 polyhedra (diagonal lines), and 1 polytope.* Note: Only the facets are listed.

1.2.3 Subdivisions

In this section, we introduce subdivisions and relate them to Laurent polynomials. We follow the constructions from GKZ [8].

Given a finite set $A \subset \mathbb{R}^{k-1}$, we obtain a polytope $Q = \text{conv}(A)$; if A consists of the vertices of Q, then we say that the pair (Q, A) is a **marked polytope**. Given a marked polytope $Q \subset \mathbb{R}^{k-1}$ with $\dim(Q) = k - 1$, a **subdivision** of (Q, A) is a collection of marked polytopes $S = \{(Q_i, A_i) \mid i \in I,\, Q_i = \text{conv}(A_i)\}$ such that

1. Each $A_i \subset A$ and $\dim(Q_i) = k - 1$
2. $Q_i \cap Q_j \preceq Q_i, Q_j$ and $A_i \cap (Q_i \cap Q_j) = A_j \cap (Q_i \cap Q_j)$ for all $i, j \in I$
3. $\bigcup_{i \in I} Q_i = Q$

A particularly special subcase of subdivisions is that of triangulations. Given a marked polytope (Q, A), a **triangulation** of (Q, A) is a subdivision such that each Q_i is a simplex and the A_i are (still) vertices of Q_i. So, for example, if S is a subdivision of (Q, A) where there is some $i \in I$ such that A_i contains interior points, then S is not a triangulation.

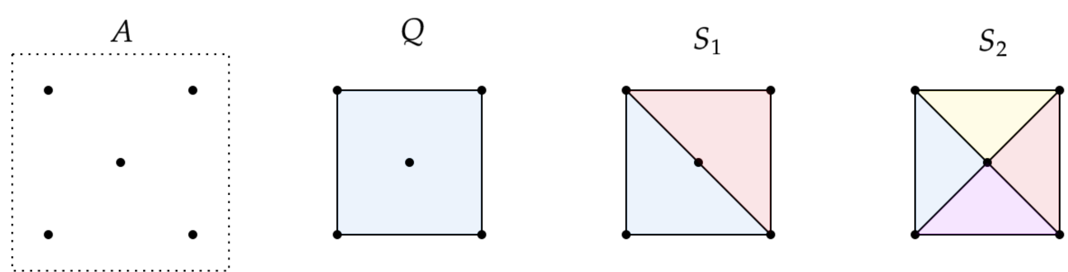

Figure 1.8: *A marked polytope (Q, A) with two subdivisions S_1, S_2, where S_1 is not a triangulation, but S_2 is. The faces $Q_i = \text{conv}(A_i)$ are the distinctly-colored simplices. The subdivisions are implied to be sets containing each of the shaded polytopes together with their vertices, along with Q.*

Next, we require a way to generate subdivisions by computing with linear functions; this leads us to the topic of regular/coherent subdivisions. In GKZ [8], these are historically referred to as **coherent subdivisions**, but are now also frequently called **regular subdivisions**—the terms are interchangeable. The following shows how to produce a regular subdivision S_ψ relative to a fixed linear function $\psi : \mathbb{R}^{k-1} \to \mathbb{R}$.

Given a marked polytope (Q, A), with $Q = \operatorname{conv}(A) \subset \mathbb{R}^{k-1}$, $\dim(Q) = k-1$, and a linear functional $\psi : \mathbb{R}^{k-1} \to \mathbb{R}$, we lift Q to $\mathbb{R}^{(k-1)+1}$ by adjoining an additional coordinate to each of the vertices in A and then use the k-tuples as the vertices for the boundary of a (unbounded) polyhedra; more precisely

$$Q_\psi := \operatorname{conv}\{(a,\, \psi(a)) \in \mathbb{R}^k \mid y \leq \psi(a),\, a \in A \subset \mathbb{R}^{k-1},\, y \in \mathbb{R}\}.$$

Next, we take the **upper hull**, or, **upper boundary** of the polyhedron Q_ψ by realizing the boundary as the graph G_ψ of the function $g_\psi(x) := \max_{(x,y)\in G_\psi}(y)$. With the upper hull constructed, we can then define the faces of the subdivision S_ψ as the **domains of linearity** of g_ψ, which are the subsets $Q_i \subset \mathbb{R}^{k-1}$ such that $g_\psi|_{Q_i}$ is linear. The domains of linearity also coincide with the projections of facets $F \preceq G_\psi \subset \mathbb{R}^k$ to $\pi(F) \subset \mathbb{R}^{k-1}$ where $\pi : \mathbb{R}^k \to \mathbb{R}^{k-1}$ is the projection map onto the first $k-1$ coordinates. And we define the vertices $a \in A_i \subset A$ as those such that $a \in A \cap Q_i$ and $g_\psi(a) = \psi(a)$. The resulting subdivision $S_\psi = \{(Q_i, A_i) \mid Q_i = \operatorname{conv}(A_i),\, i \in I\}$ is called a **regular subdivision induced by ψ.** If ψ is chosen sufficiently generically, then S_ψ is a triangulation [8].

Not all subdivisions are regular; a famous example is given by a projection of a capped pyramid in $\mathbb{R}^3$ to $\mathbb{R}^2$.

1.2.4 Duality

Recall that given a $\mathbb{R}$-vector space V, its **dual space** is the $\mathbb{R}$-vector space of linear functionals $\psi : V \to \mathbb{R}$, denoted $V^\vee := \hom(V, \mathbb{R})$. The map $\varphi : V \to V^\vee$ is given by $\varphi(\psi) = \varphi_\psi$ for any $\psi \in V$ and $\varphi_\psi(u) = \langle \psi, x\rangle$ for all $x \in V$. If V is finite-dimensional, then the Riesz Representation Theorem gives an isomorphism between V and $V^\vee$. What this means in

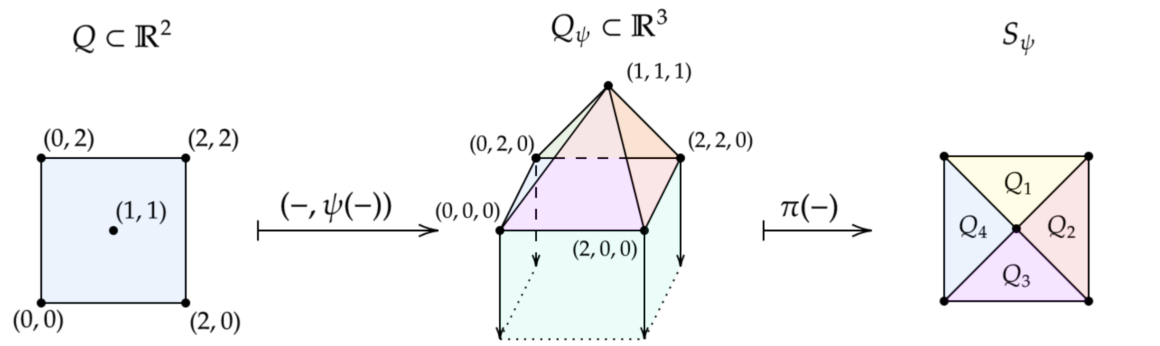

Figure 1.9: *A regular subdivision induced by a (nonlinear) convex function* $\psi(x,y) = -\frac{1}{2}\|(x,y)-(1,1)\|_{\ell_1} + 1$ *produces a triangulation* S_ψ.

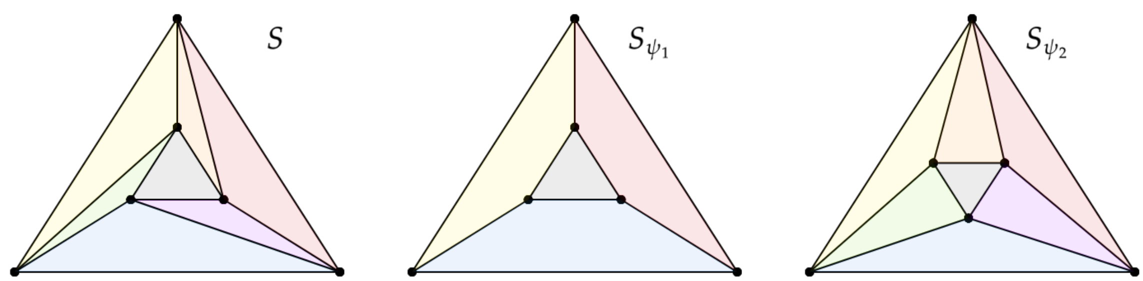

Figure 1.10: (L) *A subdivision* S *of a polytope that is not regular.* (M) *A non-generic choice of* ψ_1 *results in this subdivision.* (R) *Translating the inner vertices results in a polytope with regular subdivision* S_{ψ_2} *for generic* ψ_2.

particular is that there is a natural correspondence between points and linear functions on $\mathbb{R}^k$, which we can apply to polyhedra.

Given a polyhedra P, its faces $F \preceq P$ have the description

$$\mathrm{face}_w(P) = \{x \in P \mid w(x) = \min_{y \in P} w(y)\}$$

and we define the **inner normal fan** of P as the fan

$$\begin{aligned}\mathcal{F}_P^{\downarrow} &= \left\{\operatorname{cone}\left\{w \in (\mathbb{R}^d)^\vee \mid \operatorname{face}_w(P) = F\right\} \mid F \preceq P\right\} \\ &= \left\{\operatorname{cone}\left\{w \in (\mathbb{R}^d)^\vee \mid w|_F = \min w|_P\right\} \mid F \preceq P\right\}\end{aligned}$$

which consists of **inner normal cones** $\mathcal{N}_F^{\downarrow} := \operatorname{cone}\left\{w \in (\mathbb{R}^d)^\vee \mid w|_F = \min w|_P\right\}$ to the faces $F \preceq P$ generated by inward-pointing normal vectors $w \in (\mathbb{R}^d)^\vee$. It should be noted that the polytope P and its corresponding normal fan $\mathcal{F}_P$ can be drawn in the same plane via the isomorphism $\mathbb{R}^d \cong (\mathbb{R}^d)^\vee$, but one should always keep in mind that $P \subset \mathbb{R}^d$ and $\mathcal{F}_P \subset (\mathbb{R}^d)^\vee$ reside in different spaces, which can *then* be identified via this isomorphism.

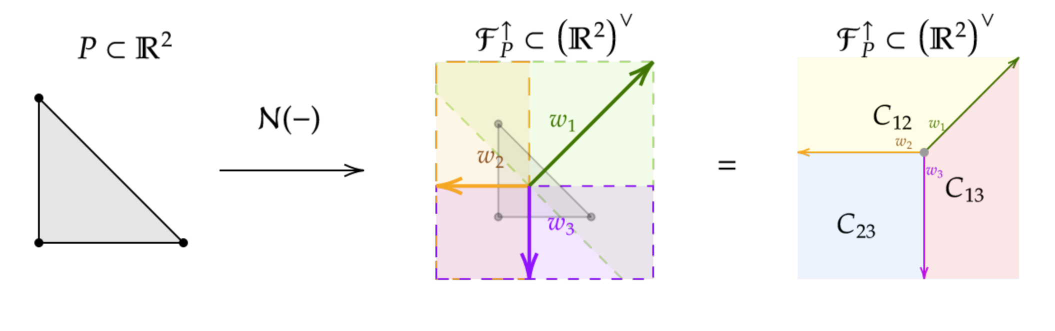

Figure 1.11: *The outer normal fan to a simplex.*

We similarly define the **outer normal fan** of P as

$$\begin{aligned}\mathcal{F}_P^{\uparrow} &= -\mathcal{F}_P^{\downarrow} \\ &= \left\{\operatorname{cone}\left\{-w \in (\mathbb{R}^d)^\vee \mid -w|_F = -\min w|_P\right\} \mid F \preceq P\right\} \\ &= \left\{\operatorname{cone}\left\{-w \in (\mathbb{R}^d)^\vee \mid -w|_F = \max(-w|_P)\right\} \mid F \preceq P\right\}\end{aligned}$$

and the **outer normal cones** as $\mathcal{N}_F^{\uparrow} = -\mathcal{N}_F^{\downarrow}$ where $\mathcal{F}_P^{\downarrow}$ and $\mathcal{N}_F^{\downarrow}$ denote the inner normal fan and inner normal cone; if the construction type is irrelevant or clear from context, we will omit the arrows. Therefore, in practice, the use of outer normals or inner normals is entirely up to personal preference, as the two are isomorphic via the map $C \mapsto -C$.

Next, we briefly describe the process for associating a dual complex to a polyhedral complex Σ composed of polyhedra $P \in \Sigma$. Each polyhedra has an associated dual fan $\mathcal{F}_P$ as described above, so all that is required is a way to glue the fans together to form a dual polyhedral complex. Suppose that $P, Q \in \Sigma$ share a facet $R \preceq P \cap Q$. Then in the dual space, $\mathcal{F}_P$ and $\mathcal{F}_Q$ will share generators $w \in \mathcal{F}_R$ but with opposite sign. Since P and Q are assumed to be (convex) polyhedra, this generator is unique, so we may glue $\mathcal{F}_P$ and $\mathcal{F}_Q$ by identifying $w \in \mathcal{F}_P$ and $-w \in \mathcal{F}_Q$. Proceeding in this manner for all facets $P \in \Sigma$, we obtain a polyhedral complex called the **dual complex** of Σ, denoted $\mathcal{F}_\Sigma$. Note that this construction does not indicate where the coordinates for the vertices of $\mathcal{F}_P$ are, but instead gives a complex that is combinatorially equivalent.

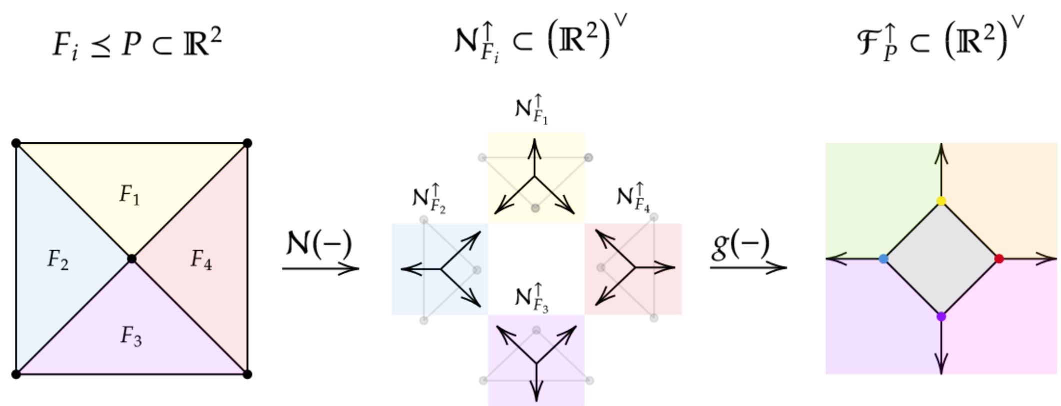

Figure 1.12: *The dual complex to a subdivision of a polytope formed by taking dual (outer) fans to facets via $\mathcal{N}(-)$, then gluing the generators with opposing signs via $g(-)$.*

1.2.5 Tropical Hypersurfaces

An important application of this in Tropical Geometry is the duality between a tropical hypersurface and regular subdivisions of Newton polytopes. For a Laurent polynomial

$f \in k[\vec{x}^{\pm}]$ represented as

$$f(x) = \sum_{u \in S \subset \mathbb{Z}^n} a_u x^u$$

over an algebraically-closed valued field k, where the vectors u are taken from a finite set of points $S \subset \mathbb{Z}^n$, and every coefficient $a_u \in k^*$ is nonzero, the **Newton polytope** of f is the convex polytope

$$\mathrm{Newt}(f) := \mathrm{conv}\{u \mid u \in \mathcal{A}\} \subset \mathbb{R}^n$$

and the **ψ-lifted Newton polytope** is the $(\dim \mathrm{Newt}(f) + 1)$-dimensional polytope created by appending an extra coordinate consisting of function values

$$\widetilde{\mathrm{Newt}}_\psi(f) := \mathrm{conv}\{(u, \psi(a_u)) \mid u \in \mathcal{A}, a_u \in k^*\} \subset \mathbb{R}^{n+1}\,.$$

It can be checked that $\widetilde{\mathrm{Newt}}_\psi \subset Q_\psi$, where Q_ψ is the GKZ analogue [8] construction of regular subdivisions, and that the lower faces of both coincide; so projecting the lower faces of either polytope back to $\mathbb{R}^n$ induces a ψ-regular subdivision S_ψ on the marked polytope $(\mathrm{Newt}(f), \mathcal{A})$. If $\psi = \mathrm{val}$, then the weights are the valuations of the coefficients $a_u \in S$, so we may obtain a val-induced regular subdivision of of the Newton polytope of f, S_{val}, by the process described earlier in Section 1.2.3. Following the duality procedure from Section **??**, we may obtain a dual complex Σ of S_{val}. This construction is summarized by the following theorem.

Theorem 1.2. *[10] Let $f \in k[\vec{x}^{\pm}]$ be a Laurent polynomial. The tropical hypersurface* $\mathrm{trop}(\mathcal{V}(f))$ *is the support of a pure Γ_{val}-rational polyhedra complex of dimension $n-1$ in $\mathbb{R}^n$. It is the $(n-1)$-skeleton of the polyhedral complex Σ dual to the* val*-regular subdivision of the Newton polytope of $f = \sum_{u \in \mathcal{A}} a_u x^u$ given by the weights* $\mathrm{val}(a_u)$ *on the lattice points in* $\mathrm{Newt}(f)$.

So, what this means is that we can effectively compute a tropical variety through duality, without having to solve each pairwise system of linear equations in $\mathrm{trop}(f)$. This polyhedral geometry approach is what we will use in applications of tropical geometry.

1.2.6 Secondary Polytopes & Lafforgue Polytopes

Next, we describe a few operations we can apply to our study of polyhedra.

Given two subsets $A, B \subset \mathbb{R}^n$, their **Minkowski sum** is defined as the set

$$A + B = \{a + b \in \mathbb{R}^n \mid a \in A,\, b \in B\}.$$

If A and B are polyhedra, then $A + B$ is also a polyhedra; in fact, if $A = \mathrm{conv}\{u_i\}_{i=1}^{m_1}$ and $B = \mathrm{conv}\{v_j\}_{j=1}^{m_2}$ are polytopes, then

$$A + B = \mathrm{conv}\{u_i + v_j \mid i \in [m_1],\, j \in [m_2]\}.$$

In fact, all polyhedra can be decomposed as a Minkowski sum of a polytope and cone (possibly $C = \{0\}$) [15]. Visually, a Minkowski sum can be viewed as 'smearing' the first factor by the second factor (or vice-versa).[15].

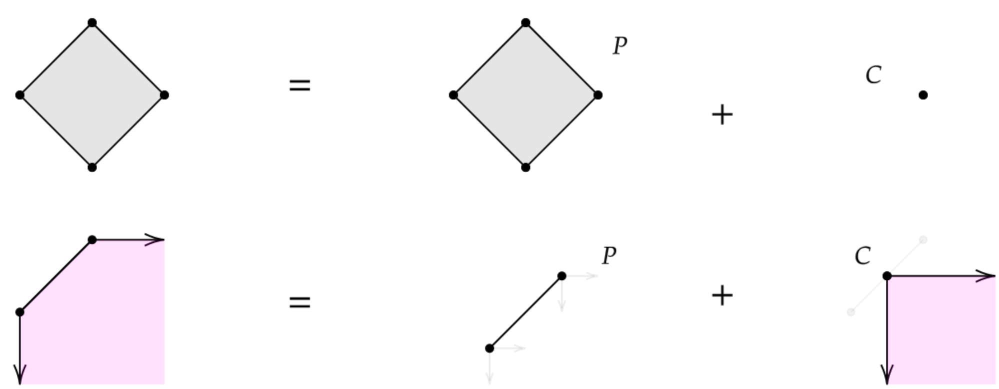

Figure 1.13: *Every polyhedron may be decomposed as a Minkowski sum of a polytope and a cone.*

When given two polyhedral complexes $\Sigma_1, \Sigma_2 \subset \mathbb{R}^n$ with a common support, their **com-**

mon refinement is the collection of polyhedra

$$\Sigma_1 \wedge \Sigma_2 := \{P_1 \cap P_2 \mid P_1 \in \Sigma_1,\ P_2 \in \Sigma_2\}.$$

When given two dual fans $\mathcal{F}_P, \mathcal{F}_Q$ for some polytopes $P, Q \subset \mathbb{R}^n$, it can be shown that [15]

$$\mathcal{F}_{P+Q} = \mathcal{F}_P \wedge \mathcal{F}_Q.$$

Furthermore, Minkowski sums behave well under Newton polytopes; for Laurent polynomials $f, g \in k[x_1^{\pm}, \ldots, x_m^{\pm}]$, we have

$$\mathrm{Newt}(f \cdot g) = \mathrm{Newt}(f) + \mathrm{Newt}(g).$$

It can be easily seen that the refinement defines a partial ordering on the set of regular subdivisions of a polytope, where $\Sigma_1 \preceq \Sigma_2$ if and only if $\Sigma_1 \wedge \Sigma_2 = \Sigma_2$. So for a marked polytope $(Q, \mathcal{A})$ where $Q = \mathrm{conv}(\mathcal{A})$ and $\mathcal{A}$ is the set of vertices of Q, the maximal elements of the poset under refinement are the triangulations of Q, while the minimal element is Q–potentially with fewer vertices in the case of interior points. This poset may then be identified with a certain polytope, called the **secondary polytope** $\Sigma(\mathcal{A})$; the details of this construction will be explained shortly.

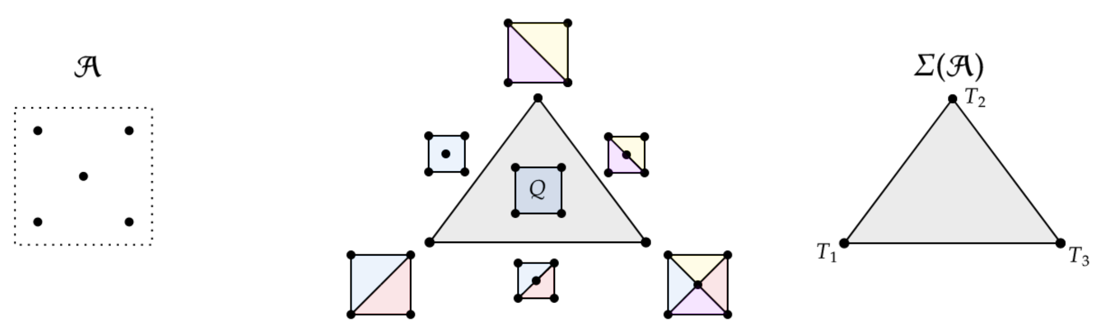

Figure 1.14: *A secondary polytope of a square with central vertex results in a triangle.*

But first, we recall the construction of the fiber polytopes from Billera and Sturmfels[4]. The fiber polytope is closely related to the secondary polytope, monotone path polytope, and other constructions. For our purposes, we require fiber polytopes in order to compute secondary polytopes, so we will include some background here. Fiber polytopes can be viewed as a generalization of previous work by Gelfand, Kapranov, and Zelevinsky[8] and provide a unifying framework for the study of polyhedra.

