

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Tropical implicitization

Permalink

<https://escholarship.org/uc/item/7d6845sr>

Author

Cueto, Maria Angelica

Publication Date

2010

Peer reviewed|Thesis/dissertation

Tropical Implicitization

by

Maria Angelica Cueto

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, BERKELEY

Committee in charge:

Professor Bernd Sturmfels, Chair

Professor Martin Olsson

Professor John P. Huelsenbeck

Fall 2010

Tropical Implicitization

Copyright 2010
by
Maria Angelica Cueto

Abstract

Tropical Implicitization

by

Maria Angelica Cueto

Doctor of Philosophy in Mathematics

University of California, BERKELEY

Professor Bernd Sturmfels, Chair

In recent years, tropical geometry has developed as a theory on its own. Its two main aims are to answer open questions in algebraic geometry and to give new proofs of celebrated classical results. The main subject of this thesis is concerned with the former: the solution of *implicitization problems via tropical geometry*. We develop new and *explicit* techniques that completely solve these challenges in four concrete examples.

We start by studying a family of challenging examples inspired by algebraic statistics and machine learning: the restricted Boltzmann machines $\mathcal{F}(n, k)$. These machines are highly structured projective varieties in tensor spaces. They correspond to a statistical model encoded by the complete bipartite graph $K_{k,n}$, by marginalizing k of the $n + k$ binary random variables. In Chapter 2, we investigate this problem in the most general setting. We conjecture a formula for the expected dimension of the model, verifying it in all relevant cases. We also study inference functions and their interplay with tropicalization of polynomial maps.

In Chapter 3, we focus on the particular case $\mathcal{F}(4, 2)$, answering a question by Drton, Sturmfels and Sullivant regarding the degree (and Newton polytope) of the homogeneous equation in 16 variables defining this model. We show that its degree is 110 and compute its Newton polytope. Along the way, we derive theoretical results in tropical geometry that are crucial for other examples in this thesis, as well as novel computational methods.

In Chapter 4, we study the first secant varieties of monomial projective curves from a tropical perspective. Our main tool is the theory of geometrical tropicalization developed by Hacking, Keel and Tevelev. Their theory hinges on computing the tropicalization of subvarieties of tori by analyzing the combinatorics of their boundary in a suitable (tropical) compactification. We enhance this theory by providing a formula for computing multiplicities on tropical varieties. We believe that this construction will give insight to understand higher secants of monomial projective curves, which are key objects in toric and birational geometry.

In Chapter 5, we answer the general question of implicitization of parametric surfaces in 3-space via geometric tropicalization. The generic case, together with its higher-dimensional analog, was studied by Sturmfels, Tevelev and Yu. We address this problem for non-generic

surfaces. This involves understanding the combinatorics of the intersection of irreducible algebraic curves in the two-dimensional torus and explicitly resolving singularities of points in curves by blow-ups. We conclude with a brief discussion on open problems, including connections to Berkovich spaces, extension of the theory to non-archimedean valued fields and applications of tropical implicitization to classifying tropical surfaces in three-space.

To my (academic) family.

Contents

List of Algorithms	iv
List of Figures	v
List of Tables	vii
1 Introduction	1
1.1 Tropical geometry	1
1.1.1 Constant coefficient case	2
1.1.2 Arbitrary coefficients case	8
1.2 Hadamard products and their tropicalization	11
1.3 Geometric tropicalization and tropical implicitization	14
1.4 Algorithms for implicitization	20
2 Geometry of the restricted Boltzmann machine	26
2.1 Introduction	26
2.2 Algebraic varieties, Hadamard product and tropicalization	28
2.3 The first secant variety of the n -cube	32
2.4 The tropical model and its dimension	37
2.5 Polyhedral geometry of parametric inference	41
3 An implicitization challenge for the restricted Boltzmann machine	47
3.1 Introduction	47
3.2 Geometry of the model	49
3.3 Tropicalizing the model	50
3.4 Newton polytope of the defining equation	59
3.4.1 Vertices and facets	59
3.4.2 Computing vertices	61
3.4.3 Implementation	65
3.4.4 Certifying facets	66
3.4.5 Completing the polytope	68

4	Tropical secant graphs of monomial curves	72
4.1	Introduction	72
4.2	The master graph	74
4.3	Combinatorics of monomial curves	79
4.4	The master graph under Hadamard products	86
4.5	The Newton polytope of the secant hypersurface in \mathbb{P}^4	95
4.6	Chow polytopes, tropical secant lines, toric arrangements and beyond	101
5	Implicitization of surfaces via geometric tropicalization	107
5.1	Introduction	107
5.2	Tropical elimination and tropical implicitization	109
5.3	Tropical implicitization for generic surfaces	116
5.4	Tropical implicitization for non-generic surfaces	127
5.5	Further remarks	134
	Bibliography	139

List of Algorithms

3.1	Ray-shooting algorithm.	62
3.2	Walking algorithm.	63
3.3	Facet certificate algorithm.	67
3.4	Approximation of a polytope by a subpolytope.	70
5.1	Tropical implicitization for generic surfaces.	138

List of Figures

1.1	A tropical surface in \mathbb{R}^3 described as a collection of two-dimensional cones in \mathbb{R}^3 or as a non-planar graph in \mathbb{S}^2	13
1.2	Intersection complex on a first compactification of $X \subset \mathbb{P}^2$ and on a resolution of this compactification with CNC.	19
1.3	Ray-shooting algorithm in dimension two.	21
2.1	Graphical representation of the restricted Boltzmann machine.	27
2.2	Quartet trees associated to the flattenings for $n = 4$	35
2.3	Partitions of the 3-cube that define non-empty cones on which Φ is linear.	37
2.4	Subdivisions of the 3-cube that represent vertices and facets of $\mathcal{T}M_3^1$	43
2.5	The polyhedral complex associated to the tropical model $\mathcal{T}M_3^1$	44
2.6	Parameterization of $\mathcal{T}M_2^1$	46
3.1	The restricted Boltzmann machine $\mathcal{F}(4, 2)$	47
3.2	Ray-shooting and walking algorithms combined.	64
3.3	Walking from vertex to vertex in the Newton polytope.	65
3.4	Approximation algorithm.	69
4.1	The building blocks of the <i>master graph</i>	75
4.2	The master graph associated to the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$	77
4.3	Compactification of Z in \mathbb{P}^n , with exponents $\{0, 30, 45, 55, 78\}$	80
4.4	Two affine charts describing the singularities of the embedding of the surface $(1 - \lambda, \omega^{30} - \lambda, \omega^{45} - \lambda, \omega^{55} - \lambda, \omega^{78} - \lambda)$	81
4.5	Resolution diagram of a binomial arrangement at the origin.	82
4.6	Resolution diagram of a binomial arrangement at infinity.	82
4.7	The tropical secant graph and the Gröbner tropical secant graph of the monomial curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$ in \mathbb{P}^4	96
4.8	The Gröbner tropical secant graphs of two rational normal curves.	103
4.9	The first tropical secant complex of the line $\mathbb{R}\langle(0, i_1, i_2, i_3, i_4)\rangle$ in \mathbb{TP}^4	105
5.1	Newton polytopes of the polynomials f_1, f_2 and f_3	122
5.2	Normal fans of the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3	122

5.3	Minkowski sum of \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}_3 , and a strictly simplicial refinement of its normal fan.	123
5.4	Pairwise Minkowski sums of the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$	124
5.5	Tropical graph of a generic surface in \mathbb{T}^3	124
5.6	Newton polytopes of the polynomials f_1, f_2 and f_3	125
5.7	Normal fans of the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3	126
5.8	Minkowski sum of \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}_3 , and a strictly simplicial refinement of its normal fan.	127
5.9	Tropical graph of a generic surface associated to three generic nodal curves. .	128
5.10	Three Newton polytopes in \mathbb{R}^2 and common refinement of their normal fans.	128
5.11	Tropical graph of the surface parameterized by (5.11).	129
5.12	Newton polytope and dual graph of a degree 3 surface in \mathbb{C}^3	129
5.13	Boundary of the closure of X in \mathbb{P}^2 and its resolution.	132
5.14	Tropical graph of a non-generic surface associated to three nodal curves. . .	133
5.15	Affine charts with $u = 1$ and $t = 1$ and the boundary divisors on each chart.	133
5.16	Resolution diagrams at $(1 : 1 : 1)$ and $(0 : 1 : 0)$	134
5.17	Newton polytope and dual graph of a non-generic surface in \mathbb{C}^3	135

List of Tables

2.1	Special cases where Conjecture 2.2.2 holds.	45
3.1	Facet orbit sizes of the Newton polytope of V_4^2	60

Acknowledgments

Hej och hopp här ska jag göra lite
roliga mattegrejer!

Alex Engström

I would like to express my gratitude to my advisor, Bernd Sturmfels, for all he has taught me, for the many interesting questions he posed, for helping me turn them into this thesis, and for his insights and encouragement along the way. He has generously afforded me a great deal of personal attention and thoughtful mentorship, taking a long-term view of my career.

Second, I would like to thank Enrique Tobis for his support and care for the past eight years, for listening and for keeping me sane throughout the decision-making process.

I would like to thank my co-authors Shaowei Lin, Jason Morton, Bernd Sturmfels, Enrique Tobis and Josephine Yu for interesting discussions, productive collaborations, and permission to include some co-authored material as part of my dissertation.

Last but not least I would like to thank the organizers and participants of the special program in Tropical Geometry that took place at the *Mathematical Sciences Research Institute* during Fall 2009. Their enthusiasm, encouragement and insight was essential for the evolution of my work. I am honored to be part of this community of great scientists and human beings. In particular, I would like to thank Alicia Dickenstein, Jan Draisma, Alex Fink, Anders Jensen, Eric Katz, Diane Maclagan, Sam Payne, Patrick Popescu-Pampu, Jena Tevelev, Josephine Yu and Piotr Zwiernik for sharing their knowledge with me and their willingness to answer my questions.

I appreciate the editors of *Contemporary Mathematics*, *Discrete Mathematics and Theoretical Computer Science* and the *Journal of Symbolic Computation* providing permission to include published material. I thank the six anonymous referees for their careful reading of this material and their fruitful comments.

I would like to thank the University of California for supporting me financially with a Chancellor's Fellowship and Bernd for supporting me during the research years. I also thank the Laboratory for Mathematical and Computational Biology of UC Berkeley and their members for their hospitality during my time there.

Finally, I would like to thank the Department of Mathematics at UC Berkeley and all its members for the stimulating and friendly environment that made my years in Graduate School a happy time. In particular, I wish to acknowledge my academic siblings Dustin Cartwright, Melody Chan, Alex Fink, Shaowei Lin, Felipe Rincon, Anne Shiu, Caroline Uhler and Cynthia Vinzant, and also the "Miller team," Alex Engström and Raman Sanyal, for their friendship, encouragement, support, fun discussions and after-seminar dinners at Jupiter.

Chapter 1

Introduction

In this chapter we describe the basics of tropical geometry, geometric tropicalization and tropical implicitization that we use throughout this dissertation. Along the way, we give an overview of the main results of our work. Most basic definitions in commutative algebra, algebraic geometry and toric geometry are taken for granted, and we refer the reader to [31, 47, 15, 39] for further details.

1.1 Tropical geometry

The term “tropical geometry” was coined in March 2002 when B. Sturmfels visited G. Mikhalkin in Salt Lake City ([90, 68, 83, 69]). It began as a framework to link amoebas, logarithmic limit sets, and (real) algebraic geometry. It synthesized and boosted the pioneering work of Bergman [2], Bieri-Groves [5] and Viro’s “patchworking” techniques to construct real algebraic varieties by “cutting and pasting” [97, 52]. The past few years have brought on truly explosive development, establishing deep connections with enumerative algebraic geometry, symplectic and analytic geometry, number theory, dynamical systems, mathematical biology, statistical physics, random matrix theory, and mathematical physics. This dissertation concerns applications of tropical geometry to solving implicitization problems in algebraic geometry with three concrete examples: the restricted Boltzmann machine $\mathcal{F}(n, k)$ (Chapters 2 and 3), which are highly structured projective varieties in tensor space, secant varieties of monomial curves with arbitrary exponents (Chapter 4), and surfaces in three-space (Chapter 5).

Tropical geometry can be viewed as a polyhedral version of algebraic geometry: algebraic varieties are replaced by weighted balanced polyhedral complexes, in order to answer open questions or to derive simpler proofs of classical results. These objects preserve just enough data about the original varieties to remain meaningful, while discarding much of their complexity.

There are many entering points to tropical geometry. In this work, we focus on two

of these perspectives: the Gröbner approach as degenerations on toric varieties [6, 56, 83, 87, 88, 90, 94] and the valuative perspective [5, 30, 45]. Chapters 2 and 3 follow the first perspective in the *constant coefficient case*, whereas Chapters 4 and 5 approach tropical varieties as images of algebraic varieties under a non-archimedean valuation.

1.1.1 Constant coefficient case

We start by defining the tropicalization of an algebraic variety over \mathbb{C} . We could replace \mathbb{C} by any algebraically closed field with trivial valuation. The precise definition of tropicalization depends on the ambient space containing our variety. In most cases we discuss here, we work with subvarieties of affine space \mathbb{C}^n , the n -dimensional algebraic torus $\mathbb{T}^n = (\mathbb{C}^*)^n$ or projective space \mathbb{P}^n . In Chapter 5 we work with subvarieties of a projective toric variety $\mathbb{P}(\mathcal{N})$ given by a complete rational polyhedral fan $\mathcal{N} \subset \mathbb{R}^n$. We adopt the *min* convention for tropical geometry inherited from the *tropical semiring* $(\mathbb{R} \cup \{\infty\}, \min, +)$ [79]. This choice arises naturally from the valuative perspective discussed in the last two chapters of the dissertation. All results also hold in the *max* convention, after changing signs accordingly.

Definition 1.1.1. Let $X \subset \mathbb{C}^n$ be an algebraic variety not contained in a coordinate hyperplane, and let $I = I(X) \subset \mathbb{C}[x_1, \dots, x_n]$ be its defining ideal. The *tropicalization* of X or I is defined as:

$$\mathcal{T}X = \mathcal{T}(I) = \{w \in \mathbb{R}^n \mid \text{in}_w(I) \text{ contains no monomial}\},$$

where $\text{in}_w(I) = \langle \text{in}_w(f) : f \in I \rangle$, and $\text{in}_w(f)$ is the sum of all nonzero terms of $f = \sum_{\alpha} c_{\alpha} x^{\alpha}$ such that the inner product $\alpha \cdot w$ is *minimum*.

Example 1.1.2. Consider an irreducible polynomial $f \in \mathbb{C}[x_1, \dots, x_n]$ and its associated *Newton polytope* in \mathbb{R}^n , i.e. the convex hull of the exponents of all monomials in f . Assume that f is not a monomial. Then, the tropicalization of $\mathcal{T}(f)$ is the set of all objective vectors in \mathbb{R}^n that pick at least two vertices of $\text{NP}(f)$. Thus, $\mathcal{T}(f)$ is the union of all codimension one cones in the inner normal fan of $\text{NP}(f)$. \diamond

Definition 1.1.3. If $X \subset \mathbb{T}^n$, with defining ideal $I \subset \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ in the Laurent polynomial ring, we define its tropicalization as

$$\mathcal{T}X = \mathcal{T}(I) = \{w \in \mathbb{R}^n \mid 1 \notin \text{in}_w(I)\}.$$

The initial ideal $\text{in}_w(I)$ with respect to a weight vector $w \in \mathbb{R}^n$ agrees with that from Definition 1.1.1, replacing the ordinary polynomial ring with the Laurent polynomial ring.

Although it may not be clear from Definition 1.1.1, the process of tropicalization is invariant under intersections with dense open subsets of affine space. More precisely, let Y be a closed subvariety of \mathbb{T}^n , also known as a *very affine variety*. Suppose $I_Y \subseteq \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$

is its defining ideal. Consider the Zariski closure \overline{Y} of Y in \mathbb{C}^n . It is easy to see that $\mathcal{T}Y$ equals $\mathcal{T}\overline{Y}$. Indeed, this follows from the fact that I_Y is the saturation ideal $(I_{\overline{Y}}\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm n}] : (x_1 \cdots x_n)^\infty)$ and $I_{\overline{Y}} = I_Y \cap \mathbb{C}[x_1, \dots, x_n]$. Therefore, if we start with an irreducible variety $X \subset \mathbb{C}^n$ not contained in a coordinate hyperplane, then the very affine variety $Y = X \cap \mathbb{T}^n$ has the same dimension as X and $\mathcal{T}Y = \mathcal{T}X$. In short, tropicalizations are “toric in nature.” We will go back and forth between these two definitions of tropicalization and mainly consider subvarieties of \mathbb{T}^n to avoid writing the required hypothesis for affine varieties every time.

We use the previous definitions to tropicalize projective subvarieties of \mathbb{P}^n :

Definition 1.1.4. Let $X \subset \mathbb{P}^n$ and let I be its homogeneous defining ideal in $\mathbb{C}[x_0, \dots, x_n]$. Then, the tropicalization of X equals $\mathcal{T}(I)$ from Definition 1.1.1.

In all three cases, the tropical variety $\mathcal{T}X$ is a rational polyhedral fan in \mathbb{R}^n . If X is projective, then $\mathcal{T}X$ is naturally a subfan of (the negative of) the Gröbner fan of the homogeneous ideal I . In the other two cases, we can equip $\mathcal{T}X$ with an inherited polyhedral fan structure by considering the Gröbner fan of its homogenized ideal [53].

Example 1.1.5. The tropical hypersurface $\mathcal{T}(f)$ from Example 1.1.2 has the inherited fan structure of the normal (Gröbner) fan of its Newton polytope. Maximal cones in $\mathcal{T}(f)$ are dual to edges in the polytope and, in general, k -dimensional cones in $\mathcal{T}(f)$ are dual to codimension k faces in the polytope. This interpretation will be key for Chapter 3. \diamond

In most cases, the inherited Gröbner fan structure will be the *coarsest* fan structure we can give to $\mathcal{T}X$. However, not all varieties admit a coarsest fan structure [91, Example 5.2]: *hübsch* varieties will have this good property [91, Remark 3.11]. The possible fan structures on $\mathcal{T}X$ form a partially ordered set under refinement, called the *tropical poset*, whose properties were studied in [91, Section 2].

If X is an irreducible subvariety of \mathbb{T}^n , then the tropical variety $\mathcal{T}X$ is a *pure* polyhedral fan: all maximal cones in $\mathcal{T}X$ have the same dimension. In addition, $\mathcal{T}X$ is connected in codimension one, an essential tool for effective computations [54]. The fan $\mathcal{T}X$ preserves an important invariant of X : both objects have the *same dimension* [5, 90]. This essential result is known in the literature as the Bieri-Groves theorem. In particular, this implies that tropical geometry can be used to compute or certify Krull dimensions of subvarieties of tori. More precisely, if we can compute the tropicalization of an irreducible variety without knowing its defining ideal we can certify its dimension by simple linear algebra. Its dimension will be the dimension of any maximal cone in a tropical fan. Using this technique, we prove:

Theorem 2.1.2. *The restricted Boltzmann machines $\mathcal{F}(n, k)$ have the expected dimension for all relevant values of n and k . In particular, the variety $\mathcal{F}(4, 2)$, is a hypersurface in \mathbb{P}^{15} .*

If X is projective, the dimension of $\mathcal{T}X \subseteq \mathbb{R}^n$ will exceed the projective (Krull) by one. We can reduce this gap applying the following definition:

Definition 1.1.6. Let $I \subset \mathbb{C}[x_1, \dots, x_n]$ be an ideal. The set $\{w \in \mathcal{T}(I) : \text{in}_w(I) = I\}$ is a linear space in \mathbb{R}^n and it is called the *lineality space* of the fan $\mathcal{T}(I)$ or the *homogeneity space* of I .

By definition, the lineality space of the tropical variety $\mathcal{T}X$ can be spanned by integer vectors, which form a primitive lattice $\Lambda \subset \mathbb{Z}^n$. This lattice encodes the action of a maximal torus on X , given by a diagonal action. That is, if $\dim_{\mathbb{Z}} \Lambda = r$, we have an r -dimensional torus action on X by

$$(t_1, \dots, t_r) \cdot (x_1, \dots, x_n) := \left(\prod_{i=1}^r t_i^{a_{i1}} x_1, \dots, \prod_{i=1}^r t_i^{a_{in}} x_n \right), \quad \underline{t} \in \mathbb{T}^r, \underline{x} \in X, \quad (1.1)$$

where $a = (a_{ij})_{i,j} \subset \mathbb{Z}^{r \times n}$ is a matrix whose rows \mathbb{Z} -span Λ . All homogeneous varieties with a non-trivial torus action will have a non-trivial lineality space in their tropicalization.

Example 1.1.7. The lineality space of a tropical hypersurface $\mathcal{T}(g)$ equals the orthogonal complement of the affine span of the Newton polytope of g , after appropriate translation to the origin. The extreme case is that of a toric variety globally parameterized by a monomial map $\alpha_A: \mathbb{T}^d \rightarrow \mathbb{T}^n$ with associated matrix $A \in \mathbb{Z}^{n \times d}$. Its tropicalization will be a classical linear space, namely, the column span of A . We will come back to this construction soon. \diamond

Example 1.1.8. Let w be a point in the relative interior of a maximal cone σ of $\mathcal{T}X$. The variety $\text{in}_w(X) \subset \mathbb{T}^n$ defined by the initial ideal $\text{in}_w(I(X))$ is a flat deformation of X and, as such, its dimension equals the dimension of X [56]. In addition, $\text{in}_w(X)$ is a homogeneous variety, with torus action defined by the lattice $\mathbb{L}_w \cap \mathbb{Z}^n$, where \mathbb{L}_w is the linear span of the maximal cone of $\mathcal{T}X$ containing w . The tropicalization of $\mathcal{T}\text{in}_w(X)$ yields \mathbb{L}_w . \diamond

All cones in any fan structure of $\mathcal{T}X$ contain the lineality space. For this reason, we can view $\mathcal{T}X$ as a pointed fan in the $(n - r)$ -dimensional vector space $\mathbb{R}^n / (\mathbb{R} \otimes_{\mathbb{Z}} \Lambda)$, with inherited fan structure (Lemma 3.3.1). Furthermore, we can intersect this fan with the unit sphere of appropriate dimension to obtain a *polyhedral (spherical) complex* representing $\mathcal{T}X$. For example, if X is a surface in n -space with no non-trivial torus action, the associated tropical polyhedral complex $\mathcal{T}X \cap \mathbb{S}^{n-1}$ is a graph, which we call the *tropical surface graph*. This identification will be extremely useful in the constructions of Chapters 4 and 5.

If X is a subvariety of \mathbb{P}^n , we know that the all-one's vector $\mathbf{1}$ lies in the lineality space of $\mathcal{T}X$, so we can view $\mathcal{T}X$ inside the *tropical projective torus* $\mathbb{TP}^n = \mathbb{R}^{n+1} / \mathbb{R} \cdot \mathbf{1}$. In this new setting, the dimension of the fan $\mathcal{T}X \subset \mathbb{TP}^n$ equals the projective dimension of X . Notice that $\mathcal{T}X \subset \mathbb{TP}^n$ need not be a pointed cone. This will be the case of the projective hypersurface studied in Chapter 3 and the three-fold secant variety from Chapter 4.

In addition to their polyhedral structure, tropical varieties are equipped with integer positive weights on all of their maximal cones, called *multiplicities*. We now explain how

these numbers can be constructed. Let $X \subset \mathbb{T}^n$. A point $w \in \mathcal{T}X$ is called *regular* if $\mathcal{T}X$ is a linear space locally near w . Alternatively, if we fix a fan structure on $\mathcal{T}X$, its regular points are the points in the relative interior of maximal cones. The *multiplicity* m_w of a regular point w is the sum of the algebraic multiplicities of all minimal associated primes of the initial ideal $\text{in}_w(I(X))$ [6, 24]. More precisely, if P is one such minimal prime ideal, we define the multiplicity of P in $\text{in}_w(I(X))$ as

$$m(P, \text{in}_w(I(X))) = \dim_{\frac{S_P}{\overline{PS}_P}} S_P / \text{in}_w(I(X))S_P = \text{length}_{S_P}(S_P / \text{in}_w(I(X))S_P), \quad (1.2)$$

where $S = \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ and S_P is the localization of S at the prime ideal P . The rightmost term in (1.2) is the length of the quotient $S_P / \text{in}_w(I(X))S_P$ as an S_P -module. This definition of multiplicity follows the approach of [31, Section 3.6] and [24] and will be used throughout Section 3.3.

The multiplicity of a maximal cone $\sigma \subset \mathcal{T}X$ is defined as m_w for any $w \in \sigma$ in its relative interior. It can be shown that this assignment does not depend on the choice of w : multiplicities are locally constant functions on regular points [91, Corollary 3.8]. With these multiplicities, the tropical variety satisfies the *balancing condition* [91, Corollary 3.4]. For example, if X is an irreducible hypersurface, the multiplicity of a maximal cone equals the lattice length of the corresponding edge in its Newton polytope [24]. The balancing condition says that the sum of all edges (clockwise-oriented) in a two-face of this polytope equals zero. Tropical multiplicities of globally parameterized toric varieties have constant value 1.

In addition, tropical multiplicities allow two alternative interpretations. They can be thought of as *intersection numbers* [87, Chapter 2], [56, Section 9], [74, Section 2], or as *Minkowski weights*, i.e. *operational Chow cohomology classes* on an appropriate toric variety [91, Remark 3.5]. This alternative perspective will be used in Chapters 4 and 5, but we give a short introduction in Section 1.1.2.