Suppose that $P \subset \mathbb{R}^n$ and $Q \subset \mathbb{R}^d$ are polytopes with $P = \operatorname{conv} \mathcal{A}_P = \operatorname{conv}\{p_1, \ldots, p_m\}$ and $Q = \operatorname{conv} \mathcal{A}_Q = \operatorname{conv}\{q_1, \ldots, q_m\}$ where $\mathcal{A}_P, \mathcal{A}_Q$ are possibly multisets. Also, suppose that $\pi : \mathbb{R}^n \to \mathbb{R}^d$ is an affine map with $\pi(p_i) = q_i \, \forall i = 1, \ldots, m$. A **projection of polytopes** is an affine projection $\pi : \mathbb{R}^n \to \mathbb{R}^d$ such that $\pi(P) = Q$; in this case, we may define a **projection bundle** $\mathcal{B} := \pi^{-1}$ assigning a point $x \in P$ to the polytope $\mathcal{B}_x = \pi^{-1}(x)$. The set $\Gamma(\mathcal{B})$ consists of all sections $\gamma : Q \to \mathbb{R}^n$ with $\gamma(x) \in \mathcal{B}_x$. The **Minkowski integral** $\int \mathcal{B} \subset \mathbb{R}^n$ of the polytope bundle $\mathcal{B}$ is defined as

$$\int \mathcal{B} = \int_Q \mathcal{B}_x \, dx = \left\{ \int_Q \gamma(x) \, dx \mid \gamma \in \Gamma(\mathcal{B}) \right\}$$

and can be thought of as the continuous analogue of a Minkowski sum of polytopes. With this, we can then define the **fiber polytope** of P over Q as the normalized Minkowski integral over the projection bundle $\mathcal{B} = \pi^{-1}$, denoted by

$$\Sigma(P, Q) = \frac{1}{\operatorname{vol}(Q)} \int \mathcal{B}.$$

However, this description is too computationally difficult to use, as a Minkowski integral is essentially an infinite Minkowski sum. The following result implies that a finite Minkowski sum of polyhedra is sufficient to calculate the dual fan to a polyhedra resulting from a Minkowski integral.

Theorem 1.3. *[4] Let $\mathcal{B}$ be a piecewise-linear polytope bundle. Then the Minkowski integral $\int \mathcal{B} \subset \mathbb{R}^n$ is then a (convex) polytope. Furthermore, there exists some finite subset of points* *(not necessarily vertices)* $\{x_1, \ldots, x_m\} \subset Q$ *such that* $\mathcal{F}_{\int \mathcal{B}} = \mathcal{F}_{\mathcal{B}_{x_1} + \cdots + \mathcal{B}_{x_m}}$*; or rather, that the polyhedra* $\int \mathcal{B}$ *and* $\mathcal{B}_{x_1} + \cdots + \mathcal{B}_{x_m}$ *have the same dual fans.*

The fiber polytope $\Sigma(P, Q)$ has dimension $\dim(P) - \dim(Q)$ and its face poset is isomorphic to the poset of coherent subdivisions of Q ordered under refinement.

Example 1.6. There are a few accessible examples using this description of fiber polytopes, though more will come later with a discrete representation.

1. If $P = Q$ and π is the identity map, then

$$\Sigma(P, Q) = \frac{1}{\mathrm{vol}(Q)} \int_Q x \, dx = x_Q$$

   results in the centroid of Q.

2. If $\pi : \mathbb{R}^{n+d} \to \mathbb{R}^d$ is the projection onto the last d coordinates, then $\Sigma(P \times Q, Q) = \pi^{-1}(x_Q)$ results in the centroid fiber of x_Q, so is just a translation of P.

3. If $Q = \{q\}$, then $\Sigma(P, Q) = P$.

□

It was also shown that fiber polytopes are a generalization of secondary polytopes via the following:

Theorem 1.4. *The secondary polytope is a scaled fiber polytope;*

$$\Sigma(Q) = \Sigma(\mathcal{A}_Q) = (d+1)\mathrm{vol}(Q)\, \Sigma(\Delta_{n-1}, Q).$$

However, the Minkowski integral description is insufficient for computations. Fortunately, Billera and Sturmfels showed that this Minkowski integral can be replaced by a finite Minkowski sum of polytopes. With Theorems 1.4 and 1.3, we get the following (much nicer) representation of secondary polytopes.

Theorem 1.5. *[4] Let $P = \operatorname{conv}(\mathcal{A}_P) \subset \mathbb{R}^n$ and $Q = \operatorname{conv}(\mathcal{A}_Q) \subset \mathbb{R}^d$ with $\mathcal{A}_P = \{p_1, \ldots, p_m\}$ and $\mathcal{A}_Q = \{q_1, \ldots, q_m\}$ potentially containing repeated vertices. Furthermore, let $\pi : \mathbb{R}^d \to \mathbb{R}^n$ be an affine map with $\pi(P) = Q$ and $\pi(p_i) = q_i$ $\forall i = 1, \ldots, m$. For any regular subdivision S of P and face $F \in S$ of P, denote the vertex set indexed by F by $V(F)$. Then*

$$\Sigma(P, Q) = \operatorname{conv}\{\Phi_\Delta \in \mathbb{R}^n \mid T \text{ is a triangulation of } Q\}$$

where the vectors Φ_T are given by

$$\Phi_T = \frac{1}{d+1} \sum_{F \in T} \left(\frac{\operatorname{vol}(F)}{\operatorname{vol}(Q)} \sum_{p_i \in V(F)} p_i \right) \in \mathbb{R}^n .$$

There are a few special cases of fiber polytopes that are worth mentioning briefly as particularly nice examples of fiber polytopes. In each example, the choice of P relative to Q essentially changes the combinatorics that the fiber polytope $\Sigma(P, Q)$ measures with respect to Q. Note that in figures 1.15, 1.16 the regular subdivisions use the 'min' convention for $\text{face}_\psi(F)$ and specify the proper face of F where ψ is minimized.

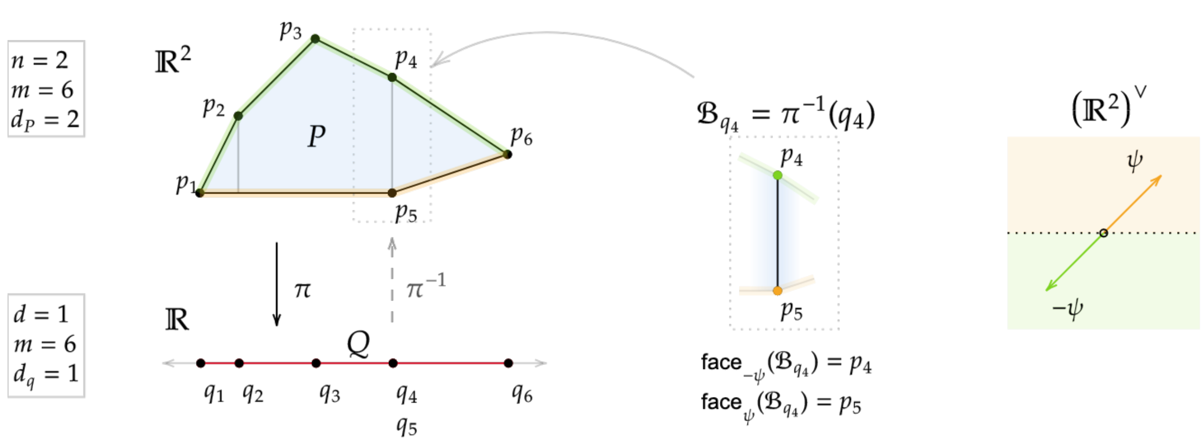

Figure 1.15: *The vertices of the fiber polytope $\Sigma(P, Q)$ correspond to the only two π-coherent subdivisions of Q. $\Sigma(P, Q)$ has dimension $\dim(P) - \dim(Q) = 1$, resides in $\mathbb{R}^m = \mathbb{R}^6$, and has two vertices. So $\Sigma(P, Q)$ is a line segment.*

Example 1.7. If $\dim(P) = \dim(Q) + 1$, then $\Sigma(P, Q)$ is a line segment, since only two π-coherent subdivisions of Q correspond to the vertices of the fiber polytope. Let $d_Q = \dim(Q)$ and embed P and Q into $\mathbb{R}^{d_Q+1}$ and $\mathbb{R}^{d_Q}$, respectively, so that their embeddings are full-dimensional. Without loss of generality, there is an induced projection map $\tilde{\pi} : \mathbb{R}^{d_Q+1} \to \mathbb{R}^{d_Q}$ onto the first d_Q coordinates with $\tilde{\pi}(i(P)) = i(Q)$. Label the last coordinate of each vertex p_i by ψ_i, so that the vertices of P are $p_i = (q_{d_Q}, \psi_i) \in \mathbb{R}^{d_Q+1}$. The vectors $\psi = (\psi_1, \dots, \psi_n)$ and $-\psi$ in $(\mathbb{R}^n)^\vee$ result in the two regular subdivisions of Q. This is illustrated in Figure 1.15 and is essentially a generalization of regular subdivisions. $\square$

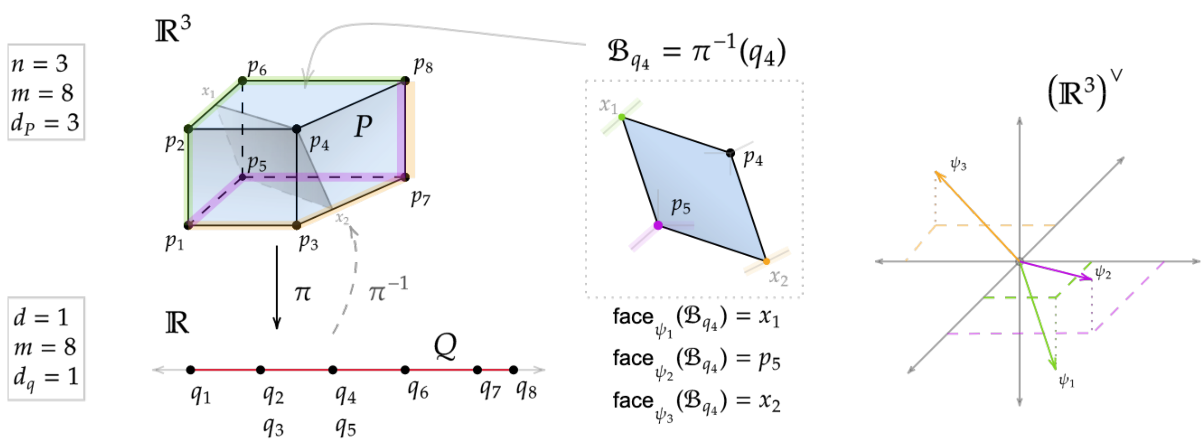

Figure 1.16: *The vertices of the fiber polytope* $\Sigma(P, Q) = \Sigma_\gamma(P)$ *correspond to* γ*-monotone paths on* P*; 3/6 paths are shown above.* $\Sigma(P, Q)$ *has dimension* $\dim(P) - \dim(Q) = 2$*, resides in* $\mathbb{R}^m = \mathbb{R}^8$*, and consists of 6 vertices. So* $\Sigma_\gamma(P)$ *is a hexagon.*

Example 1.8. If $\dim(Q) = 1$, then $\dim(\Sigma(P, Q)) = \dim(P) - 1$ and the π-coherent subdivisions of Q correspond to monotone edge paths induced by linear functionals on Q. Note that $\mathcal{A}_P$ and $\mathcal{A}_Q$ are both assumed to be multisets with m vertices, so $\mathcal{A}_Q$ will contain many repeats in this case; this can be avoided by adding additional parameters.

Assume that the vertices of P are indexed in nondecreasing order relative to $\pi \equiv \gamma \in (\mathbb{R}^n)^\vee$, so that $\gamma(p_i) \leq \gamma(p_{i+1})$ for $i = 1, \dots, m$. Any linear functional $\psi \in (\mathbb{R}^n)^\vee$ will then induce a regular subdivision on $Q = [q_1, q_m]$ by selecting the vertices of P in the direction $\psi + \alpha\gamma$ for some $\alpha \in \mathbb{R}$ via $p_{i_j} = \mathrm{face}_{\psi+\alpha\gamma}(\pi^{-1}(q_j)) \in V(P)$ for $j = 1, \dots, m$; the ordered list of vertices $V_\psi = (p_{i_1}, \ldots, p_{i_m})$ may then have repeated entries removed and this results in a **γ-monotone edge path** $\hat{V}_\psi$ of length $\hat{m} \leq m$.

As ψ ranges, the resulting γ-monotone edge paths on P can be identified with vertices of $\Sigma(P,Q) = \Sigma_\gamma(P)$ and has the description

$$\Sigma_\gamma(P) = \mathrm{conv}\{\Phi_\psi \mid \hat{V}_\psi \text{ is a } \gamma\text{-monotone edge path of length } \hat{m}\}$$

where the vertices Φ_ψ are computed by applying Theorem 1.5

$$\Phi_\psi = \sum_{j=1}^{\hat{m}} \frac{\langle \gamma, p_{i_j} - p_{i_{j-1}} \rangle}{2\,\langle \gamma, p_m - p_1 \rangle} (p_{i_{j-1}} + p_{i_j}).$$

This example is illustrated in Figure 1.16. □

Lastly, we introduce the dual counterparts of secondary and fiber polytopes; the dual fans will be of partially greater interest due to residing in functional spaces. The **secondary fan** of P is the dual fan of $\Sigma(\mathcal{A})$, denoted $\mathcal{F}_{\Sigma(\mathcal{A})}$. The secondary fan, in particular, parametrizes the functions that induce regular subdivisions on Q of a certain type, with the full-dimensional chambers representing the regular triangulations. This is in contrast with how the secondary polytope represents the (more important) regular triangulations with vertices (measure zero sets). Traversing the dual fan allows us to explore how regular triangulations change with respect to a chosen path.

Example 1.9. Let $\mathcal{A}$ consist of the five corners of a pentagon, so that $Q = \mathrm{conv}(\mathcal{A})$ is a pentagon. In coordinates, the vertices are

$$\mathcal{A} = \{a_1, a_2, a_3, a_4, a_5\} = \{(0,10), (-10,3), (-6,8), (6,8), (10,3)\}.$$

There are a total of 5 triangulations of Q, each of which is regular (this can be confirmed by drawing a picture). The five triangulations, with only their facets listed, with subscripts denoting vertices, are:

$$S_1 = \{Q_{\mathbf{1}23}, Q_{\mathbf{1}34},\ Q_{\mathbf{1}45}\}$$

$$S_2 = \{Q_{\mathbf{4}32},\ Q_{\mathbf{4}21},\ Q_{\mathbf{4}51}\} \qquad S_5 = \{Q_{\mathbf{3}12},\ Q_{\mathbf{3}51},\ Q_{\mathbf{3}45}\}$$

$$S_3 = \{Q_{\mathbf{2}51},\ Q_{\mathbf{2}45},\ Q_{\mathbf{2}34}\} \qquad S_4 = \{Q_{\mathbf{5}12},\ Q_{\mathbf{5}23},\ Q_{\mathbf{5}34}\}$$

By applying Theorem 1.5, along with Theorem 1.4 (with $d = 2, n = 5$) we get

$$\Phi_{S_1} = \frac{1}{2+1}\left(\frac{\mathrm{vol}(Q_{123})}{\mathrm{vol}(Q)}(e_1+e_2+e_3) + \frac{\mathrm{vol}(Q_{134})}{\mathrm{vol}(Q)}(e_1+e_3+e_4) + \frac{\mathrm{vol}(Q_{145})}{\mathrm{vol}(Q)}(e_1+e_4+e_5)\right)$$
$$\Phi_{S_2} = \frac{1}{2+1}\left(\frac{\mathrm{vol}(Q_{432})}{\mathrm{vol}(Q)}(e_1+e_2+e_3) + \frac{\mathrm{vol}(Q_{421})}{\mathrm{vol}(Q)}(e_4+e_2+e_1) + \frac{\mathrm{vol}(Q_{451})}{\mathrm{vol}(Q)}(e_4+e_5+e_1)\right)$$
$$\Phi_{S_3} = \frac{1}{2+1}\left(\frac{\mathrm{vol}(Q_{251})}{\mathrm{vol}(Q)}(e_2+e_5+e_1) + \frac{\mathrm{vol}(Q_{245})}{\mathrm{vol}(Q)}(e_2+e_4+e_5) + \frac{\mathrm{vol}(Q_{234})}{\mathrm{vol}(Q)}(e_2+e_3+e_4)\right)$$
$$\Phi_{S_4} = \frac{1}{2+1}\left(\frac{\mathrm{vol}(Q_{512})}{\mathrm{vol}(Q)}(e_5+e_1+e_2) + \frac{\mathrm{vol}(Q_{523})}{\mathrm{vol}(Q)}(e_5+e_2+e_3) + \frac{\mathrm{vol}(Q_{534})}{\mathrm{vol}(Q)}(e_5+e_3+e_4)\right)$$
$$\Phi_{S_5} = \frac{1}{2+1}\left(\frac{\mathrm{vol}(Q_{312})}{\mathrm{vol}(Q)}(e_3+e_1+e_2) + \frac{\mathrm{vol}(Q_{351})}{\mathrm{vol}(Q)}(e_3+e_5+e_1) + \frac{\mathrm{vol}(Q_{345})}{\mathrm{vol}(Q)}(e_3+e_4+e_5)\right)$$

After computing the volumes (in this case, areas), we get the vertices

$$\Phi_{S_1} = (0.32665971,\ 0.1031557,\ 0.22350401,\ 0.22350401,\ 0.1031557)$$

$$\Phi_{S_2} = (0.22449512,\ 0.22693529,\ 0.10492707,\ 0.32942219,\ 0.1024869)$$

$$\Phi_{S_3} = (0.10577776,\ 0.33175752,\ 0.10337372,\ 0.22597976,\ 0.2283838)$$

$$\Phi_{S_4} = (0.22449512,\ 0.1024869,\ 0.32942219,\ 0.10492707,\ 0.22693529)$$

$$\Phi_{S_5} = (0.10577776,\ 0.2283838,\ 0.22597976,\ 0.10337372,\ 0.33175752)$$

The secondary polytope is then $\Sigma(\mathcal{A}) = \mathrm{conv}(\Phi_{S_1},\ \Phi_{S_2},\ \Phi_{S_3},\ \Phi_{S_4},\ \Phi_{S_5})$ and has dimension

$$\dim(\Sigma(\mathcal{A})) = \dim(\Delta_4) - \dim(Q) = 4 - 2 = 2.$$

A 2-dimensional polytope with 5 vertices is a pentagon residing in $\mathbb{R}^5$, which can be embedded in $\mathbb{R}^2$ as shown in Figure 1.17. The vertices, edges, and convex hull of $\Sigma(A)$ correspond to regular triangulations, regular (non-triangulated) subdivisions, and the trivial subdivision, respectively.

Taking the dual fan of the projection of $\Sigma(A)$ results in the secondary fan $\mathcal{F}_{\Sigma(\mathcal{A})}$ with five polyhedral sectors (one for each vertex) and spans $(\mathbb{R}^2)^\vee$. Each sector in $\mathcal{F}_{\Sigma(\mathcal{A})}$ parameterizes the linear functionals ψ that induce a ψ-regular subdivision S_ψ. Alternatively, taking the dual fan of $\Sigma(A)$ in the ambient space $\mathbb{R}^5$ results in a dual fan with a 3-dimensional lineality space. If we take a quotient by this lineality space, then we get the same 2-dimensional dual fan $\mathcal{F}_{\Sigma(\mathcal{A})}$.

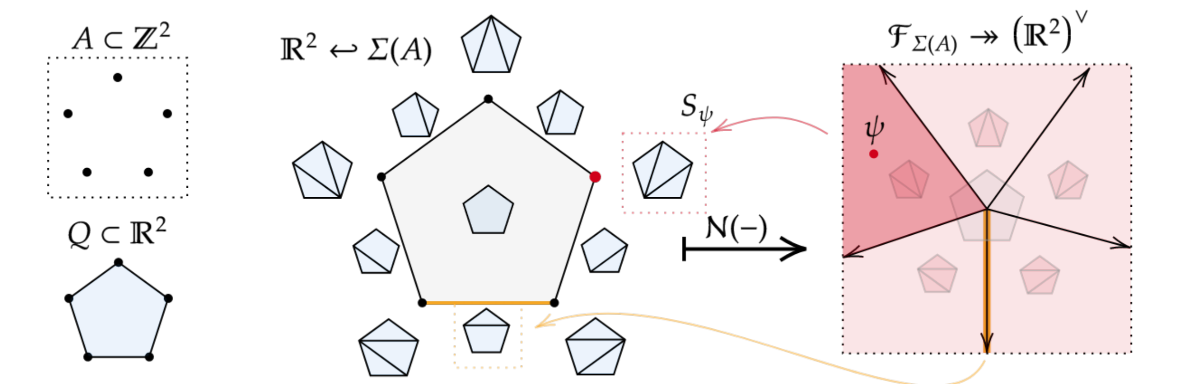

Figure 1.17: *The secondary polytope of the vertices of a hexagon.*

□

Theorem 1.4 was used to refine the secondary fan to a higher-dimensional class of fans related to fiber polytopes, which are called Lafforgue fans, named after the toric varieties associated with them. We will follow the definition in Proposition 2.3 of [6].

Let $\Delta^{\mathcal{A}} \subset \mathbb{R}^{\mathcal{A}}$ be the unit simplex, then the **Lafforgue polytope** is the Minkowski sum $\Theta(\mathcal{A}) := \Sigma(\mathcal{A}) + \Delta^{\mathcal{A}}$. The normal fan of $\Sigma(\mathcal{A})$ is then called the **Lafforgue fan**, denoted $\mathcal{F}_{\Theta(\mathcal{A})}$. The Lafforgue fan is a refinement of the secondary fan, in that each cone in $\mathcal{F}_{\Theta(A)}$ contains the dual functionals inducing certain coherent subdivisions, but now also stratifies them by subsets of points where the minimum is achieved.

The Lafforge polytope is a codimension-one polytope in $\mathbb{R}^{\mathcal{A}}$ and enumerates the same objects as the secondary polytope, while adding more information. In particular, vertices now correspond to triangulations with a marked point of the original configuration (T, a_i).

Furthermore, the Lafforgue fan is full-dimensional in $(\mathbb{R}^{\mathcal{A}})^{\vee}$ with a one-dimensional lineality space.

Example 1.10. Continuing with 1.9, we computed $\Sigma(\mathcal{A})$ as a pentagon; the secondary polytope vertex structure simply says 'there are five chambers in the dual fan, over which triangulations of a certain type occur.' Conversely, the Lafforgue polytope tells you, for every chamber in $\mathcal{F}_{\Sigma(\mathcal{A})}$, which of the five a_i's hit the wall first, which hits next, and how those choices chain together as one traverses a path through the chambers. If we want to compute the Lafforgue polytope explicitly, then we need to compute

$$\begin{aligned}
\Sigma(\mathcal{A}) + \Delta_4 &= \operatorname{conv}\{\Phi_{S_1}, \Phi_{S_2}, \Phi_{S_3}, \Phi_{S_4}, \Phi_{S_5}\} + \operatorname{conv}\{e_1, e_2, e_3, e_4, e_5\} \\
&= \operatorname{conv}\{\Phi_{S_i} + e_j \mid i, j \in [5]\} \\
&= \operatorname{conv}\{\Phi_{S_1} + e_1, \Phi_{S_2} + e_1, \Phi_{S_3} + e_1, \Phi_{S_4} + e_1, \Phi_{S_5} + e_1, \\
&\qquad \Phi_{S_1} + e_2, \Phi_{S_2} + e_2, \Phi_{S_3} + e_2, \Phi_{S_4} + e_2, \Phi_{S_5} + e_2, \\
&\qquad \Phi_{S_1} + e_3, \Phi_{S_2} + e_3, \Phi_{S_3} + e_3, \Phi_{S_4} + e_3, \Phi_{S_5} + e_3, \\
&\qquad \Phi_{S_1} + e_4, \Phi_{S_2} + e_4, \Phi_{S_3} + e_4, \Phi_{S_4} + e_4, \Phi_{S_5} + e_4, \\
&\qquad \Phi_{S_1} + e_5, \Phi_{S_2} + e_5, \Phi_{S_3} + e_5, \Phi_{S_4} + e_5, \Phi_{S_5} + e_5\}
\end{aligned}$$

In coordinates, the vertices for the convex hull are:

$$
\begin{aligned}
\{&(1.32665971, 0.1031557, 0.22350401, 0.22350401, 0.1031557),\\
&(1.22449512, 0.22693529, 0.10492707, 0.32942219, 0.1024869),\\
&(1.10577776, 0.33175752, 0.10337372, 0.22597976, 0.2283838),\\
&(1.22449512, 0.1024869, 0.32942219, 0.10492707, 0.22693529),\\
&(1.10577776, 0.2283838, 0.22597976, 0.10337372, 0.33175752),\\
&(0.32665971, 1.1031557, 0.22350401, 0.22350401, 0.1031557),\\
&(0.22449512, 1.22693529, 0.10492707, 0.32942219, 0.1024869),\\
&(0.10577776, 1.33175752, 0.10337372, 0.22597976, 0.2283838),\\
&(0.22449512, 1.1024869, 0.32942219, 0.10492707, 0.22693529),\\
&(0.10577776, 1.2283838, 0.22597976, 0.10337372, 0.33175752),\\
&(0.32665971, 0.1031557, 1.22350401, 0.22350401, 0.1031557),\\
&(0.22449512, 0.22693529, 1.10492707, 0.32942219, 0.1024869),\\
&(0.10577776, 0.33175752, 1.10337372, 0.22597976, 0.2283838),\\
&(0.22449512, 0.1024869, 1.32942219, 0.10492707, 0.22693529),\\
&(0.10577776, 0.2283838, 1.22597976, 0.10337372, 0.33175752),\\
&(0.32665971, 0.1031557, 0.22350401, 1.22350401, 0.1031557),\\
&(0.22449512, 0.22693529, 0.10492707, 1.32942219, 0.1024869),\\
&(0.10577776, 0.33175752, 0.10337372, 1.22597976, 0.2283838),\\
&(0.22449512, 0.1024869, 0.32942219, 1.10492707, 0.22693529),\\
&(0.10577776, 0.2283838, 0.22597976, 1.10337372, 0.33175752),\\
&(0.32665971, 0.1031557, 0.22350401, 0.22350401, 1.1031557),
\end{aligned}
$$

$$(0.22449512, 0.22693529, 0.10492707, 0.32942219, 1.1024869),$$

$$(0.10577776, 0.33175752, 0.10337372, 0.22597976, 1.2283838),$$

$$(0.22449512, 0.1024869, 0.32942219, 0.10492707, 1.22693529),$$

$$(0.10577776, 0.2283838, 0.22597976, 0.10337372, 1.33175752)\}$$

In this case, $\Theta(\mathcal{A})$ is 3-dimensional in $\mathbb{R}^5$ and has a total of $25 = 5 \cdot 5$ vertices—each triangulation of $\mathcal{A}$ was refined into 5 marked triangulations (T, a_i) for $a_i \in \mathcal{A}$. $\Theta(\mathcal{A})$ also has a total of 246 3-dimensional facets. Dually, the Lafforgue fan is a 5-dimensional fan in $(\mathbb{R}^5)^\vee$ with a 1-dimensional lineality space.

We also see that there are $25 + 246 - 2 = 269$ edges on the Lafforgue polytope. Edges of $\Theta(\mathcal{A})$ encode labelled flips, telling not only which diagonal in the coherent subdivision changes but also which vertex lies on the lower side of every height function that realises the flip, and the edges record how such flips commute. In other words, $\mathcal{F}_{\Theta(\mathcal{A})}$ strictly refines the secondary fan: cones that were previously identical split according to the subset of points at which the lifted height attains its minimum, capturing the order in which different vertices of $\mathcal{A}$ become coplanar under coherent subdivisions.

□

Chapter 2

Applications to Ising Systems

2.1 Classical Ising Problems

The Ising Model, named after physicists Ernst Ising and Wilhelm Lenz, is a mathematical model of ferromagnetism in statistical mechanics. The model consists of discrete variables $\sigma_1, \ldots, \sigma_d$ that represent magnetic dipole moments of atomic spins that can be in one of two states; $\sigma_i = 1$ or $\sigma_i = -1$. These spins are arranged in a graph (usually a one-, two-, or three-dimensional lattice), allowing each spin to interact with its neighbors. In most models, the interactions are limited to two-nearest-neighbor interactions, but more general models allow for more interactions. Ising systems have a tendency to evolve to the lowest energy state, but heat may prevent this from occurring.

2.1.1 Forward Ising Problem

The forward Ising problem refers to problems in which the graph has a predetermined configuration, but the lowest energy state is unknown. The forward Ising problem is heavily involved in designing adiabatic quantum optimization algorithms. More precisely, for the forward Ising problem, we are given

- a lattice Λ with d states (vertices) σ_i carrying values in ± 1

- a matrix $J \in \mathbb{R}^{d\times d}$ consisting of pairwise interactions (weighted edges) between states σ_i and σ_j
- a vector $h \in \mathbb{R}^d$ consisting of external forces on states/vertices σ_i

And with these, we can construct the **Hamiltonian function**

$$H_{J,h}(\sigma) := -\sum_{i<j\leq d} \sigma_i\sigma_j J_{ij} - \sum_{i=1}^{d} \sigma_i h_i$$

and the **energy function**

$$E_{J,h} := H_{J,h}(\sigma^*) = \min_{\sigma\in\{\pm 1\}^d} H_{J,h}(\sigma).$$

The goals of the forward Ising problem are to compute the minimizer for the Hamiltonian, this is called the **ground state** σ^*, along with the minimal value for the Hamiltonian function, called the **ground state energy** $E_{J,h}$ for the Hamiltonian $H_{J,h}(\sigma)$. It should be noted that in the care where the interactions have the same value, $J = \mathbf{1}$, the Hamiltonian becomes

$$H_{J,h}(\sigma) := -J \sum_{i<j\leq d} \sigma_i\sigma_j - \sum_{i=1}^{d} \sigma_i h_i$$

which results in the classical Ising model of ferromagnetism (or antiferromagnetism when $J = -\mathbf{1}$). In this case, the ground state consists of spin states that are all aligned, $\sigma = \pm\mathbf{1}$.

Example 2.1.

If for every pair of interactions $J_{\sigma_i\sigma_j} = 1$ or $J_{\sigma_i\sigma_j} = -1$ (i.e. if J represents ferromagnetic or antiferromagnetic interactions) then the induced subdivision is determined by the sign of the spin states. Uniformity of the couplings means

$$\langle J, \gamma(\sigma)\rangle = \pm \sum_{\sigma_i,\sigma_j\in E(\Gamma)} \sigma_i\sigma_j.$$

The resulting val-coherent subdivision is then completely determined by the structure of the Ising graph.

□

If J is assumed to be non-uniform, then we have the **Ising spin glass** model, which is more interesting mathematically due to the presence of **geometric frustration**. Geometric frustration occurs due to the presence of opposing interactions J_{ij}, J_{jk} between neighbors $\sigma_i, \sigma_j, \sigma_k$ on Γ, each corresponding to a spin state with the same energy level; these correspond to local minima of the Hamiltonian. In a spin glass, the interactions are a mixture of ferromagnetic and antiferromagnetic bonds; these patterns of aligned and misaligned states create **frustrated interactions**. The ground states for spin glass problems are more complex, and often there are more than two ground states for any given Ising model configuration as a consequence of geometric frustration. The ground states in a spin glass model tend to be stuck within local minima surrounded by configurations with much higher energy states, or **free energy barriers**, making the dynamics of iterative methods very challenging to work with and the complexity of the forward Ising spin glass problem notably difficult. An interesting characteristic of spin glass models is the presence of **degenerate** ground state energy levels which are ground state energy levels $E_{J,h}$ that occur at multiple spin state configurations.

2.1.2 Inverse Ising Problem

In contrast to the forward Ising problem, the inverse Ising problem involves working backwards from the spin states in order to determine the Ising graph parameters. More precisely, we are given the following.

- a lattice Λ with d states (vertices) σ_i carrying values in ± 1
- a set of m spin state configurations (vertex values) $\sigma^{(1)}, \ldots, \sigma^{(m)} \in \{\pm 1\}^d$ generated from some distribution

Compute the pairwise interactions $J^* \in \mathbb{R}^{d\times d}$ and external forces $h^* \in \mathbb{R}^d$ corresponding to the given spin state configurations $\{\sigma^{(i)}\}_{i=1}^m \subset \{\pm 1\}^d$

$$H_\sigma(J,h) := -\sum_{i<j\leq d} \sigma_i\sigma_j J_{ij} - \sum_{i=1}^{d} \sigma_i h_i$$

$$E(J,h) := H_{\sigma^*}(J^*,h^*) = \min_{\sigma\in\{\pm 1\}^d} H_\sigma(J,h).$$

Additional assumptions may also be made on the given parameters, just as in the forward Ising problem. For example, in the spin glass model, we have the presence of geometric frustration, which results in degenerate spin states, so in this case the set of spin states may be assumed to be degenerate ground states. Due to geometric frustration, there are multiple ground states that correspond to a minimal energy configuration.

2.2 An Tropical Perspective

We begin by recalling the field of Puiseaux series $K = \mathbb{C}\{\{t\}\}$ as the field consisting of all formal power series of the form

$$x(t) = \sum_{n=n_0}^{\infty} a_n t^{m_n} \in K$$

where $a_n \in \mathbb{C}^*$ and $m_n \in \mathbb{Q}\ \forall n \geq n_0 \in \mathbb{Z}$. Since $\mathbb{C}$ is algebraically closed, by a theorem of Serre [14], K is algebraically closed, so Theorem 1.1 holds. The valuation $\mathrm{val} : \mathbb{C}\{\{t\}\} \to \mathbb{R}$ determines the order of vanishing

$$\mathrm{val}(x(t)) = \begin{cases} \min\{m_n \mid n \geq n_0\} & x(t) \in K^* \\ \infty & x(t) = 0 \end{cases}.$$

We then use the field of Puiseaux series as the coefficient field for the Laurent polynomials $K[\vec{x}^{\pm}, \vec{y}^{\pm}] := K[x_1^{\pm 1}, \ldots, x_d^{\pm 1}, y_1^{\pm 1}, \ldots, y_e^{\pm 1}]$ defined as

$$f(x) = \sum_{(u,v) \in S \subsetneq \mathbb{Z}^{d+e}} a_{(u,v)} x^u y^v$$

where $u \in \mathbb{Z}^d, v \in \mathbb{Z}^e$ are taken from a finite set $S \subsetneq \mathbb{Z}^{d+e}$, and $a_{(u,v)} \in K$ is a Puiseux series. Then, there exist projection maps $\pi_x : y^{\vec{v}} \mapsto t^{\langle J,v \rangle}$ and $\pi_y : x^{\vec{u}} \mapsto t^{\langle h,u \rangle}$ such that the diagram

$$\begin{tikzcd}
K[\vec{x}^{\pm}, \vec{y}^{\pm}] \arrow[r, "\pi_y"] \arrow[d, "\pi_x"] & K[\vec{y}^{\pm}] \arrow[r, "\mathrm{trop}"] & \hom(\overline{\mathbb{R}}_{\mathrm{trop}}, \mathbb{R}) \arrow[d, "\mathrm{ev}_J"] \\
K[\vec{x}^{\pm}] \arrow[r, "\mathrm{trop}"] & \hom(\overline{\mathbb{R}}_{\mathrm{trop}}, \mathbb{R}) \arrow[r, "\mathrm{ev}_h"] & \mathbb{R}
\end{tikzcd}$$

commutes for every choice of parameters $J \in \mathbb{R}^e$, $h \in \mathbb{R}^d$.

For the sake of genericity, we consider Ising systems over a generic graph, though this work can also be considered over lattices. Let Γ be a simple weighted graph with $\#E(\Gamma) = e$ edges and $\#V(\Gamma) = d$ vertices. Set $\mathcal{A} = \{\sigma^{(i)}\}_{i=1}^{2^d} = \{\pm 1\}^d \subset \mathbb{Z}^d$ to be the set of all binary sequences of ± 1 with length d so that $Q = \mathrm{conv}(\mathcal{A}) \subset \mathbb{R}^d$ is the d-cube. Note that there are 2^d states $\sigma \in \mathcal{A}$, each with d-coordinates σ_i specifying weights for each vertex in $V(\Gamma)$.

Let $J \in \mathbb{R}^{d \times d}$ be the interaction matrix and let $h \in \mathbb{R}^d$ be the external force vector from the Ising model. Since Γ is a simple undirected graph, we may assume without loss of generality that J can be embedded in the much lower dimensional space $\mathbb{R}^e$. Define the **Ising polynomial** of Γ as the Laurent polynomial

$$\mathcal{E}(\vec{x}, \vec{y}) := \sum_{\sigma \in \mathcal{A}} x^{-\sigma} y^{-\gamma(\sigma)}, \quad \gamma(\sigma) := (\sigma_i \sigma_j)_{(\sigma_i, \sigma_j) \in E(\Gamma)} \in \mathbb{R}^e$$

where each monomial in $\mathcal{E}$ can be seen to represent a spin state configuration $\sigma \in \mathcal{A}$, and each $\gamma(\sigma)$ is the edge parity vector of $\mathcal{E}(\Gamma)$. The x^{σ} terms represent the external forces for the Ising system, while the $y^{\gamma(\sigma)}$ terms represent the state interactions of the Ising system.

Define the image of $\mathcal{E}$ under the projection maps as the pair of Laurent polynomials

$$\mathcal{E}_J(\vec{x}) := \pi_y(\mathcal{E}(\vec{x},\vec{y})) = \sum_{\sigma\in\mathcal{A}} t^{-\langle J,\gamma(\sigma)\rangle} x^{-\sigma} \in K[\vec{x}^{\pm}]$$
$$\mathcal{E}_h(\vec{y}) := \pi_x(\mathcal{E}(\vec{x},\vec{y})) = \sum_{\sigma\in\mathcal{A}} t^{-\langle h,\sigma\rangle} y^{-\gamma(\sigma)} \in K[\vec{y}^{\pm}]$$

as the **field polynomial** and **coupling polynomial**, respectively. Note that under the tropicalization map, both the interaction and the field polynomials are mapped to the **field energy** $E_J(h)$ and **coupling energy** $E_h(J)$ tropical functions which both evaluate to the **ground state energy** $E_{J,h}$ with respect to $J \in \mathbb{R}^e$ and $h \in \mathbb{R}^d$

$$\begin{aligned}
\mathrm{trop}_{\mathrm{val}}(\mathcal{E}_J(x)) &= \min_{\sigma\in\mathcal{A}} \left(-\langle J,\gamma(\sigma)\rangle - \langle \sigma,\vec{x}\rangle\right)\\
E_J(x) &= \min_{\sigma\in\mathcal{A}} \left(-\sum_{(\sigma_i,\sigma_j)\in E(\Gamma)} \sigma_i\sigma_j J_{ij} - \sum_{i=1}^{d} \sigma_i x_i\right)\\
\Rightarrow E_J(h) &= \min_{\sigma\in\mathcal{A}}(H_{J,\sigma}(h))\\
\Rightarrow E_{J,h} &= \min_{\sigma\in\mathcal{A}}(H_{J,h,\sigma})\\
\Rightarrow E_h(J) &= \min_{\sigma\in\mathcal{A}}(H_{h,\sigma}(J))\\
E_h(y) &= \min_{\sigma\in\mathcal{A}} \left(-\sum_{(\sigma_i,\sigma_j)\in E(\Gamma)} \sigma_i\sigma_j y_{ij} - \sum_{i=1}^{d} \sigma_i h_i\right)\\
\mathrm{trop}_{\mathrm{val}}(\mathcal{E}_h(x)) &= \min_{\sigma\in\mathcal{A}} \left(-\langle h,\sigma\rangle) - \langle \gamma(\sigma),\vec{y}\rangle\right)
\end{aligned}$$

where the middle equalities result from the evaluation maps.

By construction, the energy functions E_J, E_h are tropical functions, so they have a locus of nonlinearity $\mathcal{V}(E_J), \mathcal{V}(E_h)$. These tropical varieties consist of points where multiple spin states achieve the ground state energy, so we call these tropical varieties the **degeneracy loci** of $E(J,h)$. We also have the following conjecture.

The field and interaction polynomials determine the tropical varieties in different ambient spaces, so $\mathbb{R}^e$ is the **J-space** and $\mathbb{R}^d$ is the **h-space**.

Since each field polynomial $\mathcal{E}_J(h)$ induces a val-coherent subdivision with respect to a choice of J, the interaction matrix parameterizes the coherent subdivisions of the unit cube $Q = \mathrm{conv}(\mathcal{A}) = [-1,1]^d$. By Theorem 1.2, each coherent subdivision is dual to the tropical hypersurface $\mathcal{V}(E_J)$, so the field degeneracy locus $\mathcal{V}(E_J) \subset \mathbb{R}^d$ has its combinatorial type determined by the corresponding face $F_J \preceq \Sigma(\mathcal{A})$ of the secondary polytope $\Sigma(\mathcal{A})$. But since the regular triangulations correspond to (measure-zero) vertices, we want to enlarge the moduli space for the field degeneracy locus by first projecting $\Sigma(\mathcal{A})$ to $\mathbb{R}^{2^d-d-1}$, then taking the dual fan in order to obtain the secondary fan $\mathcal{F}_{\Sigma(\mathcal{A})}$. So we may now view the collection of field degeneracy loci as naturally residing inside of the moduli space $\mathcal{F}_{\Sigma(\mathcal{A})}$. Thus, specifying a choice for J fixes the combinatorial type of the degeneracy locus. This is summarized in the following result.

Theorem 2.1. *Let Γ be a simple undirected weighted graph on d vertices and e edges with $J \in \mathbb{R}^e$ representing the spin-state interactions and $h \in \mathbb{R}^d$ representing the external forces on the spin-states. Let $\mathcal{A} = \{\pm 1\}^d \subset \mathbb{Z}^d$ be the set of vertices of the d-cube $Q = \mathrm{conv}(\mathcal{A})$. Then the cones of the secondary fan $\mathcal{F}_{\Sigma(\mathcal{A})}$ parameterize the space of field degeneracy loci in terms of the couplings J. This moduli space has dimension 2^d, with a lineality space of dimension $d+1$ and is dual to the* val*-induced coherent subdivision of the field polynomial $\mathcal{E}_J$.*

Proof. The discussion above covers the parameterization of the space of field degeneracy loci $\{\mathcal{V}(E_J) \mid J \in \mathbb{R}^e\} \cong \mathbb{R}^d \times \mathbb{R}^e$. The only part that requires justification is the dimension count. By Theorem 1.4, we can represent the secondary polytope as a fiber polytope $\Sigma(\mathcal{A}) = (d+1)\,\mathrm{vol}(Q)\Sigma(\Delta_{2^d-1}, Q)$. By the discussion in Theorem 1.5 in [8], we have that

$$\begin{aligned}\dim \Sigma(\mathcal{A}) &= \dim (d+1)\,\mathrm{vol}(Q)\Sigma(\Delta_{2^d}, Q)\\ &= \dim \Sigma(\Delta_{2^d}, Q)\\ &= \dim \Delta_{2^d} - \dim Q\\ &= 2^d - 1 - d.\end{aligned}$$

Since the dual fan is complete, $\mathcal{F}_{\Sigma(\mathcal{A})}$ is 2^d-dimensional with $(d+1)$-dimensional lineality space. $\square$

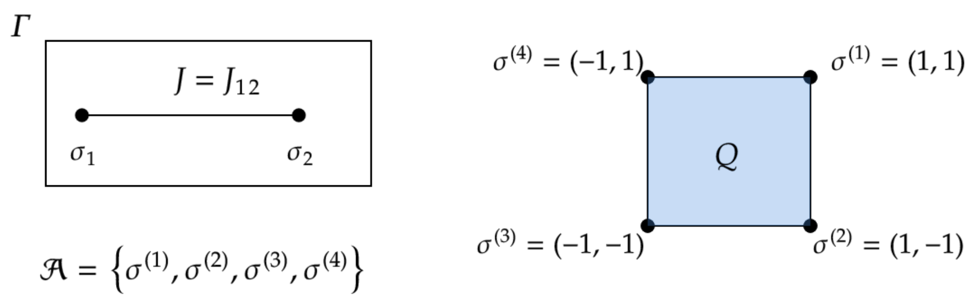

Figure 2.1: *The case when $d = 2$ and $\Gamma = P_1$ results in the following polyhedra Q.*

Example 2.2. Let Γ be the path graph on two vertices. Then $\sigma = (\sigma_1, \sigma_2) \subset \{\pm 1\}^2$, $J \in \mathbb{R}^1$, $h = (h_1, h_2) \in \mathbb{R}^2$. The coupling polynomial $\mathcal{E}_J \in K[x]$ is

$$\begin{aligned}\mathcal{E}_J(\vec{x}) &= \sum_{\sigma \in \mathcal{A}} t^{\langle J, \gamma(\sigma)\rangle} x^{\sigma} \\ &= t^{-J(1)(1)} x^{-(1,1)} + t^{-J(1)(-1)} x^{-(1,-1)} + t^{-J(-1)(-1)} x^{-(-1,-1)} + t^{-J(-1)(1)} x^{-(-1,1)} \\ &= t^{-J} x^{-1} y^{-1} + t^{J} x^{-1} y + t^{-J} xy + t^{J} xy^{-1}\end{aligned}$$

A choice of J completely determines the weights $\langle J, \gamma(\sigma)\rangle$ and, consequently, the val-coherent subdivision.

Tropicalizing $\mathcal{E}_J$ gives the field energy function

$$\begin{aligned}E_J(h) &= \min_{\sigma \in \mathcal{A}} (H_\sigma(J, h)) \\ &= \min\left(-(h_1 + h_2 + J),\, -(h_1 - h_2 - J),\, -(-h_1 - h_2 + J),\, -(-h_1 + h_2 - J)\right)\end{aligned}$$

$\square$

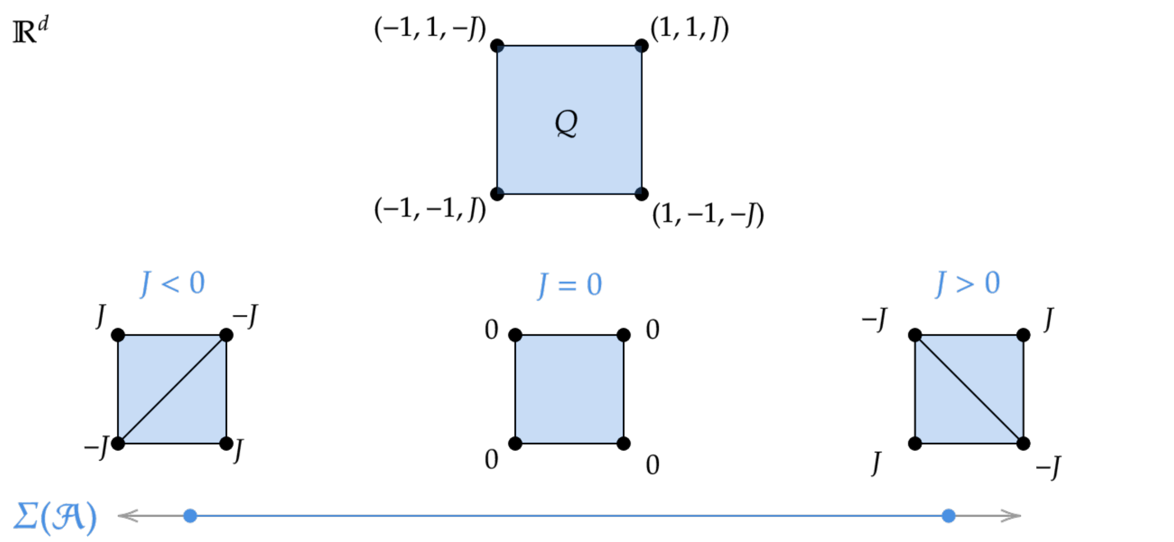

Figure 2.2: *The* val-*induced subdivisions of Q for various choices of faces of $\Sigma(\mathcal{A})$ are pictured.*

Example 2.3. Let Γ be the 2-regular graph on three vertices. Then $\sigma = (\sigma_1, \sigma_2, \sigma_3) \subset \{\pm 1\}^3$, $J \in \mathbb{R}^{3\times 3}$, $h = (h_1, h_2, h_3) \in \mathbb{R}^3$. The interaction polynomial $\mathcal{E}_J$ has 2^3 monomials; one for each vertex on the 3-cube.