Just as we tropicalize varieties, we can also tropicalize polynomial maps. To do so, we just need to define the tropicalization of a single polynomial. We replace the usual arithmetic operations of sum and product in \mathbb{R} by the *tropical* operations \oplus and \odot in the tropical semiring $(\mathbb{R} \cup \{\infty\}, \oplus, \odot)$, where $a \oplus b = \min\{a, b\}$ and $a \odot b = a + b$. In addition, we endow \mathbb{C} with the trivial valuation, i.e. $\text{val}(\mathbb{C}^*) = 0$ and $\text{val}(0) = \infty$ [88]. The resulting *tropical polynomial* is a continuous piecewise linear map:

Definition 1.1.9. Let $f = \sum_{\alpha} a_{\alpha} x^{\alpha} \in \mathbb{C}[x_1, \dots, x_n]$. Then, the *tropicalization* of f is the piecewise linear map $\text{trop}(f): \mathbb{R}^n \rightarrow \mathbb{R}$ that is defined by

$$\text{trop}(f)(w) = \bigoplus_{\alpha} (\text{val}(a_{\alpha}) \odot \alpha \cdot w) = \min_{\alpha} \{\alpha \cdot w\}. \quad (1.3)$$

The tropicalization $\text{trop}(\mathbf{f})$ of a polynomial map $\mathbf{f}: \mathbb{C}^n \rightarrow \mathbb{C}^d$ is defined as the map obtained by tropicalizing each coordinate.

As expected, these definitions extend naturally to algebraically closed fields with a non-archimedean valuation, the subject of Section 1.1.2. In particular, if $X \subset \mathbb{T}^n$ is a very affine variety parameterized by a Laurent polynomial map $\mathbf{f}: \mathbb{T}^d \rightarrow \mathbb{T}^n$, then it follows that the image of the tropical map $\text{trop}(\mathbf{f})$ is contained in $\mathcal{T}X$. Varieties coming from statistical models often admit such a polynomial parameterization and it is worth studying the associated tropical maps to understand their properties, including their inference functions [75]. Recent applications of this “min-plus” algebra include solutions of continuous-time, continuous-state optimal control problems [67].

Chapter 2 discusses this methodology for the *restricted Boltzmann machines* $\mathcal{F}(n, k)$. These “machines” are highly-structured algebraic varieties in tensor space. They correspond to graphical statistical models given by a complete bipartite graph $K_{k,n}$ with k binary hidden nodes and n observed binary nodes. From the statistical point of view, their inference functions are in one-to-one correspondence with the regions of lineality of $-\text{trop}(f)$, where f is the parameterization of the variety $\mathcal{F}(n, k)$. Such regions will correspond to certain slicings on the n -cube (Proposition 2.5.1). This novel approach yields a new interpretation for the space of explanations as a subfan of the secondary fan of the n -cube, which connects nicely with combinatorial and geometric work of Develin and Draisma on tropical secants of toric varieties [22, 25]. The tropical point of view allows us to *organize the geometric information* of the space of inference functions into the image of $-\text{trop}(f)$, which can then be analyzed with the tools of tropical and polyhedral geometry. This approach is explained in Section 2.5.

One of the weaknesses of tropical geometry is the lack of functoriality. Even if we had a polynomial map between two varieties $\mathbf{f}: X \rightarrow Y$, and we could tropicalize both varieties and polynomial maps, we cannot a priori guarantee that the image of the piecewise linear map $\text{trop}(\mathbf{f}): \mathcal{T}X \rightarrow \mathbb{R}^d$ lies in $\mathcal{T}Y$. However, if we restrict the class of varieties and maps to subvarieties of tori and homomorphisms of tori (i.e. monomial maps), then functoriality does hold. We now give a precise statement. Let A be a $d \times n$ integer matrix defining a monomial map $\alpha: \mathbb{T}^n \rightarrow \mathbb{T}^d$ and a linear map $A: \mathbb{R}^n \rightarrow \mathbb{R}^d$ defined by left multiplication.

Theorem 1.1.10. [91, Theorem 3.12] *Let $X \subset \mathbb{T}^n$ be a subvariety. Then*

$$\mathcal{T}(\alpha(X)) = A(\mathcal{T}X).$$

Moreover, if α induces a generically finite morphism of degree δ on X , then the multiplicity of $\mathcal{T}(\alpha(X))$ at a regular point w is

$$m_w = \frac{1}{\delta} \sum_v m_v \text{index}(\mathbb{L}_w \cap \mathbb{Z}^d, A(\mathbb{L}_v \cap \mathbb{Z}^n)), \quad (1.4)$$

where we sum over all points $v \in \mathcal{T}X$ with $Av = w$. We also assume that the number of such v is finite, all of them are regular in $\mathcal{T}X$, and $\mathbb{L}_v, \mathbb{L}_w$ are linear spans of neighborhoods of $v \in \mathcal{T}X$ and $w \in A(\mathcal{T}X)$ respectively.

We refer to (1.4) as the *push-forward formula* for multiplicities.

If the fibers of the map α restricted to X are infinite, then the sum in (1.4) is infinite and, a priori, there is no formula connecting multiplicities on the source and the target of the linear map A . However, in numerous examples, subvarieties of tori admit torus actions that allow us to think of α as if it were finite. More precisely, after taking the quotient of the domain by a maximal torus action, we build a canonical generically finite monomial map $\bar{\alpha}$ in the quotient space. Thus, we can freely use the push-forward formula (1.4) to compute tropical multiplicities on the target. Using this small remark, we can extend Theorem 1.1.10 to maps with finite generic fibers modulo a torus action. Our main technical tool is the following lemma, that also justifies the heuristics for tropicalizing homogeneous varieties:

Lemma 3.3.1. *Let $X \subset \mathbb{T}^n$ and let L be a subspace of the lineality space of the tropical variety $\mathcal{T}X$ generated by integer vectors. Consider the action of $\mathbb{T}^{\dim L}$ on X induced by the lattice $L \cap \mathbb{Z}^n$ and the induced quotient X' . Endow the quotient \mathcal{T}/L with induced fan structure and multiplicity. Then, $\mathcal{T}X/L$ is the tropicalization of X' .*

Theorem 3.3.2. *Let $\alpha: \mathbb{T}^n \rightarrow \mathbb{T}^d$ be a monomial map with associated integer matrix A and let $X \subset \mathbb{T}^n$ be a closed subvariety. Then,*

$$\mathcal{T}(\alpha(X)) = A(\mathcal{T}X).$$

Suppose X has a torus action given by a rank l lattice $\Lambda \subset \mathbb{Z}^n$. Let X' be the quotient by this torus action and π be its quotient map. Let $\bar{\alpha}: X' \rightarrow \mathbb{T}^d$ be the induced monomial map, with associated integer matrix A' .

Suppose $\Lambda' = A(\Lambda)$ is a primitive sublattice of \mathbb{Z}^d and that $\bar{\alpha}$ induces a generically finite morphism of degree δ on X' . Let π_2 be the projection from \mathbb{R}^n to the quotient vector space $\mathbb{R}^n/\Lambda' \otimes_{\mathbb{Z}} \mathbb{R}$. Then, the multiplicity of a regular point w in $\mathcal{T}(\alpha(X))$ can be computed as:

$$m_w = \frac{1}{\delta} \cdot \sum_{\substack{\pi(v) \\ A \cdot v = w}} m_v \cdot \text{index}(\mathbb{L}_w \cap \mathbb{Z}^d, A(\mathbb{L}_v \cap \mathbb{Z}^n))$$

where the sum is over any set of representatives of points $\{v' = \pi(v) \in \mathcal{T}X' \mid \pi_2(A'v') = \pi_2(w')\}$. We also assume that the number of such v' is finite, all of them are regular in $\mathcal{T}X'$ and $\mathbb{L}_v, \mathbb{L}_w$ are linear spans of neighborhoods of $v \in \mathcal{T}X$ and $w \in A(\mathcal{T}X)$ respectively.

In Remark 3.3.3, we extend the previous result to the case where Λ' is not a primitive lattice. The corresponding formula involves an extra factor, namely the index of the lattice Λ' with respect to its saturation $\Lambda'^{\text{sat}} \subset \mathbb{Z}^d$. These results are extensively discussed in Chapter 3, particularly in Section 3.3.

1.1.2 Arbitrary coefficients case

In Section 1.1.1, we discussed the tropicalization of varieties over \mathbb{C} , a field with trivial valuation. As expected, the geometry becomes richer if we allow algebraically closed fields with non-archimedean, non-trivial valuation. This is the topic we cover in this section.

First, we recall the definition of such valued fields:

Definition 1.1.11. A *non-archimedean valued field* is a pair (\mathbb{K}, val) where $\mathbb{K} = \overline{\mathbb{K}}$ is a field and $\text{val}: \mathbb{K}^* \longrightarrow \mathbb{Q}$ satisfies:

- (i) $\text{val}(x + y) \geq \min\{\text{val}(x), \text{val}(y)\}$,
- (ii) $\text{val}(xy) = \text{val}(x) + \text{val}(y)$,
- (iii) $\text{val}(0) = \infty, \text{val}(1) = 0$.

For simplicity, we will assume that val is a surjective map, i.e. there exists an element $t \in \mathbb{K}$ with valuation 1. The field \mathbb{K} contains a discrete valuation ring $\mathcal{R} = \text{val}^{-1}(\mathbb{R}_{\geq 0} \cup \{\infty\})$, with maximal ideal $\mathcal{M} = \text{val}^{-1}(\mathbb{R}_{> 0} \cup \{\infty\})$ and residue field $\mathcal{K} = \mathcal{R}/\mathcal{M}$.

Our archetypical example of such valued fields is the *field of Puiseux series* $\mathbb{C}\{\{t\}\} = \bigcup_{n \in \mathbb{N}} \mathbb{C}((t^{1/n}))$, where the valuation is the order in t . In this case, $\mathcal{R} = \bigcup_{n \in \mathbb{N}} \mathbb{C}[[t^{1/n}]]$ and \mathcal{M} is the ideal of all formal power series in \mathcal{R} with zero constant term. In addition, $\mathcal{K} = \mathbb{C}$ and embeds naturally in \mathbb{K} .

Just as we tropicalized subvarieties of the algebraic torus \mathbb{T}^n over \mathbb{C} , we can tropicalize subvarieties of the n -dimensional algebraic torus $\mathbb{T}_{\mathbb{K}}^n$ over the valued field \mathbb{K} . As a result we obtain weighted balanced rational polyhedral complexes. If $X \subset \mathbb{T}_{\mathbb{K}}^n$ is irreducible, then $\mathcal{T}X$ is a pure complex of dimension $\dim X$ and connected in codimension one [87, 56, 91]. In what follows, we explain this construction in analogy to the constant coefficient case.

Definition 1.1.12. Let $X \subset \mathbb{T}_{\mathbb{K}}^n$ be a subvariety with defining ideal $I \subset \mathbb{K}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. We define the tropicalization of X as

$$\mathcal{T}X = \mathcal{T}I = \overline{\{w \in \mathbb{Q}^n \mid 1 \notin \text{in}_w(I)\}} \subseteq \mathbb{R}^n,$$

where $\text{in}_w(I) \subset \mathcal{K}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ is the ideal generated by all $\text{in}_w(f)$ with $f \in I \setminus \{0\}$, where $\text{in}_w(f)$ is defined as the class of $t^{-\text{trop}(f)(w)} f(t^{w_1} x_1, \dots, t^{w_n} x_n) \in \mathcal{R}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ modulo \mathcal{M} , and $\text{trop}(f)$ is the tropicalization of f as in (1.3).

The set $\mathcal{T}X$ is a rational polyhedral complex. It consists of the closure of all weights $w \in \mathbb{Q}^n$ such that the variety $\text{in}_w(X) \subset \mathcal{K}^n$ determined by $\text{in}_w(I)$ intersects the big open torus $\mathbb{T}_{\mathcal{K}}^n \subset \mathcal{K}^n$. In the homogeneous case, each weight vector $w \in \mathbb{Q}^n$ determines a relatively open polyhedron in \mathbb{R}^n in which the initial ideal $\text{in}_w(I(X))$ is constant. Thus, $\mathcal{T}X$ inherits a polyhedral complex structure from (the negative of) this *Gröbner complex*. In the non-homogeneous case, as we saw in Section 1.1.1, the structure is obtained by choosing a homogeneization of the defining ideal of the variety and cutting the resulting complex in \mathbb{R}^{n+1} at height one.

Example 1.1.13 (Examples 1.1.2 and 1.1.18 revisited). Let $f = \sum_{\alpha} a_{\alpha} x^{\alpha} \in \mathbb{K}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. We define the *extended Newton polytope* of f as $\widetilde{\text{NP}}(f) := \text{Conv}\{(\alpha, b) : b \leq -\text{val}(a_{\alpha}) \forall a_{\alpha} \neq 0\} \subseteq \mathbb{R}^{n+1}$. The upper hull of this convex set is a polyhedron parameterized by the *Legendre transform* of f , which defines a regular polyhedral subdivision of the Newton polytope of f by projection. The tropical hypersurface $\mathcal{T}(f)$ is the union of all codimension one cells in the dual complex of this subdivision. Maximal cells in $\mathcal{T}(f)$ are dual to edges in the subdivision of $\text{NP}(f)$ and their weight agrees with the lattice length of the associated edges.

If $\mathbb{K} = \mathbb{C}$, the polytope $\widetilde{\text{NP}}(f)$ equals $\text{NP}(f) \times \{0\}$ and induces the trivial subdivision with a single cell, namely, the polytope $\text{NP}(f)$. The dual subdivision is the inner normal fan of $\text{NP}(f)$ and we recover the characterization from Examples 1.1.2 and 1.1.18. \diamond

Tropical varieties over arbitrary coefficients have an explicit characterization in terms of the valuation map: these polyhedral complexes will only remember the valuation of each point in the corresponding algebraic varieties. In the Puiseux series case this means that we only remember the order of a power series. It is in this spirit that tropical varieties are *combinatorial shadows* of algebraic varieties [65].

Theorem 1.1.14 (Fundamental theorem of tropical geometry). *Let $X \subset \mathbb{T}_{\mathbb{K}}^n$ be a closed subvariety. Then, the tropicalization of X is the closure of the image of X under the valuation map with the usual Euclidean topology, that is:*

$$\mathcal{T}X = \overline{\{(\text{val}(x_1), \dots, \text{val}(x_n)) \mid x \in X\}} \subseteq \mathbb{R}^n.$$

The previous result allows us to compute tropicalizations of varieties over \mathbb{C} (the constant coefficient case) by base-changing to the field of Puiseux series.

Even though the fundamental theorem has a long history that began with Kapranov *et al.* [30], with its latest installment by Payne [78], the valuative approach lacks an interpretation of multiplicities on tropical varieties. As we will see in Section 1.4 these numbers are essential for recovering information about algebraic varieties from their tropical counterparts. For this reason, we will use Theorem 1.1.14 only for set-theoretic descriptions of tropical varieties.

As in the constant coefficient case, we can tropicalize polynomial maps between algebraic varieties using (1.3). Functoriality will hold for subvarieties of tori and monomial maps:

Theorem 1.1.15. *Let $X \subset \mathbb{T}_{\mathbb{K}}^n$ be a closed subvariety and $\alpha: \mathbb{T}_{\mathbb{K}}^n \rightarrow \mathbb{T}_{\mathbb{K}}^d$ a monomial map associated to an integer matrix $A \in \mathbb{Z}^{d \times n}$. Then:*

$$\mathcal{T}(\alpha(X)) = A(\mathcal{T}X).$$

Notice that the previous statement does not provide a push-forward formula for multiplicities. An analogous formula to (1.4) is expected to hold since the local geometry of

tropical varieties can be reduced to the constant coefficient case [87, Prop 2.2.3]. Recent results by Osserman and Payne suggest that this is indeed the case [74].

We conclude this section with a brief discussion on tropicalization of subvarieties of toric varieties and tropical multiplicities from the degenerations perspective. This approach will be key for deriving the main result of Chapter 5: a combinatorial formula for computing tropical multiplicities of rational varieties.

Our starting point is an irreducible subvariety X of $\mathbb{T}_{\mathbb{K}}^n$ of dimension d . As in the constant coefficient case, given a tropical complex structure of $\mathcal{T}X$, we define the regular points of $\mathcal{T}X$ as the set of all points in the relative interior of maximal cells of $\mathcal{T}X$. Any regular point $w \in \mathcal{T}X$ defines a proper ideal $\text{in}_w(I(X))$ and, hence, a subvariety $\text{in}_w(X) \subset \mathbb{T}_{\mathcal{K}}^n$. This subvariety is a flat deformation of $X \subset \mathbb{T}_{\mathbb{K}}^n$ of the same dimension as X .

Let σ be the maximal cell of $\mathcal{T}X$ containing the point w , and assume $\dim X = d$. Then, the tropical variety of $\text{in}_w(X)$ is the d -dimensional subspace spanned by the primitive lattice $\Lambda_w := (\mathbb{R} \otimes (\sigma - w)) \cap \mathbb{Z}^n$. In particular, $\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n$ is a d -dimensional variety with a d -dimensional torus action defined by the lattice Λ_w [56]. Therefore, $\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n$ is supported on finitely many d -dimensional torus orbits. Thus,

$$\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n = \bigsqcup_{i=1}^r (\mathbb{T}_{\mathcal{K}}^d \cdot p_i), \quad (1.5)$$

for some points $p_i \in \mathbb{T}_{\mathcal{K}}^n$ and some $r \in \mathbb{Z}_{>0}$. This phenomenon extends Example 1.1.8 and is key to the interpretation of tropical multiplicities.

If w is a point in the relative interior of a cone $\tilde{\tau}$ in $\mathcal{T}X$, we can compute m_w from the variety X or its initial degeneration $\text{in}_x(X)$. More precisely,

$$m_w = \deg([X] \cdot [V(\tilde{\tau})]) = \deg([\text{in}_w(X)] \cdot [V(\tilde{\tau})]). \quad (1.6)$$

The right-most equality holds because intersection product commutes with specialization.

Definition 1.1.16. [56, Definition 6.10] Let σ be a maximal cell in $\mathcal{T}X$ of dimension d , and let w be a point in its relative interior. Following (1.5), we decompose the underlying cycle of $\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n$ as

$$[\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n] = \sum_{i=1}^r m_i [\mathbb{T}_{\mathcal{K}}^d \cdot p_i],$$

for some $m_i \in \mathbb{Z}_{>0}$. The multiplicity m_σ is defined as $m_\sigma = \sum_{i=1}^r m_i$.

We extend the operation of tropicalization to subvarieties of a complete toric variety $Y(\Sigma)$ defined over \mathcal{K} . We will use this construction in Chapter 5.

Definition 1.1.17. Let $X \subseteq Y(\Sigma)$ be a d -dimensional subvariety, and $w \in \mathbb{Q}^n$. Consider the subvariety $\overline{t^{-w} \cdot X} \subset Y(\Sigma)$. Then the *initial degeneration of X* with respect to w is defined as the variety

$$\text{in}_w(X) := \overline{t^{-w} \cdot X} \times_{\text{Spec } \mathcal{R}} \text{Spec } \mathcal{K} \subseteq Y(\Sigma) \times_{\text{Spec } \mathcal{R}} \text{Spec } \mathcal{K}.$$

This definition gives a locally constant map on \mathbb{Q}^n , which we extend to \mathbb{R}^n by continuity. As before, we define $\mathcal{T}X$ as the set of all points w with $\text{in}_w(X) \cap \mathbb{T}_{\mathcal{K}}^n \neq \emptyset$.

Example 1.1.18. If X is a hypersurface defined by a polynomial $f \in \mathbb{K}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, then $\overline{t^{-w} \cdot X}$ is defined by the polynomial $f(t^{w_1}x_1, \dots, t^{w_n}x_n)$. The order in t of this polynomial is $\text{trop}(f)(w)$ so $\text{in}_w(X) \subset \mathbb{T}_{\mathcal{K}}^n$ coincides with Definition 1.1.12. \diamond

We conclude by extending the definition of tropical multiplicities to subvarieties of toric varieties defined over \mathbb{C} that properly intersect torus orbits. This proper intersection condition ensures that the tropicalization of the subvariety $X \subseteq Y(\Sigma)$ is supported on all cones of Σ whose dimension equals the dimension of X [94, Lemma 2.2]. Notice that if the fan is strictly simplicial then the intersections are automatically proper [56, Lemma 8.10].

Definition 1.1.19. [56, Definition 9.3] Let $Y(\Sigma)$ be a complete toric variety defined over \mathcal{K} , and let $\tilde{Y} = Y(\tilde{\Sigma})$ be a smooth toric resolution of Y defined by a strictly simplicial refinement $\tilde{\Sigma}$ of the fan Σ . Let d be the dimension of X and $\tilde{\tau}$ a d -dimensional cone in $\tilde{\Sigma}$. Then the multiplicity of $\tilde{\tau}$ in $\mathcal{T}X$ equals $\deg([X] \cdot [V(\tilde{\tau})])$, where $V(\tilde{\tau})$ is the closure of the torus orbit of $\tilde{\tau}$ in \tilde{Y} .

1.2 Hadamard products and their tropicalization

A central construction in Chapters 2 through 4 is the *Hadamard product* of algebraic varieties. These products arise naturally in many problems posed in algebraic statistics, including restricted Boltzmann machines. They also play an important role in the construction of secant varieties of homogeneous spaces, such as monomial projective curves (Proposition 4.4.1). Hadamard products have an important feature: they behave well with respect to tropicalization. In this section, we study these algebraic varieties and explain their connection to tropical geometry.

We start with the definition of Hadamard products for subvarieties of tori:

Definition 1.2.1. Let $X, Y \subset \mathbb{T}_{\mathbb{K}}^n$ be two closed subvarieties of tori. The *Hadamard product* of X and Y is

$$X \cdot Y = \overline{\{x \cdot y = (x_1y_1, \dots, x_ny_n) \mid x \in X, y \in Y\}} \subset \mathbb{T}_{\mathbb{K}}^n.$$

Note that each coordinate of $x \cdot y$ is bihomogeneous of degree $(1, 1)$. Therefore, the definition of Hadamard products carries over to projective varieties as well: the extra condition required is that $x \cdot y \neq \mathbf{0}$. With this definition, $X \cdot Y$ is a projective subvariety of \mathbb{P}^n . Note that if X and Y are not contained in a coordinate hyperplane, the same will hold for their Hadamard product. This small remark will be important when tropicalizing this construction in the projective setting.

A special instance of Hadamard products, treated in Chapters 2 through 3, corresponds to the case when $X = Y$. The Hadamard product $X \cdot X$ and, in general, of r copies of a variety X are called *Hadamard powers* and denoted by $X^{[r]}$. The Hadamard square $X^{[2]}$ corresponds to the diagonal part of the *Segre embedding* $X \subset \mathbb{P}_{\mathbb{K}}^n \hookrightarrow \mathbb{P}_{\mathbb{K}}^{n^2+2n}$ and its defining ideal can be obtained from this description.

The next result shows how to tropicalize Hadamard products:

Corollary 3.3.6. *Given $X, Y \subset \mathbb{T}_{\mathbb{K}}^n$ two irreducible subvarieties of tori, we can consider the associated irreducible variety $X \cdot Y \subset \mathbb{T}_{\mathbb{K}}^n$. Then as sets:*

$$\mathcal{T}(X \cdot Y) = \mathcal{T}X + \mathcal{T}Y,$$

where the sum on the right-hand side denotes the Minkowski sum in \mathbb{R}^n .

This set-theoretic result is motivated by, and is a direct consequence of, Theorem 1.1.14 and the fact that valuations turn products into sums. It extends naturally to projective varieties, provided they do not lie in coordinate hyperplanes.

As we explained before, $\mathcal{T}(X \cdot Y)$ has a canonical polyhedral complex structure (a fan, in the constant coefficient case). However, the previous theorem gives no information about this structure. Moreover, we are not claiming that $\mathcal{T}(X \cdot Y)$ inherits a fan structure from $\mathcal{T}X$ and $\mathcal{T}Y$. In general, it might happen that maximal cells (cones) in the Minkowski sum get subdivided to give maximal cells (cones) in $\mathcal{T}(X \cdot Y)$ or, furthermore, the union of several cells (cones) in the Minkowski sum gives a maximal cell (cone) in $\mathcal{T}(X \cdot Y)$. We should not expect a natural description of the structure of $\mathcal{T}(X \cdot Y)$. We illustrate this undesirable property with an example.

Example 1.2.2. It may seem surprising at first that the combinatorial structure (e.g. f -vector) of the Newton polytope does not follow easily from the description of the tropical hypersurface as a Minkowski sum of two fans. Moreover, the number of edges of the polytope (and even the number of vertices) may exceed the number of maximal cones of the tropical hypersurface *given as a set*. To see this in a small example, consider the tropical curve in \mathbb{R}^3 whose six rays are columns of the following matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & -5 \\ 0 & 0 & 1 & 1 & 2 & -4 \\ 0 & 1 & 0 & 2 & 1 & -4 \end{pmatrix},$$

and consider the Minkowski sum of the fan with itself. This tropical hypersurface is described as a union of 15 cones (or as a non-planar graph in \mathbb{S}^2 with 6 nodes and 15 edges), but the dual Newton polytope has 16 vertices, 25 edges, and 11 facets. If we intersect the tropical hypersurface with a sphere around the origin, we would see the planar graph in Figure 1.1.

The planar regions correspond to the 16 vertices. The black dots correspond to the columns in the above matrix and the arcs between them correspond to cones generated by




Figure 1.1: A tropical surface in \mathbb{R}^3 described as a collection of two-dimensional cones in \mathbb{R}^3 or as a non-planar graph in \mathbb{S}^2 .

them. The nodes in the graph correspond to the facets in the Newton polytope. Six of these facets correspond to the black dots in Figure 1.1 and they are the six nodes in the non-planar graph description of the tropical hypersurface. The remaining five facets correspond to the five missing intersection points between the edges of the non-planar graph in the picture. Adding these five nodes to the graph will give us a planar graph with 11 nodes and 25 edges that encodes the fan structure of the tropical variety and the combinatorics of the Newton polytope.

If we had started instead with a tropical curve whose six rays are $\pm e_i$ for $i = 1, 2, 3$, then the dual polytope would be a cube with f -vector $(8, 12, 6)$. \diamond

As we saw in the previous example, knowledge of a fan structure for $\mathcal{T}(X \times Y)$ is useful when studying combinatorial properties of $X \times Y$. In the hypersurface case, it allows us to obtain the f -vector of the corresponding Newton polytope. In Chapter 4 we show an instance where the intrinsic Gröbner tropical fan structure can be derived (Theorem 4.5.3). The corresponding Hadamard product is the first secant variety of a monomial curve $(1 : t^{i_1} : t^{i_2} : t^{i_3} : t^{i_4})$ in \mathbb{P}^4 , with arbitrary set of exponents.