Let Γ be the 2-regular graph on three vertices. Then $\sigma = (\sigma_1, \sigma_2, \sigma_3) \in \{\pm 1\}^3$, $J \in \mathbb{R}^{3\times 3}$, $h = (h_1, h_2, h_3) \in \mathbb{R}^3$. The interaction polynomial $\mathcal{E}_J$ has 2^3 monomials; one for each vertex on the 3-cube.

$$
\begin{aligned}
\mathcal{E}_J(\vec{x}) &= \sum_{\sigma \in \mathcal{A}} t^{\gamma_J(\sigma)} x^{\vec{\sigma}} \\
&= t^{-\gamma_J(1,1,1)} x^{-(1,1,1)} + t^{-\gamma_J(1,-1,1)} x^{-(1,-1,1)} + t^{-\gamma_J(-1,-1,1)} x^{-(-1,-1,1)} \\
&\quad + t^{-\gamma_J(-1,1,1)} x^{-(-1,1,1)} + t^{-\gamma_J(1,1,-1)} x^{-(1,1,-1)} + t^{-\gamma_J(1,-1,-1)} x^{-(1,-1,-1)} \\
&\quad + t^{-\gamma_J(-1,-1,-1)} x^{-(-1,-1,-1)} + t^{-\gamma_J(-1,1,-1)} x^{-(-1,1,-1)}
\end{aligned}
$$

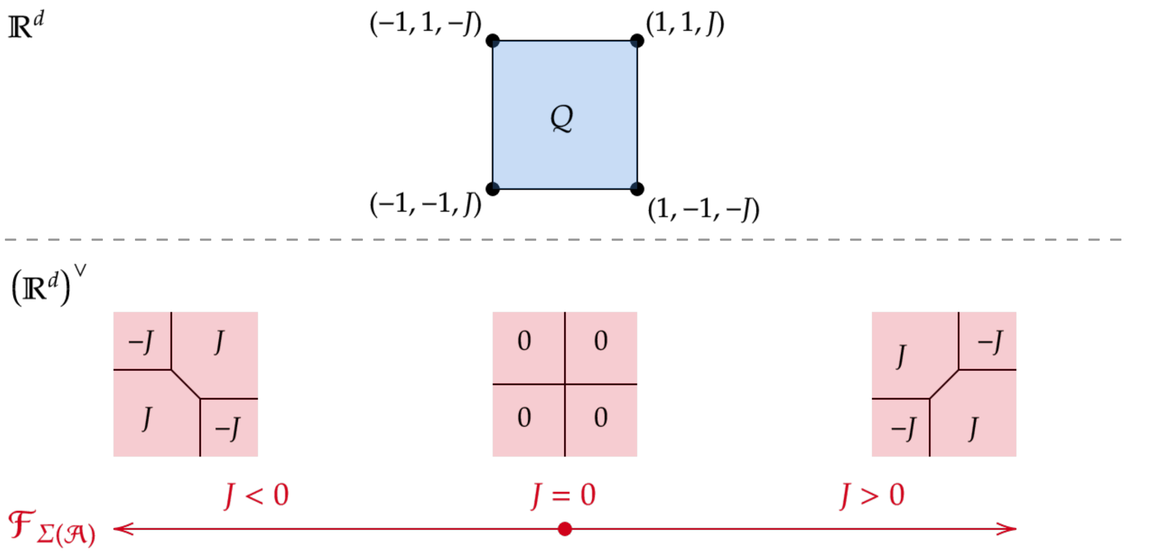

Figure 2.3: *The* val*-induced subdivisions of Q for various choices of faces of $\mathcal{F}_{\Sigma(\mathcal{A})}$ are pictured.*

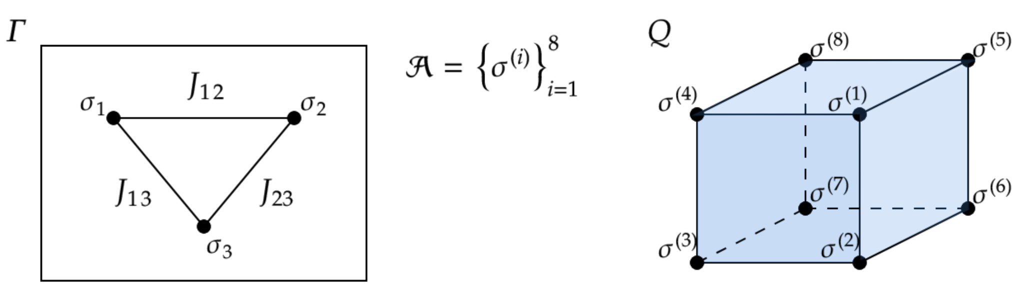

Figure 2.4: *The case when $d = 3$ and $\Gamma = C_3$ results in the following polyhedra Q.*

The tropicalization of $\mathcal{E}_J$ gives

$$
\begin{aligned}
E_J(h) &= \min_{\sigma\in\mathcal{A}}(H_\sigma(J,h)) \\
&= \min(\; -(J_{12}+J_{23}+J_{13}+h_1+h_2+h_3), \quad -(J_{12}-J_{23}+J_{13}-h_1+h_2+h_3), \\
&\qquad -(J_{12}-J_{23}-J_{13}+h_1-h_2+h_3), \quad -(J_{12}-J_{23}-J_{13}+h_1+h_2-h_3), \\
&\qquad -(J_{12}-J_{23}-J_{13}-h_1-h_2+h_3), \quad -(-J_{12}-J_{23}+J_{13}-h_1+h_2-h_3), \\
&\qquad -(-J_{12}+J_{23}-J_{13}+h_1-h_2-h_3), \quad -(J_{12}+J_{23}+J_{13}-h_1-h_2-h_3)\;)
\end{aligned}
$$

However, the secondary fan for the 3-vertex case has $\dim \Sigma(\mathcal{A}) = 2^3 - 3 - 1 = 4$ and is embedded in $\mathbb{R}^8$. So even the projection to $\mathbb{R}^4$ cannot be drawn; only the resulting slices in the h-space can be visualized since $\mathcal{V}(E_J) \subset \mathbb{R}^3$. A computer calculation shows that $\Sigma(\mathcal{A})$ has a total of 74 triangulations, of which all 74 are regular triangulations. Another lengthy calculation can show that all but 6 of these can be generated by rotations of other regular triangulations.

Taking the slice at $J = \mathbf{0}$ gives a degenerate subdivision (at a wall) of $(Q, \mathcal{A})$ and the tropical variety below.

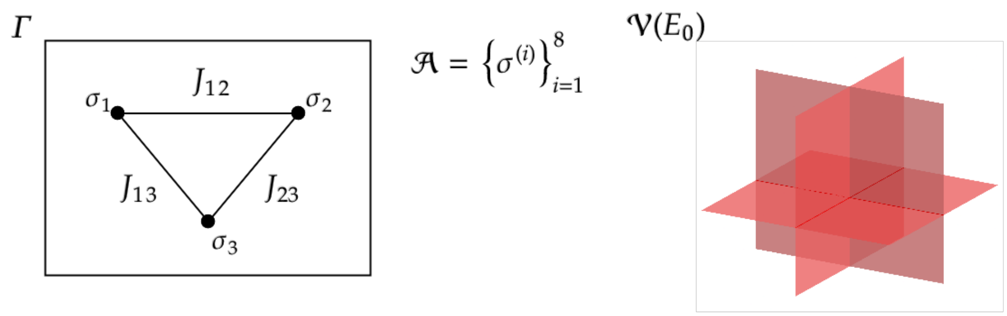

Figure 2.5: *The case when $d = 2$ and $\Gamma = P_1$ results in the following polyhedra Q. A degenerate tropical variety is pictured on the right.*

□

2.2.1 Computational Analysis

On the topic of enumerating the number of regular triangulations $\#R$ of the n-cube, Billera, Filliman, and Sturmfels [3] came up with an algorithm that computes $\#R$ by first computing the secondary fan through gale transforms, then applying a hyperplane arrangement algorithm from [7]; this approach has complexity $\mathcal{O}(n^{(d+1)(n-d-2)})$. De Loera [5] then came up with a breadth-first search enumeration algorithm, but this had complexity $\Omega(\#R)$ simply due to it being a breadth-first search method. Maseda [11] approached it from a reverse-search method that was adapted from an approach from [1, 2]; this resulted in a more efficient algorithm that was then improved on by [9] to get an algorithm with complexity $\mathcal{O}\left(d^2s^2\mathrm{LP}(n-d-1,s)\,\#R\right)$ and memory $\mathcal{O}(ds)$, where s is the upper bound of the number of simplices of all dimensions contained in a regular triangulation and $\mathrm{LP}(n-d-1,s)$ is the time required to solve a linear programming problem in s strict inequalities and $n-d-1$ variables.

The main idea used in [9] is that the vertices of the secondary polytope are connected by edges which represent edge-flips, or, bistellar flips. In short, traversing edges on the secondary polytope is almost equivalent to swapping edges; it is important to note that although the flip graph can always be viewed as a subset of vertices and edges of $\Sigma(\mathcal{A})$, they are not isomorphic [5]. In dimension two, the graph of triangulations of every point configuration is connected [5]. However, in dimensions 5 and higher, the statement is false. The relevant counterexamples were constructed by Santos [12, 13]. In 5 dimensions, the point configuration has 17 points in general position, and in 6 dimensions, the point configuration has 324 points and its triangulation is an isolated vertex in the flip graph of the point configuration. The situation in dimensions 3 and 4 is still unknown.

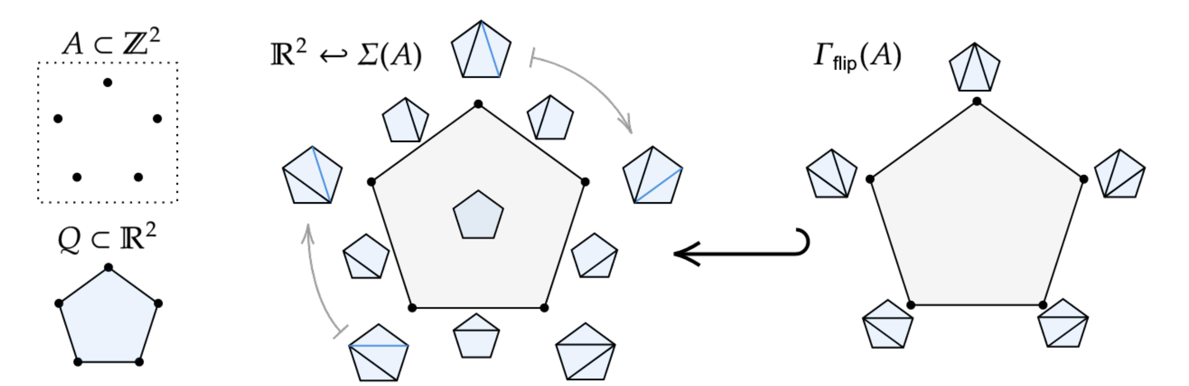

Figure 2.6: *The flip graph of a point configuration $\Gamma_{flip}(A)$ can always be embedded in 2-skeleton of the secondary polytope $\Sigma(A)$. The converse is false and the lowest-dimensional counterexample currently known is 5-dimensional.*

Possibly more concerning is the growth rate of the number of regular triangulations of the

d-cube. Using the algorithm from [9], one can show that the 3-cube has 74 triangulations, of which all are regular and can be grouped into 6 equivalence classes modulo rotational symmetries of S_4 acting on the 3-cube. Similarly, the 4-cube has 92,487,256 triangulations, of which 87,959,448 are regular and can be grouped into 235,277 modulo S_8 equivalence classes. Recall that the 3-cube and 4-cube correspond to Ising graphs on 3 and 4 vertices, respectively, are the smallest non-trivial examples of this tropical approach. So, one would first need to get a handle on how to utilize the S_{2^d} symmetries to simplify computing the secondary polytope/fan in order for this approach to remain somewhat feasible for small Ising system.

These computational observations are what led us to explore alternative methods for determining the coherent subdivisions of a point configuration.

2.2.2 Reducing Complexity via Fiber Polytopes

When a polytope $Q \subset \mathbb{R}^n$ is not acted on directly, but rather appears as the image of a higher-dimensional polytope $P \xrightarrow{\pi} Q$, the regular subdivisions of Q that can actually occur are constrained: a height function $\psi : P \to \mathbb{R}$ is coherent on Q only if it is constant along every fibre $\pi^{-1}(q)$. Billera–Sturmfels capture precisely those fibre-wise coherent height functions in the fiber polytope as previously defined

$$\Sigma(P, Q) = \frac{1}{\mathrm{vol}(Q)} \int_{q \in Q} \pi^{-1}(q)\, dq \subset \mathbb{R}^m,$$

whose vertices correspond bijectively to the π-coherent subdivisions of Q. Computing $\Sigma(P, Q)$ therefore only produces the triangulations that can be realised by lifting to P, so one never has to enumerate the exponentially larger secondary polytope of Q and discard the incompatible subdivisions after the fact.

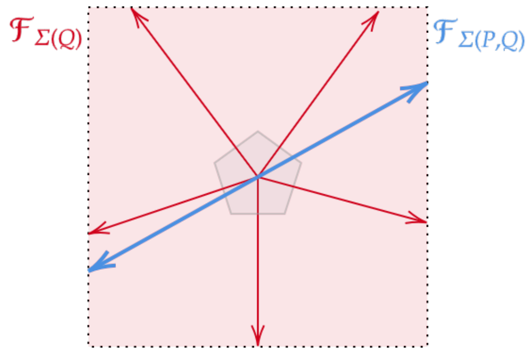

Geometrically, the fiber fan $\mathcal{F}_{\Sigma(P,Q)}$ sits inside the full secondary fan $\mathcal{F}_{\Sigma(Q)}$ of Q as a linear subspace. Intersecting $\mathcal{F}_{\Sigma(Q)}$ with the subspace $\pi^*((\mathbb{R}^d)^\vee) \subset (\mathbb{R}^A)^\vee$ consisting of those weight vectors that are constant on each fiber of π. In this way the fiber fan cuts out exactly the chambers of the secondary fan compatible with the face lattice structure of P, organising them into a much smaller and more tractable arrangement. From an computational perspective, one now works in a $(\dim P - \dim Q)$-dimensional space, rather than an $(|\mathcal{A}| - 1)$-dimensional space, so convex-hull and subdivision enumeration computations become more feasible by several orders of magnitude.

This viewpoint motivates the J-cubes used for the Tropical Ising perspective; a suitable polytope $P = J_\Gamma \subset \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$ should carry all spin and coupling information, while $Q = \operatorname{conv}(\mathcal{A})$ records the spin states. If J_Γ projects linearly onto the cube, then every choice of interaction matrix J will yield a height function on J_Γ, which will determine a point in the fiber polytope $\Sigma(J_\Gamma, Q)$. Studying this much smaller polytope and its fan structure, will provide us with a cleaner polyhedral description of all ground-state subdivisions that the Hamiltonian can induce on Q. In subsequent sections we will analyse the face structure and symmetries of J-cubes to obtain explicit and computable models of Ising system degeneracies.

Chapter 3

A Second Approach

3.1 Constructing the J-Cube

Let Γ be a simple (undirected) graph with vertex set $\mathrm{V}(\Gamma)$ of size n and edge set $\mathrm{E}(\Gamma)$ of size m. We denote by $\mathcal{A} = \{\pm 1\}^n$ the set of all states (i.e. assignments of ± 1 to each of the n vertices). Equivalently, $\mathcal{A}$ can be identified with the set of vertices of the standard n-cube $Q = \operatorname{conv}\big(\{\pm 1\}^n\big) \subset \mathbb{R}^n$.

We will construct a linear map $\varphi : \mathbb{R}^{\mathcal{A}} \longrightarrow \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$, such that the images of the basis vectors u_σ (for $\sigma \in \mathcal{A}$) form the vertices of a polytope

$$J_\Gamma = \operatorname{conv}\big\{\varphi(u_\sigma) = p_\sigma \mid \sigma \in \mathcal{A}\big\} \subset \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}\,.$$

We will also show why φ arises in a canonical way from the sign assignments, and how a $\mathbb{Z}_2^n$-symmetry (sign-flip action) makes it sufficient to study the facets incident to one chosen vertex (e.g. $p_{-\mathbf{1}}$).

To understand how φ fits into the bigger picture of quotient and inclusion maps, we can represent the various spaces via two commutative diagrams (a primal diagram and a dual diagram).

Primal Diagram

$$\begin{tikzcd}
 & \mathbb{R}^{E(\Gamma)} \arrow[r] & \mathbb{R}^{V(\Gamma)+E(\Gamma)} \\
\mathbb{R}^{2^n-n} \arrow[ur, two heads, "i^*"] \arrow[r, hook] & \mathbb{R}^{\mathcal{A}} \arrow[u, two heads, "\gamma^*"'] \arrow[ur, "\varphi"'] \arrow[d, two heads, "\pi"] & \\
\mathbb{R}^{2^n-n-m} \arrow[u] \arrow[ur] & \mathbb{R}^{V(\Gamma)} &
\end{tikzcd}$$

Dual Diagram

$$\begin{tikzcd}
 & (\mathbb{R}^{E(\Gamma)})^{\vee} \arrow[dl, hook', "i"'] \arrow[d, hook, "\gamma"] & (\mathbb{R}^{V(\Gamma)+E(\Gamma)})^{\vee} \\
(\mathbb{R}^{2^n-n})^{\vee} & (\mathbb{R}^{\mathcal{A}})^{\vee} \arrow[l] & \\
(\mathbb{R}^{2^n-n-m})^{\vee} & (\mathbb{R}^{V(\Gamma)})^{\vee} \arrow[u, hook, "\pi^*"'] &
\end{tikzcd}$$

Here, i^*, γ^*, π are quotient maps or injections, and φ is the top-right arrow in the primal diagram. Dually, i, γ, π^* are embeddings in $(\mathbb{R}^{\cdot})^{\vee}$. These diagrams affirm that φ is built out of canonical sums involving vertex and edge contributions.

We now verify that φ is well-defined and faithful (injective in the necessary sense). The essential step is to show that no nontrivial linear combination of vertex or edge basis elements can vanish on all states.

Lemma 3.1. *If for some $a \in \mathbb{R}^{\mathrm{V}(\Gamma)}$, $\sum_{v\in \mathrm{V}(\Gamma)} \sigma(v)a_v = 0$ for every choice of state $\sigma \in \mathcal{A}$, then $a = 0$.*

Proof. Choose any vertex $k \in \mathrm{V}(\Gamma)$. Then for each state $\sigma \in \mathcal{A}$, there exists another state $\sigma' \in \mathcal{A}$ such that

$$\sigma'(v) = \begin{cases} \sigma'(k) & v = k \\ \sigma(v) & v \neq k \end{cases}$$

Then

$$0 = \sum_{v\in \mathrm{V}(\Gamma)} \sigma(v)a_v - \sum_{v\in \mathrm{V}(\Gamma)} \sigma'(v)a_v = 2\sigma(k)a_k$$

so we have $a_k = 0$. We may repeat this argument for each of the remaining $v(i) \in \mathrm{V}(\Gamma) \setminus k$ to conclude $a = 0$. $\square$

Lemma 3.2. *If for some $b \in \mathbb{R}^{\mathrm{E}(\Gamma)}$, $\sum_{(i,j)\in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j)b_{(i,j)} = 0$ for every choice of state $\gamma \in \mathcal{A}$, then $b = 0$.*

Proof. Choose any edge cover of Γ and fix a vertex $k \in \mathrm{V}(\Gamma)$. Consider its set of neighbors of $N(k) \subset \mathrm{V}(\Gamma)$ on Γ. Since Γ is a simple weighted graph, the induced subgraph $\Gamma(N(k))$ of Γ is a rooted tree. So we may then choose a state $\sigma^{(1)} \in \mathcal{A}$ such that $\sigma^{(1)}(k)\sigma^{(1)}(j) = \operatorname{sign}(b_{(k,j)})$ for all $j \in N(k)$. Define $\sigma^{(2)} \in \mathcal{A}$ by

$$\sigma^{(2)}(j) := \begin{cases} -\sigma^{(1)}(k) & j = k \\ \sigma^{(1)}(j) & j \neq k \end{cases}.$$

Then we get the following

$$\begin{aligned}
0 &= \sum_{(i,j)\in \mathrm{E}(\Gamma)} \sigma^{(1)}(i)\sigma^{(1)}(j)b_{(i,j)} - \sum_{(i,j)\in \mathrm{E}(\Gamma)} \sigma^{(2)}(i)\sigma^{(2)}(j)b_{(i,j)} \\
&= \sum_{(i,j)\in \mathrm{E}(\Gamma)} \left(\sigma^{(1)}(i)\sigma^{(1)}(j) - \sigma^{(2)}(i)\sigma^{(2)}(j)\right) b_{(i,j)} \\
&= \sum_{\substack{(i,j)\in \mathrm{E}(\Gamma) \\ i,j\neq k}} \left(\sigma^{(1)}(i)\sigma^{(1)}(j) - \sigma^{(2)}(i)\sigma^{(2)}(j)\right) b_{(i,j)} \\
&\qquad + \sum_{(k,j)\in \mathrm{E}(\Gamma)} \left(\sigma^{(1)}(k)\sigma^{(1)}(j) - \sigma^{(2)}(k)\sigma^{(2)}(j)\right) b_{(k,j)} \\
&= \sum_{\substack{(i,j)\in \mathrm{E}(\Gamma) \\ i,j\neq k}} \left(\sigma^{(1)}(i)\sigma^{(1)}(j) - \sigma^{(1)}(i)\sigma^{(1)}(j)\right) b_{(i,j)} \\
&\qquad + \sum_{(k,j)\in \mathrm{E}(\Gamma))} \left(\sigma^{(1)}(k)\sigma^{(1)}(j) + \sigma^{(1)}(k)\sigma^{(1)}(j)\right) b_{(k,j)} \\
&= \sum_{(k,j)\in \mathrm{E}(\Gamma)} 2\operatorname{sign}(b_{(k,j)})b_{(k,j)} \\
&= 2 \sum_{(k,j)\in \mathrm{E}(\Gamma))} |b_{(k,j)}|
\end{aligned}$$

Thus, $b_{(k,j)} = 0$ for all edges $(k,j) \in \mathrm{E}(\Gamma)$. Repeating this argument as k ranges over the edge cover of Γ shows that $b = 0$. So $\ker \gamma = 0$. $\square$

Lemma 3.3. *The map $\gamma : (\mathbb{R}^{\mathrm{E}(\Gamma)})^\vee \to (\mathbb{R}^{\mathcal{A}})^\vee$ defined as*

$$\gamma(b) := \sum_{\sigma \in \mathrm{Hom}(\mathrm{E}(\Gamma), \{\pm 1\})} \left(\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} \right) f_\sigma$$

acting on the basis elements $f_{(i,j)} \in (\mathbb{R}^{\mathrm{E}(\Gamma)})^\vee$ and $f_\sigma \in (\mathbb{R}^{\mathcal{A}})^\vee$ is an injective mapping.