From Definition 1.2.1 we see that Hadamard products are images of monomial maps $\alpha: \mathbb{T}_{\mathbb{K}}^{2n} \rightarrow \mathbb{T}_{\mathbb{K}}^n$ associated to matrices of the form $A = (I_n \mid I_n)$, restricted to the product subvariety $X \times Y$. Applying Theorem 1.1.15 and knowing that $\mathcal{T}(X \times Y) = \mathcal{T}X \times \mathcal{T}Y$ (Theorem 3.3.4), we recover the set-theoretic tropicalization of $X \cdot Y$. As we mentioned in Section 1.1.2, up to now there is no push-forward formula for multiplicities in the arbitrary coefficient setting described in the literature. Such a formula would be extremely useful for computing weights on tropicalizations of Hadamard products over valued fields.

In the constant coefficient case (i.e. $\mathbb{K} = \mathbb{C}$), we can derive formulas for tropical multiplicities on Hadamard products thanks to the push-forward formula (1.4) and an explicit description of a canonical fan structure and weights in $\mathcal{T}(X \times Y)$. Theorem 3.3.4 and Corollary 3.3.7 provide such formulas under suitable finiteness conditions of our ‘‘Hadamard’’ monomial map $\alpha: X \times Y \subset \mathbb{T}^{2n} \rightarrow X \cdot Y \subset \mathbb{T}^n$. The examples studied in Chapters 3 and 4 possess this finiteness property.

1.3 Geometric tropicalization and tropical implicitization

In this section we present the basics of *geometric tropicalization* and its consequences for tropical implicitization. The crux of geometric tropicalization is to read off the tropicalization of an algebraic variety directly from the combinatorics of its boundary in a suitable compactification: a tropical compactification in the sense of Tevelev [94]. As one may suspect, computing the tropicalization of an algebraic variety without any information about its defining ideal is not an easy task. This point of view was pioneered by the work of Kapranov and his collaborators [30], and further developed by Hacking, Keel and Tevelev [45]. The starting point of this construction is a parametric representation of the variety and the characterization of its tropicalization in terms of *divisorial valuations*, following the spirit of [5]. This can be quite difficult if the variety is non-generic, which explains the small sample of computational examples available in the literature. Chapters 4 and 5 contribute to the theory by constructing weighted graphs encoding tropical surfaces using geometric tropicalization techniques.

We now discuss the main result in the theory of geometric tropicalization. Starting from a variety $X \subset \mathbb{T}^n$ and a suitable compactification, we construct an abstract simplicial complex: the intersection complex of its boundary. This complex encodes the combinatorics of the tropical fan $\mathcal{T}X$. This fan is the cone over a *realization* of this complex in \mathbb{R}^n , obtained by assigning an integer vector to every vertex of intersection complex. This vector encodes the divisorial valuation of the divisor associated to this vertex. We choose the word “realization” rather than embedding because this map need not be injective, as we show in Chapter 4. The support of this fan will be independent of the chosen tropical compactification.

Theorem 1.3.1 (Geometric Tropicalization [45, §2]). *Let χ_1, \dots, χ_n be the basis of characters of the n -dimensional torus \mathbb{T}^n , and let X be a closed subvariety of \mathbb{T}^n of dimension d . Suppose X is smooth and $\overline{X} \supset X$ is any compactification whose boundary $D = \overline{X} \setminus X$ is a smooth divisor with simple normal crossings. Let D_1, \dots, D_m be the irreducible components of D , and write $\Delta_{X,D}$ for the intersection complex of the boundary divisor D , i.e. the simplicial complex on $\{1, \dots, m\}$ defined by*

$$\{k_1, \dots, k_d\} \in \Delta_{X,D} \iff D_{k_1} \cap \dots \cap D_{k_d} \neq \emptyset.$$

Define the integer vectors $[D_k] := (\text{val}_{D_k}(\chi_1), \dots, \text{val}_{D_k}(\chi_n)) \in \mathbb{Z}^n$ ($k = 1, \dots, m$) where $\text{val}_{D_k}(\chi_j)$ is the order of zero/pole of χ_j along D_k . For any $\sigma \in \Delta_{X,D}$, let $[\sigma]$ be the semigroup spanned by $\{[D_k] : k \in \sigma\} \subset \mathbb{Z}^n$ and let $\mathbb{R}_{\geq 0}[\sigma]$ be the cone in \mathbb{R}^n spanned by the same integer vectors. Then,

$$\mathcal{T}X = \bigcup_{\sigma \in \Delta_{X,D}} \mathbb{R}_{\geq 0}[\sigma]. \quad (1.7)$$

Remark 1.3.2. The proof of this result in [45] shows that $\mathcal{T}X$ is contained in the right-hand side of (1.7) if we assume X to be normal [91, Remark 2.7]. Notice that if more than $\dim X$ divisors pass through a point, then the right hand side of (1.7) could (and in general will) contain cones of dimension greater than $\dim X$, violating the Bieri-Groves theorem. Thus, an expected necessary condition to achieve equality is to have a compactification X whose boundary has *combinatorial normal crossings* (CNC), i.e. where any of its k components intersect in codimension k . In particular, the complex $\Delta_{X,D}$ will be simplicial of pure dimension $\dim X - 1$.

We now explain how to obtain the divisorial valuations val_{D_k} . Consider the m irreducible components of the boundary of X , the basis of characters $\{\chi_1, \dots, \chi_n\}$ of \mathbb{T}^n and pullback the basis of characters along the inclusion $i: X \hookrightarrow \mathbb{T}^n$. By construction, the pullback $i^*(\chi_j)$ is a unit of \mathcal{O}_X and a rational function on \overline{X} . Therefore, for every j , the pullback $i^*(\chi_j)$ has zeros and poles *only* along the boundary of X . In particular, we can write $i^*(\chi_j)$ as an integer linear combination of the prime divisors D_k for all $k = 1, \dots, m$. In symbols,

$$i^*(\chi_j) = \sum_{k=1}^m a_{jk} D_k.$$

The divisorial valuation val_{D_k} satisfies $val_{D_k}(i^*(\chi_j)) = a_{jk}$. Thus, the integer vectors $[D_k]$ in the theorem correspond to the columns of the matrix of coefficients $(a_{jk})_{j,k}$.

To compute $\mathcal{T}X$ using Theorem 1.3.1, we need a method to construct a compactification $\overline{X} \supset X$ whose boundary has simple normal crossings (SNC). In words, we require all the components of the divisor D to be smooth and to intersect “as transversally as possible:” they must behave locally as an intersection of coordinate hyperplanes. The best compactification \overline{X} we can pick will be the *tropical* one, obtained by taking the closure $X(\Sigma)$ of X in a suitable projective toric variety $\mathbb{P}(\Sigma)$ associated to a fan $\Sigma \subset \mathbb{R}^n$:

$$\begin{array}{ccc} X & \hookrightarrow & \mathbb{T}^n \\ \downarrow & & \downarrow \\ X(\Sigma) & \hookrightarrow & \mathbb{P}(\Sigma) \end{array} \tag{1.8}$$

The fan Σ will be supported on the *tropical fan* $\mathcal{T}X$ [94].

One method for producing such a compactification is to take the closure \overline{X} of X in \mathbb{P}^n and resolve the singularities of the boundary $\overline{X} \setminus X$ to fulfill the SNC condition, for example, by blow-ups. The latter step can be difficult.

If X is a surface, we can relax the resolution process to give a boundary with CNC, i.e. such that no three of its irreducible components intersect at a point [91]. Roughly speaking, a full resolution will give us several extra (exceptional) divisors that yield bivalent nodes in the intersection complex $\Delta_{X,D}$. If we contract these curves with negative self-intersection,

we obtain a singular surface whose boundary divisor has CNC. This weaker condition will suffice to construct tropical surfaces using the algorithm described in Theorem 1.3.1.

Despite the richness of this new “valuative” approach, characterizing tropical varieties by divisorial valuations, this geometric perspective has a major obstacle to its full development: it does not provide information about the tropical variety as a *weighted* set. Multiplicities of maximal cells are missing in the construction of [45] and they are essential for tropical implicitization methods. We overcome this difficulty by providing an explicit combinatorial formula to compute these numbers:

Theorem 5.2.4. *In the notation of Theorem 1.3.1, the multiplicity of a regular point w in the tropical variety $\mathcal{T}X$ equals*

$$m_w = \sum_{\substack{\sigma \in \Delta_{X,D} \\ w \in \mathbb{R}_{\geq 0}[\sigma]}} (D_{k_1} \cdot \dots \cdot D_{k_d}) \operatorname{index}(\mathbb{R}[\sigma] \cap \mathbb{Z}^n, \mathbb{Z}[\sigma]), \quad (1.9)$$

where $D_{k_1} \cdot \dots \cdot D_{k_d}$ denotes the intersection number of these d divisors and we sum over all $(d-1)$ -dimensional cells σ in $\Delta_{X,D}$ whose associated rational cone $\mathbb{R}_{\geq 0}[\sigma]$ contains the point w .

It is worth pointing out that similar formulas were presented in the work of Sturmfels and Tevelev for generic complete intersections [91, Section 4]. In their paper, the intersection number of d divisors is replaced by a mixed volume of d polytopes associated to the objective vector $w \in \mathbb{R}^n$ and the equations defining the divisors D_{k_j} .

Tropical implicitization was pioneered by the work of Sturmfels, Tevelev and Yu [91, 92]. Their methods are well suited for *generic* varieties, and are built on the theory of geometric tropicalization. However, real life is seldom generic, so it is crucial to attack the *non-generic* versions of these problems.

In Chapter 5, we study this question for rational surfaces. Our main contribution is the development of new, explicit techniques for tropical implicitization of non-generic surfaces. In what follows, we explain the interaction between tropical implicitization and geometric tropicalization. Our main references are [91, 92].

We begin by discussing the *classical implicitization* problem. Let f_1, \dots, f_n be Laurent polynomials in $\mathbb{C}[t_1^\pm, \dots, t_r^\pm]$ and consider the rational map $\mathbf{f}: \mathbb{T}^r \dashrightarrow \mathbb{T}^n$, $\mathbf{f} = (f_1, \dots, f_n)$. For simplicity, we assume that the generic fiber of \mathbf{f} is finite. Our goal is to compute the defining ideal of Y , the Zariski closure of the image of \mathbf{f} . This information would provide a better understanding of this variety and would allow us to study the image of this map and certify membership and dimension computations.

Classical tools to solve implicitization problems rely on Gröbner bases computations and multidimensional resultants. These methods are infeasible to implement when the dimension of Y or its ambient space becomes large. So we need a new approach to go beyond this

computational threshold. Instead of calculating the defining ideal of Y , *tropical implicitization* will seek for the tropicalization of the variety $Y \subseteq \mathbb{T}^n$, without this ideal. We wish to compute $\mathcal{T}Y$ from the geometry of Y , using the map \mathbf{f} .

Following Theorem 1.3.1, our first step is to construct a nice compactification of Y and compute the associated divisorial valuations of its boundary. As we illustrate in Sections 4.3 and 5.4, the combinatorics involved in this compactification process is by no means trivial. We should not expect them to be simple since they are the manifestation of the algebro-geometric process of resolution of singularities.

In this dissertation, we approach this task from a different angle. Using the finiteness of the map \mathbf{f} , we study the dense open set $X \subset \mathbb{T}^r$ where \mathbf{f} is well-defined, that is, the complement of the hypersurface $\bigcup_{i=1}^n (f_i = 0) \subset \mathbb{T}^r$. Next, we construct a tropical compactification of X , together with its boundary intersection complex $\Delta_{X,D}$. This abstract complex will be isomorphic to the corresponding boundary complex for some compactification of Y . We recover $\mathcal{T}Y$ from the theorem by composing the divisorial valuations on Y with the induced map of residue fields $\mathbf{f}^\#: \mathbb{C}(Y) \rightarrow \mathbb{C}(\overline{X})$. In words, we construct the realization of $\Delta_{X,D}$ in \mathbb{R}^r and we use the map \mathbf{f} to push-forward this fan in \mathbb{R}^r to a fan in \mathbb{R}^n . Here is the precise statement, which combines Theorems 5.2.2 and 5.2.4 and Corollary 5.2.5.

Theorem 1.3.3. *Let $\mathbf{f}: X \subset \mathbb{T}^r \rightarrow \mathbb{T}^n$ be a generically finite map of degree δ and let $Y \subset \mathbb{T}^n$ be the Zariski closure of its image. Consider \overline{X} a tropical compactification of X and the corresponding boundary intersection complex $\Delta_{X,D}$. For each boundary component D_k of \overline{X} , let $\tilde{D}_k := \text{val}_{D_k}(\chi \circ \mathbf{f}) = (\text{val}_{D_k}(\chi_1 \circ \mathbf{f}), \dots, \text{val}_{D_k}(\chi_n \circ \mathbf{f})) \in \mathbb{Z}^n$ and define $\tilde{\sigma}$ as the semigroup spanned by $\{\tilde{D}_k : k \in \sigma\}$ for each $\sigma \in \Delta_{X,D}$. Then, the tropical variety $\mathcal{T}Y$ equals:*

$$\mathcal{T}Y = \bigcup_{\sigma \in \Delta_{X,D}^{\text{top}}} \mathbb{R}_{\geq 0}[\tilde{\sigma}],$$

where the superscript top indicates that we only consider top dimensional cells. The multiplicity of a regular point w in $\mathcal{T}Y$ equals

$$m_w = \frac{1}{\delta} \sum_{\substack{\sigma \in \Delta_{X,D} \\ w \in \mathbb{R}_{\geq 0}[\sigma]}} (D_{k_1} \cdot \dots \cdot D_{k_d}) \text{index}(\mathbb{R}[\tilde{\sigma}] \cap \mathbb{Z}^n, \mathbb{Z}[\tilde{\sigma}]).$$

As we discussed, tropical compactifications are required to have boundaries with simple normal crossings. In the surface case, we show that the sufficient condition is a combinatorial normal crossing boundary, i.e. where no triple intersections among boundary components occur, and where pairs of divisors admit branches with no common tangent directions. Such conditions can be achieved by contracting negative curves after a full resolution.

Proposition 5.2.10. *Let \overline{X} be a compactification of a surface X whose boundary satisfies the combinatorial normal crossing condition and such that no pair of boundary components have branches with the same tangent directions. Then \overline{X} computes the tropical surface $\mathcal{T}(X)$.*

When the coefficients of f_1, \dots, f_n are generic with respect to their Newton polytopes, an explicit \overline{X} was proposed in [92, 91]. The genericity condition guarantees transverse intersection and smoothness of the hypersurfaces ($f_i = 0$) and gives a natural compactification of X in a projective toric variety $\mathbb{P}(\mathcal{N})$, whose fan \mathcal{N} is a strictly simplicial refinement of the common refinement of all $\mathcal{T}(f_i)$, $i = 1, \dots, n$. In addition to the toric divisors associated to the rays of the fan \mathcal{N} we have n extra divisors \overline{E}_i representing the closure of the n hypersurfaces ($f_i = 0$) in \mathbb{T}^n . For every ray ρ in \mathcal{N} , we let n_ρ be its primitive integer generator and D_ρ the corresponding toric divisor. The realization of the complex $\Delta_{Y,D}$ will correspond to the assignment:

$$[D_\rho] = \text{trop}(f)(n_\rho), \quad [\overline{E}_i] = e_i,$$

for all rays ρ in \mathcal{N} and indices $i = 1, \dots, n$, where e_i is the i^{th} element of the canonical basis of \mathbb{R}^n . Multiplicities are calculated in terms of mixed volumes [91]. Section 5.3 describes this algorithm for generic surfaces.

As expected, the non-generic case is intrinsically more difficult. A first attempt to solve this tropical implicitization question is to proceed as if the polynomials f_i were generic. This will give a compactification $\overline{X} \subset \mathbb{P}(\mathcal{N})$, whose boundary has bad crossings. Resolving these singularities in the toric case will not be advantageous, since they will not be torus-invariant in most cases. Therefore, *toric blow-up* methods cannot be applied.

A second approach starts by picking the naive compactification of X in \mathbb{P}^d . This adds one extra divisor to the boundary $\bigcup_{i=1}^n (f_i = 0)$: the divisor at infinity. Next, we analyze the crossings among these divisors and resolve all singularities. In principle, this can be achieved by blow-ups, and the difficulty becomes algebraic, since we need to carry all valuations along the successive blow-ups. Section 4.3 discusses this algorithm for a binomial surface in \mathbb{T}^n and Section 5.4 illustrates this procedure for non-generic surfaces in \mathbb{C}^3 .

Example 5.4.3. Consider the surface in \mathbb{T}^3 parameterized by the three polynomials

$$\begin{cases} f_1(s, t) = s^2 - s^3 - t^2, \\ f_2(s, t) = t^2 - t^3 - s^2, \\ f_3(s, t) = (s + t)^2 - (s + t)^3 - (s - t)^2 = 4st - s^3 - t^3 - 3st^2 - 3s^2t. \end{cases}$$

In the first step, we compactify the set $X = \mathbb{T}^2 \setminus \bigcup_{i=1}^3 (f_i = 0)$ inside \mathbb{P}^2 . As a result, we obtain the set \overline{X} whose boundary divisor consists of four components: $D_1 = (s^2u - s^3 - t^2u = 0)$, $D_2 = (t^2u - t^3 - s^2u = 0)$, $D_3 = (4stu - s^3 - t^3 - 3st^2 - 3s^2t = 0)$ from \mathbf{f} , and the divisor at infinity $D_\infty = (u = 0)$. However, the first three divisors intersect at $(0 : 0 : 1)$, giving a three-dimensional cone in the right hand side of (1.7).

In each step of the resolution, we add a ray in the interior of this three-dimensional cone, corresponding to an exceptional divisor. After each blow-up, the strict transform of some of the original divisors no longer intersect, and some of them will also intersect with an exceptional divisor, giving a three-dimensional subcone of the original one. In the end of the resolution, all triple intersections disappear and we are left with an abstract graph $\Delta_{\overline{X},D}$.

Its realization will be the desired tropical surface graph. In the first compactification, the divisorial valuations are obtained from the columns of the matrix of coefficients describing $\mathbf{f}^*(\chi_1, \chi_2, \chi_3)$:

$$\begin{cases} \mathbf{f}^*(\chi_1) = (f_1^h(s, t, u)) = & D_1 & -3D_\infty, \\ \mathbf{f}^*(\chi_2) = (f_2^h(s, t, u)) = & D_2 & -3D_\infty, \\ \mathbf{f}^*(\chi_3) = (f_3^h(s, t, u)) = & D_3 - 3D_\infty, \end{cases}$$

where $f_i^h(s, t, u)$ is the homogeneization of $f_i(s, t)$ with respect to the variable u , that is $f_i^h(s, t, u) = u^{\deg f_i} f_i(s/u, t/u)$.

The above linear combination gives $[D_1] = e_1$, $[D_2] = e_2$, $[D_3] = e_3$ and $[D_\infty] = (-3, -3, -3)$. In addition, the pairwise intersection numbers equal $D_i \cdot D_j = 9$, $D_\infty \cdot D_i = 3$, where $i, j = 1, 2, 3$, $i \neq j$. Figure 1.2 shows both intersection complexes and their weights, following formula (1.9). Further details, including the resolution process are given in Section 5.4 and Example 5.4.3 therein.




Figure 1.2: From left to right: first intersection complex for $\overline{X} \subset \mathbb{P}^2$ and intersection complex \mathcal{TY} obtained at the end of the resolution. The nodes $[E_i]$, $i = 1, \dots, 4$, correspond to exceptional divisors.

◇

Chapter 4 illustrates the richness of geometric tropicalization for a particular choice of surfaces: binomial surfaces representing a dehomogenization of the secant varieties of monomial projective curves with *arbitrary* sets of exponents (e.g. the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$ from Example 4.2.3). Section 4.3 describes the construction of a tropical compactification of this binomial arrangement, which resembles the recent extensions by Moci [71] of De Concini-Procesi’s *wonderful models* for hyperplane arrangements [19]. This example illustrates our philosophy: only partial resolutions are required for tropicalizing surfaces. The difference between Moci’s construction and ours is simply a matter of perspective. Consider a fixed subfamily of plane curves. In Moci’s compactification, every intersection point between members of this subfamily must be resolved independently. However, they all share

the same local behavior, so for all intents and purposes we can treat them as the same point. This is, precisely, our approach to the problem.

As a corollary of the previous construction, using Hadamard products we can compute the tropicalization of the secant variety of any monomial curve in \mathbb{P}^n and calculate its degree. The tropical approach reinterprets classical formulas for computing the degree of such secant varieties by Conca and Ranestad [13, 82]. More importantly, it provides a way of recovering their Chow polytopes, as we show in the next section.

1.4 Algorithms for implicitization

In this section, we discuss the *inverse problem*, that is, how to obtain relevant information about algebraic objects from their tropical counterparts. Our attention will be concentrated on the hypersurface case. As we mentioned in Section 1.1, tropical hypersurfaces are dual to polytopal subdivisions. These weighted polyhedral complexes are obtained as the collection of codimension one cells in the dual complex of subdivided lattice polytopes. Each maximal cell in the tropical variety is dual to an edge of the subdivided polytope, and its weight equals the lattice length of this edge. Thus, a natural question to ask is the following. If we are given a tropical hypersurface, how can we recover the corresponding polytopal subdivision?

In what follows, we assume our base field to be the complex numbers, or any algebraically closed field with trivial valuation. In this case, our question reduces to the duality between inner normal fans and lattice polytopes. Suppose our hypersurface is defined by the equation $f \in \mathbb{C}[x_1, \dots, x_n]$. A construction for the vertices of the Newton polytope $\text{NP}(f)$ from its normal fan $\mathcal{T}(f)$ equipped with multiplicities was developed in [24]. The algorithm takes as input a generic objective vector w in \mathbb{R}^n , i.e. a vector in the interior of a maximal cone (chamber) of the normal fan of $\text{NP}(f)$, and outputs an extreme monomial of f , namely, the one obtained as the initial form $\text{in}_w(f)$. Once the Newton polytope is computed, the defining equation f can be obtained using linear algebra.

The following is a special case of [24, Theorem 2.2]. Since the operation $\mathcal{T}(f)$ interprets f as a Laurent polynomial, $\text{NP}(f)$ will be determined from $\mathcal{T}(f)$ up to translation. The algorithm described in Theorem 1.4.1 computes a representative of $\text{NP}(f)$ which lies in the positive orthant and touches all coordinate hyperplanes, i.e. f is a polynomial not divisible by any non-constant monomial. We present a pseudocode implementation in Algorithm 3.1.

Theorem 1.4.1 (*Ray-shooting algorithm*, [24]). *Suppose $w \in \mathbb{R}^n$ is a generic vector so that the ray $(w + \mathbb{R}_{>0} e_i)$ intersects $\mathcal{T}(f)$ only at regular points for all i . Let \mathcal{P}^w be the vertex of the polytope $\mathcal{P} = \text{NP}(f)$ that attains the minimum of $\{w \cdot x : x \in \mathcal{P}\}$. Then the i^{th} coordinate of \mathcal{P}^w equals*

$$\sum_v m_v \cdot |l_i^v|, \tag{1.10}$$

where the sum is taken over all points $v \in \mathcal{T}(f) \cap (w + \mathbb{R}_{>0} e_i)$, m_v is the multiplicity of v in

$\mathcal{T}(f)$, and l_i^v is the i^{th} coordinate of the primitive integral normal vector l^v to the maximal cone in $\mathcal{T}(f)$ containing v .

Figure 1.3 illustrates the heart of formula (1.10). For a vertex \mathcal{P}^w , we represent its i^{th} coordinate v_i as the distance from the vertex \mathcal{P}^w to the hyperplane $H_i = (x_i = 0)$. We partition the segment joining \mathcal{P}^w and the nearest point q in H_i into pieces. Each piece will be the projection of an edge of the polytope \mathcal{P} in the edge path from \mathcal{P}^w to q obtained from the algorithm: each point in $\mathcal{T}(f) \cap (w + \mathbb{R}_{>0}e_i)$ corresponds to an edge in this path. Orient the edges in this path from \mathcal{P}^w to q . The projection of each edge is associated to a point $v \in \mathcal{T}(f) \cap (w + \mathbb{R}_{>0}e_i)$ and has length equal to $m_v |l_i^v|$, i.e. the lattice length of the edge in \mathcal{P} times the absolute value of the i^{th} coordinate of its generating vector. Figure 1.3 illustrates this algorithm for a full dimensional polytope in \mathbb{R}^2 and expresses \mathcal{P}_1^w as the sum of the contributions of two edges.




Figure 1.3: Ray-shooting algorithm: the picture shows how to obtain the first coordinate of vertex \mathcal{P}^w as the sum of the contributions of the two edges connecting this vertex to the y -axis. These two edges are dual to the points v, v' obtained by ray-shooting along the direction e_1 from a generic objective vector w in the chamber associated to vertex \mathcal{P}^w .

Theorem 1.4.1 can be generalized to projective varieties of higher codimension, replacing the ray-shooting by *orthant-shooting* [24, Theorem 2.2]. In this case, the role of the Newton polytope is played by the *Chow polytope* [55] of the variety. Knowledge of a single vertex of this polytope gives us the degree (or multidegree) of the corresponding projective variety. Geometrically, its vertices encode toric degenerations of the algebraic cycle underlying the projective variety. The connection between tropicalization and Chow polytopes began with [24] and has been further studied in [37, 38]. Example 4.6.1 illustrates this method, where we compute the Chow polytope associated to the secant variety of the rational normal curve in \mathbb{P}^n with $n \geq 5$ (Example 4.6.1).

Note that we do not need a fan structure on $\mathcal{T}(f)$ to use Theorem 1.4.1. A description of $\mathcal{T}(f)$ as a set, together with a way to compute the multiplicities of regular points gives us enough information to compute vertices of $\text{NP}(f)$ in any generic direction. This small remark is crucial for practical applications, since most characterizations of tropical varieties in the literature only give a set theoretic description of the space and multiplicities of regular points. In most cases, it is extremely un-natural to obtain any tropical fan structure from this description.