Proof. Clearly $\gamma(0) = 0$, so suppose $\gamma\left(\sum_{(i,j)\in \mathrm{E}(\Gamma)} b_{(i,j)} f_{(i,j)}\right) = 0$ for some nonzero $b = \sum_{(i,j)\in \mathrm{E}(\Gamma)} b_{(i,j)} f_{(i,j)} \in (\mathbb{R}^{\mathrm{E}(\Gamma)})^\vee$. Then

$$\sum_{\sigma \in \mathrm{Hom}(\mathrm{V}(\Gamma), \{\pm 1\})} \left(\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} \right) f_\sigma = 0$$

So for each $\sigma \in \mathrm{Hom}(\mathrm{V}(\Gamma), \{\pm 1\})$, we have

$$\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} = 0.$$

so by Lemma 3.2 $b = 0$. $\square$

Lemma 3.4. *Within $(\mathbb{R}^{\mathcal{A}})^\vee$, the images of γ and π^* intersect trivially, i.e.,* $\mathrm{im}\,(\gamma) \cap \mathrm{im}\,(\pi^*) = \{0\}$*. Therefore, there is a (non-conflicting) inclusion $i : (\mathbb{R}^{\mathrm{E}(\Gamma)})^\vee \to (\mathbb{R}^{2^n - n})^\vee$ in the dual diagram.*

Proof.

$$\begin{tikzcd}
& (\mathbb{R}^{E(\Gamma)})^\vee \arrow[dl, two heads, "i^*"'] \arrow[d, hook, "\gamma"] \\
(\mathbb{R}^{2^n-n})^\vee & (\mathbb{R}^{\mathcal{A}})^\vee \arrow[l, hook'] \\
& (\mathbb{R}^{V(\Gamma)})^\vee \arrow[u, hook, "\pi^*"']
\end{tikzcd}$$

Let $X = \operatorname{im}(\gamma)$ and $Y = \operatorname{im}(\pi)^*$. Clearly $0 \in X \cap Y$, so take $p \in X \cap Y$ to be nonzero. Then we have

$$\sum_{\sigma \in \operatorname{Hom}(\mathrm{V}(\Gamma), \{\pm 1\})} \left(\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} \right) f_\sigma = \sum_{\sigma \in \operatorname{Hom}(\mathrm{V}(\Gamma), \{\pm 1\})} \left(\sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) a_v \right) f_\sigma$$

for some a, b mapping to p,so for every $\sigma \in \operatorname{Hom}(\mathrm{V}(\Gamma), \{\pm 1\}) \cong \mathcal{A}$, we have that

$$\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} - \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) a_v = 0.$$

For any $\sigma \in \mathcal{A}$, define $\sigma' = -\sigma$. Note that for all edges $(i,j) \in \mathrm{E}(\Gamma)$,

$$\sigma'(i)\sigma'(j) = (-\sigma(i))(-\sigma(j)) = \sigma(i)\sigma(j).$$

Adding the entries from f_σ and $f_{\sigma'}$, we get

$$\begin{aligned} 0 &= \sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma'(i)\sigma'(j) b_{(i,j)} - \sum_{v \in \mathrm{V}(\Gamma)} \sigma'(v) a_v + \sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} - \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) a_v \\ &= \sum_{(i,j) \in \mathrm{E}(\Gamma)} 2\sigma(i)\sigma(j) b_{(i,j)} \end{aligned}$$

Since this is independent of choice of σ, for all states $\sigma \in \mathcal{A}$

$$\sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} = 0$$

So by Lemma 3.2, $b_{(i,j)} = 0$ for all $(i,j) \in \mathrm{E}(\Gamma)$. Thus, for every $\sigma \in \operatorname{Hom}(\mathrm{V}(\Gamma), \{\pm 1\}) \cong \mathcal{A}$, we have

$$0 = \sum_{(i,j) \in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j) b_{(i,j)} - \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) a_v = \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) a_v$$

so by Lemma 3.1, $a_v = 0$ for all $v \in \mathrm{V}(\Gamma)$. Thus, $p = 0$. $\square$

3.1.1 Definition and Naturality of φ

On basis vectors $u_\sigma \in \mathbb{R}^{\mathcal{A}}$, define

$$\varphi(u_\sigma) = \sum_{v\in \mathrm{V}(\Gamma)} \sigma(v)\, u_v + \sum_{(i,j)\in \mathrm{E}(\Gamma)} \sigma(i)\sigma(j)\, u_{(i,j)}.$$

Equivalently, the image has coordinates

$$[\varphi(u_\sigma)]_v = \sigma(v), \quad [\varphi(u_\sigma)]_{(i,j)} = \sigma(i)\,\sigma(j).$$

The previous results establish that φ is faithful: it identifies different basis vectors u_σ with distinct images in $\mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$. No extra choices are made: we simply record, for each vertex v, the sign $\sigma(v)$, and for each edge (i,j), the product $\sigma(i)\sigma(j)$. By the previous lemmas, this embedding is non-degenerate.

Diagrammatically, φ factors as an inclusion from $\mathbb{R}^{\mathcal{A}}$ onto a subspace inside $\mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$, reflecting the direct sum decomposition in the dual diagram. This is what we mean by naturality—the construction follows directly from the combinatorial structure of Γ.

3.1.2 The Coupling Polytope J_Γ and its $\mathbb{Z}_2^n$-Action Symmetry

Define $J_\Gamma := \mathrm{conv}\{\varphi(u_\sigma) \mid \sigma \in \mathcal{A}\}$. This is an $(n+m)$-dimensional polytope in $\mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$, with 2^n vertices $p_\sigma = \varphi(u_\sigma)$.

The group $\{\pm 1\}^n$ acts transitively on $\mathcal{A}$ by

$$(\tau \cdot \sigma)(v) = \tau(v)\,\sigma(v).$$

Hence any vertex p_σ of J_Γ can be sent to any other $p_{\sigma'}$ by a suitable choice of τ. This symmetry extends to J_Γ itself, permuting vertices and their incident facets.

Because τ is an affine automorphism of J_Γ, the local facet structure at p_σ is the same as that at $p_{\tau\cdot\sigma}$. Consequently, to classify all facets of J_Γ, it suffices to examine those facets that meet (for instance) $p_{-\mathbf{1}}$. We will heavily rely on this principle in the upcoming section on classifying the facets, so it should be assumed that the basepoint $p_{-\mathbb{1}}$ is chosen.

In summary, we have shown that:

- The map φ is well-defined and injective because no nontrivial linear combination of vertex- or edge-contributions can vanish on all sign assignments $\sigma \in \mathcal{A}$.
- The image of φ is a polytope J_Γ in $\mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$, encoding both per-vertex signs and per-edge products of signs.
- A natural $(\pm 1)^n$-action (sign-flips) acts transitively on vertices of J_Γ, allowing us to focus on a single "base vertex" for facet analysis.

These core ideas underlie subsequent work on the classification of edge-induced and cycle-induced facets, particularly how certain chordless cycles (or more complex 2-connected subgraphs) give rise to new facet families.

3.2 Classifying Facets of J_Γ

Fix $\sigma_0 = -\mathbb{1} \in \mathcal{A}$. Then for any other state $\sigma \in \mathcal{A}$, define the set $S_\sigma \subset \mathrm{V}(\Gamma)$ as the set of vertices in $\mathrm{V}(\Gamma)$ where σ differs in sign from $-\mathbb{1}$, and define $\sigma_S \in \mathcal{A}$ as the state corresponding to S_σ. Viewed as a weighted subgraph of Γ, each choice of $\sigma \in \mathcal{A}$ induces the sets S_σ and S_σ^c which are separated by boundary edges $\partial \mathrm{S}_\sigma \subset \mathrm{E}(\Gamma)$. In addition, define the vertex of J_Γ that corresponds to this state in $\mathcal{A}$ as p_{σ_S}.

By the discussion in the previous section on transitive actions, it is sufficient to study the facets of J_γ incident to $p_{-\mathbb{1}}$.

Lemma 3.5. *If $\psi_F \in (\mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)})^\vee$ is an inward-facing linear functional on J_Γ, defining* *a face $F \preceq J_\Gamma$ incident to $p_{-\mathbb{1}}$ then*

$$\sum_{v \in S_\sigma} a_v \geq \sum_{e \in \partial \mathrm{S}_\sigma} b_e \quad \forall S_\sigma \subset \mathrm{V}(\Gamma)$$

where $\partial \mathrm{S} \subset \mathrm{E}(\Gamma)$ denotes the set of edges e with one vertex $e_0 \in S$ and the other vertex $e_1 \in \mathrm{V}(\Gamma) \setminus S$.

Proof. For each $\sigma \in \mathcal{A}$, we have $p_{-\mathbb{1}}$ as a vertex of J_Γ. For each linear functional

$$\psi_F = \sum_{v \in \mathrm{V}(\Gamma)} a_v u_v^\vee + \sum_{e \in \mathrm{E}(\Gamma)} b_e u_e^\vee$$

that defines a facet incident to $p_{-\mathbb{1}}$, for the associated vertex to S_σ, p_{σ_S}, we have

$$\begin{aligned}
\psi_F(p_{\sigma_S}) &= \sum_{v \in \mathrm{V}(\Gamma)} a_v u_v^\vee(p_{\sigma_S}) + \sum_{e \in \mathrm{E}(\Gamma)} b_e u_e^\vee(p_{\sigma_S}) \\
&= \sum_{v \in \mathrm{V}(\Gamma)} a_v \sigma_S(v) + \sum_{e \in \mathrm{E}(\Gamma)} b_e \sigma_S(e_0)\sigma_S(e_1) \\
&= \sum_{v \in S} a_v - \sum_{v \in S^c} a_v - \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0 \in S \\ e_1 \in S^c}} b_e + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0, e_1 \in S}} b_e + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0, e_1 \in S^c}} b_e
\end{aligned}$$

Similarly, for $p_{-\mathbb{1}}$ we have

$$\begin{aligned}
\psi_F(p_{-\mathbb{1}}) &= \sum_{v \in \mathrm{V}(\Gamma)} a_v u_v^\vee(p_{-\mathbb{1}}) + \sum_{e \in \mathrm{E}(\Gamma)} b_e u_e^\vee(p_{-\mathbb{1}}) \\
&= \sum_{v \in \mathrm{V}(\Gamma)} a_v - \mathbb{1}(v) + \sum_{e \in \mathrm{E}(\Gamma)} b_e - \mathbb{1}(e_0) - \mathbb{1}(e_1) \\
&= -\sum_{v \in S} a_v - \sum_{v \in S^c} a_v + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0 \in S \\ e_1 \in S^c}} b_e + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0, e_1 \in S}} b_e + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e_0, e_1 \in S^c}} b_e
\end{aligned}$$

Since ψ_F defines a facet, we have that $\psi_F(p_\sigma) \geq \psi_F(p_{-\mathbb{1}})$ for any vertex p_σ of J_Γ. Applying this inequality to p_{σ_S} we get the following,

$$\psi_F(p_{\sigma_S}) \geq \psi_F(p_{-\mathbb{1}}) \quad \Rightarrow \quad 2\sum_{v\in S} a_v \geq 2 \sum_{\substack{e\in \mathrm{E}(\Gamma)\\ e_0\in S\\ e_1\in S^c}} b_e.$$

□

So for any subset of vertices $S_\sigma \subset \mathrm{V}(\Gamma)$ corresponding to the state $\sigma \in \mathcal{A}$, define the ray in the direction of p_{σ_S} relative to $p_{-\mathbb{1}}$ as

$$u_{S_\sigma} := \sum_{v\in S} u_v - \sum_{e\in\partial \mathrm{S}} u_e \in \ \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}\,.$$

The local cone at $\sigma_0 = -\mathbb{1}$, denoted $\mathrm{cone}(u_{S_\sigma} \mid \forall\sigma \in \mathcal{A})$ is the cone at $p_{-\mathbb{1}}$ consisting of points p_{σ_S} induced by the sets S_σ with respect to $-\mathbb{1}$, taken over all states of $\mathcal{A}$. The local cone at $-\mathbb{1}$ contains a translate of J_Γ.

Lemma 3.6. *A vertex p_σ of J_Γ is connected to $p_{-\mathbb{1}}$ by a ridge along F if and only if $\psi_F(u_{S_\sigma}) = 0$ for each facet-defining linear functional $\psi_F = \sum_{v\in \mathrm{V}(\Gamma)} a_v u_v^\vee + \sum_{e\in \mathrm{E}(\Gamma)} b_e u_e^\vee$ where F is incident to $p_{-\mathbb{1}}$.*

Proof. If two points share a ridge on a facet F of J_Γ, for all $\lambda \in [0,1]$,

$$\begin{aligned}
\psi_F(p_{-\mathbb{1}}) &= \psi(\lambda p_{-\mathbb{1}} + (1-\lambda)p_\sigma)\\
&= \lambda\left(\sum_{v\in \mathrm{V}(\Gamma)} a_v - \mathbb{1}(v) + \sum_{e\in \mathrm{E}(\Gamma)} b_e - \mathbb{1}(e)\right) + (1-\lambda)\left(\sum_{v\in \mathrm{V}(\Gamma)} a_v\sigma(v) + \sum_{e\in \mathrm{E}(\Gamma)} b_e\sigma(e)\right)\\
&= -\lambda\sum_{v\in S} a_v + \lambda\sum_{e\in\partial \mathrm{S}_\sigma} b_e + (1-\lambda)\sum_{v\in S} a_v - (1-\lambda)\sum_{e\in\partial \mathrm{S}_\sigma} b_e\\
&\quad -\lambda\sum_{v\in S^c} a_v + \lambda\sum_{e\in\partial \mathrm{S}_\sigma^c} b_e - (1-\lambda)\sum_{v\in S^c} a_v + (1-\lambda)\sum_{e\in\partial \mathrm{S}_\sigma^c} b_e\\
&= (1-2\lambda)\left(\sum_{v\in S_\sigma} a_v - \sum_{e\in\partial \mathrm{S}_\sigma} b_e\right) - \sum_{v\in S^c} a_v + \sum_{e\in\partial \mathrm{S}_\sigma^c} b_e
\end{aligned}$$

if and only if $\sum_{v \in S_\sigma} a_v = \sum_{e \in \partial \mathrm{S}_\sigma} b_e$. $\square$

We now explore the first classification of facets of J_Γ. We begin in the tree case, so that the only combinatorial structures are edges; this is due to the fact that if the defining set $S_\sigma \subset \mathrm{V}(\Gamma)$ for ψ induces a disconnected subgraph $\Gamma[S_\sigma] = G_1 \cup G_2$, then ψ can be reduced to $\psi_{G_1} + \psi_{G_2}$.

Lemma 3.7. *For every edge $e = (v_0, v_1)$ of a tree Γ, the following three functionals define the only inward-facing normals to facets of J_Γ incident to $p_{-\mathbb{1}}$:*

$$\begin{aligned}
\psi_{e_{++}} &:= u_{v_0}^\vee + u_{v_1}^\vee + u_e^\vee \\
\psi_{e_{+-}} &:= u_{v_0}^\vee - u_{v_1}^\vee - u_e^\vee \\
\psi_{e_{-+}} &:= -u_{v_0}^\vee + u_{v_1}^\vee - u_e^\vee
\end{aligned}$$

Proof. Suppose ψ is the unique linear functional that defines a facet $F \preceq J_\Gamma$. We characterize faces by the subsets S_σ whose corresponding rays lie in them. Specifically, for a functional ψ, we consider

$$F_\psi^1 = \left\{ S_\sigma \subset V(\Gamma) \middle| \psi(u_{S_\sigma}) = \sum_{v \in S_\sigma} a_v - \sum_{e \in \partial \mathrm{S}} b_e = 0,\ \Gamma[S_\sigma] \text{ is connected} \right\}$$

to be the maximal set of S_σ's with respect to containment. Fix an edge $e' = (v_0, v_1) \in \mathrm{E}(\Gamma)$ and consider F_ψ^1 under several cases.

Case 1: If e' is a bridge in Γ, and $b_{e'} > 0$. Suppose $\{v_0, v_1\} \subset S_\sigma \in F_\psi^1$. Since e' is a bridge, removing it disconnects the graph into two components S_+ and S_- where:

$$\begin{aligned}
S_- &:= \{v \in S_\sigma \mid \text{there is a path from } v_0 \text{ to } v \text{ not traversing } e'\} \\
S_+ &:= \{v \in S_\sigma \mid \text{there is a path from } v_1 \text{ to } v \text{ not traversing } e'\}
\end{aligned}$$

The boundary edges split as $\partial \mathrm{S}_\sigma = \partial \mathrm{S}_+ \setminus \{e'\} \sqcup \partial \mathrm{S}_- \setminus \{e'\}$ and the ray u_{S_σ} then decomposes as:

$$
\begin{aligned}
u_S &= \sum_{v \in S_\sigma} a_v u_v - \sum_{e \in \partial \mathrm{S}_\sigma} b_e u_e \\
&= \sum_{v \in S_+} a_v u_v - \sum_{e \in \partial \mathrm{S}_+} b_e u_e + \sum_{v \in S_-} a_v u_v - \sum_{e \in \partial \mathrm{S}_-} b_e u_e + 2b_{e'} u_{e'} \\
&= (u_{S_+} + b_{e'} u_{e'}) + (u_{S_-} + b_{e'} u_{e'})
\end{aligned}
$$

Since ψ defines a facet, $\psi(u_{S_\sigma}) \geq 0$ for all S_σ.Thus $\psi(u_{S_+}) \geq 0$ and $\psi(u_{S_-}) \geq 0$. But then:

$$
\psi(u_S) = \psi(u_{S_+}) + \psi(u_{S_-}) + 2b_{e'} > 0
$$

contradicting $S_\sigma \in F^1_\psi$. Therefore, any S_σ containing both endpoints of e' cannot be in F^1_ψ.

Now consider any $T_\sigma \in F^1_\psi$. By the above argument, T_σ cannot contain both endpoints of e'. Moreover, since ψ defines a facet of J_Γ, T_σ must be in $F_{\psi_{e'_{++}}}$, as these are the only sets giving rays in J_Γ that do not contain both endpoints of e'.

Thus $F^1_\psi \subseteq F_{\psi_{e'_{++}}}$. Since $F_{\psi_{e'_{++}}}$ is maximal among such collections of sets (as it defines a facet), we must have $F^1_\psi = F_{\psi_{e'_{++}}}$.

Case 2: Next, suppose that e' is a twig and v_0 is a leaf. Define

$$
\mathcal{C}_\Gamma := \{S_\sigma \subset \mathrm{V}(\Gamma) \mid \Gamma[S_\sigma] \text{ is connected}\}
$$

be the indices of cone rays incident to $p_{-\mathbb{1}}$ in J_Γ. If $a_{v_0} \neq b_{e'}$, then $\psi(u_{\{v_0\}}) = a_{v_0} - b_{e'} > 0$, so $a_{v_0} > b_{e'}$. Then $\{v_0\} \notin F^1_\psi$.

But note that for any $S_\sigma \in \mathcal{C}_\Gamma \setminus \{e'\}$ we have $\psi_{e'_{+-}}(u_{S_\sigma}) = 0$, thus $\mathcal{C}_\Gamma \setminus \{e'\} \subset F^1_{\psi_{e'_{+-}}}$. By maximality, $\mathcal{C}_\Gamma \setminus \{e'\} = F_{\psi_{e'_{+-}}}$, and $F^1_\psi = F_{\psi_{e'_{+-}}}$.

Case 3: Lastly, suppose that e' is a twig, v_0 is a leaf, and $a_{v_0} = b_{e'} < 0$. Let $S_\sigma \in \mathcal{C}_\Gamma$ contain $\{v_0, v_1\}$. Removing a leaf leaves the subgraph connected, so $S_\sigma \setminus \{v_0\} \in \mathcal{C}_\Gamma$. Then:

$$
\psi(u_{S \setminus \{e_0\}}) > \psi(u_{S \setminus \{e_0\}}) + a_{v0} + b_{e'} = \psi(u_S) \geq 0.
$$

Therefore $S_\sigma \setminus \{v_0\} \notin F^1_\psi$, but $\{v_0\} \in F^1_\psi$, so $\mathbb{R}_{\geq 0}\,\psi = \mathbb{R}_{\geq 0}\,\psi_{e_{-+}}$.

In all cases, the functionals induced by e' define facets of J_Γ incident to $p_{-\mathbb{1}}$, completing the proof. $\square$

So we have that if Γ is a tree, then each edge induces three facets of J_Γ that are incident to $p_{-\mathbb{1}}$. Moreover, these are the only facets that appear when Γ is a tree. These facets will be referred to as edge-facets, and when specificity isn't required, we refer to one of the facet-defining functions as simply $\psi_{e_{\pm\pm}}$, without specifying a sign pattern.

In the following, we extend this classification by adding a single edge to a tree in order to introduce a single chordless cycle. To begin, we classify the kernel of these cycle-facet functionals.

Lemma 3.8. *For every chordless cycle $\gamma = (e_1, e_2, \ldots, e_\ell)$ of Γ with length ℓ, the following ℓ functionals*

$$\psi_{\gamma_k} := u^\vee_{e_k} - \sum_{\substack{i=1\\ i\neq k}}^{\ell} u^\vee_{e_i}$$

for $k = 1, \ldots \ell$ vanish precisely on the set

$$F^1_{\gamma_k} := \{S_\sigma \subset V(\gamma) \mid \Gamma[S_\sigma] \text{ a path on } \gamma \text{ with at least one endpoint on } e_k \text{ and not crossing } e_k\}$$

Proof. Recall that the ray corresponding to $S_\sigma \subset V(\gamma)$ is:

$$u_{S_\sigma} = \sum_{v \in S_\sigma} u_v - \sum_{e \in \partial \mathrm{S}_\sigma} u_e,$$

where $\partial\mathrm{S}_\sigma$ consists of edges in γ that have exactly one endpoint in S_σ.

Since ψ_{γ_k} acts only on the edge components u_e, this reduces to:

$$\psi_{\gamma_k}(u_{S_\sigma}) = \sum_{\substack{i=1\\ i\neq k}}^{\ell} \delta_{e_i} - \delta_{e_k}$$

where δ_{e_i} indicates whether $e_i \in \partial \mathrm{S}_\sigma$. We now analyze when ψ_{γ_k} vanishes.

1. If $\Gamma[S_\sigma]$ is a vertex path that includes an endpoint on e_k and does not cross e_k, then $e_k \in \partial \mathrm{S}_\sigma$ because it has exactly one endpoint in S_σ. Furthermore, no other edges $e_i \neq e_k$ contribute to $\partial \mathrm{S}_\sigma$, since they are internal to the path $\Gamma[S_\sigma]$. Hence $\psi_{\gamma_k}(u_{S_\sigma}) = 1-1 = 0$.

2. If S_σ does not satisfy the conditions of $F^1_{\gamma_k}$, then either: $e_k \notin \partial \mathrm{S}_\sigma$, or, other edges $e_i \neq e_k$ contribute to $\partial \mathrm{S}_\sigma$. In these cases $\psi_{\gamma_k}(u_{S_\sigma}) = 2$.