By picking an objective vector in each chamber of the normal fan and running the ray-shooting algorithm we obtain a single vertex of the polytope. A natural goal in any implementation is a method for picking objective vectors by passing exactly once through each chamber. This problem is harder to solve if we lack a fan structure description, i.e. if we define the chambers as the complement of a collection of codimension one cones. Moreover, if this collection is large, it is very ineffective to use ray-shooting to obtain vertices one at a time.

In Chapter 3 we give an answer to these questions inspired by a concrete implicitization task for the restricted Boltzmann machine. We develop three algorithms to successfully compute the Newton polytope of a tropical hypersurface, starting from a single generic objective vector in \mathbb{R}^n . We call them the *walking algorithm* (Algorithm 3.2), the *facet certificate algorithm* (Algorithm 3.3) and the *approximation algorithm* (Algorithm 3.4).

The walking algorithm performs the following operation: starting from an objective vector and the corresponding associated vertex, we use the principle of ray-shooting to recursively pick an objective vector from a neighboring chamber together with its associated vertex, each time we cross a codimension one cone. These crossings are recorded by the points of $\mathcal{T}(f) \cap (w + \mathbb{R}_{>0}e_i)$, $i = 1, \dots, n$.

Theorem 1.4.2 (Walking algorithm). *Suppose $w \in \mathbb{R}^n$ is a generic vector so that the ray $(w + \mathbb{R}_{>0}e_i)$ intersects $\mathcal{T}(f)$ only at regular points of $\mathcal{T}(f)$, for all i . Let \mathcal{P}^w be the vertex of the polytope $\mathcal{P} = \text{NP}(f)$ that attains the minimum of $\{w \cdot x : x \in \mathcal{P}\}$. For each coordinate, call v_1, \dots, v_r the intersection points of $w + \mathbb{R}_{>0}e_i$ and $\mathcal{T}(f)$, ordered increasingly by their distance to w . Call $t_1, \dots, t_r \in \mathbb{R}_{>0}$ the corresponding parameter of each intersection point and pick $t_j < s_j < t_{j+1}$ for each $j = 1, \dots, r - 1$. Then, vertices and random vectors along the direction e_i can be computed as:*

$$w_j = w + s_j \cdot e_i, \quad \mathcal{P}^{w_j} = \mathcal{P}^w + m_{v_j} \cdot l^{v_j}, \quad j = 1, \dots, r - 1.$$

Algorithm 3.2 describes a pseudocode implementation of the previous theorem. In practice, we start from a random objective vector w , and we apply the ray-shooting algorithm to obtain both \mathcal{P}^w and the intersection points v_1, \dots, v_r . Then, we use each new random vector obtained by the walking algorithm as the input for a new iteration of ray-shooting, and we repeat this process. Figures 3.2 and 3.3 show the interface between the ray-shooting and walking algorithms. Once the combination of ray-shooting and walking algorithms stops

giving new vertices, we need a certificate of completion of the polytope $\text{NP}(f)$. This is done by combining the facet certificate and approximation algorithms.

The *facet certificate algorithm* provides a way of determining if a given linear inequality defines a facet of a polytope by looking at its tropicalization. By duality, each facet direction is a ray in the Gröbner fan structure of the tropical variety. Since the tropical hypersurface is given as a collection of weighted codimension one cones with no fan structure, we need a computational certificate for rays in the tropical hypersurface. This is provided by the following lemma, discussed in Section 3.4.

Lemma 3.4.3. *Let $w \in \mathbb{R}^n$ and $\mathcal{T}(f)$ be a tropical hypersurface given by a collection of cones. Let d be the dimension of its lineality space. Let $\mathcal{H} = \{\sigma_1, \dots, \sigma_k\}$ be the list of cones containing w . Let l^{σ_i} be the normal vector to cone σ_i for $i = 1, \dots, k$. Then, w is a ray of $\mathcal{T}(f)$ equipped with the Gröbner fan structure if and only if $\{l^{\sigma_1}, \dots, l^{\sigma_k}\}$ generates a $(n - d - 1)$ -dimensional vector space if and only if w is a facet direction of $\text{NP}(f)$.*

Given a ray w in the tropical variety, the corresponding facet in the polytope \mathcal{P} will be determined by a constant. This number can be obtained by taking the inner product between w and any vertex contained in the facet \mathcal{P}^w . An objective vector in any neighboring chamber of w will provide such vertex. The facet certificate algorithm outputs such objective vector, by a slight modification of ray-shooting and walking algorithms. The core of the algorithm lies in the following result:

Theorem 1.4.3 (Facet certificate). *Let $w \in \mathbb{R}^n$ and $i = 1, \dots, n$. Let $t_i > 0$ be such that $w + t_i e_i$ is the closest point to w in the nonempty set $\mathcal{T}(f) \cap (w + \mathbb{R}_{>0} e_i)$, or set t_i to be infinity otherwise. Fix $0 < s_i < t_i$ and define $\tilde{w} = w + s_i e_i$. Then \tilde{w} lies in a cone in the normal fan of $\text{NP}(f)$ of dimension at least $\min\{n, 1 + \text{codim } \mathcal{P}^w\}$, where \mathcal{P}^w is the face of $\text{NP}(f)$ defined by w .*

In particular, we can obtain a generic vector on a neighboring chamber of a ray in $\mathcal{T}(f)$ by starting with this ray, the index $i = 1$, and iterating the process by increasing the index by one on each step and using the output vector \tilde{w} as the input for the next iteration until we reach a chamber. Algorithm 3.3 contains a pseudocode implementation of these ideas.

Finally, we discuss the *approximation algorithm* and its geometric meaning. The main objects in this algorithm are the edges of $\text{NP}(f)$ normal to the cones in $\mathcal{T}(f)$ and the *tangent cones* at each vertex of the polytope $\text{NP}(f)$. For each vertex v in a polytope \mathcal{P} we define the tangent cone by:

$$\mathcal{T}_v^{\mathcal{P}} := v + \mathbb{R}_{\geq 0} \langle w - v : w \in \mathcal{P} \rangle = v + \mathbb{R}_{\geq 0} \langle e : e \text{ edge of } \mathcal{P} \text{ adjacent to } v \rangle.$$

The approximation algorithm is based on the following technical lemma from Section 3.4.5:

Lemma 3.4.5. *Let \mathcal{P} be a polytope and $\mathcal{Q} \subseteq \mathcal{P}$ be the convex hull of a subset of the vertices in \mathcal{P} . If all facets of \mathcal{Q} are facets of \mathcal{P} , then $\mathcal{Q} \supseteq \mathcal{P}$, so $\mathcal{Q} = \mathcal{P}$.*

Given a subset of the vertices of a polytope \mathcal{P} we construct its convex hull. This polytope \mathcal{Q} will be a subpolytope of \mathcal{P} and we aim to recover \mathcal{P} from \mathcal{Q} by successive approximations, using the list of known edge directions of \mathcal{P} . We start by computing the tangent cones of all vertices of \mathcal{Q} and we compare them with the corresponding tangent cones in \mathcal{P} . This comparison is done by constructing a cone $C_v^{\mathcal{Q},\mathcal{P}}$ with apex v a vertex in \mathcal{Q} and whose generating rays are all the differences $w - v$ where w runs along the set of all vertices of \mathcal{Q} and $w - v$ is *parallel* to an edge in \mathcal{P} .

Theorem 1.4.4 (Approximation algorithm). *Let $\mathcal{Q} \subseteq \mathcal{P}$ be a subpolytope spanned by some vertices of \mathcal{P} , and let v be a vertex of \mathcal{Q} . Consider the cone $C_v^{\mathcal{Q},\mathcal{P}}$ as above. If all facet inequalities of this cone are facets of \mathcal{P} then we have $C_v^{\mathcal{Q},\mathcal{P}} = \mathcal{T}_v^{\mathcal{Q}} = \mathcal{T}_v^{\mathcal{P}}$. Moreover, if this equality holds for every vertex in \mathcal{Q} , then $\mathcal{Q} = \mathcal{P}$.*

If $C_v^{\mathcal{Q},\mathcal{P}} \not\subseteq \mathcal{T}_v^{\mathcal{Q}}$ or a facet of $C_v^{\mathcal{Q},\mathcal{P}}$ is not a facet of \mathcal{P} , we obtained a new vertex adjacent to v by perturbing an objective vector giving v , adding this new vertex of \mathcal{P} to the list of vertices of \mathcal{Q} and running the algorithm again until we certify the hypothesis of the theorem. Methods to perform this perturbation in a systematic way are discussed in Section 3.4. Algorithm 3.4 gives a pseudocode implementation of Theorem 1.4.4. Figure 3.4 shows a pictorial description of this result in dimension two.

Even though there is no condition on the cones in the collection $\mathcal{T}(f)$, our practical implementations of all four algorithms described in this section assume our cones are simplicial. This strong condition is used to simplify the computation of intersection points between rays and cones, and to test membership of vectors to $\mathcal{T}(f)$. If the input data contains non-simplicial cones, we preprocess these bad cones. We subdivide them into simplicial pieces with standard methods in geometric combinatorics and replace each one with this finite list of simplicial cones. Each simplicial piece will inherit its weight from the original non-simplicial cone.

In analogy with the generalized *orthant-shooting* algorithm [24, Theorem 2.2], the other three algorithms can also be extended to higher codimension irreducible projective varieties and their Chow polytopes. The four algorithms described in this section give a complete answer to the task of recovering the Newton polytope (or Chow polytope) from its associated tropical variety. Implementations of these algorithms are discussed extensively in Section 3.4.3.

Chapter 3 illustrates these powerful tropical implicitization tools with a concrete example: the restricted Boltzmann machine $\mathcal{F}(4, 2)$. This “machine” is an algebraic variety obtained from the undirected graphical model of the complete bipartite graph $K_{2,4}$ by marginalizing two of the six binary random variables (Figure 3.1). It defines a degree 110 hypersurface in \mathbb{P}^{15} with a four-dimensional torus action and a natural symmetry given by B_4 , the symmetry group of the 4-cube (Theorem 3.4.2). The sheer size of this computational challenge forced us to derive, and efficiently implement, the techniques described above to go from a tropical hypersurface $\mathcal{T}(f)$ to the Newton polytope of f . After months of strenuous computations,

and using the symmetry of this object, we were able to describe all vertices and facets of this polytope:

Theorem 3.4.1. *The Newton polytope of the restricted Boltzmann machine $\mathcal{F}(4, 2)$ has 17 214 912 vertices in 44 938 symmetry classes and 70 646 facets in 246 symmetry classes under the group B_4 .*

Chapter 2

Geometry of the restricted Boltzmann machine

This chapter is joint work with Jason Morton and Bernd Sturmfels. It has been published in *Algebraic Methods in Statistics and Probability* (M. Viana and H. Wynn, eds.), American Mathematical Society, Contemporary Mathematics **516** (2010), 135–153, under the same title. The present version incorporates some minor changes, largely for consistency with other chapters.

2.1 Introduction

In recent years, a fruitful interaction between (computational) algebraic geometry and statistics has emerged, under the form of algebraic statistics. A primary focus in algebraic statistics is the study of statistical models that can be represented by polynomials in the model parameters. This class of algebraic statistical models includes graphical models for both Gaussian and discrete random variables [28, 27]. In this chapter we study a family of binary graphical models with hidden variables. The underlying graph is the complete bipartite graph $K_{k,n}$.

The k white nodes in the top row of Figure 2.1 represent hidden random variables. The n black nodes in the bottom row represent observed random variables. The *restricted Boltzmann machine* (RBM) is the undirected graphical model for binary random variables specified by this bipartite graph. Each node represents a binary random variable and each edge represents a dependency between two random variables. In other words, if there is no edge between two random variables, then they are conditionally independent given the rest of the variables. We identify the model with the set M_n^k of its joint distributions inside the probability simplex Δ_{2^n-1} .

The graphical model for Gaussian random variables represented by Figure 2.1 is the *factor analysis* model, whose algebraic properties were studied in [7, 25, 26]. Thus, the restricted Boltzmann machine is the binary undirected analog of factor analysis. Our aim here is to




Figure 2.1: Graphical representation of the restricted Boltzmann machine.

study this model from the perspectives of algebra and geometry. Unlike in the factor analysis study [26], an important role will now be played by *tropical geometry* [75]. This will be seen for $n = 4$ and $k = 2$ in the solution of the implicitization challenge in [27, Problem 7.7], describe in Chapter 3.

The restricted Boltzmann machine has been the subject of a recent resurgence of interest due to its role as the building block of the deep belief network [48, Section 17.4.4]. Deep belief networks are designed to learn feature hierarchies to automatically find high-level representations for high-dimensional data. A deep belief network comprises a stack of restricted Boltzmann machines. Given a piece of data (state of the lowest visible variables), each layer's most likely hidden states are treated as data for the next layer. A new effective training methodology for deep belief networks, which begins by training each layer in turn as an RBM using contrastive divergence, was introduced by Hinton et al. [49]. This method led to many new applications in general machine learning problems including object recognition and dimensionality reduction [50]. While promising for practical applications, the scope and basic properties of these statistical models have only begun to be studied. For example, Le Roux and Bengio [61] showed that any distribution with support on r visible states may be arbitrarily well approximated provided there are at least $r + 1$ hidden nodes. Therefore, any distribution can be approximated with $2^n + 1$ hidden nodes.

The natural starting question is whether the restricted Boltzmann machine model is *identifiable*, i.e. whether the parametrization of the model is locally one-to-one. The dimension of the fully observed binary graphical model on $K_{k,n}$ is equal to $nk + n + k$, the number of nodes plus the number of edges. We conjecture that this dimension is preserved under the projection corresponding to the algebraic elimination of the k hidden variables. Here is the precise statement:

Conjecture 2.1.1. *The restricted Boltzmann machine has the expected dimension, i.e. M_n^k is a semialgebraic set of dimension $\min\{nk + n + k, 2^n - 1\}$ in $\Delta_{2^n - 1}$.*

This conjecture is shown to be true in many special cases. In particular, it holds for all k when $n + 1$ is a power of 2. This is a consequence of the following:

Theorem 2.1.2. *The restricted Boltzmann machine has the expected dimension $\min\{nk + n + k, 2^n - 1\}$ when $k \leq 2^{n - \lceil \log_2(n+1) \rceil}$ and when $k \geq 2^{n - \lfloor \log_2(n+1) \rfloor}$.*

We note that Theorem 2.1.2 covers most cases of restricted Boltzmann machines as used in practice, as those generally satisfy $k \leq 2^{n - \lceil \log_2(n+1) \rceil}$. In particular, we conclude that the model is identifiable in these cases. The case of large k is primarily of theoretical interest and has been studied recently in [61].

This chapter is organized as follows. In Section 2.2 we introduce four geometric objects, namely, the RBM model, the RBM variety, the tropical RBM model, and the tropical RBM variety, and we formulate a strengthening of Conjecture 2.1.1. Section 2.3 is concerned with the case $k = 1$. Here the RBM variety is the variety of secant lines of the Segre variety $(\mathbb{P}^1)^n \subset \mathbb{P}^{2^n - 1}$. The general case $k > 1$ arises from that secant variety by way of a construction we call the *Hadamard product of projective varieties*, introduced in Section 1.2. In Section 2.4 we analyze the tropical RBM model, we establish a formula for its dimension (Theorem 2.4.2), and we draw on results from coding theory to derive Theorem 2.1.2 and Table 2.1. In Section 2.5 we study the piecewise-linear map that parameterizes the tropical RBM model. The *inference functions* of the model (in the sense of [32, 75]) are k -tuples of *linear threshold functions*. We discuss the number of these functions. Figures 2.5 and 2.6 shows the combinatorial structure of the tropical RBM model for $k = 1$ and $n=3$ and 2, respectively.

2.2 Algebraic varieties, Hadamard product and tropicalization

We begin with an alternative definition of the restricted Boltzmann machine. This “machine” is a statistical model for binary random variables where n of the variables are visible and k of the variables are hidden. The states of the hidden and visible variables are written as binary vectors $h \in \{0, 1\}^k$ and $v \in \{0, 1\}^n$ respectively. We introduce $nk + n + k$ model parameters, namely, the entries of a real $k \times n$ matrix W and the entries of two vectors $b \in \mathbb{R}^n$ and $c \in \mathbb{R}^k$, and we set

$$\psi(v, h) = \exp(h^\top W v + b^\top v + c^\top h). \quad (2.1)$$

The probability distribution on the visible random variables in our model equals

$$p(v) = \frac{1}{Z} \cdot \sum_{h \in \{0, 1\}^k} \psi(v, h), \quad (2.2)$$

where $Z = \sum_{v, h} \psi(v, h)$ is the *partition function*. We denote by M_n^k the subset of the open probability simplex $\Delta_{2^n - 1}$ consisting of all such distributions $(p(v) : v \in \{0, 1\}^n)$ as the parameters W, b and c run over $\mathbb{R}^{k \times n}$, \mathbb{R}^n and \mathbb{R}^k respectively.

In what follows we refer to M_n^k as the *RBM model* with n visible nodes and k hidden nodes. It coincides with the binary graphical model associated with the complete bipartite graph $K_{k,n}$ as described in the introduction. This is indicated in Figure 2.1 by the labeling with the states v, h and the model parameters c, W, b .

The parameterization in (2.1) is not polynomial because it involves the exponential function. However, it is equivalent to the polynomial map obtained by replacing each model parameter by its value under the exponential function:

$$\gamma_i = \exp(c_i), \quad \omega_{ij} = \exp(W_{ij}), \quad \beta_j = \exp(b_j)^{1/k}.$$

This coordinate change translates (2.1) into the squarefree monomial

$$\psi(v, h) = \prod_{i=1}^k \gamma_i^{h_i} \cdot \prod_{i=1}^k \prod_{j=1}^n \omega_{ij}^{h_i v_j} \cdot \prod_{j=1}^n \beta_j^{v_j \cdot k},$$

and we see that the probabilities in (2.2) can be factored as follows:

$$p(v) = \frac{1}{Z} \prod_{i=1}^k \left(\beta_1^{v_1} \beta_2^{v_2} \cdots \beta_n^{v_n} \cdot (1 + \gamma_i \omega_{i1}^{v_1} \omega_{i2}^{v_2} \cdots \omega_{in}^{v_n}) \right) \quad \text{for } v \in \{0, 1\}^n. \quad (2.3)$$

The RBM model M_n^k is the image of the polynomial map $\mathbb{R}_{>0}^{nk+k+n} \rightarrow \Delta_{2^n-1}$ whose v^{th} coordinate equals (2.3). The Tarski-Seidenberg Theorem from real algebraic geometry implies that M_n^k is a semialgebraic subset of Δ_{2^n-1} .

When faced with a high-dimensional semialgebraic set in statistics, it is often useful to simplify the situation by disregarding all inequalities and by replacing the real numbers \mathbb{R} by the complex numbers \mathbb{C} . This leads us to considering the Zariski closure V_n^k of the RBM model M_n^k . This is the algebraic variety in the complex projective space \mathbb{P}^{2^n-1} parameterized by (2.3). We call V_n^k the *RBM variety*.

For any projective variety X , we may consider its Hadamard square $X^{[2]} = X \cdot X$ and its higher Hadamard powers $X^{[k]} = X \cdot X^{[k-1]}$, as in Definition 1.2.1. If M is a subset of the open simplex Δ_{m-1} then its Hadamard powers $M^{[k]}$ are also defined by componentwise multiplication followed by rescaling so that the coordinates sum to one. This construction is compatible with taking Zariski closures, i.e. we have $\overline{M^{[k]}} = \overline{M}^{[k]}$.

In the next section we shall take a closer look at the case $k = 1$, and we shall recognize V_n^1 as a secant variety and M_n^1 as a phylogenetic model. Here, we prove that the case of $k > 1$ hidden nodes reduces to $k = 1$ using Hadamard powers.

Proposition 2.2.1. *The RBM variety and model factor as Hadamard powers:*

$$V_n^k = (V_n^1)^{[k]} \quad \text{and} \quad M_n^k = (M_n^1)^{[k]}.$$

Proof. A strictly positive vector p with coordinates $p(v)$ as in (2.3) admits a componentwise factorization into similar vectors for $k = 1$, and, conversely, the componentwise product of k probability distributions in M_n^1 becomes a distribution in M_n^k after division by the partition function. Hence $M_n^k = (M_n^1)^{[k]}$ in Δ_{2^n-1} . The equation $V_n^k = (V_n^1)^{[k]}$ follows by passing to the Zariski closure in \mathbb{P}^{2^n-1} . \square

The emerging field of *tropical mathematics* is predicated on the idea that $\log(\exp(x) + \exp(y))$ is approximately equal to $\max(x, y)$ when x and y are quantities of different scale. The process of passing from ordinary arithmetic to the max-plus algebra is known as *tropicalization*. The same approximation motivates the definition of the *softmax* function in the neural networks literature. A statistical perspective is offered in work by Pachter and Sturmfels [76, 75].

If $q(v)$ approximates $-\log(p(v))$ in the sense of tropical mathematics, and if we disregard the global additive constant $\log Z$ and change signs of our parameters W, b, c , then (2.2) translates into the formula

$$q(v) = \min\{h^\top Wv + b^\top v + c^\top h : h \in \{0, 1\}^k\}. \quad (2.4)$$

This expression is a piecewise-linear concave function $\mathbb{R}^{nk+n+k} \rightarrow \mathbb{R}$ on the space of model parameters (W, b, c) . As v ranges over $\{0, 1\}^n$, there are 2^n such concave functions, and these form the coordinates of a piecewise-linear map

$$\Phi : \mathbb{R}^{nk+n+k} \rightarrow \mathbb{TP}^{2^n-1}. \quad (2.5)$$

The image of the map Φ is denoted $\mathcal{T}M_n^k$ and is called the *tropical RBM model*. The map Φ is the *tropicalization* of the given parameterization of the RBM model. It is our objective to investigate its geometric properties.

This situation fits precisely into the general scheme of parametric maximum a posteriori (MAP) inference introduced in [75] and studied in more detail by Elizalde and Woods [32]. In Section 2.5 below, we discuss the statistical relevance of the map Φ and we examine its geometric properties. Of particular interest are the domains of linearity of Φ , and how these are mapped onto the cones of the model $\mathcal{T}M_n^k$.

Finally, we define the *tropical RBM variety* $\mathcal{T}V_n^k$ to be the tropicalization of the RBM variety V_n^k . As explained Section 1.1.1, the tropical variety $\mathcal{T}V_n^k$ is the intersection in \mathbb{TP}^{2^n-1} of all the tropical hypersurfaces $\mathcal{T}(f)$ where f runs over *all* polynomials that vanish on V_n^k (or on M_n^k). By definition, $\mathcal{T}(f)$ is the union of all codimension one cones in the inner normal fan of the Newton polytope of f . If the homogeneous prime ideal of the variety V_n^k were known then the tropical variety $\mathcal{T}V_n^k$ could in theory be computed using the algorithms in [6] which are implemented in the software `Gfan` ([54]). However, this prime ideal is not known in general. In fact, even for small instances, its computation is very hard and relies primarily on tropical geometry techniques such as the ones developed in Chapter 3. For instance, the main result in that chapter states that the RBM variety V_4^2 is a hypersurface of degree 110 in \mathbb{P}^{15} ,

and it remains a challenge to determine a formula for the defining irreducible polynomial of this hypersurface. To appreciate this challenge, note that the number of extreme monomials in this polynomial equals 17 214 912 (Theorem 3.4.1), whereas the number of monomials in the relevant multidegree equals 5 529 528 561 944.

Here is a brief summary of the four geometric objects we have introduced:

- The semialgebraic set $M_n^k \subset \Delta_{2^n-1}$ of probability distributions represented by the restricted Boltzmann machine. We call M_n^k the *RBM model*.
- The Zariski closure V_n^k of the RBM model M_n^k . This is an algebraic variety in the complex projective space \mathbb{P}^{2^n-1} . We call V_n^k the *RBM variety*.
- The image $\mathcal{T}M_n^k$ of the tropicalized parameterization Φ . This is the subset of \mathbb{TP}^{2^n-1} consisting of all optimal score value vectors in the MAP inference problem for the RBM. We call $\mathcal{T}M_n^k$ the *tropical RBM model*.
- The tropicalization $\mathcal{T}V_n^k$ of the variety V_n^k . This is a tropical variety in the tropical projective torus \mathbb{TP}^{2^n-1} . We call $\mathcal{T}V_n^k$ the *tropical RBM variety*.

We have inclusions $M_n^k \subset V_n^k$ and $\mathcal{T}M_n^k \subset \mathcal{T}V_n^k$. The latter inclusion is the content of the second statement in [75, Theorem 2]. We shall see that both inclusions are strict even for $k = 1$. For example, M_3^1 is a proper subset of $V_3^1 \cap \Delta_7 = \Delta_7$ since points in this set must satisfy the inequality $\sigma_{12}\sigma_{13}\sigma_{23} \geq 0$ as indicated in Theorem 2.3.4 below. Likewise, $\mathcal{T}M_3^1$ is a proper subfan of $\mathbb{TP}^7 = \mathcal{T}V_3^1$. This subfan will be determined in our discussion of the secondary fan structure in Example 2.5.2.

The dimensions of our four geometric objects satisfy the following chain of equations and inequalities:

$$\dim(\mathcal{T}M_n^k) \leq \dim(\mathcal{T}V_n^k) = \dim(V_n^k) = \dim(M_n^k) \leq \min\{nk + n + k, 2^n - 1\}. \quad (2.6)$$

Here, the tropical objects $\mathcal{T}M_n^k$ and $\mathcal{T}V_n^k$ are polyhedral fans, and by their dimension we mean the dimension of any cone of maximal dimension in the fan. When speaking of the dimension of V_n^k we mean the Krull dimension of the projective variety, and for the model M_n^k we mean its dimension as a semialgebraic set.

The leftmost inequality in (2.6) holds because $\mathcal{T}M_n^k \subset \mathcal{T}V_n^k$. The left equality holds by the Bieri-Groves Theorem (cf. [25, Theorem 4.5]) which ensures that every irreducible variety has the same dimension as its tropicalization (see also Section 1.1.1).