Therefore, $\psi_{\gamma_k}(u_{S_\sigma}) = 0$ if and only if $S_\sigma \in F^1_{\gamma_k}$, i.e., $\Gamma[S_\sigma]$ is a path on γ with at least one endpoint on e_k and does not cross e_k. And for all other subsets S_σ, $\psi_{\gamma_k}(u_{S_\sigma}) \neq 0$. $\square$

Lemma 3.9. *For every chordless cycle $\gamma = (e_1, e_2, \dots, e_\ell)$ of Γ with length ℓ, the following ℓ functionals define inward-facing normals to a facet incident to $p_{-\mathbb{1}}$:*

$$\psi_{\gamma_k} := u^\vee_{e_k} - \sum_{\substack{i=1 \\ i \neq k}}^{\ell} u^\vee_{e_i}$$

for $k = 1, \dots \ell$.

Proof. Let ψ be the unique linear functional defining a facet $F \preceq J_\Gamma$, and let

$$F^1_\psi = \left\{ S_\sigma \subset V(\Gamma) \;\middle|\; \psi(u_{S_\sigma}) = \sum_{v \in S_\sigma} a_v - \sum_{e \in \partial \mathrm{S}_\sigma} b_e = 0,\ \Gamma[S_\sigma] \text{ connected} \right\}$$

be the maximal set of S_σ's whose rays u_{S_σ} lie on F.

Fix a chordless cycle $\gamma = (e_1, \dots, e_\ell)$ in Γ, and let $V(\gamma) = \{v_1, \dots, v_\ell\}$ be the vertices of γ, indexed modulo ℓ. We will show that if $F^1_\psi \supseteq F_{\psi_{\gamma_k}}$ for some $k = 1, \dots, \ell$, then $\psi = \psi_{\gamma_k}$ for some k. Fix an edge $e_k = (v_{k-1}, v_k)$ of the cycle γ and suppose that F^1_ψ contains $F^1_{\gamma_k}$ as defined in Lemma 3.8:

$$F^1_{\gamma_k} = \{S_\sigma \subset V(\gamma) \mid \Gamma[S_\sigma] \text{ a path on } \gamma \text{ with at least one endpoint on } e_k \text{ and not crossing } e_k\}.$$

The facet-defining property of ψ ensures that for any $S_\sigma \in F^1_{\gamma_k} \subset F^1_\psi$, we have $\psi(u_{S_\sigma}) = 0$.

First, we consider two specific vertex paths in $F^1_{\gamma_k}$:

$$S_{\sigma_1} = \{v_{k-1}, v_{k-2}, \ldots, v_{i+1}\}, \text{ with boundary edges } \{e_k, e_i\}$$
$$S_{\sigma_2} = \{v_k, v_{k+1}, \ldots, v_i\}, \text{ with boundary edges } \{e_k, e_i\}$$

If γ's vertices are ordered 'counterclockwise', then S_{σ_1} runs clockwise from v_{k-1} along $\gamma \setminus e_k$, while S_{σ_2} runs counterclockwise from v_k along $\gamma \setminus e_k$. Applying ψ to their respective rays gives:

$$\sum_{v \in S_{\sigma_1}} a_v = b_{e_k} + b_{e_i} \quad \text{and} \quad \sum_{v \in S_{\sigma_2}} a_v = b_{e_k} + b_{e_i}.$$

Next, note that the path $S_{\sigma_1} \cup S_{\sigma_2} \in F_{\psi_{\gamma_k}}$ and has no boundary edges in γ. Hence

$$\sum_{v \in S_{\sigma_1} \cup S_{\sigma_2}} a_v = 0 \quad \Rightarrow \quad -\sum_{v \in S_{\sigma_2}} a_v = \sum_{v \in S_{\sigma_1}} a_v.$$

Adding the two expressions for $\psi(u_{S_{\sigma_i}}$, we get:

$$0 = 2b_{e_k} + 2b_{e_i} \quad \Rightarrow \quad -b_{e_k} = b_{e_i}.$$

Since i resulted from an arbitrary choice of vertex with $i \neq k$, we have that $b_{e_i} = -b_{e_k}$ for all $i \neq k$. As a consequence, we also have that $b_{e_i} = b_{e_j}$ for all $i, j \neq k$.

Next, to show that all vertex coefficients for ψ vanish, we pick an arbitrary vertex v_i, with $i \neq k-2, k-1, k$ and consider two overlapping counterclockwise paths that differ in boundary edges and in their inclusion of v_i:

$$S_{\sigma_1} = \{v_k, v_{k+1}, \ldots, v_{i-1}, v_i\} \quad \text{and} \quad S_{\sigma_2} = \{v_k, v_{k+1}, \ldots, v_{i-1}\}$$

where the indices are again taken modulo ℓ. Applying ψ yields

$$\sum_{v\in S_{\sigma_1}} a_v = b_{e_k} + b_{e_{i+1}} \quad \text{and} \quad \sum_{v\in S_{\sigma_2}} a_v = b_{e_k} + b_{e_i}$$

Subtracting these equations, we get:

$$\left(\sum_{v\in S_\sigma} a_v\right) - \left(\sum_{v\in S'_\sigma} a_v\right) = (b_{e_k} + b_{e_{i+1}}) - (b_{e_k} + b_{e_i})$$
$$a_{v_{i-1}} = b_{e_{i+1}} - b_{e_i}$$

which simplifies to $a_{v_i} = 0$ since $i, i+1 \neq k-1, k$ modulo ℓ. Hence, we get that $a_{v_i} = 0$ for all $i \neq k-2, k-1, k$.

Repeating this argument with two clockwise paths and $i \neq k-1, k, k+1$,

$$S_{\sigma_1} = \{v_{k-1}, v_{k-2}, \ldots, v_{i+1}, v_i\} \quad \text{and} \quad S_{\sigma_2} = \{v_{k-1}, v_{k-2}, \ldots, v_{i+1}\}$$

we get

$$\sum_{v\in S_{\sigma_1}} a_v = b_{e_k} + b_{e_i} \quad \text{and} \quad \sum_{v\in S_{\sigma_2}} a_v = b_{e_k} + b_{e_{i+1}}$$

which again simplifies to $a_{v_i} = 0$ as $i+1 \neq k-1, k$ modulo ℓ. Combining both types of paths, we get that $a_{v_i} = 0$ for all $i \neq k-1, k$. Finally, note that $u_{\{v_{k-1}\}}, u_{\{v_{k-1}\}} \in F^1_{\gamma_k}$, so

$$\psi(u_{\{v_{k-1}\}}) = a_{v_{k-1}} - b_{e_k} - b_{e_{k-1}} = 0 \quad \Rightarrow \quad a_{v_{k-1}} = 0$$
$$\psi(u_{\{v_k\}}) = a_{v_k} - b_{e_k} - b_{e_{k+1}} = 0 \quad \Rightarrow \quad a_{v_k} = 0$$

Hence we have that $\mathbb{R}_{\geq 0}\, \psi = \mathbb{R}_{\geq 0}\, \psi_{\gamma_k}$. □

So we have that edges induce facets on tree graphs, and chordless cycles induce facets on cycle graphs, but it remains to see whether the edge facets persist in the transition from Γ a tree, to Γ a chordless cycle.

Moreover, each facet-defining dual functional $\psi_F = \sum_{v \in \mathrm{V}(\Gamma)} a_v u_v + \sum_{e \in \mathrm{E}(\Gamma)} b_e u_e$ can be identified with a weighted subgraph $\Gamma(\sigma_F)$ by assigning a_v, b_e to the vertex and edge weights, respectively. Here we allow the edge and vertex weights to be zero, but still include the edges and vertices with zero weight.

Example 3.1. Consider the graph $\Gamma = T_5$, by 3.7, there are $3 \cdot |\,\mathrm{E}(\Gamma)\,| = 18$ facets of J_Γ incident to $p_{-\mathbb{1}}$. Indeed, computing the J-cube directly and embedding the coefficients of each ψ as a subgraph of Γ, gives the following:

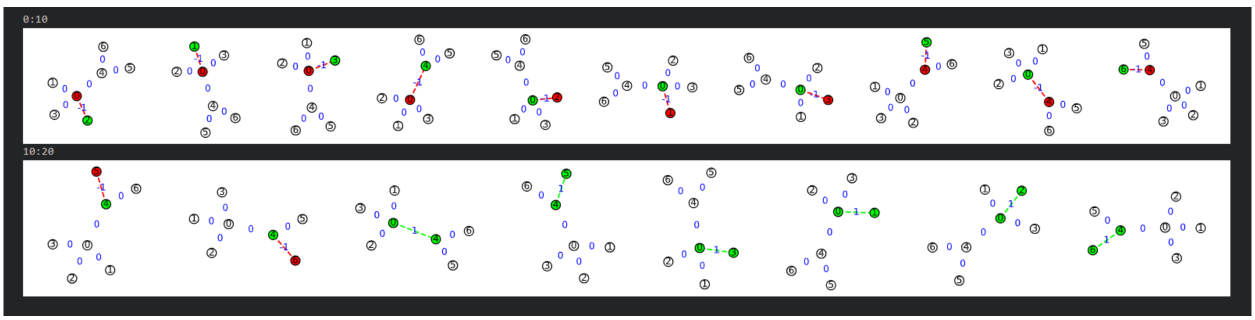

Similarly, if we consider the cycle graph $\Gamma = C_5$, then by 3.9, there are $\mathrm{len}(C_5) = 5$ facets of the J-cube that are incident to $p_{-\mathbb{1}}$. However, if we compute the facets of J_Γ directly, we observe $5 + 3|\,\mathrm{E}(\Gamma)\,| = 20$ facets—the edge-facets seemingly carried over.

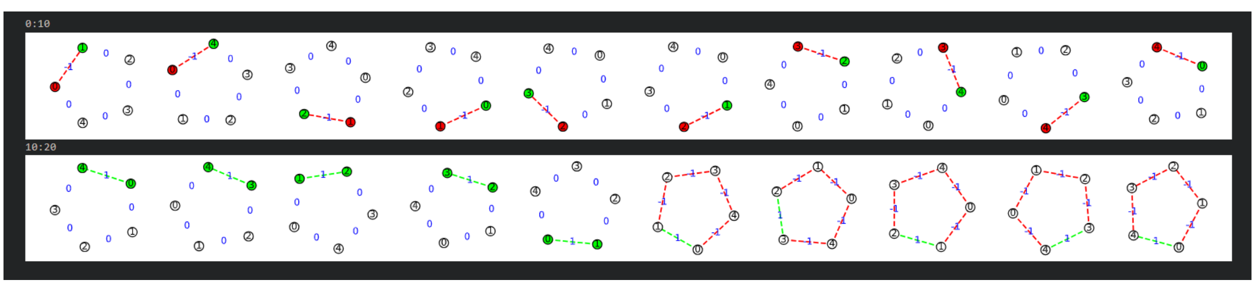

□

We will see that this phenomenon carries over into all graphs; in other words, edge-facets seem to persist among J-cubes of more complex graphs. Naturally, adding more edges to a cycle graph will affect the cycle structure, so we will also see that while edge-facets persist, cycle-facets are more delicate, so can decompose if a chord is added.

3.2.1 Local Classification of Facets

Lemma 3.7 classifies edge-induced facets by examining sets of vertices S_σ whose rays u_{S_σ} lie in a facet. In a tree, every edge naturally falls into one of the categories (bridge or twig), and the lemma constructs three facet-defining functionals for each edge from these configurations. Now, if we add a single edge to a tree to form a chordless cycle γ, some edges may no longer be bridges or twigs. Instead, they lie on the cycle.

The crux is to confirm that the presence of the cycle does not invalidate the original construction of the three edge facets. The newly formed cycle adds more structure, but it does not erase or merge the original edge-defined configurations. Instead, it provides a new layer of facet structure on top.

Theorem 3.1. *For $\Gamma = T \cup \{e'\}$ a graph with a single chordless cycle γ such that T is a tree, the tree edges and cycle edges classify all of the facets of J_Γ incident to $p_{-\mathbb{1}}$. Additionally, there are $3|E(\Gamma)| + \ell(\gamma)$ of these facets.*

Before tackling Theorem 3.1, we prove the following necessary lemma on edge-facet stability.

Lemma 3.10. *Let Γ be a generic graph on n vertices. Then the linear functionals of the form $\psi_{e_{\pm\pm}}$ define edge-facets for any J_Γ.*

Moreover, the canonical projection $J_{K_n} \twoheadrightarrow J_\Gamma$ sends edge-facets of J_{K_n} to edge-facets of J_Γ, so long as $\mathrm{V}(K_n) = \mathrm{V}(\Gamma)$.

Proof. Fix an edge $e' = (v_1, v_2) \in \mathrm{E}(\Gamma)$. Then $J_\Gamma \subset \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$. We will show that $\dim(\ker(\psi_{e'_{\alpha,\beta}})) = |\,\mathrm{V}(\Gamma)\,| + |\,\mathrm{E}(\Gamma)\,| - 1$ for any sign pattern $\alpha, \beta \in \{\pm 1\}$. We will show that

$$B = \operatorname{span}\{\{u_v \mid v \neq v_1, v_2\} \cup \{u_e \mid e \neq e'\} \cup \{u_{e'} - \beta u_{v_1},\, u_{e'} - \alpha u_{v_2}\}\}$$

defines a basis for the kernel of $\psi_{e'_{\pm\pm}}$. Clearly, for any $u_e \in V_{K_n}$ with $e \neq e'$, we have

$u_e^\vee(u_e') = 0$, so

$$\psi_{e'_{\alpha,\beta}}\left(\alpha u_{v_1}^\vee + \beta u_{v_2}^\vee + \alpha\beta u_{e'}^\vee\right)(u_e) = 0.$$

Similarly, for any $u_v \in V_{K_n}$ with $v \neq v_1, v_2$, we have $u_v^\vee(u_{v_1}) = u_v^\vee(u_{v_2}) = 0$, so

$$\psi_{e'_{\alpha,\beta}}\left(\alpha u_{v_1}^\vee + \beta u_{v_2}^\vee + \alpha\beta u_{e'}^\vee\right)(u_v) = 0.$$

Lastly, if $\alpha = \beta$, then $\alpha^2 = \alpha$ and

$$\begin{aligned}
\psi_{e'_{\alpha,\beta}}(u_{e'} - \alpha\beta u_{v_1}) &= \left(\alpha u_{v_1}^\vee + \beta u_{v_2}^\vee + \alpha\beta u_{e'}^\vee\right)(u_{e'} - \beta u_{v_1}) \\
&= \alpha\beta - \alpha\beta \\
&= 0.
\end{aligned}$$

Similarly, we also have

$$\begin{aligned}
\psi_{e'_{\alpha,\beta}}(u_{e'} + \alpha\beta u_{v_2}) &= \left(\alpha u_{v_1}^\vee + \beta u_{v_2}^\vee + \alpha\beta u_{e'}^\vee\right)(u_{e'} - \alpha\beta u_{v_2}) \\
&= 0.
\end{aligned}$$

Thus the set

$$B = \{u_v \mid v \neq v_1, v_2\} \cup \{u_e \mid e \neq e'\} \cup \{u_{e'} - \beta u_{v_1},\, u_{e'} - \alpha u_{v_2}\}$$

is contained in the kernel of $\psi_{e'_{\pm\pm}}$ for any sign pattern. Note that this set is linearly independent and maximal; if $\psi_{e'_{\alpha,\beta}}(x) = 0$, then $x \in u_{e'}$, but then that implies $\alpha\beta = 0$. Thus, $\operatorname{span}(B) = \ker(\psi_{e'_{\pm\pm}})$ is codimension 1.

Next, we will show that edge-facets descend to edge-facets. Let $V_\Gamma := \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$ and $V_{K_n} := \mathbb{R}^{\mathrm{V}(K_n)} \oplus \mathbb{R}^{\mathrm{E}(K_n)}$.

Let $\pi : V_{K_n} \to V_\Gamma$ be the canonical linear map sending $u_v \mapsto u_v$ and $u_e \mapsto u_e$ if $e \in \mathrm{E}(\Gamma)$ and $u_e \mapsto 0$ otherwise. On the vertex set side π restricts to the identity on $\mathbb{R}^{\mathrm{V}(\Gamma)} = \mathbb{R}^{\mathrm{V}(K_n)}$, since K_n and Γ are assumed to have the same vertex set. We claim that π is a linear projection of polytopes. Indeed, for every vertex $p_\sigma \in J_{K_n}$,

$$\pi(p_\sigma) = \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) u_v + \sum_{(v_1,v_2)=e \in \mathrm{E}(\Gamma)} \sigma(v_1)\sigma(v_2) u_e = q_\sigma \in J_\Gamma.$$

Now, note that under the projection π,

$$\pi(\ker(\psi_{e'_{\pm\pm}})) = \operatorname{span}\left\{\{u_v \mid v \neq v_1, v_2\} \cup \{u_e \mid e \neq e'\} \cup \{u_{e'} - \beta u_{v_1},\, u_{e'} - \alpha u_{v_2}\}\right\} \subset V_\Gamma$$

is a basis for the kernel of the pullback of $\psi_{e'_{\pm\pm}}$ and has dimension $|\,\mathrm{V}(\Gamma)\,| + |\,\mathrm{E}(\Gamma)\,| - 1$. $\square$

This lemma shows us that edge-facets always exist and behave nicely under projections. The same cannot be said about more generic facet types.

Theorem 3.2. *For $\Gamma = T \cup \{e'\}$ a graph with a chordless cycle γ such that T is a tree, the tree edges and cycle edges classify* ***all*** *of the facets of J_Γ incident to $p_{-\mathbb{1}}$. Additionally, there are $3|E(\Gamma)| + \ell(\gamma)$ of these facets.*

Proof. In the tree scenario, for each edge $e = (v_0, v_1)$, lemma 3.7 constructs three functionals:

$$\psi_{e_{++}} := u_{v_0}^\vee + u_{v_1}^\vee + u_e^\vee, \quad \psi_{e_{+-}} := u_{v_0}^\vee - u_{v_1}^\vee - u_e^\vee, \quad \psi_{e_{-+}} := -u_{v_0}^\vee + u_{v_1}^\vee - u_e^\vee.$$

Each of these functionals corresponds to a facet constructed by considering subsets S_σ that reflect different sign patterns at v_0 and v_1. This classification ensures every edge induces these three facets incident to $p_{-\mathbb{1}}$, with $F^1_{\psi_{e_{\pm\pm}}}$ characterizing the maximal sets of S_σ determining each facet.

Now form Γ by adding a single edge to a tree T, creating exactly one chordless cycle $\gamma = (e_1, \ldots, e_\ell)$. By Lemma 3.7, each edge in Γ induces three edge-facets $F \preceq J_\Gamma$.

The chordless cycle γ yields new cycle facets due to Lemma 3.9, each corresponding to functionals ψ_{γ_k}. These $\ell(\gamma)$ functionals define fresh facets in J_Γ, extending the fan with exactly $\ell(\gamma)$ new rays (one per cycle edge).

Note that the construction of J_Γ depends only on the vertices and edges of Γ (not the order in which we construct Γ). Because a single-cycle graph has no additional chordless cycles or substructures beyond those accounted for by the newly introduced edge, no further facet-defining functionals arise.

Hence, we have the old rays from edges in T, plus three new rays for the new edge e', plus $\ell(\gamma)$ rays for the cycle facets. Thus, the total is $3|E(\Gamma)| + \ell(\gamma)$ facets of J_Γ incident to the point $p_{-\mathbb{1}}$. $\square$

So what we have is a classification of facets for graphs consisting of a single chordless cycle γ with tree graphs attached to γ—in other words, a classification of facets of J_Γ for hairy cycle graphs. Note that this theorem does not say that we can continue to inductively add edges and maintain an accurate facet classification. A natural next step is to determine how these types of polytopes behave when connected by an edge.

Theorem 3.3. *For G_1, G_2 hairy cycle graphs each containing a single chordless cycle, define $\Gamma = G_1 \cup_{e'} G_2$ as the graph obtained by connecting the two graphs by an edge e'. The edges in Γ and cycle paths on γ_1, γ_2 classify all of the facets of J_Γ, with there being a total of*

$$3|E(G_1)| + \mathrm{len}(\gamma_1) + 3|E(G_2)| + \mathrm{len}(\gamma_2) + 1$$

facets containing $p_{-\mathbb{1}}$.

Using induction, we can extend this result to a more general statement about links of graphs.

Theorem 3.4. *For $G_1, \ldots, G_k$ be hairy cycle graphs each containing a single chordless cycle, define $\Gamma = G_1 \cup_{e_1} G_2 \cup_{e_2} \cdots \cup_{e_{k-1}} G_k$ as the graph obtained by connecting the k graphs by a sequence of $k-1$ edges such that they do not form any loops. The edges in Γ and cycle* *paths on $\gamma_1, \gamma_2, \dots \gamma_k$ classify all of the facets of J_Γ, with there being a total of*

$$\sum_{i=1}^{k} 3|E(G_i)| + \operatorname{len}(\gamma_i) + 3(k-1)$$

facets containing $p_{-\mathbb{1}}$.

Example 3.2. Consider $\Gamma = G_1 \cup_{e'} G_2$ where G_1 contains a C_4 and G_2 contains a C_5, each cycle also containing trees along the cycles, with a lone edge e' connecting the two. In this case, we have

$$3|E(G_1)| + \operatorname{len}(C_4) + 3|E(G_2)| + \operatorname{len}(C_3) + 3 \cdot 1 = 64$$

facets in total. Twelve of them are shown below.

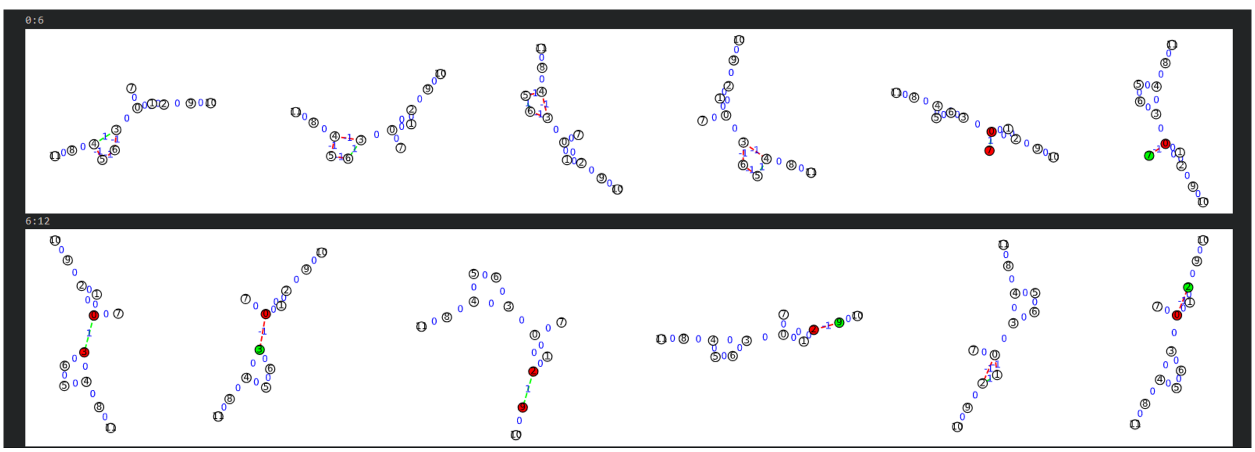

3.2.2 Global Classification of Facets

If we want to compute **all** of the facets of J_Γ (not just those incident to the basepoint $p_{-\mathbb{1}}$) then there are some minor considerations. We know that since the vertices of J_Γ are transitive under the group action $\tau \in (\mathbb{Z}/2\mathbb{Z})^n$, for the facets that are not incident to the basepoint, there is some τ such that the 'missing' facet-defining functional ψ^* is generated by one of the $\psi^{(i)}$'s in our list; i.e., $\psi^* = \psi^{(i)} \circ \tau^{-1}$ for some $i = 1, \dots, 3|E(\Gamma)| + \ell(\gamma)$.