Every polynomial function that vanishes on the image of the map p in (2.3) also vanishes on V_n^k . This means that the model M_n^k is *Zariski dense* in the variety V_n^k . From this we conclude the validity of the second equality in (2.6). Finally, the rightmost inequality in (2.6) is seen by counting parameters in the definition (2.1)–(2.2) of the RBM model M_n^k , and by bounding its dimension by the dimension of the ambient space Δ_{2^n-1} .

We conjecture that both of the inequalities in (2.6) are actually equalities:

Conjecture 2.2.2. *The tropical RBM model has the expected dimension, i.e. $\mathcal{T}M_n^k$ is a polyhedral fan of dimension $\min\{nk + n + k, 2^n - 1\}$ in $\mathbb{TP}^{2^n - 1}$.*

In light of the inequalities (2.6), Conjecture 2.2.2 implies Conjecture 2.1.1. In Section 2.4 we shall prove some special cases of these conjectures, including Theorem 2.1.2.

2.3 The first secant variety of the n -cube

We saw in Proposition 2.2.1 that the RBM for $k \geq 2$ can be expressed as the Hadamard power of the RBM for $k = 1$. Therefore, it is crucial to understand the model with one hidden node. In this section we fix $k = 1$ and we present an analysis of that case. In particular, we shall give a combinatorial description of the fan $\mathcal{T}M_n^1$ which shows that it has dimension $2n + 1$, as stated in Conjecture 2.2.2.

We begin with a reparameterization of our model that describes it as a secant variety. Let $\lambda, \delta_1, \dots, \delta_n, \epsilon_1, \dots, \epsilon_n$ be real parameters which range over the open interval $(0, 1)$, and consider the polynomial map $p : (0, 1)^{2n+1} \rightarrow \Delta_{2^n - 1}$ whose coordinates are given by

$$p(v) = \lambda \prod_{i=1}^n \delta_i^{1-v_i} (1 - \delta_i)^{v_i} + (1 - \lambda) \prod_{i=1}^n \epsilon_i^{1-v_i} (1 - \epsilon_i)^{v_i} \quad \text{for } v \in \{0, 1\}^n. \quad (2.7)$$

Proposition 2.3.1. *The image of p coincides with the RBM model M_n^1 .*

Proof. Recall the parameterization (2.3) of the RBM model M_n^1 from Section 2.2:

$$p(v) = \frac{1}{Z} \beta_1^{v_1} \beta_2^{v_2} \cdots \beta_n^{v_n} (1 + \gamma \omega_1^{v_1} \omega_2^{v_2} \cdots \omega_n^{v_n}) \quad \text{for } v \in \{0, 1\}^n. \quad (2.8)$$

We define a bijection between the parameter spaces $\mathbb{R}_{>0}^{2n+1}$ and $(0, 1)^{2n+1}$ as follows:

$$\beta_i = \frac{1 - \delta_i}{\delta_i} \quad \text{and} \quad \omega_i = \frac{\delta_i}{1 - \delta_i} \frac{1 - \epsilon_i}{\epsilon_i} \quad \text{for } i = 1, 2, \dots, n,$$

$$\gamma = Z(1 - \lambda)\epsilon_1\epsilon_2 \cdots \epsilon_n \quad \text{where} \quad Z = (\lambda\delta_1\delta_2 \cdots \delta_n)^{-1}.$$

This substitution is invertible and it transforms (2.8) into (2.7). \square

Proposition 2.3.1 shows that M_n^1 is the first mixture of the independence model for n binary random variables. In phylogenetics, it coincides with the *general Markov model on the star tree* with n leaves. A semi-algebraic characterization of that model follows as a special case from recent results of Zwiernik and Smith [100]. We shall present and discuss their characterization in Theorem 2.3.4 below.

First, however, we remark that the Zariski closure of a mixture of an independence model is a secant variety of the corresponding Segre variety. This fact is well-known (see e.g. [27, §4.1]) and is here easily seen from (2.7). We conclude:

Corollary 2.3.2. *The first RBM variety V_n^1 coincides with the first secant variety of the Segre embedding of the product of projective lines $(\mathbb{P}^1)^n$ into \mathbb{P}^{2^n-1} , and the first tropical RBM variety \mathcal{TV}_n^1 is the tropicalization of that secant variety.*

We next describe the equations defining the first secant variety V_n^1 . The coordinate functions $p(v)$ are the entries of an n -dimensional table of format $2 \times 2 \times \cdots \times 2$. For each set partition $\{1, 2, \dots, n\} = A \sqcup B$ we can write this table as an ordinary two-dimensional matrix of format $2^{|A|} \times 2^{|B|}$, with rows indexed by $\{0, 1\}^A$ and columns indexed by $\{0, 1\}^B$. These matrices are the *flattenings* of the $2 \times 2 \times \cdots \times 2$ -table. Pachter and Sturmfels [75, Conjecture 13] conjectured that the homogeneous prime ideal of the projective variety $V_n^1 \subset \mathbb{P}^{2^n-1}$ is generated by the 3×3 -minors of all the flattenings of the table $(p(v))_{v \in \{0,1\}^n}$. This conjecture has been verified computationally for $n \leq 5$, and recently proven by Raicu [81]. A more general form of this conjecture was stated in [41, Section 7]. The set-theoretic version of that general conjecture was proved by Landsberg and Manivel in [58, Theorem 5.1]. Their results imply:

Theorem 2.3.3 (Landsberg-Manivel, [59, 58]). *The projective variety $V_n^1 \subset \mathbb{P}^{2^n-1}$ is the common zero set of the 3×3 -minors of all the flattenings of the table $(p(v))_{v \in \{0,1\}^n}$.*

We now come to the inequalities that determine M_n^1 among the real points of V_n^1 . For any pair of indices $i, j \in \{1, 2, \dots, n\}$ we write σ_{ij} for the covariance of the two random variables X_i and X_j obtained by marginalizing the distribution, and we write $\Sigma = (\sigma_{ij})$ for the $n \times n$ -covariance matrix. We regard Σ as a polynomial map from the simplex Δ_{2^n-1} to the space $\mathbb{R}^{\binom{n+1}{2}}$ of symmetric $n \times n$ -matrices. The off-diagonal entries of the covariance matrix Σ are the 2×2 -minors obtained by marginalization from the table $(p(v))$. For example, for $n = 4$ the covariances are

$$\begin{aligned} \sigma_{12} &= \det \begin{pmatrix} p_{0000} + p_{0001} + p_{0010} + p_{0011} & p_{0100} + p_{0101} + p_{0110} + p_{0111} \\ p_{1000} + p_{1001} + p_{1010} + p_{1011} & p_{1100} + p_{1101} + p_{1110} + p_{1111} \end{pmatrix}, \\ \sigma_{13} &= \det \begin{pmatrix} p_{0000} + p_{0001} + p_{0100} + p_{0101} & p_{0010} + p_{0011} + p_{0110} + p_{0111} \\ p_{1000} + p_{1001} + p_{1100} + p_{1101} & p_{1010} + p_{1011} + p_{1110} + p_{1111} \end{pmatrix}, \quad \text{etc.} \end{aligned}$$

Zwiernik and Smith [100] gave a semi-algebraic characterization of the general Markov model on a trivalent phylogenetic tree in terms of covariances and moments. The statement of their characterization is somewhat complicated, so we only state a weaker necessary condition rather than the full characterization. Specifically, applying [100, Theorem 4.2] to the star tree on n leaves implies the following result.

Corollary 2.3.4. *If a probability distribution $p \in \Delta_{2^n-1}$ lies in the first RBM model M_n^1 then all its matrix flattenings (as in Theorem 2.3.3) have rank ≤ 2 and*

$$\sigma_{ij}\sigma_{ik}\sigma_{jk} \geq 0 \quad \text{for all distinct triples } i, j, k \in \{1, 2, \dots, n\}.$$

These inequalities follow easily from the parameterization (2.8), which yields

$$\sigma_{ij} = \lambda(1 - \lambda)(\delta_i - \epsilon_i)(\delta_j - \epsilon_j) \frac{\delta_i \delta_j}{\prod_{s=1}^n \delta_s} \frac{\epsilon_i \epsilon_j}{\prod_{s=1}^n \epsilon_s}.$$

This factorization also shows that the binomial relations $\sigma_{ij}\sigma_{kl} = \sigma_{il}\sigma_{jk}$ hold on M_n^1 . These same binomial relations are valid for the covariances in factor analysis [26, Theorem 16], thus further underlining the analogies between the Gaussian case and the binary case. Theorem 20 in [100] extends the covariance equations $\sigma_{ij}\sigma_{kl} = \sigma_{il}\sigma_{jk}$ to a collection of quadratic binomial equations in all tree-cumulants, which in turn can be expressed in terms of higher order correlations. For the star tree, these equations are equivalent on Δ_{2^n-1} to the rank ≤ 2 constraints. However, for general tree models, the binomial equations in the tree-cumulants are necessary conditions for distributions to lie in these models.

We now turn to the tropical versions of the RBM model for $k = 1$. The variety V_n^1 is cut out by the 3×3 -minors of all flattenings of the table $(p(v))_{v \in \{0,1\}^n}$. It is known that the 3×3 -minors of **one** fixed two-dimensional matrix form a tropical basis. Recall (e.g. from [6, Section 2]) that a *tropical basis* of a polynomial ideal is a generating set with the property that the intersection of the corresponding tropical hypersurfaces equals the tropical variety of the ideal. The tropical basis property of the 3×3 -minors is equivalent to [23, Theorem 6.5].

It is natural to ask whether this property continues to hold for the set of **all** 3×3 -determinants in Theorem 2.3.3. Since each flattening of our table corresponds to a non-trivial edge split of a tree on n taxa (i.e. a partition of the set of taxa into two sets each of cardinality ≥ 2), our question can be reformulated as follows:

Question 2.3.5. *Is the tropical RBM variety \mathcal{TV}_n^1 equal to the intersection of the tropical rank 2 varieties associated to non-trivial edge splits on a collection of trees on n taxa?*

The tropical rank two varieties associated to each of the edge splits have been studied by Markwig and Yu [66]. They endow this determinantal variety with a simplicial fan structure that has the virtue of being shellable. The cones of this simplicial fan correspond to weighted bicolored trees on 2^{n-1} taxa with no monochromatic cherries. The interior points in a cone can be viewed as a matrix encoding the distances between leaves with different colors in the associated weighted bicolored tree.

Question 2.3.5 is void for $n \leq 3$, so the first relevant case concerns $n = 4$ taxa. We were surprised to learn that the answer is negative already in this case:

Example 2.3.6. The prime ideal of the variety V_4^1 is generated by the sixteen 3×3 -minors of the three flattenings of the $2 \times 2 \times 2 \times 2$ -table p . As a statistical model, each one of the three flattenings corresponds to the graphical model associated to each one of the quartet trees (12|34), (13|24) and (14|23), as depicted in Figure 2.2.

Algebraically, each flattening corresponds to the variety cut out by the sixteen 3×3 -minors of a 4×4 -matrix of unknowns. These minors form a tropical basis. The tropical variety they




Figure 2.2: From left to right: quartet trees associated to the flattenings $(12|34)$, $(13|24)$ and $(14|23)$ for $n = 4$.

define is a pure fan of dimension 11 in \mathbb{TP}^{15} with a 6-dimensional lineality space. The simplicial fan structure on this variety given by [66] has the f -vector $(98, 1152, 4248, 6072, 2952)$. Combinatorially, this object is a shellable 4-dimensional simplicial complex which is the bouquet of 73 spheres. However, this determinantal variety admits a different fan structure, induced from the Gröbner fan as in Section 1.1.1, or from the fact that the sixteen 3×3 -minors form a tropical basis. Its f -vector is $(50, 360, 1128, 1680, 936)$.

The tropical variety \mathcal{TV}_4^1 is a pure fan of dimension 9 in \mathbb{TP}^{15} . Its lineality space has dimension 4, and the cones of various dimensions are tallied in its f -vector

$$f(\mathcal{TV}_4^1) = (382, 3436, 11236, 15640, 7680).$$

Question 2.3.5 asks whether the 9-dimensional tropical variety \mathcal{TV}_4^1 is the intersection of the three 11-dimensional tropical determinantal varieties associated with the three trees in Figure 2.2. The answer is “no”. Using the software `Gfan` [54], we computed the tropical prevariety cut out by the union of all forty-eight 3×3 -minors. The output is a *non-pure* polyhedral fan of dimension 10 with a 4-dimensional lineality space (the same one as of \mathcal{TV}_4^1), having f -vector $(298, 2732, 9440, 13992, 7304, 96)$. The tropical variety \mathcal{TV}_4^1 is a triangulation of a proper subfan, and each of the 96 10-dimensional maximal cones lies in the prevariety but not in the variety. An example of a such a vector in the relative interior of a maximal cone is

$$q = (59, 1, 80, 86, 102, 108, 107, 113, 109, 115, 100, 106, 78, 84, 21, 43),$$

where coordinates are indexed in lexicographic order $p_{0000}, p_{0001}, \dots, p_{1111}$. Given the weights q , the initial form of each 3×3 -minor of each flattening is a binomial, however, the initial form of the following polynomial in the ideal of V_4^1 is the underlined monomial:

$$\begin{aligned} & \underline{p_{0000}p_{0110}p_{1010}p_{1101}} - p_{0010}p_{0100}p_{1000}p_{1111} + p_{0010}p_{0100}p_{1001}p_{1110} \\ & - p_{0000}p_{0110}p_{1001}p_{1110} - p_{0001}p_{0110}p_{1010}p_{1100} + p_{0000}p_{0010}p_{1100}p_{1111} \\ & - p_{0000}p_{0010}p_{1101}p_{1110} + p_{0001}p_{0110}p_{1000}p_{1110}. \end{aligned}$$

Anders Jensen performed another computation, using `Gfan` and `SoPlex` [99], which verified that we get a tropical basis by augmenting the 3×3 -minors with the above quartic and its images under the symmetry group of the 4-cube. This is a non-trivial computation because the corresponding fan structure on \mathcal{TV}_4^1 has the f -vector

$$(37442, 321596, 843312, 880488, 321552).$$

Using the language of [23], we may conclude from our computational results that the notions of tropical rank and Kapranov rank disagree for $2 \times 2 \times 2 \times 2$ -tensors. \diamond

Last but not least, we examine the tropical model $\mathcal{T}M_n^1$. This is a proper subfan of the tropical variety $\mathcal{T}V_n^1$, namely, $\mathcal{T}M_n^1$ is the image of the tropical morphism $\Phi : \mathbb{R}^{2n+1} \rightarrow \mathbb{TP}^{2^n-1}$ which is the specialization of (2.5) for $k = 1$. Equivalently, Φ is the tropicalization of the map (2.8), and its coordinates are written explicitly as

$$q(v) = b^\top v + \min\{0, \omega v + c\}. \quad (2.9)$$

This convex function is the minimum of two linear functions. The $2n + 1$ parameters are given by a column vector $b \in \mathbb{R}^n$, a row vector $\omega \in \mathbb{R}^n$, and a scalar $c \in \mathbb{R}$. A different tropical map which has the same image as Φ can be derived from (2.7). As v ranges over $\{0, 1\}^n$, there are 2^n such convex functions, and these form the coordinates of the tropical morphism Φ . We note that Φ made its first explicit appearance in [75, Equation (10)], where it was discussed in the context of *ancestral reconstruction* in statistical phylogenetics. Subsequently, Develin [22] and Draisma [25, Section 7.2] introduced a tropical approach to secant varieties of toric varieties, and our model fits well into the context developed by these two authors.

Remark 2.3.7. The first tropical RBM model $\mathcal{T}M_n^1$ is the image of the tropical secant map for the Segre variety $(\mathbb{P}^1)^n$ in the sense of Develin [22] and Draisma [25]. The linear space for their constructions has basis $\{\sum_{\alpha \in \{0,1\}^n, \alpha_i=1} e_\alpha : i = 1, \dots, n\}$, and the underlying point configuration consists of the vertices of the n -cube.

In light of Example 2.3.6, it makes sense to say that the $2 \times \dots \times 2$ -tensors in the tropical variety $\mathcal{T}V_n^1$ are precisely those that have *Kapranov (tensor) rank* ≤ 2 . This would be consistent with the results and nomenclature in [22, 23]. A proper subset of the tensors of Kapranov rank ≤ 2 are those that have *Barvinok (tensor) rank* ≤ 2 . These are precisely the points in the first tropical RBM model $\mathcal{T}M_n^1$.

We close this section by showing that $\mathcal{T}M_n^1$ has the expected dimension:

Proposition 2.3.8. *The dimension of the tropical RBM model $\mathcal{T}M_n^1$ is $2n + 1$.*

Proof. Each region of linearity of the map Φ is defined by a partition C of $\{0, 1\}^n$ into two disjoint subsets C^- and C^+ , according to the condition $\omega v + c < 0$ or $\omega v + c > 0$. Thus, the corresponding region is an open convex polyhedral cone, possibly empty, in the parameter space \mathbb{R}^{2n+1} . It consists of all triples (b, ω, c) such that $\omega v + c < 0$ for $v \in C^-$ and $\omega v + c > 0$ for $v \in C^+$. Assuming $n \geq 3$, we can choose a partition C of $\{0, 1\}^n$ such that this cone is non-empty and both C^- and C^+ affinely span \mathbb{R}^n . The image of the cone under the map Φ spans a space isomorphic to the direct sum of the images of $b \mapsto (b^\top v : v \in C)$ and $(\omega, c) \mapsto (\omega v + c : v \in C^-)$. Hence this image has dimension $2n + 1$, as expected. \square

An illustration of the proof of Proposition 2.3.8 is given in Figure 2.3. The technique of partitioning the vertices of the cube will be essential in our dimension computations for general k in the next section. In Section 2.5 we return to the small models $\mathcal{T}M_n^1$ and take a closer look at their geometric and statistical properties.




Figure 2.3: Partitions of $\{0, 1\}^3$ that define non-empty cones on which Φ is linear. Here C^+ and C^- are indicated by black (\bullet) and white (\circ) vertices of the 3-cube. The slicing on the right represents a cone in the parameter space whose image under Φ is full-dimensional, while the one on the left does not.

2.4 The tropical model and its dimension

This section is concerned with Conjecture 2.2.2, which states that the tropical RBM model has the expected dimension. Namely, our aim is to show that

$$\dim(\mathcal{T}M_n^k) = kn + k + n \quad \text{for} \quad k \leq \frac{2^n - 1 - n}{n + 1}.$$

For $k = 1$ this is Proposition 2.3.8, and we now consider the general case $k \geq 2$. Our main tool towards this goal is the dimension formula in Theorem 2.4.2 below. As in the previous section, we study the regions of linearity of the tropical morphism Φ .

Let A denote the matrix of format $2^n \times n$ whose rows are the vectors in $\{0, 1\}^n$. A subset C of the vertices of the n -cube is a *slicing* if there exists a hyperplane that has the vertices in C on the negative side and the remaining vertices of the n -cube on the other side. In the notation in the proof of Proposition 2.3.8, the subset C was denoted by C^- . Two examples of slicings for $n = 3$ are shown in Figure 2.3.

For any slicing C of the n -cube, let A_C be the $2^n \times (n+1)$ -matrix whose rows v indexed by the vertices in C are $(1, v) \in \{0, 1\}^{n+1}$, and whose other rows are all identically zero. The following result extends the argument used for Proposition 2.3.8.

Lemma 2.4.1. *On each region of linearity, the tropical morphism Φ in (2.5) coincides with the linear map represented by a $2^n \times (nk + n + k)$ -matrix of the form*

$$\mathcal{A} = (A \mid A_{C_1} \mid A_{C_2} \mid \cdots \mid A_{C_k}),$$

for some slicings C_1, C_2, \dots, C_k of the n -cube.

Proof. The tropical map $\Phi : \mathbb{R}^{nk+n+k} \rightarrow \mathbb{TP}^{2^n-1}$ can be written as follows:

$$\Phi(W, b, c) = \left(\min_{h \in \{0,1\}^k} \{h^\top (Wv + c), 0\} + b^\top v \right)_{v \in \{0,1\}^n}.$$

Consider a parameter vector θ with coordinates

$$\theta := (b_1, b_2, \dots, b_n, c_1, \omega_{11}, \dots, \omega_{1n}, c_2, \omega_{21}, \dots, \omega_{2n}, \dots, c_k, \omega_{k1}, \dots, \omega_{kn}).$$

We associate to this vector the k hyperplanes $H_i(\theta) = \{v \in \mathbb{R}^n : \omega_{i1}v_1 + \dots + \omega_{in}v_n + c_i = 0\}$ for $i = 1, 2, \dots, k$. Let us assume that θ is chosen generically. Then, for each index i , we have $\{0, 1\}^n \cap H_i(\theta) = \emptyset$, and we obtain a slicing of the n -cube with $C_i(\theta) := \{v \in \{0, 1\}^n : \sum_{j=1}^n \omega_{ij}v_j + c_i < 0\}$. The generic parameter vector θ lies in a unique open region of linearity of the tropical morphism Φ . More precisely, this region corresponds to the cone of all θ' in \mathbb{R}^{nk+n+k} such that $C_i(\theta) = C_i(\theta')$ for $i = 1, 2, \dots, k$. By construction, the map $\Phi : \mathbb{R}^{nk+n+k} \rightarrow \mathbb{R}^{2^n}$ is linear on this cone. Following the definition of Φ we see that this linear map is left multiplication of the vector θ by a matrix whose rows are indexed by the observed states v and columns indexed by the coordinates of θ . This matrix is precisely the matrix \mathcal{A} above, where $C_i = C_i(\theta)$ for $i = 1, 2, \dots, k$. The result follows by continuity of the map Φ . \square

As an immediate consequence of Lemma 2.4.1 we obtain the following result:

Theorem 2.4.2. *The dimension of the tropical RBM model \mathcal{TM}_n^k equals the maximum rank of any matrix of size $2^n \times (nk + n + k)$ of the form*

$$\mathcal{A} = (A \mid A_{C_1} \mid A_{C_2} \mid \cdots \mid A_{C_k}),$$

where $\{C_1, C_2, \dots, C_k\}$ is any set of k slicings of the n -cube.

Theorem 2.4.2 furnishes a tool to attack Conjecture 2.2.2. What remains is the combinatorial problem of finding a suitable collection of slicings of the n -cube. In what follows we shall apply existing results from coding theory to this problem.

There are two quantities from the coding theory literature [4, 12, 14, 51] that are of interest to us. The first one is $A_2(n, 3)$, the size (number of codewords) of the largest binary code on n bits with each pair of codewords at least Hamming distance (number of bit flips) 3 apart. The second one is $K_2(n, 1)$, the size of the smallest *covering code* on n bits. In other words, $K_2(n, 1)$ is the least number of codewords such that every string of n bits lies within Hamming distance one of some codeword. We obtain:

Corollary 2.4.3. *The dimension of the tropical RBM model satisfies*

- $\dim \mathcal{T}M_n^k = nk + n + k$ for $k < A_2(n, 3)$,
- $\dim \mathcal{T}M_n^k = \min\{nk + n + k, 2^n - 1\}$ for $k = A_2(n, 3)$,
- $\dim \mathcal{T}M_n^k = 2^n - 1$ for $k \geq K_2(n, 1)$.

Proof. For the first statement, let $k \leq A_2(n, 3) - 1$ and fix a code in n bits of size $k + 1$ with minimum distance ≥ 3 . For each codeword let C_j denote its Hamming neighborhood, that is, the codeword together with all strings that are at Hamming distance 1. These $k + 1$ sets C_j are pairwise disjoint, and each of them corresponds to a slicing of the cube as in Theorem 2.4.2. The disjointness of the $k + 1$ neighborhoods means that $nk + n + k \leq 2^n - 1$. Elementary row and column operations can now be used to see that the corresponding $2^n \times (nk + n + k)$ matrix $\mathcal{A} = (A|A_{C_1}|\cdots|A_{C_k})$ has rank $nk + n + k$. This is because, after such operations, \mathcal{A} consists of a block of format $n \times n$ and k blocks of format $(n + 1) \times (n + 1)$ along the diagonal. The first block has rank n and the remaining k blocks have rank $n + 1$ each. The same reasoning is valid for $k = A_2(n, 3)$ except that it may now happen that $nk + k + n \geq 2^n$. In this case, the k blocks have total rank $k(n + 1)$ and together with the first $n \times n$ block they give a matrix of maximal rank $\min\{nk + n + k, 2^n - 1\}$.

For the third statement, we suppose C_1, \dots, C_k are slicings with subslicings $C'_i \subseteq C_i$ such that the C'_i are disjoint and no $n + 1$ of the vertices in a given C_i lie in a hyperplane. Then $\text{rank}(\mathcal{A}) \geq n + \sum_{i=1}^k |C'_i|$ by similar arguments. This is because we may construct the C'_i by pruning neighbors from codewords, and we are left with a lower-dimensional Hamming neighborhood which is a slicing. \square

The computation of $A_2(n, 3)$ and $K_2(n, 1)$, both in general and for specific values of n , has been an active area of research since the 1950s. In Table 2.1 we summarize some of the known results for specific values of n . This table is based on [12, 62]. For general values of n , the following bounds can be obtained.

Proposition 2.4.4. *For binary codes with $n \geq 3$, the Varshamov bound*

$$A_2(n, 3) \geq 2^{n - \lceil \log_2(n+1) \rceil}$$

holds, whereas for covering codes,

$$K_2(n, 1) \leq 2^{n - \lceil \log_2(n+1) \rceil}.$$

For $n = 2^\ell - 1$ with $\ell \geq 3$, we have the equality $A_2(n, 3) = K_2(n, 1) = 2^{2^\ell - \ell - 1}$.

Proof. A proof of the Varshamov bound on $A_2(n, 3)$ may be found in [51]. The last statement holds because $A_2(n, 3) = K_2(n, 1)$ for perfect Hamming codes: for every $\ell \geq 3$ there is a perfect $(2^\ell - 1, 2^\ell - \ell - 1, 3)$ Hamming code (i.e. a perfect Hamming code on $2^\ell - 1$ bits, of size $2^\ell - \ell - 1$, and with Hamming distance 3). For a proof of this result, see [14]. Additionally, we have $K_2(2^m - 1, 1) = 2^{2^m - m - 1}$ for $m \geq 3$; see [12].

The simple upper bound on $K_2(n, 1)$ can be obtained by using overlapping copies of the next smallest Hamming code. Suppose $n \neq 2^\ell - 1$ for any ℓ , i.e. n is strictly between two integers of the form $2^\ell - 1$ (*Hamming integer numbers*). Let \underline{n} be the largest Hamming integer smaller than n , with $\ell = \lfloor \log_2(n + 1) \rfloor$, so $\underline{n} = 2^\ell - 1$. The number of hidden nodes needed to cover the \underline{n} -cube is exactly $K_2(\underline{n}, 1) = 2^{2^\ell - \ell - 1}$. We may use the \underline{n} codes to cover each of the $2^{n - \underline{n}}$ faces of the n -cube with $2^{\underline{n}}$ vertices, although we will have overlaps. That is,

$$K_2(n, 1) \leq K_2(\underline{n}, 1) \cdot 2^{n - \underline{n}}. \quad (2.10)$$

Taking \log_2 in the inequality (2.10), we obtain

$$\log_2 K_2(n, 1) \leq \log_2(K_2(\underline{n}, 1)2^{n - \underline{n}}) = n - \lfloor \log_2(n + 1) \rfloor.$$

This implies $K_2(n, 1) \leq 2^{n - \lfloor \log_2(n + 1) \rfloor}$. □

Our method results in the following upper and lower bounds for arbitrary values of n . Note that the bound is tight if $n + 1$ is a power of 2. Otherwise there might be a multiplicative gap of up to 2 between the lower and upper bound. In addition to these general bounds, we have the specific results recorded in Table 2.1.

Corollary 2.4.5. *The coding theory argument leads to the following bounds:*

- If $k < 2^{n - \lceil \log_2(n + 1) \rceil}$, then $\dim \mathcal{T}M_n^k = nk + n + k$.
- If $k = 2^{n - \lceil \log_2(n + 1) \rceil}$, then $\dim \mathcal{T}M_n^k = \min\{nk + n + k, 2^n - 1\}$.
- If $k \geq 2^{n - \lceil \log_2(n + 1) \rceil}$, then $\dim \mathcal{T}M_n^k = 2^n - 1$.

Proof of Theorem 2.1.2. This is now easily completed by combining Corollary 2.4.5 with the inequalities in (2.6). □

We close this section with the remark that the use of Hamming codes is a standard tool in the study of dimensions of secant varieties. We learned this technique from Tony Geramita and his collaborators [10]. For a review of the relevant literature see Draisma's paper [25]. It is important to note that, in spite of the combinatorial similarities, the varieties we study here are different from and more complicated than higher secant varieties of Segre varieties. This may be because the varieties here involve both the secant construction and Hadamard products.

2.5 Polyhedral geometry of parametric inference

The tropical model $\mathcal{T}M_n^k$ is not just a convenient tool for estimating the dimension of the statistical model M_n^k . It is also of interest as the geometric object that organizes the space of inference functions which the model can compute. This statistical interpretation of tropical spaces was introduced in [75] and further developed in [32, 76]. We shall now discuss this perspective for the RBM model.

Given an RBM model with fixed parameters learned by some estimation procedure and an observed state v , we want to infer which value \hat{h} of the hidden data maximizes $\text{Prob}(h | v)$. The inferred string \hat{h} might be used in classification or as the input data for another RBM in a deep architecture. Such a vector of hidden states is called an *explanation* of the observation v . Each choice of parameters $\theta = (b, W, c)$ defines an *inference function* I_θ sending $v \mapsto \hat{h}$. The value $I_\theta(v)$ equals the hidden string $h \in \{0, 1\}^k$ that attains the minimum in the negative of the tropical polynomial

$$\max_{h \in \{0, 1\}^k} \{h^\top W v + c^\top h + b^\top v\} = -((-b)^\top v + \min_{h \in \{0, 1\}^k} \{h^\top (-W)v + (-c)^\top h\}). \quad (2.11)$$

In order for the inference function I_θ to be well-defined, it is necessary (and sufficient) that $-\theta = -(b, W, c)$ lies in an open cone of linearity of the tropical morphism Φ . In that case, the maximum in equation (2.11) is attained for a unique value of h . That h can be recovered from the expression of Φ as we vary the parameters in the fixed cone of linearity. Thus, up to sign, the inference functions are in one-to-one correspondence with the regions of linearity of the tropical morphism Φ .

The RBM model grew out of work on artificial neurons modeled as linear threshold functions [70, 84]. We pause our geometric discussion to offer remarks about these functions and the types of inference functions that our model can represent.

A *linear threshold function* is a function $\{0, 1\}^n \rightarrow \{0, 1\}$ defined by choosing a weight vector $\omega \in \mathbb{R}^n$ and a target weight $\pi \in \mathbb{R}$. For any point $v \in \{0, 1\}^n$ we compute the value ωv , we test if this quantity is at most π or not, and we assign value 0 or 1 to V depending on $\pi \geq \omega v$ or $\pi < \omega v$. The weights ω, π define a hyperplane in \mathbb{R}^n such that the vertices of the n -cube lie on the “true” or “false” side of the hyperplane. Using the linear threshold functions, we construct a k -valued function $\{0, 1\}^n \rightarrow \{0, 1\}^k$ where we replace the weight vector ω by a $k \times n$ matrix W and the target weight π by a vector $\pi \in \mathbb{R}^k$. More precisely, the function assigns a vertex of the k -cube where the i^{th} coordinate equals 0 if $(Wv)_i \leq \pi_i$ and 1 if not. Our discussion of slicings of the n -cube in Section 2.4 implies the following observation:

Proposition 2.5.1. *The inference functions for the restricted Boltzmann machine model M_n^k are precisely those Boolean functions $\{0, 1\}^n \rightarrow \{0, 1\}^k$ for which each of the k coordinate functions $\{0, 1\}^n \rightarrow \{0, 1\}$ is a linear threshold function.*

Most Boolean functions are not linear threshold functions, that is, are not inference functions for the model M_n^1 . For example, the parity function cannot be so represented. To

be precise, while the number of all Boolean functions is 2^{2^n} , it is known [73] that for $n \geq 8$ the number $\lambda(n)$ of linear threshold functions satisfies

$$2^{\binom{n}{2}+16} < \lambda(n) \leq 2^{n^2}.$$

The exact number $\lambda(n)$ of linear threshold functions has been computed for up to $n = 8$. The *On-Line Encyclopedia of Integer Sequences* [86, A000609] reveals

$$\lambda(1 \dots 8) = 4, 14, 104, 1882, 94572, 15028134, 8378070864, 17561539552946. \quad (2.12)$$

Combining k such functions for $k \geq 2$ yields $\lambda(n)^k = 2^{\Theta(kn^2)}$ possible inference functions for the RBM model M_n^k . This number grows exponentially in the number of model parameters. This is consistent with the result of Elizalde and Woods in [32] which states that the number of inference functions of a graphical model grows polynomially in the size of the graph when the number of parameters is fixed.

In typical implementations of RBMs using IEEE 754 doubles, the size in bits of the representation is $64(nk+n+k)$. Thus the number $2^{\Theta(kn^2)}$ of inference functions representable by a theoretical RBM M_n^k will eventually outstrip the number $2^{64(nk+n+k)}$ representable in a fixed-precision implementation; for example with $k = 100$ hidden nodes, this happens at $n \geq 132$. As a result, the size of the regions of linearity will shrink to single points in floating point representation. This is one possible contributor to the difficulties that have been encountered in scaling RBMs.

The tropical point of view allows us to organize the geometric information of the space of inference functions into the tropical model $\mathcal{T}M_n^k$, which can then be analyzed with the tools of tropical and polyhedral geometry. We now describe this geometry in the case $k = 1$. Geometrically, we can think of the linear threshold functions as corresponding to the vertices of the $(n + 1)$ -dimensional zonotope corresponding to the n -cube. This zonotope is the Minkowski sum in \mathbb{R}^{n+1} of the 2^n line segments $[(1, 0, \dots, 0), (1, v)]$ where v ranges over the set $\{0, 1\}^n$.

The quantity $\lambda(n)$ is the number of vertices of these zonotopes, and their facet numbers were computed by Aichholzer and Aurenhammer [1, Table 2]. They are

$$4, 12, 40, 280, 6508, 504868, 142686416, 172493511216, \dots \quad (2.13)$$

For example, the second entry in (2.12) and (2.13) refers to a 3-dimensional zonotope known as the *rhombic dodecahedron*, which has 12 facets and $\lambda(2) = 14$ vertices. Likewise, the third entry in (2.12) and (2.13) refers to a 4-dimensional zonotope with 40 facets and $\lambda(3) = 104$ vertices. The normal fan of that zonotope is an arrangement of eight hyperplanes, indexed by $\{0, 1\}^3$, which partitions \mathbb{R}^4 into 104 open convex polyhedral cones. That partition lifts to a partition of the parameter space \mathbb{R}^7 for M_3^1 whose cones are precisely the regions on which the tropical morphism Φ is linear. The image of that morphism is the first non-trivial tropical RBM model $\mathcal{T}M_3^1$. This model has the expected dimension 7 and it happens to be a pure fan.




Figure 2.4: Subdivisions of the 3-cube that represent vertices and facets of $\mathcal{T}M_3^1$. The first two from the left correspond to the special subdivisions D_{1357} and V_6 .

Example 2.5.2. The tropical RBM model $\mathcal{T}M_3^1$ is a 7-dimensional fan whose lineality space is 3-dimensional. It is a subfan of the secondary fan of the 3-cube [22, Corollary 2.2]. The secondary fan of the 3-cube can be represented as a 3-dimensional polyhedral sphere with f -vector $(22, 100, 152, 74)$. The 74 facets of that 3-sphere correspond to triangulations of the 3-cube. The tropical model $\mathcal{T}M_3^1$ consists of all regular subdivisions of the 3-cube with two regions covering all eight vertices. It sits inside the polyhedral 3-sphere as a *simplicial* subcomplex with f -vector $(14, 40, 36, 12)$. Its 12 facets (tetrahedra) correspond to a single triangulation type of the 3-cube as depicted in the right of Figure 2.4. The 14 vertices of $\mathcal{T}M_3^1$ come in two families: six vertices D_j corresponding to diagonal cuts, as in the left side of Figure 2.4, and eight vertices V_i representing corner cuts, as in the center of Figure 2.4. The edges come in three families: four edges V_iV_j corresponding to pairs of corner cuts at antipodal vertices of the cube, twenty-four edges V_iD_j , and twelve edges D_iD_j . Finally, of the four possible triangles, only two types are present: the ones with two vertices of different type. Thus, they are 12 triangles $V_iV_jD_k$ and 24 triangles $V_iD_jD_k$.

Figure 2.5 depicts the simplicial complex $\mathcal{T}M_3^1$ which is pure of dimension 3. The six vertices D_i and the twelve edges D_jD_k form the edge graph of an octahedron. The four nodes interior to the shaded triangles represent pairs of antipodal vertices V_i that are joined by an edge. Each of the shaded triangles represents three tetrahedra that are glued together along a common edge V_iV_j . Thus the twelve tetrahedra in $\mathcal{T}M_3^1$ come as four triangulated bipyramids. The four bipyramids are then glued into four of the triangles in the octahedron graph. Our analysis shows that the simplicial complex $\mathcal{T}M_3^1$ has reduced homology concentrated in degree 1 and it has rank 3. \diamond

The previous example is based on the fact that the image of the tropical map $\Phi : \mathbb{R}^{2n+1} \rightarrow \mathbb{R}^{2^n}$ is a subfan of the secondary fan of the n -cube. However, it is important to note that Φ is **not** a morphism of fans with respect to the natural fan structure on the parameter space \mathbb{R}^{2n+1} given by the slicings of the n -cube.




Figure 2.5: The tropical model $\mathcal{T}M_3^1$ is glued from four triangulated bipyramids. In this octahedron graph, each of the bipyramids is represented by a shaded triangle.

Example 2.5.3. Consider the case $n = 2$. Here M_2^1 equals \mathbb{R}^4 with its secondary fan structure coming from the two triangulations of the square. Modulo lineality, this fan is simply the standard fan structure $\{\mathbb{R}_{\leq 0}, \{0\}, \mathbb{R}_{\geq 0}\}$ on the real line. The fan structure on the parameter space \mathbb{R}^7 has 14 maximal cones. Modulo lineality, this is the normal fan of the rhombic dodecahedron, i.e. a partition of \mathbb{R}^3 into 14 open convex cones by an arrangement of four planes through the origin. Ten of these 14 open cones are mapped onto cones, namely, four are mapped onto $\mathbb{R}_{\leq 0}$, two are mapped onto $\{0\}$, and four onto $\mathbb{R}_{\geq 0}$. The remaining four cones are mapped onto \mathbb{R}^1 , so Φ does not respect the fan structures relative to these four cones, as Figure 2.6 reveals.

The situation is analogous for $n = 3$ but more complicated. The tropical map Φ is injective on precisely eight of the 104 maximal cones in the parameter space. These eight cones are the slicings shown on the right side of Figure 2.4. The map Φ is injective on such a cone, but the cone is divided into three subcones by the secondary fan structure on M_3^1 . The resulting $24 = 3 \cdot 8$ maximal cells in the parameter space are mapped in a 2-to-1 fashion onto the 12 tetrahedra in Figure 2.5. It would be worthwhile to study the combinatorics of the graph of Φ for $n \geq 3$. \diamond

n	$k \leq$	$k \geq$	n	$k \leq$
5	2^2	7		
6	2^3	12	35	$2^{23} \cdot 83$
7	2^4	2^4	37	$2^{26} \cdot 41$
8	$2^2 \cdot 5$	2^5	39	$2^{31} \cdot 5$
9	$2^3 \cdot 5$	62	47	$2^{38} \cdot 9$
10	$2^3 \cdot 9$	120	63	2^{57}
11	$2^4 \cdot 9$	192	70	$2^{43} \cdot 1657009$
12	2^8	380	71	$2^{63} \cdot 3$
13	2^9	736	75	$2^{63} \cdot 41$
14	2^{10}	1408	79	$2^{70} \cdot 5$
15	2^{11}	2^{11}	95	$2^{85} \cdot 9$
16	$2^5 \cdot 85$	2^{12}	127	2^{120}
17	$2^6 \cdot 83$	2^{13}	141	$2^{113} \cdot 1657009$
18	$2^8 \cdot 41$	2^{14}	143	$2^{134} \cdot 3$
19	$2^{12} \cdot 5$	31744	151	$2^{138} \cdot 41$
20	$2^{12} \cdot 9$	63488	159	$2^{149} \cdot 5$
21	$2^{13} \cdot 9$	122880	163	$2^{151} \cdot 19$
22	$2^{14} \cdot 9$	245760	191	$2^{180} \cdot 9$
23	$2^{15} \cdot 9$	393216	255	2^{247}
24	2^{19}	786432	270	$2^{202} \cdot 1021273028302258913$
25	2^{20}	1556480	283	$2^{254} \cdot 1657009$
26	2^{21}	3112960	287	$2^{277} \cdot 3$
27	2^{22}	6029312	300	$2^{220} \cdot 3348824985082075276195$
28	2^{23}	12058624	303	$2^{289} \cdot 41$
29	2^{24}	23068672	319	$2^{308} \cdot 5$
30	2^{25}	46137344	327	$2^{314} \cdot 19$
31	2^{26}	2^{26}	383	$2^{371} \cdot 9$
32	$2^{20} \cdot 85$	2^{27}	511	2^{502}
33	$2^{21} \cdot 85$	2^{28}	512	$2^{443} \cdot 1021273028302258913$

Table 2.1: Special cases where Conjecture 2.2.2 holds, based on [12, 62] and Corollary 2.4.3. Bold entries show improvements made by various researchers on the bounds provided by Corollary 2.4.5. For example, for $n = 19$, TM_n^k has the expected dimension if $k \leq 2^{12} \cdot 5 = 20480$ and dimension $2^n - 1$ if $k \geq 31744$, while Corollary 2.4.5 bounds are $2^{14} = 16384$ and $2^{15} = 32768$, respectively. The “ $k \leq$ ” columns list lower bounds on $A_2(n, 3)$ while the “ $k \geq$ ” column lists upper bounds on $K_2(n, 1)$.




Figure 2.6: The tropical map Φ parameterizing $\mathcal{T}M_2^1$ does not respect the fan structure of the secondary fan of the square.

Chapter 3

An implicitization challenge for the restricted Boltzmann machine

This chapter is joint work with Enrique Tobis and Josephine Yu. It is to appear in the *Journal of Symbolic Computation* - Special Issue MEGA'09 under the title “An implicitization challenge for binary factor analysis”, as [doi:10.1016/j.jsc.2010.06.011](https://doi.org/10.1016/j.jsc.2010.06.011). The present version incorporates some minor changes, largely for consistency with other chapters.

3.1 Introduction

In this chapter, we focus our attention on a special discrete graphical model, introduced in Chapter 2: the *restricted Boltzmann machine* $\mathcal{F}(4, 2)$ with two hidden nodes H_1, H_2 and four observed nodes X_1, X_2, X_3, X_4 (cf. Figure 3.1). This machine was introduced and studied in Chapter 2. As we discussed there, the main invariant of interest in these models is the expected dimension, and, furthermore, lower bounds on k (the number of hidden nodes) such that the probability distributions are a dense subset of the probability simplex Δ_{2^n-1} . By direct computation, it is easy to show that $\mathcal{F}(2, 2)$ and $\mathcal{F}(3, 2)$ are dense subsets of Δ_3 and Δ_7 respectively, so $\mathcal{F}(4, 2)$ is the first interesting example worth studying.




Figure 3.1: The model $\mathcal{F}(4, 2)$. Each node represents a binary random variable.

The set of all possible joint probability distributions (X_1, X_2, X_3, X_4) that arise in this way

forms a semialgebraic set V_4^2 in the probability simplex Δ_{15} . To simplify our construction, we consider the Zariski closure V_4^2 of the joint probability distributions in \mathbb{P}^{15} , following the notation of Chapter 2. As stated in Theorem 2.1.2, this variety is expected to have codimension one and be defined by a homogeneous polynomial in 16 variables.

Problem 3.1.1. (*An Implicitization Challenge, [27, Chapter VI, Problem 7.7]*) Find the degree and the defining polynomial of the model V_4^2 .

Our main results state that the variety V_4^2 is a hypersurface of degree 110 in \mathbb{P}^{15} (Theorem 3.4.2) and explicitly enumerate all vertices and facets of the polytope (Theorem 3.4.1). Our methods are based on tropical geometry. Since the polynomial is multihomogeneous, we get its *multidegree* from just one vertex. Interpolation techniques will allow us to compute the corresponding irreducible homogeneous polynomial in 16 variables, using the lattice points in the Newton polytope. However, this polytope will turn out to be too big for interpolation to be practically feasible.

The chapter is organized as follows. In Section 3.2 we describe the parametric form of our model and we express our variety as the Hadamard square of the first secant of the Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$. In Section 3.3 we present the tropical interpretation of our variety. By means of the nice interplay between the Hadamard squares and their tropicalization (Corollary 3.3.6), we compute this tropical variety as a collection of cones with multiplicities. We should remark that we do not obtain a fan structure, but, nonetheless, our characterization is sufficient to solve Problem 3.1.1. The key ingredient is the computation of multiplicities by the so called push-forward formula (Theorem 1.1.10) of Sturmfels and Tevelev [91], which we generalize to match our setting (Theorem 3.3.2). We finish Section 3.3 by describing the effective computation of the tropical variety and discussing some of the underlying combinatorics.

In Section 3.4 we compute the multidegree of our model with respect to a natural 5-dimensional grading, which comes from the tropical picture in Section 3.3. Once this question is answered, we shift gears and move to the study of the Newton polytope of our variety. We present two algorithms that compute vertices of this polytope by “shooting rays” (Algorithm 3.1) and “walking” from vertex to vertex in the Newton polytope (Algorithm 3.2). Using these methods, introduced in Section 1.4, and also taking advantage of the B_4 symmetry of the polynomial and the Newton polytope, we compute all 17 214 912 vertices our polytope (in 44 938 orbits under B_4), which shows the intrinsic difficulties of this “challenging” problem. Along the way, we also compute the tangent cones at each symmetry class of vertices and certify the facet normal directions by looking at the local behavior of the tropical variety around these vectors (after certifying they belong to the tropical variety). In particular, by computing dimensions of a certain linear space (Algorithm 3.3) we can check if the vector is a ray of the tropical variety. In this way, we certify all 246 facets of the polytope modulo symmetry. We believe these methods will pave the way to attack combinatorial questions about high dimensional polytopes with symmetry as the one analyzed in this chapter.

3.2 Geometry of the model

We start this section by describing the parametric representation of the model we wish to study. Recall that all our six random variables are binary, with four observed nodes and two hidden ones. Since the model comes from an undirected graph (see [27, 76]), we can parameterize it by a map $p: \mathbb{R}^{32} \rightarrow \mathbb{R}^{16}$, where

$$p_{ijkl} = \sum_{s=0}^1 \sum_{r=0}^1 a_{si} b_{sj} c_{sk} d_{sl} e_{ri} f_{rj} g_{rk} h_{rl} \quad \text{for all } (i, j, k, l) \in \{0, 1\}^4.$$

Notice that our coordinates are homogeneous of degree 1 in the subset of variables corresponding to each edge of the graph. Therefore, there is a natural interpretation of this model in projective space. On the other hand, by the distributive law we can write down each coordinate as a product of two points in the model corresponding to the 4-claw tree, which is the first secant variety of the Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$ (see Corollary 2.3.2). Namely,

$$p: (\mathbb{P}^1 \times \mathbb{P}^1)^8 \rightarrow \mathbb{P}^{15} \quad p_{ijkl} = \left(\sum_{s=0}^1 a_{si} b_{sj} c_{sk} d_{sl} \right) \left(\sum_{r=0}^1 e_{ri} f_{rj} g_{rk} h_{rl} \right) \quad \forall (i, j, k, l) \in \{0, 1\}^4.$$

The next proposition follows from the construction and Proposition 2.2.1:

Proposition 3.2.1. *The algebraic variety of the model $\mathcal{F}(4, 2)$ is $V_4^2 = (V_4^1)^{[2]}$, where V_4^1 is the first secant variety V_1^4 of the Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$.*

Notice that the binary nature of our random variables enables us to define a natural \mathbb{S}_2 -action by permuting the values 0 and 1 on each index in our 4-tuples. Combining this with the \mathbb{S}_4 -action on the 4-tuples of indices, we see that our model comes equipped with a natural $\mathbb{S}_4 \times (\mathbb{S}_2)^4$ -action. In other words, the 16 coordinates p_{ijkl} of \mathbb{P}^{15} , for $i, j, k, l \in \{0, 1\}$, are in natural bijection with the vertices of a 4-dimensional cube. Since V_4^2 is a hypersurface (as we know from Theorem 2.1.1 and will reprove in Section 3.3), its defining polynomial is invariant under the group B_4 of symmetries of the 4-cube, which has order 384. This group action will be *extremely* helpful for our computations in the next two sections.

We now describe the ideal associated to the secant variety $V_4^1 = \text{Sec}(\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1)$, discussed in Section 2.3. The Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$ has a monomial parameterization $p_{ijkl} = u_i \cdot v_j \cdot w_k \cdot x_l$ for $i, j, k, l \in \{0, 1\}$. Its defining prime ideal is generated by the 2×2 -minors of all three 4×4 -flattenings, together with some 2×2 -minors of the 2×8 -flattenings [34, Section 3]:

$$F_{(12|34)} := \begin{pmatrix} p_{0000} & p_{0001} & p_{0010} & p_{0011} \\ p_{0100} & p_{0101} & p_{0110} & p_{0111} \\ p_{1000} & p_{1001} & p_{1010} & p_{1011} \\ p_{1100} & p_{1101} & p_{1110} & p_{1111} \end{pmatrix}, \quad F_{(13|24)} := \begin{pmatrix} p_{0000} & p_{0001} & p_{0100} & p_{0101} \\ p_{0010} & p_{0011} & p_{0110} & p_{0111} \\ p_{1000} & p_{1001} & p_{1100} & p_{1101} \\ p_{1010} & p_{1011} & p_{1110} & p_{1111} \end{pmatrix},$$

$$F_{(14|23)} := \begin{pmatrix} p_{0000} & p_{0010} & p_{0100} & p_{0110} \\ p_{0001} & p_{0011} & p_{0101} & p_{0111} \\ p_{1000} & p_{1010} & p_{1100} & p_{1110} \\ p_{1001} & p_{1011} & p_{1101} & p_{1111} \end{pmatrix}.$$

From Theorem 2.3.3 we know that the secant variety $V_4^1 = \text{Sec}(\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1) \subset \mathbb{P}^{15}$ is the nine-dimensional *irreducible* subvariety consisting of all $2 \times 2 \times 2 \times 2$ -tensors of tensor rank at most 2. In turn, its defining ideal can be computed from the previous three 4×4 -flattening matrices: it is generated by all the 3×3 -minors of the three flattenings (Theorem 2.3.3). This information will be crucial when turning to the tropical geometry framework.

3.3 Tropicalizing the model

In this section we compute the tropicalization of the model V_4^2 by means of the nice interplay between Hadamard squares and tropicalization (Corollary 3.3.6). We refer the reader to Section 1.1 and references therein for more details about tropical varieties. Since part of our computations are done with `Gfan`, which follows the `max` convention for tropical geometry, we will need to change the sign of the output of any computation we perform with this software to match our `min` convention.

As we discussed in the previous section (Proposition 3.2.1) our variety V_4^2 is expressed as the Hadamard square of the well-known variety V_4^1 . This Hadamard square has a dense set which can be parameterized in terms of the *Hadamard monomial map* (the coordinatewise product of two points). The integer matrix of exponents corresponding to this map is $(I_{16} \mid I_{16}) \subset \mathbb{Z}^{16 \times 32}$. Therefore, we can compute the tropicalization of V_4^2 as a set from Theorem 1.1.10.

We now treat the computation of multiplicities on $\mathcal{T}V_4^2$. At first sight, the hypothesis of the second part of that theorem is not satisfied by our variety because the restrict of the map α to the cartesian product $V_4^1 \times V_4^1$ is not generically finite. However it is very close to having the required finiteness behavior. Namely, after taking the quotient $(V_4^1)'$ of V_4^1 by a maximal torus action, and a choice of a suitable monomial map $\bar{\alpha}$, this new monomial map becomes generically finite on $(V_4^1)' \times (V_4^1)'$ and we can apply Theorem 1.1.10 to compute the tropical hypersurface $\mathcal{T}V_4^2$. We now explain this reduction process.