So from a combinatorial classification viewpoint, the 'missing' facet is already included in the "(base) facet + sign-flip" dictionary we have, without having to explicitly enumerate all flips in practice.

So the brute-force approach of:

1. Classify the facets incident to

2. For each such facet, apply all sign-flips.

3. Collect the new facets at all other vertices.

is clearly exponential—and typically not the style of classification or algorithm that we want.

Instead, we will approach the previously classified facets separately and note that they result from similar facets.

Lemma 3.11. *For each edge $e \in \mathrm{E}(\Gamma)$, four edge-facets of J_Γ are defined by $\psi_{e,++}$, $\psi_{e,+-}$, $\psi_{e,-+}$, $\psi_{e,--}$.*

Proof. Only the existence of the $\psi_{e,++}$ edge-facet is required here. To see that this does define a facet of the J-cube, note that if equality is achieved by $\psi_{e,--}$ on a set of vertices $p_\sigma \in S_{++}$, then $\psi_{e,--} = \psi_{e,++} \circ \tau^{-1}$ achieves the minimum on the same set with opposite parity $p_{-\sigma} \in S_{--}$. □

Lemma 3.12. *For each chordless cycle γ, there are $2^{\mathrm{len}(\gamma)-1}$ cycle-facets, whose coefficient patterns are defined by frustration patterns on γ.*

Proof. To see the connection to geometric frustration on γ, recall the first definition of the Hamiltonian from 2.1 is the negative of our working definition

$$H_\sigma(J,h) = \sum_{v \in \mathrm{V}(\Gamma)} a_v \sigma_v + \sum_{e \in \mathrm{E}(\Gamma)} b_e \sigma_e.$$

So when connecting the embedded sugraphs representing each of the cycle-facets, the positive couplings represent edges where adjacent vertices disagree in spin, while negative couplings are where adjacent nodes agree in their spin. Viewed in this way, we can attain these frustration configurations on a cycle by assigning consecutive blocks of an odd number of edges to be frustrated (graphically, red paths of odd length separated by green paths of even length).

So for any chordless cycle γ, there are

$$\binom{\mathrm{len}(\gamma)}{1} + \binom{\mathrm{len}(\gamma)}{3} + \cdots + \binom{\mathrm{len}(\gamma)}{\mathrm{len}(\gamma)-1} = 2^{\mathrm{len}(\gamma)-1}$$

many such configurations. $\square$

In summary, we have the following global facet classification.

Theorem 3.5. *For Γ a hairy cycle graph, J_Γ consists of only these $4|\operatorname{E}(\Gamma)| + 2^{\mathrm{len}(\gamma)-1}$ edge-facets and cycle-facets.*

3.3 A Gluing Approach

3.3.1 A First Approach

Consider a polytope $Q \subset \mathbb{R}^n$ in the vector space $W = \mathbb{R}^n$, and suppose there is an involution $\nu\colon \mathbb{R}^n \to \mathbb{R}^n$ satisfying $\nu^2 = \mathrm{id}$. We form two parallel slices of Q in $\mathbb{R}^{n+1}$ by setting

$$Q_+ := Q \oplus 1 = \{(x,1) \mid x \in Q\}, \quad Q_- := \nu(Q) \oplus (-1) = \{(\nu(x),-1) \mid x \in Q\},$$

so that Q_+ is a copy of Q at the affine hyperplane $z = 1$ and Q_- is a copy of $\nu(Q)$ at $z = -1$. Taking the convex hull of these two sets then defines

$$Q_\nu = \mathrm{conv}\big(Q_+ \cup Q_-\big) \subset \mathbb{R}^{n+1}.$$

A linear functional ψ on $\mathbb{R}^{n+1}$ can be written as $\psi = \psi_W \oplus c$, where ψ_W lies in $W^\vee$ (the dual space of $\mathbb{R}^n$) and $c \in \mathbb{R}$. For any point (x, z) in $\mathbb{R}^n \oplus \mathbb{R}$,

$$\psi((x, z)) = \psi_W(x) + cz.$$

When this is restricted to Q_+, each point there has the form $(x, 1)$, so

$$\psi\big((x, 1)\big) = \psi_W(x) + c.$$

Restricting ψ to Q_- yields

$$\psi\big((\nu(x), -1)\big) = \psi_W\big(\nu(x)\big) - c = \nu^*\psi_W(x) - c,$$

where ν^* denotes the pullback under ν; so we may focus soley on restricting to Q_+.

A face $F_\psi \preceq Q_\nu$ is determined by minimizing ψ on Q_ν. The face F_ψ then consists of all points in Q_ν where ψ attains this minimal value $\min_{y \in Q_\nu} \psi(y)$. It follows that if ψ achieves its minimum exclusively on the top slice Q_+, then the resulting face lies entirely in Q_+. Alternatively, if the minimum is strictly lower on the bottom slice Q_-, the face is contained in Q_-. In the event that ψ takes the same minimal value on both slices, one obtains a face spanning both Q_+ and Q_-.

To characterize this more concretely, suppose that there is a face $F_+ \subseteq Q_+$ on which ψ_W is constant and minimized; let $c^+ = \min_{x \in F_+} \psi_W(x)$ stand for this boundary value. In Q_+, the value of ψ is therefore $c^+ + c$, when restricted to this face.

There is also a face $F_- \subseteq Q_-$ where ψ_W is also constant and minimized (equivalently, a face of Q_+ where $\nu^*\psi_W$ is constant and minimized). So let

$$c^- = \min_{y \in F_-}(\psi_W)(y) = \min_{y \in F_+}(\nu^*\psi_W)(y)$$

be this boundary value. Then in Q_-, the value of ψ is $c^- - c$ on this face.

Then we compare $2c$ with $c^- - c^+$.

- If $2c < c^- - c^+$, then $c^+ + c < c^- - c$ and $\psi_W|_{Q_+} < \psi_W|_{Q_-}$, thus the minimum occurs strictly in Q_+.

- If $2c > c^- - c^+$, then $c^+ + c > c^- - c$ and $\psi_W|_{Q_+} > \psi_W|_{Q_-}$, thus the minimum is strictly in Q_-.

- Only when $2c = c^- - c^+$, hence $c^+ + c = c^- - c$ and $\psi_W|_{Q_+} = \psi_W|_{Q_-}$, do the points in both slices simultaneously realize the same minimum, causing the face $F_\psi := \operatorname{conv}(F_+, F_-) \preceq Q_\nu$, induced by ψ, to be one of the newly generated faces in the convex hull of the top and bottom faces. Note however, that this may not necessarily be a *facet* of Q_ν.

In this way, one sees that when we extend ψ_W to ψ, it slices through Q_ν in precisely three ways, depending on how its values compare at $z = 1$ and $z = -1$. These comparisons directly determine whether the induced face lies entirely in one slice or straddles both. In order to answer this question as to when extending ψ_W to ψ induces a facet, let's recall some information.

Let A and B be two affine subsets in some ambient space. Consider $\operatorname{conv}(A \cup B)$, the convex hull of $A \cup B$. Generically, one can find the dimension formula

$$\dim\bigl(\operatorname{conv}(A \cup B)\bigr) \;=\; \dim \operatorname{Aff}(A) \;+\; \dim \operatorname{Aff}(B) \;-\; \dim\,(\,\operatorname{Aff}(A) \cap \operatorname{Aff}(B)\,) \;+\; \delta,$$

where $\delta = 0$ or 1. Concretely:

If A and B lie in the same affine hyperplane, then $\delta = 0$. In that degenerate case, $\operatorname{conv}(A \cup B)$ doesn't gain an extra dimension because its convex hull remains in a single hyperplane.

If A and B lie in nonparallel affine hyperplanes, and the intersection $\operatorname{Aff}(A) \cap \operatorname{Aff}(B)$ has dimension at most $\dim(A) + \dim(B) - 1$, then $\delta = 1$. Intuitively, the line segment joining a point in A to a point in B adds an extra direction beyond any contained in either affine set alone.

Connecting this back to our current setup above, since the faces F_+, F_- lie in separate slices, if F_ψ indeed defines a facet of Q_ν, then the pair of faces (F_+, F_-) must satisfy

$$\dim\bigl(Q_\nu\bigr) \;=\; \dim \operatorname{Aff}(F_+) + \dim \operatorname{Aff}(F_-).$$

This generic gluing strategy is how we will approach computing generic J-cubes of graphs by expanding along an edge via 'splitting' a vertex. However, note that in this process we went from $\mathbb{R}^n$ to $\mathbb{R}^{n+1}$ in a predictable manner, which suggests that we may be able to perform these computations in $\mathbb{R}^n$ instead. Indeed, we will make some adjustments in order to apply this approach, but first we will illustrate the above constructions more concretely with an example.

Example 3.3. Consider a graph Γ with just one vertex (and no edges), viewed as a contraction of P_2, the path graph on two vertices. Its state set $\{\pm 1\}^1$ has exactly two vertices: $\sigma = 1$ or $\sigma = -1$. In the usual J-cube construction, each state ± 1 maps into a point in $\mathbb{R}^{V(\Gamma)} \oplus \mathbb{R}^{E(\Gamma)}$. But since there are no edges and only one vertex, we effectively just get two points spanning a line segment. J_Γ is the convex hull of those two points, so it is a 1D segment $\operatorname{conv}(\pm(u_{v_1} + u_{v_2}))$. In coordinates,

$$Q := J_\Gamma = \{x \in \mathbb{R} \mid \lambda(1,1) + (1-\lambda)(-1,-1), \lambda \in [0,1]\}.$$

Note that we use the coordinates for the vertices v_1, v_2; the lone edge in P_2 contracts these two vertices along the edge e to a single vertex $v_\pm$. The uncontracting map sends $v_\pm \mapsto u_{v_+} + u_{v_-}$, here denoted v_1, v_2. Hence our 'glued' polytope will reside in in $\mathbb{R}^3$ after uncontracting.

Next, we define the involution $\nu : \mathbb{R}^2 \to \mathbb{R}^2$ as the linear map that changes sign on the u_{v_2}-coordinate

$$\nu(u_{v_1}) = u_{v_1} \text{ and } \nu(u_{v_2}) = -u_{v_2}$$

or in coordinates, $\nu(x,y) = (x,-y)$.

We next move up one dimension, to $\mathbb{R}^2 \oplus \mathbb{R} = \mathbb{R}^3$ with coordinates (x,y,z). Define

$$\begin{aligned} Q_+ &:= Q \oplus 1 \\ &= J_\Gamma \oplus 1 \\ &= \mathrm{conv}(\pm(u_{v_1} + u_{v_2})) \oplus 1 \\ &= \{(x,x,1) \in \mathbb{R}^3 \mid x \in [-1,1]\}. \end{aligned}$$

placing a copy of the J-cube at the plane $z = 1$. Simultaneously, define

$$\begin{aligned} Q_- &:= \nu(Q) \oplus (-1) \\ &= \nu(J_\Gamma) \oplus (-1) \\ &= \mathrm{conv}(\pm(\nu(u_{v_1}) + \nu(u_{v_2})) \oplus (-1) \\ &= \mathrm{conv}(\pm(u_{v_1} - u_{v_2}) + u_e) \\ &= \{(x,-x,-1) \in \mathbb{R}^3 \mid x \in [-1,1]\}. \end{aligned}$$

Thus Q_+ is a horizontal line segment at the upper $z = 1$ slice, running from $(1,1,1)$ to $(-1,-1,1)$ and Q_- is a vertical line segment at the lower $z = -1$ slice, running from $(1,-1,-1)$ to $(-1,1,-1)$.

Viewing these two slices together in $\mathbb{R}^3$, we get four distinct corner points: $(-1,0,1)$ and $(1,0,1)$ across the top slice, $(0,-1,-1)$ and $(0,1,-1)$ across the bottom slice.

Taking the convex hull of those four points yields a 3-dimensional polytope with four vertices. One can check that they do not all lie in a single plane, so their convex hull is a tetrahedron. Hence:

$$Q_\nu = \mathrm{conv}\big(Q_+ \cup Q_-\big) = (\text{a tetrahedron in } \mathbb{R}^3).$$

Finally, recall that the J-cube of the path graph P_2 on two vertices (and one edge) is known to be a tetrahedron. Indeed, P_2 has state space $\{\pm 1\}^2$ of four distinct vertices, and those vertices map into $\mathbb{R}^2 \oplus \mathbb{R} \cong \mathbb{R}^3$ in such a way that their convex hull is a tetrahedron. The four corner points we see above match the typical sign configurations $\sigma\colon \{\text{two vertices}\} \to \{\pm 1\}$.

Thus, by starting with the one-vertex J-cube (a line segment) in the plane $\mathbb{R}^2$ and applying the involution ν to produce two slices in $\mathbb{R}^3$, we exactly reconstruct the J-cube of P_2. This demonstrates how we can 'glue' two copies of a 1-dimensional J-cube into the 3-dimensional J-cube of a larger graph, by 'expanding' P_1 along the edge $e \in E(P_2)$.

Hence we have:

$$J_{P_1} \cong J_{\Gamma/e} \quad \xrightarrow[\nu]{\text{glue}} \quad Q_\nu = J_{P_2} = J_\Gamma,$$

showing how we can start from the J-cube of P_1, and produce the J-cube associated to P_2 by splitting a vertex into an edge. $\square$

We summarize the gluing process in full generality below.

Given a graph Γ' which contracts along an edge $e' = (v_+, v_-)$ to $\Gamma := \Gamma'/e'$, with the vertices of e' contracting to the vertex $v_\pm$, define the inclusion $i : J_\Gamma \to \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$ by

$$i := \begin{cases} u_{v_\pm} \mapsto u_{v_+} + u_{v_-} & \\ u_v \mapsto u_v & v \neq v_\pm \\ u_e \mapsto u_e & e \neq e' \end{cases}$$

Under this embedding, we 'expand' the vertex $v_\pm$ to the vertices it originated from, and we obtain a copy of J_Γ in a one-higher dimension space

$$Q := i(J_\Gamma) \subset W := \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}.$$

Next, we define the involution $\nu : \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)} \to \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)}$ on this space by flipping the sign on the vertex v_-, along with a subset of its neighboring edges $S_- \subset\neq \partial\{v_-\}$, chosen in advance:

$$\nu := \begin{cases} u_v \mapsto u_v & v \neq v_- \\ u_{v_-} \mapsto -u_{-v} & \\ u_e \mapsto u_e & e \notin S_- \\ u_e \mapsto -u_e & e \in S_- \end{cases}$$

Intuitively, the choice of S_- determines which edges were incident to v_- prior to contracting along e' on Γ'. The choice of S_- will determine the expanded graph we obtain, and there is no clear notion of independence of S_-; for example, a star graph with n leaves can be expanded into several nonisomorphic graphs, by splitting the root vertex in $\binom{n}{2}$ different ways, resulting in a total of $\lfloor \frac{n}{2} \rfloor$ isomorphism classes of graphs. We'll revisit this question in later work.

Next, we embed the J-cube of the contracted graph Γ in the ambient space for its expansion along e', or Γ'. Thus we obtain two isomorphic copies of J_Γ at the slices corresponding to $\sigma(e') = \pm 1$:

$$Q_+ = Q \oplus 1 \subset W \oplus \mathbb{R} \subset \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma')}$$

$$Q_- := \nu(Q) \oplus (-1) \subset W \oplus \mathbb{R} \subset \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma')}$$

Finally, we take the convex hull of the two slices to obtain the J-cube for the uncontracted graph Γ'

$$Q_\nu := \mathrm{conv}(Q \cup \nu(Q)) = J_{\Gamma'} \subset V_{\Gamma'} := \mathbb{R}^{\mathrm{V}(\Gamma')} \oplus \mathbb{R}^{\mathrm{E}(\Gamma')}.$$

Now, given any linear functional $\psi_W \in W^\vee$ with representation

$$\psi_W = \sum_{\substack{v \in \mathrm{V}(\Gamma) \\ v \neq v_-, v_+}} a_v u_v^\vee + \sum_{e \in \Gamma} b_e u_e^\vee + a_{v_\pm}(u_{v_+}^\vee + u_{v_-}^\vee)$$

we extend it to a functional $\psi = \psi_W \oplus b_{e'} \in V_{\Gamma'}^\vee$ by defining the edge constant $b_{e'}$ in the following manner:

$$\begin{aligned} b_{e'} &:= \frac{1}{2}\left(\min_{x\in Q_-}\psi_W(x) - \min_{x\in Q_+}\psi_W(x)\right) \\ &= \frac{1}{2}\left(\min_{x\in Q}\nu^*\psi_W(x) - \min_{x\in Q}\psi_W(x)\right) \\ &:= \frac{1}{2}\left(\Delta_Q\right). \end{aligned}$$

Recall that this choice of constant resulted from the requirement '$2c = c^- - c^+$', from our more generic discussion.

Extending ψ in this manner, if the functional ψ_W defines a pair of faces $(F_-, F_+) \preceq (Q_-, Q_+)$ of Q in the usual manner, then ψ defines a face $F_\nu := \mathrm{conv}(F_- \cup F_+) \preceq Q_\nu$ which is not fully contained in either $W \oplus 1$ or $W \oplus (-1)$.

Example 3.4. Revisiting the setup from Example 3.3, recall that we have the following bounds on the dimensions for the facets of J_Γ:

$$\begin{aligned} \dim\big(Q_\nu\big) - 1 &= \dim \mathrm{Aff}(F_+) + \dim \mathrm{Aff}(F_-) - \dim\left(\mathrm{Aff}(F_+) \cap \mathrm{Aff}(F_-)\right) + 1 \\ 2 &= \dim \mathrm{Aff}(F_+) + \dim \mathrm{Aff}(F_-) - 0 + 1 \\ 1 &= \dim \mathrm{Aff}(F_+) + \dim \mathrm{Aff}(F_-) \end{aligned}$$

One can readily see that the only pairs of faces (F_+, F_-) that satisfy this are a line segment and a point in Q_ν. We summarize the facet structure of Q_ν below, in terms of pairs of faces (F_+, F_-):

F_+	F_-	$\dim \mathrm{Aff}(F_+, F_-)$
$u_{v_+} + u_{v_-} + u_{e'}$	$\mathrm{conv}(\pm(u_{v_+} - u_{v_-}) - u_{e'})$	$(0, 1)$
$-u_{v_+} - u_{v_-} + u_{e'}$	$\mathrm{conv}(\pm(u_{v_+} - u_{v_-}) - u_{e'})$	$(0, 1)$
$\mathrm{conv}(\pm(u_{v_+} + u_{v_+}) + u_{e'})$	$u_{v_+} - u_{v_-} - u_{e'}$	$(1, 0)$
$\mathrm{conv}(\pm(u_{v_+} + u_{v_+}) + u_{e'})$	$-u_{v_+} + u_{v_-} - u_{e'}$	$(1, 0)$

The linear functions that induce each of these facets are listed below:

Next, we'll verify the construction by computing ψ from $\psi_W \in Q^\vee$.

□

$\psi \sim F_\psi \preceq Q_\nu$	$\dim \mathrm{Aff}(F_+, F_-)$
$-u^\vee_{v_+} - u^\vee_{v_-} + u^\vee_{e'}$	$(0,1)$
$u^\vee_{v_+} + u^\vee_{v_-} + u^\vee_{e'}$	$(0,1)$
$-u^\vee_{v_+} + u^\vee_{v_-} - u^\vee_{e'}$	$(1,0)$
$u^\vee_{v_+} - u^\vee_{v_-} - u^\vee_{e'}$	$(1,0)$

ψ_W	$\nu^*\psi_W$	$\frac{1}{2}(\min \nu^*\psi_W - \min \psi_W)$	ψ
$u^\vee_{v_+} + u^\vee_{v_-}$	$u^\vee_{v_+} - u^\vee_{v_-}$	$\frac{1}{2}(0-(-2)) = 1$	$u^\vee_{v_+} + u^\vee_{v_-} + (1)u^\vee_{v'_e}$
$-u^\vee_{v_+} - u^\vee_{v_-}$	$-u^\vee_{v_+} + u^\vee_{v_-}$	$\frac{1}{2}(0-(-2)) = 1$	$-u^\vee_{v_+} - u^\vee_{v_-} + (1)u^\vee_{v'_e}$
$u^\vee_{v_+} - u^\vee_{v_-}$	$u^\vee_{v_+} + u^\vee_{v_-}$	$\frac{1}{2}(-2-0) = -1$	$u^\vee_{v_+} - u^\vee_{v_-} + (-1)u^\vee_{v'_e}$
$-u^\vee_{v_+} + u^\vee_{v_-}$	$-u^\vee_{v_+} - u^\vee_{v_-}$	$\frac{1}{2}(-2-0) = -1$	$-u^\vee_{v_+} + u^\vee_{v_-} + (-1)u^\vee_{v'_e}$

In this example, computing the edge coefficient $b_{e'}$ is time-consuming, so we provide a reduction.

Lemma 3.13. *If ψ_W attains its minimum on $F \preceq Q$, and $\nu^*\psi_W$ attains its minimum on $F^* \preceq Q$ where $F \cap F^* \neq \emptyset$, then the coefficient for $u^\vee_{e'}$ can be simply computed as a local sum at $v_\pm$*

$$b^*_{e'} = \min_{p \in F_\psi \cap F_{\nu^*\psi}} \left(-a_{v_\pm} u^\vee_{v_\pm}(p) - \sum_{e \in S_-} b_e u^\vee_e(p) \right).$$

Indeed, suppose we have $p \in F \cap F^* \neq \emptyset$, then

$$\begin{aligned}
2b_{e'} &= \min_{x \in Q} \nu^*\psi_W(x) - \min_{x \in Q} \psi_W(x) \\
&= \nu^*\psi_W(p) - \psi_W(p) \\
&= \left(\sum_{\substack{v \in \mathrm{V}(\Gamma) \\ v \neq v_+, v_-}} a_v u^\vee_v + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e \notin S_-}} b_e u^\vee_e + a_{v_+} u^\vee_{v_+} - a_{v_-} u^\vee_{v_-} - \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e \in S_-}} b_e u^\vee_e \right)(p) \\
&\quad - \left(\sum_{\substack{v \in \mathrm{V}(\Gamma) \\ v \neq v_+, v_-}} a_v u^\vee_v + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e \notin S_-}} b_e u^\vee_e + a_{v_+} u^\vee_{v_+} + a_{v_-} u^\vee_{v_-} + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e \in S_-}} b_e u^\vee_e \right)(p) \\
&= -2 \left(a_{v_-} u^\vee_{v_-} + \sum_{\substack{e \in \mathrm{E}(\Gamma) \\ e \in S_-}} b_e u^\vee_e \right)(p)
\end{aligned}$$

Thus, in this situation where ψ_W defines two adjacent faces F^*, F of Q, we obtain the reduction

$$\begin{aligned} b_{e'} &= \min_{p \in F \cap F^* \preceq Q} \left(-a_{v_\pm} u_{v_-}^\vee(p) - \sum_{e \in S_-} b_e u_e^\vee(p) \right) \\ &= \min_{p \in i^{-1}(F) \cap i^{-1}(F^*) \preceq J_\Gamma} \left(-a_{v_\pm} u_{v_\pm}^\vee(p) - \sum_{e \in S_-} b_e u_e^\vee(p) \right). \end{aligned}$$

□

So in this manner, if we know the facet structure on J_Γ, then we can attain the facet structure on $J_{\Gamma'}$. This begs the question though, is is possible to perform our calculations solely within the dual space of an initial J-cube, rather than in its imbedding as Q? This is what we answer now. The key idea is to express the iterative step abstractly, entirely within the original space $\mathbb{R}^m$, via cones and normal fans, mirroring ideas from toric geometry.