Let $X \subset \mathbb{T}^n$ a subvariety, $\alpha: \mathbb{T}^n \rightarrow \mathbb{T}^d$ a monomial map, and let $Y = \alpha(X)$. Consider the lineality space $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda \subset \mathcal{T}X$, and let $\Lambda' = A(\Lambda)$. We identify $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda$ with a \mathbb{Z} -basis of the primitive lattice Λ . Notice that Λ' need not be a primitive lattice in \mathbb{Z}^d in general. Call $(\Lambda')^{\text{sat}}$ its saturation in \mathbb{Z}^d , that is $(\Lambda')^{\text{sat}} = (\mathbb{R} \otimes_{\mathbb{Z}} \Lambda') \cap \mathbb{Z}^d$. We know by construction and Theorem 1.1.10 that $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda'$ is contained in the lineality space of $\mathcal{T}Y$. Therefore, we can consider the linear map between these tropical varieties after moding out by $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda$ and $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda'$ respectively. As we mentioned in Section 1.1 and in (1.1), the lineality space of each tropical variety determines the maximal torus action. For example, if $l = \text{rk } \Lambda$, then \mathbb{T}^s acts on X by $t \cdot (x_1, \dots, x_n) := (t^{a_1} x_1, \dots, t^{a_n} x_n)$ where \underline{a} lies in Λ .

The linear map A sends Λ onto Λ' , which is contained in the lineality space of $\mathcal{T}(\alpha(X))$. In addition, the monomial map α is compatible with the torus actions on X and $\alpha(X)$. In particular, the equality $\alpha(\Lambda \otimes_{\mathbb{Z}} \mathbb{C}^*) = \Lambda' \otimes_{\mathbb{Z}} \mathbb{C}^*$ induces an action on Y by a subtorus (the one corresponding to the primitive lattice $(\Lambda')^{\text{sat}}$). Thus, we can take the quotient of X and Y by the corresponding actions of tori H and H' . We obtain the commutative diagram:

$$\begin{array}{ccc} X & \xrightarrow{\alpha} & Y \\ \pi \downarrow & & \downarrow \pi \\ X' = X/H & \xrightarrow{\bar{\alpha}} & Y/H' = Y'. \end{array} \quad (3.1)$$

Here, $H = \Lambda \otimes_{\mathbb{Z}} \mathbb{C}^* \cong \mathbb{T}^{\dim \Lambda}$ and $H' = \Lambda' \otimes_{\mathbb{Z}} \mathbb{C}^* \cong \mathbb{T}^{\dim \Lambda'}$. Since Λ is a primitive sublattice of \mathbb{Z}^n , it admits a primitive complement in \mathbb{Z}^n . Fix one of them and call it Λ^\perp . Note that this complement need not be the usual orthogonal complement in \mathbb{Z}^n .

Assume for simplicity that Λ' is a *primitive* sublattice of \mathbb{Z}^d . Therefore, we can identify $\bar{\alpha}$ with the monomial map corresponding to the linear map:

$$A': (\mathbb{R} \otimes \Lambda)^\perp \rightarrow (\mathbb{R} \otimes \Lambda')^\perp, \quad (3.2)$$

which follows by the identifications $(\mathbb{R} \otimes \mathbb{Z}^n)/(\mathbb{R} \otimes \Lambda) \simeq \mathbb{R} \otimes \Lambda^\perp = (\mathbb{R} \otimes \Lambda)^\perp$ and $(\mathbb{R} \otimes \mathbb{Z}^d)/(\mathbb{R} \otimes \Lambda') \simeq \mathbb{R} \otimes \Lambda'^\perp = (\mathbb{R} \otimes \Lambda')^\perp$.

Since Λ is primitive, $(\mathbb{R} \otimes \Lambda)^\perp \cap \mathbb{Z}^n = \Lambda^\perp$, and likewise for Λ'^\perp .

To simplify notation, call $L := \mathbb{R} \otimes \Lambda$ and $L' := \mathbb{R} \otimes \Lambda'$. From the construction it is easy to see that $\mathcal{T}X' = \mathcal{T}X/L$ and $\mathcal{T}Y' = \mathcal{T}Y/L'$ as sets. But in fact, they agree as weighted balanced polyhedral fans. More precisely,

Lemma 3.3.1. *Let $X \subset \mathbb{T}^r$ and let L be a subspace of the lineality space of the tropical variety $\mathcal{T}X$ generated by integer vectors. Then $\mathcal{T}X/L$ is a balanced weighted polyhedral fan, where the multiplicities of regular points w' are defined as $m_{w'} = m_w$ for any w in the fiber of w' under the projection map. With these weights, $\mathcal{T}X/L$ coincides with the tropical variety $\mathcal{T}X'$, where X' is the quotient of X by the torus $(L \cap \mathbb{Z}^r) \otimes_{\mathbb{Z}} \mathbb{C}^* \cong \mathbb{T}^{\dim L}$, which is a subtorus of the maximal torus acting on X .*

Proof. By definition, we know that $\text{in}_{w+L}(I) = \text{in}_w(I)$ for any $w \in \mathbb{R}^r$. Let $l := \dim L$. Call $\Lambda := L \cap \mathbb{Z}^n$ the underlying lattice of L . Since Λ is a primitive lattice, we can extend any \mathbb{Z} -basis of Λ to a \mathbb{Z} -basis of \mathbb{Z}^n . Thus, after a linear change of coordinates (i.e. a monomial change of coordinates given by this new \mathbb{Z} -basis of \mathbb{Z}^n) we can assume $\Lambda = \mathbb{Z}\langle e_1, \dots, e_l \rangle$. And in this case, we can pick the direct summand Λ^\perp of Λ to be $\mathbb{Z}\langle e_{l+1}, \dots, e_r \rangle$. In particular, the projection map $\pi: X \rightarrow X' = X/H$ corresponds to the monomial map $\alpha: \mathbb{T}^n \rightarrow \mathbb{T}^{n-l}$ determined by the integer matrix $A = (0 \mid I_{n-l}) \in \mathbb{Z}^{(n-l) \times n}$, whose rows are a \mathbb{Z} -basis of Λ^\perp .

By construction, $I = I(X) \subset \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ is homogeneous with respect to the grading $\deg(x_i) = e_i$ for $i \leq l$ and $\deg(x_j) = \underline{0}$ for $j > l$. Since any homogeneous Laurent polynomial

is of the form $f = \underline{x}^\alpha g(x_{l+1}, \dots, x_n)$, we see that I is generated by Laurent polynomials in the variables $\{x_{l+1}, \dots, x_n\}$. Call g_1, \dots, g_s these generators. Therefore $I' = I(X') = \langle g_1(x_{l+1}, \dots, x_n), \dots, g_s(x_{l+1}, \dots, x_n) \rangle \subset \mathbb{C}[x_{l+1}^{\pm 1}, \dots, x_n^{\pm 1}]$ and $I = I'\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$.

From Theorem 1.1.10 we know that $\mathcal{T}X' = A(\mathcal{T}X) = \mathcal{T}X/L$ as sets. Moreover, since the subspace L lies in all cones of $\mathcal{T}X$, then the set $\mathcal{T}X'$ has a natural fan structure inherited from the one of $\mathcal{T}X$. By definition, if w' is a regular point in $\mathcal{T}X'$ then any lifting point in $w + L$ would be a regular point in $\mathcal{T}X$. Moreover, $\text{in}_w(I) = \text{in}_{w'}(I')\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. In particular, a primary decomposition $\text{in}_{w'}(I')$ determines a primary decomposition of $\text{in}_w(I)$ by extending each ideal to the whole Laurent polynomial ring in n variables. Therefore, to show $m_{w'} = m_w$ it suffices to show that the multiplicity of any minimal prime $P \subset \mathbb{C}[x_{l+1}^{\pm 1}, \dots, x_n^{\pm 1}]$ of $\text{in}_{w'}(I')$ equals the multiplicity of $P \subset \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ in $\text{in}_w(I)$. This claim follows from the definition of multiplicity. More precisely:

$$\begin{aligned} m(P, \text{in}_{w'}(I')) &= \dim \frac{{}_{PS}S_P}{S_P \text{in}_{w'}(I')} = \dim_{({}_{PS}S_P)[x_1^{\pm 1}, \dots, x_l^{\pm 1}]} \frac{S_P[x_1^{\pm 1}, \dots, x_l^{\pm 1}]}{S_P[x_1^{\pm 1}, \dots, x_l^{\pm 1}] \text{in}_{w'}(I')} \\ &= \dim \frac{{}_{PS[x_1^{\pm 1}, \dots, x_l^{\pm 1}]}S[x_1^{\pm 1}, \dots, x_l^{\pm 1}]_P}{S[x_1^{\pm 1}, \dots, x_l^{\pm 1}]_P \text{in}_w(I)} = m(P, \text{in}_w(I)), \end{aligned}$$

where $S = \mathbb{C}[x_{l+1}^{\pm 1}, \dots, x_n^{\pm 1}]$. □

Using the previous construction, we extend Theorem 1.1.10 to the case of monomial maps that are generically finite after taking quotients by appropriate tori. This extension fits perfectly into our setting.

Theorem 3.3.2. *Let $\alpha: \mathbb{T}^n \rightarrow \mathbb{T}^d$ be a monomial map with associated integer matrix A and let $X \subset \mathbb{T}^n$ be a closed subvariety. Then,*

$$\mathcal{T}(\alpha(X)) = A(\mathcal{T}X).$$

Suppose X has a torus action given by a rank l lattice $\Lambda \subset \mathbb{Z}^n$. Let X' be the quotient by this torus action. Let $\bar{\alpha}: X' \rightarrow \mathbb{T}^d/\alpha(\Lambda \otimes_{\mathbb{Z}} \mathbb{C}^)$ be the induced monomial map, with associated integer matrix A' as in (3.2).*

Suppose $\Lambda' = A(\Lambda)$ is a primitive sublattice of \mathbb{Z}^d and that $\bar{\alpha}$ induces a generically finite morphism of degree δ on X' . Then, the multiplicity of a regular point w in $\mathcal{T}(\alpha(X))$ can be computed as:

$$m_w = \frac{1}{\delta} \cdot \sum_{\substack{\pi(v) \\ A'v=w}} m_v \cdot \text{index}(\mathbb{L}_w \cap \mathbb{Z}^d, A(\mathbb{L}_v \cap \mathbb{Z}^n)), \quad (3.3)$$

where the sum is over any set of representatives of points $\{v' = \pi(v) \in \mathcal{T}X' \mid A'v' = w'\}$ given $w' = \pi(w) \in \mathbb{R}^d/(\mathbb{R} \otimes_{\mathbb{Z}} \Lambda') = \mathbb{R} \otimes_{\mathbb{Z}} \Lambda'^{\perp}$. We also assume that the number of such v' is finite, all of them are regular in $\mathcal{T}X'$ and $\mathbb{L}_v, \mathbb{L}_w$ are linear spans of neighborhoods of $v \in \mathcal{T}X$ and $w \in A(\mathcal{T}X)$ respectively.

Remark 3.3.3. In case Λ' is not a primitive lattice, the formula for m_w will involve an extra factor, namely, the index of Λ' with respect to its saturation Λ'^{sat} in \mathbb{Z}^d . In this case, Λ'^{\perp} will correspond to any complement of the primitive lattice Λ'^{sat} inside \mathbb{Z}^d .

Proof of Theorem 3.3.2. The equality as sets follows from Theorem 1.1.10. To prove the formula for multiplicities, we first note that the sum in (3.3) is finite. This follows because $\bar{\alpha}$ induces a generically finite morphism if and only if $\ker A' \cap \mathcal{T}X' = \{0\}$, and this holds if and only if $A(\Lambda^\perp) \cap \Lambda' = \{0\}$.

From the diagram (3.1) and the surjectivity of α and $\bar{\alpha}$, we know that the multiplicity formula (1.4) holds for $\mathcal{T}Y'$ and the morphism $\bar{\alpha}$. Pick w' a regular point of $\mathcal{T}X'$ and pick any point w in the fiber $\pi^{-1}(w') = w + (\mathbb{R} \otimes \Lambda')$. By definition, w is a regular point of $\mathcal{T}X$ and we have $m_w = m_{w'}$ by Lemma 3.3.1. We assume all v' in the fiber of A' at w' are regular in $\mathcal{T}X'$ and $\mathbb{L}_{\pi(v)}, \mathbb{L}_{\pi(w)}$ are linear spans of neighborhoods of $\pi(v) \in \mathcal{T}X'$ and $\pi(w) \in A'(\mathcal{T}X')$ respectively.

By construction, the index set in the formula (1.4) for $m_{w'}$ agrees with the index set in formula (3.3) for m_w . Therefore, our goal would be to show that each summand indexed by $\pi(v)$ in the formula for $m_{w'}$ equals its corresponding summand in formula for m_w . We know that $m_v = m_{\pi(v)}$ by Lemma 3.3.1. Hence, we only need to prove that the lattice indices on each summand are the same, i.e.

$$\text{index}(\mathbb{L}_w \cap \mathbb{Z}^d, A(\mathbb{L}_v \cap \mathbb{Z}^n)) = \text{index}(\mathbb{L}_{\pi(w)} \cap (\Lambda'^{\perp}), A'(\mathbb{L}_{\pi(v)} \cap \Lambda'^{\perp})). \quad (3.4)$$

Note that by construction, $\Lambda' \subset \mathbb{L}_w \cap \mathbb{Z}^d$, $\Lambda \subset \mathbb{L}_v$, and likewise $A(\Lambda) = \Lambda' \subset A(\mathbb{L}_v \cap \mathbb{Z}^n)$. Hence, we can consider the quotient of $\mathbb{L}_w \cap \mathbb{Z}^d$ and $A(\mathbb{L}_v \cap \mathbb{Z}^n)$ by Λ' . We obtain

$$\frac{\mathbb{L}_w \cap \mathbb{Z}^d}{A(\mathbb{L}_v \cap \mathbb{Z}^n)} \cong \frac{(\mathbb{L}_w \cap \mathbb{Z}^d)/\Lambda'}{A(\mathbb{L}_v \cap \mathbb{Z}^n)/\Lambda'}.$$

The equality in (3.4) follows from this isomorphism and the identifications $(\mathbb{L}_w \cap \mathbb{Z}^d)/\Lambda' = \mathbb{L}_{\pi(w)} \cap (\Lambda'^{\perp})$ and $A(\mathbb{L}_v \cap \mathbb{Z}^n)/\Lambda' = A'(\mathbb{L}_{\pi(v)} \cap \Lambda'^{\perp})$, via projecting to the complemented Λ'^{\perp} of the primitive lattice Λ' . \square

Theorem 3.3.4. *Given $X, Y \subset \mathbb{T}^n$ two irreducible varieties, consider the associated variety $X \times Y \subset \mathbb{T}^{2n}$. Then*

$$\mathcal{T}(X \times Y) = \mathcal{T}X \times \mathcal{T}Y$$

as weighted polyhedral complexes, with $m_{\sigma \times \tau} = m_{\sigma} m_{\tau}$ for maximal cones $\sigma \subset \mathcal{T}X, \tau \subset \mathcal{T}Y$, and $\sigma \times \tau \subset \mathcal{T}(X \times Y)$.

Proof. The equality of the sets $\mathcal{T}(X \times Y)$ and $\mathcal{T}X \times \mathcal{T}Y$ as polyhedral fans is a direct consequence of the equality $\text{in}_{(u,v)}(I + J) = \text{in}_u(I) + \text{in}_v(J)$, which follows by Buchberger's criterion [31, Theorem 15.8] and the fact that the generators of I and J involve disjoint sets of variables.

Let $u \in \mathcal{T}X$, $v \in \mathcal{T}Y$ be regular points. From the equality as polyhedral fans it is easy to check that (u, v) is a regular point in $\mathcal{T}(X \times Y)$. Our goal is to prove the multiplicity formula.

Given two primary decompositions $\text{in}_u(I) = \bigcap_i M_i \subset \mathbb{C}[\underline{x}]$, $\text{in}_v(J) = \bigcap_j N_j \subset \mathbb{C}[\underline{y}]$, we claim that $\text{in}_{(u,v)}(I + J) = \bigcap_{i,j} (M_i + N_j) \subset \mathbb{C}[\underline{x}, \underline{y}]$ is also a primary decomposition. The equality as sets follows immediately, so we only need to show that $M_i + N_j \subset \mathbb{C}[\underline{x}, \underline{y}]$ is a primary ideal. Let $P_i \subset \mathbb{C}[\underline{x}]$ and $Q_j \subset \mathbb{C}[\underline{y}]$ be associated prime ideals to M_i and N_j respectively. Since \mathbb{C} is algebraically closed, and M_i and N_j involved disjoint sets of variables, it is immediate to check that $P_i + Q_j \subset \mathbb{C}[\underline{x}, \underline{y}]$ is a prime ideal. Namely, the quotient ring $\mathbb{C}[\underline{x}, \underline{y}]/(P_i + Q_j)$ equals $(\mathbb{C}[\underline{x}]/P_i)[\underline{y}] \otimes_{\mathbb{C}} (\mathbb{C}[\underline{y}]/Q_j)[\underline{x}]$, a tensor product of two domains over \mathbb{C} , hence also a domain [60, Chapter XVI, Section 6].

Moreover, since both M_i and N_j involve disjoint sets of variables, we have

$$\text{Ann}(M_i + N_j) = \text{Ann } M_i \otimes_{\mathbb{C}} \mathbb{C}[\underline{y}] + \mathbb{C}[\underline{x}] \otimes_{\mathbb{C}} \text{Ann } N_j.$$

From this and the fact that $P_i^{s_i} \subset \text{Ann } M_i \subset P_i$ and $Q_j^{t_j} \subset \text{Ann } N_j \subset Q_j$ for suitable numbers $s_i, t_j \in \mathbb{N}$, we conclude $(P_i + Q_j)^{s_i+t_j} \subset \text{Ann}(M_i + N_j) \subset P_i + Q_j$ thus proving by definition that $M_i + N_j$ is a $(P_i + Q_j)$ -primary ideal.

With similar arguments we conclude that all minimal primes of $\text{in}_{(u,v)}(I + J)$ are sums of minimal primes of $\text{in}_u(I)$ and $\text{in}_v(J)$. This follows because, given $P, P' \subset \mathbb{C}[\underline{x}]$ and $Q, Q' \subset \mathbb{C}[\underline{y}]$ prime ideals, it is straightforward to check that $P + Q \subset P' + Q'$ if and only if $P \subset P'$ and $Q \subset Q'$.

Let σ, τ be maximal cones on $\mathcal{T}X$ and $\mathcal{T}Y$, and let u, v be regular points in σ and τ , respectively. By definition of multiplicity of a maximal cone, we have

$$m_\sigma = \sum_{\substack{P \in \text{Ass}(\text{in}_u(I)) \\ P \text{ minimal}}} m(P, \mathbb{C}[\underline{x}]/\text{in}_u(I)) = \sum_{\substack{P \in \text{Ass}(\text{in}_u(I)) \\ P \text{ minimal}}} \dim_{(\mathbb{C}[\underline{x}]/P)_P} (\mathbb{C}[\underline{x}]/\text{in}_u(I))_P ;$$

$$m_\tau = \sum_{\substack{Q \in \text{Ass}(\text{in}_v(J)) \\ Q \text{ minimal}}} \dim_{(\mathbb{C}[\underline{y}]/Q)_Q} (\mathbb{C}[\underline{y}]/\text{in}_v(J))_Q ; \quad m_{\sigma \times \tau} = \sum_{\substack{P \in \text{Ass}(\text{in}_u(I)) \\ Q \in \text{Ass}(\text{in}_v(J)) \\ P, Q \text{ minimal}}} \dim_{(\frac{\mathbb{C}[\underline{x}, \underline{y}]}{P+Q})_{P+Q}} \left(\frac{\mathbb{C}[\underline{x}, \underline{y}]}{\text{in}_u(I) + \text{in}_v(J)} \right)_{P+Q}.$$

The statement $m_{\sigma \times \tau} = m_\sigma m_\tau$ follows from the distributive law and Lemma 3.3.5. \square

Lemma 3.3.5. *Let $I \subset \mathbb{C}[\underline{x}]$, $J \subset \mathbb{C}[\underline{y}]$ be ideals and let $P \subset \mathbb{C}[\underline{x}]$, $Q \subset \mathbb{C}[\underline{y}]$ be minimal primes containing I and J , respectively. Then*

$$\dim_{(\mathbb{C}[\underline{x}, \underline{y}]/(P+Q))_{P+Q}} \left(\frac{\mathbb{C}[\underline{x}, \underline{y}]}{I + J} \right)_{P+Q} = \dim_{(\mathbb{C}[\underline{x}]/P)_P} (\mathbb{C}[\underline{x}]/I)_P \cdot \dim_{(\mathbb{C}[\underline{y}]/Q)_Q} (\mathbb{C}[\underline{y}]/J)_Q.$$

Proof. Consider the residue fields $F = (\mathbb{C}[\underline{x}]/P)_P$, $G = (\mathbb{C}[\underline{y}]/Q)_Q$, and $L = (\mathbb{C}[\underline{x}, \underline{y}]/(P + Q))_{P+Q}$. Note that $F \otimes_{\mathbb{C}} G \hookrightarrow L$ via the natural inclusion given by the multiplication

map, since \mathbb{C} is algebraically closed. Likewise, one can easily show that $\mathbb{C}[\underline{x}]/I \otimes_{\mathbb{C}} \mathbb{C}[\underline{y}]/J \cong \mathbb{C}[\underline{x}, \underline{y}]/(I+J)$, via the multiplication map. We wish to find a similar result for the localization of these quotients at the corresponding minimal primes.

For simplicity, call $M = (\mathbb{C}[\underline{x}]/I)_P \cong F^s$ and $N = (\mathbb{C}[\underline{y}]/J)_Q \cong G^r$ the corresponding finite dimensional vector spaces. Our goal is to prove that $M \otimes_{\mathbb{C}} N$ is a free L -vector space of rank sr . From the canonical isomorphisms $\mathbb{C}[\underline{x}]_P \otimes_{\mathbb{C}[\underline{x}]} \mathbb{C}[\underline{x}]/I \cong (\mathbb{C}[\underline{x}]/I)_P$, $\mathbb{C}[\underline{y}]_Q \otimes_{\mathbb{C}[\underline{y}]} \mathbb{C}[\underline{y}]/J \cong (\mathbb{C}[\underline{y}]/J)_Q$, we see that $M \otimes_{\mathbb{C}} N = (\mathbb{C}[\underline{x}]/I)_P \otimes_{\mathbb{C}} (\mathbb{C}[\underline{y}]/J)_Q \cong \mathbb{C}[\underline{x}, \underline{y}]/(I+J)[S^{-1}]$, where S is the multiplicatively closed set consisting of products of polynomials, each of which is pure in each set of variables, and which do not lie inside the prime ideals P or Q . More precisely:

$$S := (\mathbb{C}[\underline{x}] \setminus P)(\mathbb{C}[\underline{y}] \setminus Q).$$

Similarly, $F \otimes_{\mathbb{C}} G \cong \mathbb{C}[\underline{x}, \underline{y}]/(P+Q)[S^{-1}]$.

On the other hand, notice that $M \otimes_{\mathbb{C}} N$ comes with a natural $F \otimes_{\mathbb{C}} G$ -module structure via ‘‘coordinatewise action.’’ Hence,

$$(\mathbb{C}[\underline{x}, \underline{y}]/(I+J))_{(P+Q)} \cong L \otimes_{(F \otimes_{\mathbb{C}} G)} (M \otimes_{\mathbb{C}} N).$$

From the last isomorphism we see that to prove our lemma it suffices to show that $M \otimes_{\mathbb{C}} N$ is a free $F \otimes_{\mathbb{C}} G$ -module of rank sr . The original statement will follow after tensoring with L .

Let $\{f_i\}_i, \{g_j\}_j$ be bases of M and N , respectively. We claim that $\{f_i \otimes g_j\}_{i,j}$ is a basis of $M \otimes_{\mathbb{C}} N$ as an $F \otimes_{\mathbb{C}} G$ -module. It suffices to check the linear independence of the proposed basis. We proceed in an elementary way, by successively using the linear independence of the bases of the free modules M, N, F and G . Suppose $\sum_{i,j} a_{ij} f_i \otimes g_j = 0 \in M \otimes_{\mathbb{C}} N$, with $a_{ij} \in F \otimes_{\mathbb{C}} G$. Write $a_{ij} = \sum_{k,l} a_{ijkl} u_k \otimes v_l$ where $a_{ijkl} \in \mathbb{C}$ and $\{u_k\}_k, \{v_l\}_l$ are bases of the field extensions $F|\mathbb{C}$ and $G|\mathbb{C}$, respectively. Thus,

$$0 = \sum_{i,j} a_{ij} f_i \otimes g_j = \sum_{j,l} \left(\sum_{i,k} a_{ijkl} u_k f_i \right) \otimes_{\mathbb{C}} (v_l g_j). \quad (3.5)$$

To prove $a_{ij} = 0$ it suffices to show $a_{ijkl} = 0$ for all i, j, k, l . By a well-know result on tensor algebras (cf. [31, Lemma 6.4]), expression (3.5) implies the existence of elements $a_{jlt} \in \mathbb{C}$, $h_t \in M$ such that $\sum_t a_{jlt} h_t = \sum_{i,k} a_{ijkl} u_k f_i$ for all j, l and $\sum_{j,l} a_{jlt} v_l g_j = 0$ for all t . Hence, rearranging the sum we conclude that $\sum_j (\sum_l a_{jlt} v_l) g_j = 0$ in N for all t , which implies $\sum_l a_{jlt} v_l = 0 \in G$ for all j, t . This in turn implies $a_{jlt} = 0$ for all j, l, t .