3.3.2 Refinement of Edge Expansion

Given a graph Γ, we have a polytope $J_\Gamma \subset \mathbb{R}^m$ whose combinatorial structure is encoded by its normal fan $\mathcal{N}(J_\Gamma)$. Each cone in the normal fan corresponds bijectively to a face of J_Γ, with the dimension of a face equal to the codimension of the corresponding cone. The iterative step from Γ to Γ' is then described purely combinatorially by modifying the normal fan of J_Γ.

We start with a given graph Γ and a distinguished vertex $v_\pm$. To proceed inductively, we consider a new graph Γ' obtained from Γ by splitting the vertex $v_\pm$ into v_+ and v_- and add exactly one edge $e' = (v_+, v_-)$. This inductive step takes place abstractly through the geometry of cones in $\mathcal{N}(J_\Gamma) \cong (\mathbb{R}^m)^\vee$:

The cones in $\mathcal{N}(J_\Gamma)$ are defined as:

$$C_F = \{\psi \in (\mathbb{R}^m)^\vee : \psi|_F \text{ is minimized}\},$$

for each face $F \preceq J_\Gamma$. Each such cone is dual to a face in the usual manner.

The involution ν, acting on the base space $\mathbb{R}^m$, arises from the process of splitting $v_\pm$ with a labeling of edges. It is given explicitly by flipping the sign of the coordinate associated to edges in a subset $S_- \subseteq E(\Gamma)$ adjacent to $v_\pm$ and by flipping and scaling the coordinate associated to the vertex $v_\pm$. Recall that the graph Γ' obtained by this process is dependent on the choice of labeling S_-.

The involution ν on $\mathbb{R}^m$ induces a dual action ν^* on the dual space $(\mathbb{R}^m)^\vee$. This dual action naturally carries cones in the normal fan of J_Γ to new cones, possibly permuting or pairing them up. Specifically:

Given a cone $C_\psi \subseteq \mathcal{N}(J_\Gamma)$ defined by a functional ψ, the involution produces a new cone:

$$C_{\nu^*\psi} := \nu^*(C_\psi) \subseteq \mathcal{N}(J_\Gamma).$$

For example, in the case where we are adding a leaf to a graph by expanding a vertex $v_\pm$ along a leaf edge e', we will have the edge labeling $S_- = \emptyset$; as there are no edges 'beyond' $v_\pm$. So in this trivial case, the involution merely reflects the coordinate $u_{v_\pm}$.

Thus, we have pairs of cones $(C_\psi, C_{\nu^*\psi})$, reflecting the geometric involution. Each pair of cones corresponds to pairs of faces $(F_\psi, F_{\nu^*\psi})$ in the original polytope J_Γ.

At this stage, we have not yet increased the dimension, as everything occurs within the space $\mathbb{R}^m$ and $(\mathbb{R}^m)^\vee$. The dimension count of faces and cones is captured neatly by standard duality:

If $\dim F_\psi = d$, the corresponding cone $C_\psi \subseteq (\mathbb{R}^m)^\vee$ has codimension d. The involution doesn't alter the dimension, by merely flipping or reflecting coordinates and potentially permuting cones and their associated faces.

To proceed towards a geometric construction (the J-cube inductive step), we recall a dimension formula that connects the faces corresponding to pairs of cones under the involution. Given two cones C_ψ and $C_{\nu^*\psi}$, which correspond to faces F_ψ and $F_{\nu^*\psi}$, define a glued cone abstractly as:

$$C_{\psi,\nu^*\psi} := \operatorname{span}(C_\psi, C_{\nu^*\psi}),$$

which, under an appropriate embedding, corresponds to a face in the higher-dimensional J-cube $J_{\Gamma'}$.

The key dimension identity at the abstract level is the sum:

$$\dim(F_{\psi,\nu^*\psi}) = \dim(\operatorname{Aff}(F_\psi)) + \dim(\operatorname{Aff}(F_{\nu^*\psi})),$$

where $F_{\psi,\nu^*\psi} \preceq J_{\Gamma'}$ is the face obtained in the next step of the inductive construction.

This dimension formula reflects how the involution links and glues the faces, creating a new face whose dimension is precisely the sum of dimensions from the upper-lower faces.

To translate this abstract construction to the familiar geometric process from before, we now perform explicit embeddings:

First, the vertex-splitting step is realized by embedding:

$$i : \mathbb{R}^m \to \mathbb{R}^{m+1}, \quad u_{v_\pm} \mapsto u_{v_+} + u_{v_-},$$

with identity on all other coordinates. This step geometrically realizes the abstract imbedding i as a splitting of the vertex coordinate $v_\pm$, while separating the edges in S_- from v_+ (and those in $\mathrm{E}(\Gamma) \setminus S_-$ from v_-).

Second, the addition of the new edge e' introduces another coordinate which mirrors the additional edge that connects v_+ to v_-:

$$j : \mathbb{R}^{m+1} \to \mathbb{R}^{m+2}, \quad Q \mapsto Q \oplus 1, \quad \nu'(Q) \mapsto \nu'(Q) \oplus (-1),$$

by assigning coordinates $u_{e'} = \pm 1$ to the new edge coordinate e', resulting in the upper-and-lower slices $Q \oplus 1$, $\nu'(Q) \oplus (-1)$. The pair of functionals $\psi, \nu^*\psi \in \mathcal{N}(J_\Gamma)$ under the pushforwards become

$$\begin{aligned}\psi = \varphi + a_{v_\pm} u_{v_\pm}^\vee \quad &\mapsto \quad i^*\psi = \varphi + a_{v_\pm}(u_{v_+}^\vee + u_{v_-}^\vee)\\ \nu^*\psi = \nu^*\varphi + b_{v_\pm} u_{v_\pm}^\vee &\mapsto \nu'^* i^*\psi = \nu'^*\varphi + b_{v_\pm}(u_{v_+}^\vee - u_{v_-}^\vee)\end{aligned}$$

Define the perpendicular shift of $i^*\psi$ as

$$\begin{aligned}\tilde{\psi} &:= \varphi + a_{v_\pm}(u_{v_+}^\vee + u_{v_-}^\vee) + b_{v_\pm}(u_{v_+}^\vee - u_{v_-}^\vee)\\ &= \varphi + (a_{v_\pm} + b_{v_\pm})u_{v_+}^\vee + (a_{v_\pm} - b_{v_\pm})u_{v_-}^\vee.\end{aligned}$$

Note that $\tilde{\psi} \in (\mathbb{R}^{m+1})^\vee$ projects to ψ and lies in a perpendicular translate of $i(J_\Gamma) = Q \subset \mathbb{R}^{m+1}$, on which $u_{v_+}^\vee - u_{v_-}^\vee$ is zero. Pushing $\tilde{\psi}$ forward to $(\mathbb{R}^{m+2})^\vee = \mathcal{N}(J_{\Gamma'})$, we obtain

$$\Psi = \varphi + (a_{v_\pm} + b_{v_\pm})u_{v_+}^\vee + (a_{v_\pm} - b_{v_\pm})u_{v_-}^\vee + b_{e'}^* u_{e'}^\vee$$

where the real coefficient $b_{e'}^*$ is computed as

$$\begin{aligned}b_{e'}^* &:= \frac{1}{2}\left(\min_{p\in J_\Gamma} \nu^*\psi(p) - \min_{p\in J_\Gamma} \psi(p)\right)\\ &\geq \min_{p\in F_\psi \cap F_{\nu^*\psi}} \left(-a_{v_\pm} u_{v_\pm}^\vee(p) - \sum_{e\in S_-} b_e u_e^\vee(p)\right)\end{aligned}$$

where equality holds if and only if $p \in F_\psi \cap F_{\nu^*\psi} \neq \emptyset$.

The previously paired cones $(C_\psi, C_{\nu^*\psi})$ now correspond to the upper and lower faces $(F_+, F_-) \subseteq J_{\Gamma'}$. The abstract dimension formula translates directly to the geometric embedding:

$$\dim(F_{\psi,\nu^*\psi}) = \dim(\mathrm{Aff}(F_+)) + \dim(\mathrm{Aff}(F_-)),$$

reflecting precisely how these cones (faces) combine in the embedded geometry. The connection to the earlier inductive constructions is now clear:

Abstractly, in the lower-dimensional space $\mathbb{R}^m$, we chose a pairing of cones via involution, resulting in a dimension count encoded in the normal fans. Concretely, once the involution and pairing of cones are established, we can explicitly embed into higher-dimensional spaces, concretely realizing the cone-theoretic pairing as the faces of a higher-dimensional J-cube $J_{\Gamma'}$ as before.

Next, we note that in the reverse process—edge-contraction—contracting edges of a simple graph may result in multi-edges or loops. The following result shows that these edges may be 'pruned' without affecting the facet structure of the J-cube.

Theorem 3.6. *Let Γ be a simple graph and let Γ' be a non-simple graph with $\mathrm{V}(\Gamma) = \mathrm{V}(\Gamma')$ and $\mathrm{E}(\Gamma) \subsetneq \mathrm{E}(\Gamma')$ maximal. That is, Γ can be viewed as a maximal copy of Γ' with only loops and multiedges removed. Then $J_{\Gamma'} \cong J_\Gamma$.*

Proof. Note that if two vertices v_1, v_2 are connected by multiple edges $e_i = (v_1, v_2)$, with $i = 1, \ldots, \mathrm{mult}(e_1)$, then for any $\sigma \in \mathcal{A}$,

$$\sigma(e_1) = \sigma(v_1)\sigma(v_2) = \sigma(e_i)$$

for any i and $\sigma \in \mathcal{A}$. Thus, we have $(p_\sigma)_{u_{e_1}} = (p_\sigma)_{u_{e_i}}$.

Similarly, if $e = (v, v)$ is a loop in Γ', then for any $\sigma \in \mathcal{A}$, we have $\sigma(e) = \sigma(v)^2 = 1$. Thus, $(p_\sigma)_{u_e} = 1$ if e is a loop.

So in essence, the vertices of $J_{\Gamma'}$ are constant in the u_e-coordinates when e is a loop or has nontrivial multiplicity, and we can remove these extra directions without reducing the dimension of $J_{\Gamma'}$. So for any $q_\sigma \in \mathrm{vert}(J_\Gamma)$ and $p_\sigma \in \mathrm{vert}(J_{\Gamma'})$, we have

$$p_\sigma = \underbrace{\sum_{v \in \mathrm{V}(\Gamma')} \sigma(v) u_v + \sum_{\substack{e \in \mathrm{E}(\Gamma') \\ e \text{ a loop at } v}} u_e}_{\text{vertex lifts}} + \underbrace{\sum_{\substack{e=(v_1,v_2) \in \mathrm{E}(\Gamma') \\ e \text{ not a loop}}} \sigma(e) \left(\sum_{\substack{i=1 \\ e_i=(v_1,v_2)}}^{\mathrm{mult}(e)} u_{e_i} \right)}_{\text{edge lifts}}$$

$$q_\sigma = \sum_{v \in \mathrm{V}(\Gamma)} \sigma(v) u_v + \sum_{e \in \mathrm{E}(\Gamma)} \sigma(e) u_e + \tilde{p_\sigma}$$

Thus, J_Γ and $J_{\Gamma'}$ are affinely isomorphic.

$\square$

3.4 Computing the Fiber Fan

Recall the following diagram from earlier

$$\begin{tikzcd}
 & \mathbb{R}^{E(\Gamma)} \arrow[r] & \mathbb{R}^{\mathrm{V}(\Gamma)} \oplus \mathbb{R}^{\mathrm{E}(\Gamma)} \arrow[dl, bend left, two heads, "\pi_J"] \\
\mathbb{R}^{2^n-n} \arrow[ur, two heads, "i^*"] \arrow[r, hook] & \mathbb{R}^{\mathcal{A}} \arrow[u, two heads, "\gamma^*"'] \arrow[ur, two heads, "\varphi"'] \arrow[d, two heads, "\pi"] & \\
\mathbb{R}^{2^n-n-m} \arrow[u, hook] \arrow[ur, hook] & \mathbb{R}^{V(\Gamma)} &
\end{tikzcd}$$

All arrows are affine, surjective maps of polytopes. Lemma 2.3 of [4] therefore gives

$$\Sigma(J_\Gamma, Q) \;=\; \varphi\big(\Sigma(\Delta_{\mathcal{A}}, Q)\big) = \varphi\left(\Sigma(Q)\right).$$

In other words, the fiber polytope is a projection of the secondary polytope. Dually, the fiber fan carves out a linear subspace of the secondary fan over which the π_J-coherent subdivisions are realized by a Hamiltonian.

A weight vector $\omega \in (\mathbb{R}^{\mathcal{A}})^*$ lifts the n-cube to $\{(\sigma, \omega(\sigma)) : \sigma \in \mathcal{A}\} \subset \mathbb{R}^{\mathcal{A}} \oplus \mathbb{R}$. The lower convex hull defines a coherent subdivision of Q; varying ω sweeps the whole secondary fan of Q. But $(\mathbb{R}^{\mathcal{A}})^*$ has 2^n coordinates, and a generic weight vector yields a full 6-tetrahedra triangulation of the 3-cube, 384 for the 4-cube, and astronomically many in higher dimension—an intractable search space.

When Γ is of the form "cycle+trees" (or, more generally, when every 2-connected block of Γ is a single chordless cycle) we showed in Section 3 that every facet normal is one of

- $4|E(\Gamma)|$ edge-normals $\psi_{e,\pm\pm} = a\, u^*_{v_i} + b\, u^*_{v_j} + ab\, u^*_e$, where $a, b = \pm 1$
- $2^{\ell(\gamma)-1}$ cycle-normals $\psi_{\gamma,k} = \sum_{i \in \mathcal{I}} u^*_{e_i} - \sum_{i \notin \mathcal{I}} u^*_{e_i}$, where $\mathcal{I} \subset \mathcal{P}(E(\gamma))$ contains disjoint,

odd-length edge paths.

That is, the entire normal fan $\mathcal{N}(J_\Gamma)$ is spanned by only

$$4|E(\Gamma)| + \sum_{\text{chordless cycles } \gamma} 2^{\text{len}(\gamma)-1}$$

primitive vectors in $(\mathbb{R}^{\mathcal{A}})^*$. Working in this subfan collapses the exponential search to something linear in the size of Γ.

In general, an arbitrary functional ψ will induce a coherent lifting of Q into a higher-dimensional polyhedra. Constructing the lower convex hull of the lifted points and projecting back down yields a well-defined regular (coherent) subdivision of the cube Q.

Applying techniques from Section 1.2.5, we vertically lift vertices $\sigma \in Q$ to height $h_\psi(\sigma)$ and project the lower convex hull of this lifted configuration back down to Q in order to produce a regular (coherent) subdivision whose combinatorial structure is dictated explicitly by the facet-defining linear functional ψ. The resulting subdivision structure is straightforward to compute and enumerate explicitly.

1. Pick a facet normal $\psi = \sum_{v \in V} a_v \, u_v^* + \sum_{e \in E} b_e \, u_e^*$ of J_Γ.
2. Induce a height on the cube: for each vertex $\sigma \in \{\pm 1\}^n$ set $h_\psi(\sigma) = \psi\,(\varphi(\sigma)) = \sum_{v \in \mathrm{V}(\Gamma)} a_v \sigma_v + \sum_{e=(v_i,v_j) \in \mathrm{E}(\Gamma)} b_e \, \sigma_{v_i} \sigma_{v_j}$.
3. Lift Q to $\{(\sigma, h_\psi(\sigma))\} \subset \mathbb{R}^{n+1}$ and take the lower hull—these can be computed in higher dimensions by the process described in Chapter 5 of [5]. The maximal cells then project back down to $\mathbb{R}^n$ to obtain a coherent subdivision $\mathcal{S}_\psi$.
4. The set of all $\lambda \cdot \psi$ with $\lambda > 0$ that induce the same subdivision are one–dimensional cone. Its ray $\mathbb{R}_{>0}\, \psi$ is a generator of the secondary fan $\mathcal{N}(\Sigma(Q))$.
5. Repeating for every facet normal of J_Γ lists all rays of the secondary fan. Higher-dimensional cones arise as common refinements.

In practice we need only compare which heights coincide—two vertices ending up at the same value means the edge between them is not broken.

3.4.1 Why So Few Coherent Triangulations Appear

For the 3–cycle example we listed the 15 facet–normals of the J-cube and found only six distinct cones of the secondary fan that correspond to full triangulations of the cube $Q = [-1,1]^3$—not the 74 regular triangulations we expected. This was a consequence of the π_J-coherent subdivisions residing in the fiber fan $\Sigma(J_\Gamma, Q)$ which resides as a linear subspace of $\Sigma(Q)$.

In fact, this can be seen from a more basic example, as the Hamiltonian structure imposes fairly strict conditions on the valuations that are allowed.

Example 3.5. Concretely, for the 3-cycle one has

$$H_\sigma(J,0) = J_{12}\sigma_1\sigma_2 + J_{23}\sigma_2\sigma_3 + J_{13}\sigma_1\sigma_3.$$

If $J_{12} = J_{23} = J_{13}$ the eight energies $H_{\sigma^{(i)}}(J,0)$ occur in only three distinct values (vertices of Q), so no functional with that choice of edge coefficients can produce a generic triangulation.

□

3.4.2 Regular triangulations vs. Hamiltonian degeneracy

Denote the zero-field Hamiltonian as

$$H_\sigma(J,0) = \sum_{e=(v_i,v_j)\in \mathrm{E}(\Gamma)} J_e\,\sigma(v_i)\sigma(v_j).$$

The even part of h_ψ is literally the valuation $\sigma \mapsto \langle b, \gamma(\sigma)\rangle$ of the monomial $t^{H_\sigma(J,0)}x^\sigma$. If two spin configurations σ, σ' have the same energy then $E(\sigma) = E(\sigma')$. Hence degeneracies of the Hamiltonian collapse heights in the lifting construction, forcing coarse subdivisions which exclude these valuations from appearing in the moduli space.

For a general graph with n vertices and m edges

- The dimension of the Ising weight space $\dim W = n+m$ rather than the full $\dim(\mathbb{R}^{\mathcal{A}})^\vee = 2^n$ parameter space
- codimension in secondary fan $2^{n-1} - n - e$ linear relations $w(\sigma) + w(-\sigma) = \text{const.}$
- expected number of cones grows polynomially in n (empirically $O(n^2)$) instead of exponentially.

Hence the fiber-polytope method stays practical even when n is large, which is much better than what direct enumeration of all 2^n weight vectors could handle—graphs with $m + n \geq 15$ were approaching computational limits on a test computer with the direct enumeration approach.

The spin-glass symmetry and the edge-product constraints force Ising valuation functionals to live in a low-dimensional subspace of the full height space. Geometrically this means that the secondary fan of the cube, when intersected with that subspace, contains only a handful of maximal cones; algebraically it reflects the degeneracy pattern of the zero-field Hamiltonian.

For the 3-cycle this yields exactly six coherent triangulations—precisely the vertices of the fiber polytope $\Sigma(J_{C_3}, Q)$. The same mechanism explains the sparcity of regular triangulations for arbitrary graphs Γ: every ground-state symmetry translates into linear relations among heights and collapses whole regions of the secondary fan, leaving a computationally manageable family of cones.

3.4.3 Connection to Tropical Varieties

This explicit subdivision construction reconnects directly to the tropical geometry framework introduced earlier in the dissertation. Recall the interaction polynomial $\mathcal{E}_J$ associated to the

Ising model on our graph Γ; with the functionals $\psi \in \mathcal{F}_{J_\Gamma}$, we may obtain

$$\mathcal{E}_{\pi^*\psi}(\vec{x}) = \sum_{\sigma \in \mathcal{A}} t^{h_\psi(\sigma)} x^{\vec{\sigma}} \in K[x_1, \dots, x_n],$$

with Newton polytope $Q = \operatorname{conv}(\mathcal{A})$ and weights generated by h_ψ. Tropicalizing this polynomial yields:

$$\operatorname{trop}(\mathcal{E}_J)(w) = \min_{\sigma \in \mathcal{A}} (h_\psi(\sigma) + w \cdot \sigma),$$

whose zero locus defines a tropical variety $V(\operatorname{trop}(\mathcal{E}_J)) \subset \mathbb{R}^{\mathrm{V}(\Gamma)}$.

This tropical variety is dual to the coherent subdivision induced by the heights $h_\psi(\sigma)$. Each functional ψ selects a particular combinatorial type of tropical hypersurface associated with a specific energy degeneracy locus from the original Ising Hamiltonian setup. Thus, the facets of the J-cube provide a direct bridge back to the degenerate ground states of Γ that originally motivated this work.

Due to the explicit classification of facet normals for J_Γ, the normal fan $\mathcal{N}(J_\Gamma)$ is completely described by a finite, known set of cones, which may subsequently be used to describe the secondary fan $\Sigma(Q)$. Each of these cones corresponds exactly to one of the coherent subdivisions described above. Instead of exponentially many possible subdivisions, we explicitly enumerate a finite set of cones, substantially simplifying computations and making the enumeration feasible even for larger cases.

This approach avoids the computational explosion inherent in direct subdivision enumeration, while also explicitly reconnecting the construction to the tropical geometry framework motivating this dissertation. The key simplifications come from the explicit facet classification of the polytope J_Γ for classified graphs Γ, ensuring that the computation remains manageable even as the number of vertices of Γ grows. It remains to be seen whether this classification can be expanded to higher-connectivity graphs.

Bibliography

[1] D. Avis and K. Fukuda, *A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra*, Discrete & Computational Geometry, 8 (1992), p. 295–313.

[2] ——, *Reverse search for enumeration*, Discrete Applied Mathematics, (1996).

[3] L. J. Billera, P. Filliman, and B. Sturmfels, *Constructions and complexity of secondary polytopes*, Advances in Mathematics, 83 (1990), p. 155–179.

[4] L. J. Billera and B. Sturmfels, *Fiber polytopes*, Annals of Mathematics, 135 (1992), pp. 527–549.

[5] J. A. De Loera, J. Rambau, and F. Santos, *Triangulations in Mathematics*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 1–41.

[6] C. Diemer, L. Katzarkov, and G. Kerr, *Compactifications of spaces of landau–ginzburg models*, Izvestiya: Mathematics, 77 (2013), p. 487.

[7] H. Edelsbrunner, *Algorithms in combinatorial geometry*, EATCS monographs on theoretical computer science, Springer, Berlin Heidelberg, 1987.

[8] I. M. Gel'fand, M. Kapranov, and A. V. Zelevinsky, *Discriminants, Resultants, and Multidimensional Determinants*, Springer, 1994.

[9] I. Hiroshi, T. Masada, F. Takeuki, and K. Imai, *Enumerating triangulations in general dimensions*, International Journal of Computational Geometry &; Applications, (2002).

[10] D. Maclagan and B. Sturmfels, *Introduction to tropical geometry*, A.M.S., 2015.

[11] T. Masada, *An algorithm for the enumeration of regular triangulations*, Master's thesis, University of Tokyo, Mar. 1995.

[12] F. Santos, *A point set whose space of triangulations is disconnected*, Journal of the American Mathematical Society, 13 (2000), p. 611–637.

[13] ——, *Triangulations of oriented matroids*, Memoirs of the American Mathematical Society, 156 (2002), p. 0–0.

[14] J.-P. Serre, *Local Fields (trans. Marvin J. Greenberg)*, vol. Graduate Texts in Math. 67, Springer, 1980.

[15] G. Ziegler, *Lectures on Polytopes*, Graduate texts in mathematics, Springer-Verlag, 1995.