Using the condition $\sum_i (\sum_k a_{ijkl} u_k) f_i = \sum_t a_{jlt} h_t = 0$ for all j, l , we have $\sum_k a_{ijkl} u_k = 0$ for all i, j, l . Therefore, $a_{ijkl} = 0$ for all i, j, k, l , as we wanted to show. \square

Corollary 3.3.6. *Given $X, Y \subset \mathbb{T}^n$ two irreducible subvarieties of tori, we can consider the associated irreducible subvariety variety $X \cdot Y \subset \mathbb{T}^n$. Then as sets:*

$$\mathcal{T}(X \cdot Y) = \mathcal{T}X + \mathcal{T}Y,$$

where the sum on the right-hand side denotes the Minkowski sum in \mathbb{R}^n .

The same result holds for the projectivization of X and Y in \mathbb{P}^{n-1} .

We complement this result with a formula for computing multiplicities in $\mathcal{T}(X \cdot Y)$, which follows from Lemma 3.3.1 and Theorems 3.3.2 and 3.3.4.

Corollary 3.3.7. *Let X, Y be as in Corollary 3.3.6.*

- (i) *If $\dim X + \dim Y = \dim(X \cdot Y)$, then the multiplicity at a regular point w of $\mathcal{T}(X \cdot Y)$ equals:*

$$m_w = \frac{1}{\delta} \sum_{\substack{u \in \mathcal{T}X, \\ v \in \mathcal{T}Y, \\ w = u+v}} m_u m_v \text{ index } ((\mathbb{L}_u + \mathbb{L}_v) \cap \mathbb{Z}^n, (\mathbb{L}_u \cap \mathbb{Z}^n) + (\mathbb{L}_v \cap \mathbb{Z}^n)),$$

where δ is the degree of the generically finite monomial map $X \times Y \rightarrow X \cdot Y$. We assume that the number of such u, v is finite, all of them are regular in $\mathcal{T}X$ and $\mathcal{T}Y$, and $\mathbb{L}_u, \mathbb{L}_v$ are linear spans of neighborhoods of $u \in \mathcal{T}X$ and $v \in \mathcal{T}Y$ respectively.

- (ii) *Let X' and Y' denote the quotient of X and Y by a maximal torus given by primitive sublattices Λ and Λ' of \mathbb{Z}^n . If $\dim X' + \dim Y' = \dim(X' \cdot Y')$, multiplicities at a regular point w of $\mathcal{T}(X \cdot Y)$ can be computed as:*

$$m_w = \frac{1}{\delta} \sum_{\substack{\pi(u) \\ \pi(v)}} m_u m_v \text{ index } ((\mathbb{L}_u + \mathbb{L}_v) \cap \mathbb{Z}^n, (\mathbb{L}_u \cap \mathbb{Z}^n) + (\mathbb{L}_v \cap \mathbb{Z}^n)) \text{ index } ((\Lambda + \Lambda')^{\text{sat}}, \Lambda + \Lambda'),$$

where we sum over all classes $\pi(u) \in \mathcal{T}X/(\mathbb{R} \otimes \Lambda)$, $\pi(v) \in \mathcal{T}Y/(\mathbb{R} \otimes \Lambda')$ with $u + v \equiv w \pmod{\mathbb{R} \otimes (\Lambda + \Lambda')}$. Here, δ is the degree of the generically finite monomial map $X' \times Y' \rightarrow X' \cdot Y'$, and $(\Lambda + \Lambda')^{\text{sat}}$ is the primitive lattice $(\Lambda + \Lambda') \otimes_{\mathbb{Z}} \mathbb{R} \cap \mathbb{Z}^n$.

We assume that the number of classes of such points u, v is finite, all of them are regular in $\mathcal{T}X$ and $\mathcal{T}Y$, and $\mathbb{L}_u, \mathbb{L}_v$ are linear spans of neighborhoods of $u \in \mathcal{T}X$ and $v \in \mathcal{T}Y$ respectively.

As we mentioned in Section 1.4, the set-theoretic description in Corollary 3.3.6 is motivated by and is a direct consequence of Theorem 1.1.14 and the fact that valuations turn products into sums. The novelty of our approach is that under suitable finiteness conditions of the monomial map defining Hadamard products, we can effectively compute multiplicities of regular points in $\mathcal{T}(X \cdot Y)$ from multiplicities of $\mathcal{T}X$ and $\mathcal{T}Y$, as stated in Corollary 3.3.7. It is important to mention that this finiteness condition holds for the example we are studying in this chapter, after taking the quotient by a 5-dimensional torus action. Moreover, we are **not** claiming that $\mathcal{T}(X \cdot Y)$ inherits a fan structure from $\mathcal{T}X$ and $\mathcal{T}Y$. In general, it might happen that maximal cones in the Minkowski sum get subdivided to give maximal

cones in $\mathcal{T}(X \cdot Y)$ or, moreover, the union of several cones in the Minkowski sum gives a maximal cone in $\mathcal{T}(X \cdot Y)$. Example 1.2.2 illustrates this undesirable behavior.

Due to the lack of a fan structure in our description of the projective variety $X \cdot Y$, Corollary 3.3.6 gives no estimate for the number of maximal cones in the tropical variety $X \cdot Y$, where the fan structure is inherited from the Gröbner fan structure of the homogeneous defining ideal of $X \cdot Y$ (Section 1.1.1). Moreover, this fan structure is infeasible to obtain in general. Hence, in the hypersurface case we have no estimate on the number of edges of the dual polytope to the tropical variety $\mathcal{T}(X \cdot Y)$ and, as a consequence, no estimate on the number of vertices of the polytope. As Example 1.2.2 illustrates, the description of $\mathcal{T}(X \cdot Y)$ as a collection of weighted cones of maximal dimension contains less combinatorial information than the fan structure does and hence, the computation of the dual polytope becomes more challenging, as we show in Section 3.4.

We now describe the computation of the tropical variety $\mathcal{T}V_4^2$ associated to our model V_4^2 . By our discussions in Section 3.2, we know that the defining ideal of $V_4^1 = \text{Sec}(\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1) \subset \mathbb{P}^{15}$ is generated by the 3×3 minors of the three flattenings of $2 \times 2 \times 2 \times 2$ matrix of variables (p_{ijkl}) , for a total of 48 generators. Since V_4^1 is irreducible, we can use **Gfan** [54] to compute the tropical variety $\mathcal{T}V_4^1$.

The ideal $I(V_4^1)$ of $\mathbb{C}[p_{0000}, \dots, p_{1111}]$ is invariant under the action of B_4 , and **Gfan** can exploit the symmetry of a variety determined by an action of a subgroup of the symmetric group \mathbb{S}_{16} . For this, we need to provide a set of generators as part of the input data. The output groups cones together according to their orbits.

The tropical variety $\mathcal{T}V_4^1 \in \mathbb{R}^{16}$ has a lineality space spanned by the rows of the following integer matrix:

$$\Lambda = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad (3.6)$$

where the columns correspond to variables p_{ijkl} , for $i, j, k, l \in \{0, 1\}$, ordered lexicographically. As we explained already in this section, we can identify this linear space with the maximal torus acting on the variety V_4^1 and hence on V_4^2 . A set of generators of the corresponding lattice giving this action can be read-off from the parameterization. More precisely, consider the morphism of tori $\beta: \mathbb{T}^5 \rightarrow \mathbb{T}^{16}$ sending $(t_0, \dots, t_4) \mapsto (t^{m_1}, \dots, t^{m_{16}})$, where each m_i is of the form $(1, v_i)$, where v_i runs over all sixteen vertices of the 4-cube. Then, one can check that the closure of the image of β in \mathbb{C}^{16} is the affine cone over the Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$. More precisely, given a generic point in the image of β , we have $t_0 t_1^i t_2^j t_3^k t_4^l = \lambda x_i y_j z_k w_l$, where $(x_0 : x_1) = (1 : t_1)$, $(y_0 : y_1) = (1 : t_2)$, $(z_0 : z_1) = (1 : t_3)$, $(w_0 : w_1) = (1 : t_4) \in \mathbb{P}^1$ and $\lambda = t_0 \in \mathbb{R}$.

The **Gfan** computation confirms that the tropical variety $\mathcal{T}V_4^1$ inside \mathbb{R}^{16} has dimension 10 and a 5-dimensional lineality space. After moding out by this linear subspace, the f -vector

of the resulting pointed fan is:

$$(382, 3436, 11236, 15640, 7680).$$

Regarding the orbit structure, there are 13 rays and 49 maximal cones in \mathcal{TV}_4^1 up to symmetry and all maximal cones have multiplicity 1.

According to Corollary 3.3.6, the tropical variety of the model is (as a set)

$$\mathcal{TV}_4^2 = \mathcal{TV}_4^1 + \mathcal{TV}_4^1.$$

Since we know that this will result in a pure polyhedral fan, we only need to compute all Minkowski sums between pairs of cones of *maximal* dimension. For this step we use the B_4 group action. There is a natural (coordinatewise) action of $B_4 \times B_4$ on $\mathcal{TV}_4^1 \times \mathcal{TV}_4^1$ that translates to a B_4 -action on $\mathcal{TV}_4^1 + \mathcal{TV}_4^1$. Therefore, to compute the Minkowski sum of maximal cones, we first consider $49 \cdot 7680 = 376320$ pairs (σ_1, σ_2) , where σ_1 is taken from a set of representatives of the 49 orbits of maximal cones, and σ_2 is taken from the set of all maximal cones. We discard the pairs (σ_1, σ_2) for which $\sigma_1 + \sigma_2$ is not of maximal dimension 15. After this reduction, the total number of maximal cones computed is 92469. By construction, this list of 92469 cones contains all representatives of the orbits of maximal cones in \mathcal{TV}_4^2 . But they do not form distinct orbits. Some cones appear twice in the list as $\sigma + \tau$ and $\tau + \sigma$, and this the only possibility except for 4512 cones which arise from two different pairs, plus their flips. That is, $\sigma_1 + \tau_1 = \sigma_2 + \tau_2$ where both pairs differ only by an interchange of a *single* pair of extremal rays $(r_1, r_2) \in (\sigma_1, \tau_1)$: i.e. $\sigma_2 = (\sigma_1 \setminus \{r_1\}) \cup \{r_2\}$ and $\tau_2 = (\tau_1 \setminus \{r_2\}) \cup \{r_1\}$. Some cones σ have non-trivial stabilizers in B_4 , so there are cones $\sigma + \tau_1$ and $\sigma + \tau_2$ in the same orbit. The dimension of the maximal cones in \mathcal{TV}_4^2 confirms that V_4^2 is a hypersurface, as predicted by Theorem 2.1.2.

The total number of orbits of maximal cones is 18972, and each orbit has size 96, 192, or 384. We then let the group B_4 act on each orbit and obtain 6865824 cones of dimension 15, the union of which is the tropical variety \mathcal{TV}_4^2 , as predicted by Corollary 3.3.6. We do not have a fan structure of \mathcal{TV}_4^2 . Nonetheless, we can compute the multiplicity of any regular point in \mathcal{TV}_4^2 using Corollary 3.3.7. After taking quotients by the respective maximal torus acting on each space, the map $(V_4^1)' \times (V_4^1)' \rightarrow ((V_4^1)')^{[2]}$, is generically finite of degree two. In practice, the lattice indices in (2) are computed via greatest common divisors (gcd) of maximal minors of integer matrices whose rows span the corresponding maximal cones in \mathcal{TV}_4^1 and $\mathcal{T}(V_4^1 \times V_4^1)$. More precisely,

Lemma 3.3.8. *Given a lattice $D \subset \mathbb{Z}^n$, and an integer matrix $A \in \mathbb{Z}^{d \times n}$ with $\text{rk} A(D) = \text{rk} D$, the index $(\mathbb{R} \otimes_{\mathbb{Z}} A(D) \cap \mathbb{Z}^d, A(D))$ can be computed as follows. Pick $\{w_1, \dots, w_s\}$ a \mathbb{Z} -basis of D , and let $B := (w_1 \mid \dots \mid w_s) \in \mathbb{Z}^{n \times s}$. Then, the index of the lattice $A(D)$ in its saturation equals the quotient of the gcd of the maximal $s \times s$ -minors of the matrix $A \cdot B \in \mathbb{Z}^{d \times s}$ by the gcd of the maximal $s \times s$ -minors of the matrix B .*

Proof. Since $\mathbb{R} \otimes_{\mathbb{Z}} A(D) = \mathbb{R} \otimes_{\mathbb{Z}} A(\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n)$, the $\text{index}(\mathbb{R} \otimes_{\mathbb{Z}} A(D) \cap \mathbb{Z}^d, A(D))$ equals the product

$$\text{index}(\mathbb{R} \otimes_{\mathbb{Z}} A(D) \cap \mathbb{Z}^d, A(\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n)) \cdot \text{index}(A(\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n), A(D)).$$

By construction $\text{index}(\mathbb{R} \otimes_{\mathbb{Z}} A(D) \cap \mathbb{Z}^d, A(D))$ is the gcd of the maximal minors of the matrix $A \cdot B$. To prove the lemma, it suffices to show that $\text{index}(A(\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n), A(D))$ equals the gcd of the maximal minors of the matrix B in the statement.

Since $\text{rk } A(D) = \text{rk } D$, this implies that $\ker A \cap D = \ker A \cap (\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n) = \{0\}$. Then, $A(\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n)/A(D) \cong (\mathbb{R} \otimes_{\mathbb{Z}} D \cap \mathbb{Z}^n)/D$, which equals the gcd of the maximal minors of the matrix B , as we wanted to show. \square

In our case, B is spanned by twenty integer vectors: five from each cone $\sigma \times \mathbf{0}, \mathbf{0} \times \tau \in \mathcal{TV}_4^1 \times \mathcal{TV}_4^1$ plus the lattices $\Lambda \times \mathbf{0}, \mathbf{0} \times \Lambda$ coming from the lineality space. Call C_σ and C_τ each list of five vectors of σ and τ . Then, the matrix B in the previous lemma equals the block diagonal matrix $B = \text{diag}(B_\sigma, B_\tau)$, where $B_\sigma = (C_\sigma | \Lambda)$, $B_\tau = (C_\tau | \Lambda)$ and $A \cdot B = (C_\sigma | \Lambda | C_\tau | \Lambda)$. Thus, the index equals the quotient of $\text{gcd}(15 \times 15\text{-minors of } (C_\sigma | C_\tau | \Lambda))$ by the product $\text{gcd}(10 \times 10\text{-minors of } (C_\sigma | \Lambda)) \cdot \text{gcd}(10 \times 10\text{-minors of } (C_\tau | \Lambda))$. Each gcd calculation is done via the Hermite (alt. Smith) normal form of these matrices [72]. After computing all multiplicities using `Macaulay 2` we obtain only values one or two.

3.4 Newton polytope of the defining equation

In this section, we focus our attention on the *inverse problem*. That is, given the tropical fan of our irreducible hypersurface, we wish to compute the *Newton polytope* of the defining equation $f = \sum_a c_a \underline{x}^a$ of the hypersurface V_4^2 , i.e. the convex hull of all vectors $a \in \mathbb{Z}^{16}$ such that \underline{x}^a appears with a nonzero coefficient in f . Since we chose to work with the `min` convention, the duality between polytopes and tropical hypersurfaces will come from the *inner* normal fan of the polytope. This can be seen, for example, in Figure 3.4.

3.4.1 Vertices and facets

We will first present the results of our computation before discussing algorithms and implementation in the following subsections. Here is the ultimate result:

Theorem 3.4.1. *The Newton polytope of the defining equation of V_4^2 has 17 214 912 vertices in 44 938 orbits and 70 646 facets in 246 orbits under the symmetry group B_4 .*

Among the 44 938 orbits of vertices, 215 have size 192 and 44 723 has size 384. The maximum coordinate of a vertex ranges between 14 and 20, and the minimum coordinate is either 0 or 1. All but 46 orbits have a zero-coordinate. A vertex can have up to seven zero-coordinates. Each vertex is contained in 11 to 62 facets. There are 11 800 symmetry

classes of *simple* vertices, that is, those contained in exactly 11 facets. The following is the unique symmetry class of vertices contained in 62 facets each:

$$(0, 0, 1, 17, 13, 6, 17, 1, 17, 1, 6, 13, 1, 17, 0, 0).$$

This vertex has orbit size 192.

We index the coordinates of \mathbb{P}^{15} by $\{0, 1\}^4$ and order them lexicographically, following the convention in Chapter 2. Since our polynomial is (multi)-homogeneous, knowing even a single point in the Newton polytope gives the multidegree. We now describe the multidegree of the hypersurface V_4^2 :

Theorem 3.4.2. *The hypersurface V_4^2 has multidegree $(110, 55, 55, 55, 55)$ with respect to the grading defined by the matrix in (3.6).*

The following table lists the orbit sizes of all 246 facet orbits:

orbit size	2	8	12	16	24	32	48	64	96	192	384
number of facet orbits	1	2	1	3	1	1	7	3	15	67	145

Table 3.1: Classification of facet orbits of the Newton polytope of V_4^2 under B_4 -symmetry by sizes.

The coordinates x_{ijkl} are naturally indexed by bit strings $ijkl \in \{0, 1\}^4$. The two facet inequalities in the size-2 orbit say that the sum of x_{ijkl} such that $i + j + k + l$ is even (or odd) is at least 32. Each facet contains between 210 and 3 907 356 vertices. The unique symmetry class of facets containing the most vertices consist of coordinate hyperplanes.

Using Algorithm 3.3, we certified that out of the 13 orbits of rays of the 9-dimensional tropical variety $\mathcal{T}V_4^1$ of the first secant variety of the Segre embedding $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^{15}$, only the following eight are facet directions of $\mathcal{T}V_4^2$:

$-(1, 0, 0, 1, 0, 1, 1, 2, 2, 1, 1, 0, 1, 0, 0, 1)$
 $-(1, 3, 3, 1, 3, 1, 1, 3, 1, 3, 3, 1, 3, 1, 1, 3)$
 $-(2, 1, 1, 0, 1, 0, 0, 0, 2, 1, 1, 0, 1, 0, 0, 0)$
 $-(2, 1, 1, 2, 1, 2, 2, 1, 1, 2, 2, 1, 2, 1, 1, 2)$
 $-(3, 2, 2, 1, 2, 1, 1, 0, 2, 1, 1, 0, 1, 0, 0, 0)$
 $-(3, 3, 3, 3, 3, 3, 3, 3, 1, 3, 3, 1, 3, 1, 1, 3)$
 $-(-1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)$
 $-(-1, -1, -1, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, -1, -1, -1, -1).$

A complete list of vertices and facets, together with the scripts used for all computations, are available at

<http://math.berkeley.edu/~macueto/implicitizationChallenge.html>

3.4.2 Computing vertices

We now discuss how we obtained the Newton polytope of V_4^2 . In Section 1.4 we presented the duality between Newton polytopes and tropical hypersurfaces, together with an algorithm, known as *ray-shooting*, to go from the tropical hypersurface to the corresponding polytope. Since the operation of tropicalization interprets a polynomial as a Laurent polynomial, the polytope will be defined up to translation. The algorithm described in Theorem 1.4.1 computes a representative of $\text{NP}(f)$ which lies in the positive orthant and touches all coordinate hyperplanes, i.e. f is a polynomial not divisible by any non-constant monomial. We describe the pseudocode in Algorithm 3.1.

From the statement of Theorem 1.4.1 we see that we do not need a fan structure on $\mathcal{T}(f)$ to apply the ray-shooting algorithm. A description of the tropical hypersurface as a set, together with a way to compute the multiplicities of regular points, gives us enough information to compute vertices of the corresponding polytope in any generic direction. This observation fits into the description of the tropical hypersurface $\mathcal{T}V_4^2$ given in Section 3.3. The multiplicity of a regular point will be the sum of the multiplicities of all maximal cones containing this point.

In Section 3.3 we computed $\mathcal{T}(f)$ as a union of 6 865 824 cones. For each of those cones, we calculated the lattice index in Theorem 3.3.2 and the primitive vector which is the direction of the edge of $\text{NP}(f)$ normal to the cone. There are 15 788 distinct edge directions in $\text{NP}(f)$. We then pick a random vector $w \in \mathbb{R}^{16}$ and go through the list of 6 865 824 cones, recording the cones that meet any of the rays $w + \mathbb{R}_{>0} e_i$. For each i , we sum the numbers $m_v \cdot |l_i^v|$ over all the intersection points v and obtain the i^{th} coordinate of the vertex.

To obtain the multidegree, we only need one vertex. We computed the first vertex using `Macaulay 2` [44] in a few days. Our ultimate goal was to compute the Newton polytope $\text{NP}(f)$, a much more difficult computational problem that took us many more months to complete. As a first attempt, we bounded the number of lattice points in the polytope by the number of nonnegative lattice points of the given multidegree. Using the software `LatTE` [20], we found that the number of monomials in 16 variables with multidegree $(110, 55, 55, 55, 55)$ is 5 529 528 561 944.

By construction, it is clear that the bottleneck of Algorithm 3.1 is in going through the list \mathcal{F} of 6 865 824 cones. We can modify the algorithm to produce more than one vertex for each pass through the list. We do this in two ways. One is to process multiple objective vectors at once and save time by reducing the number of file readings and reusing the linear algebra computations for checking whether a cone meets a ray or not. Another way to produce more vertices is to keep track of the cones that we meet while ray-shooting, and use them to walk from chamber to chamber in the normal fan of $\text{NP}(f)$. This is described in Algorithm 3.2. On the polytope \mathcal{P} , this means walking from vertex P^{w+te_i} to $P^{w+t'e_i}$ for scalars $t' > t > 0$ corresponding to points between three consecutive intersection points, along an edge whose i^{th} coordinate is positive. If the vector w is generic, then we can assume that two maximal cones are parallel whenever they share an intersection point obtained by shooting from w in

Input: The list \mathcal{F} of maximal cones, with multiplicities, whose union is the codimension one cones in the outer normal fan of a polytope $\mathcal{P} \subset \mathbb{R}^n$. An objective vector $w \in \mathbb{R}^n$.

Assumption: The objective vector w does not lie in any cone in \mathcal{F} , i.e. the face P^w is a vertex. For each $i = 1, 2, \dots, n$ the ray $w + \mathbb{R}_{>0}e_i$ does not meet the boundary of any cone in \mathcal{F} .

Output: The vertex \mathcal{P}^w that maximizes the scalar product with the objective vector w .

```

 $P^w \leftarrow 0$ 
for each cone  $\sigma$  in  $\mathcal{F}$  do
  for  $i = 1, 2, \dots, n$  do
    if  $\sigma \cap (w + \mathbb{R}_{>0}e_i) \neq \emptyset$  then
       $\mathcal{P}_i^w \leftarrow \mathcal{P}_i^w + m_\sigma \cdot \ell_{\sigma,i}$ , where  $m_\sigma$  is the multiplicity of  $\sigma$  and  $\ell^\sigma$  is the
      primitive integral normal vector to  $\sigma$  such that  $\ell_i^\sigma > 0$ .
  return  $\mathcal{P}^w$ .

```

Algorithm 3.1: Ray-shooting: computing a vertex of a polytope from its normal fan.

a fixed coordinate direction. So we can use any of the cones in a parallel class to compute the edge direction of the wall we walk across. By adding up the multiplicities of the cones in each class, we get the lattice length of the corresponding edge of \mathcal{P} . This allows us to compute the coordinates of the vertices dual to the chambers we walk into. For each of these vertices found by walking from a known vertex, we also get an objective vector in the process. For example, in the notation of Algorithm 3.2, any vector of the form $w + te_i$, where $t_{k_j}^i < t < t_{k_j+1}^i$, is an objective vector for the j^{th} vertex found in the walk in direction e_i . We take t_{m+1}^i as ∞ . For numerical stability, we use exact arithmetic over the rational numbers. In particular, we always choose the new objective vectors to be integral.

Using a new vertex, with its associated objective vector, we can repeat the ray-shooting (Algorithm 3.1) and walking (Algorithm 3.2) again. The picture one should have in mind is that walking from chamber to chamber in the tropical side corresponds to walks from vertex to vertex in $\text{NP}(f)$ along edges normal to the codimension one cones traversed in the tropical hypersurface.

The combination of Algorithms 1 and 2 is illustrated in Figures 3.2 and 3.3. Starting from chamber \mathcal{C}_0 and an objective vector w_0 , we shoot rays in the coordinate axes directions. The intersection points are indicated by their defining parameters t_j^i (note that superscripts are omitted in the notation of Algorithm 3.2). As we explain below, to speed up the computation of Algorithm 3.1 we first precompute the inverses of all suitable matrices of the form $M_\sigma := (-e_i | r_1 | \dots | r_{15})$ where $\{r_1, \dots, r_{10}\}$ are generators of the cone σ and $\{r_{11}, \dots, r_{15}\}$ are the rows of the matrix A from (3.6) spanning the lineality space of \mathcal{TV}_4^1 . Using this, the condition

Input: A generic objective vector $w \in \mathbb{R}^n$, the vertex \mathcal{P}^w , and the set $\mathcal{S} := \{(\sigma, i, t) \in \mathcal{F} \times \{1, 2, \dots, n\} \times \mathbb{R}_{>0} : \sigma \cap (w + \mathbb{R}_{>0}e_i) = \{w + te_i\}\}$.
(This input is typically obtained from Algorithm 3.1.)

Output: The set of all vertices of \mathcal{P} with objective vectors of the form $w + te_i$ for some $t \in \mathbb{R}_{>0}$ and $i \in \{1, 2, \dots, n\}$.

for $i = 1, 2, \dots, n$ **do**

Let $\sigma_1, \dots, \sigma_m$ be the cones that intersect the ray $w + \mathbb{R}_{>0}e_i$ transversely.

Let $t_1, \dots, t_m \in \mathbb{R}_{>0}$ be such that $(\sigma_k, i, t_k) \in \mathcal{S}$ for $k = 1, 2, \dots, m$.

Order $\sigma_1, \dots, \sigma_m$ so that

$$t_1 = \dots = t_{k_1} < t_{k_1+1} = \dots = t_{k_2} < \dots < t_{k_l+1} = \dots = t_m := t_{k_{l+1}}.$$

$v \leftarrow \mathcal{P}^w$

for $j = 1, 2, \dots, l + 1$ **do**

$\ell^{\sigma_{k_j}} \leftarrow$ primitive integral normal vector to σ_{k_j} with $\ell_i^{\sigma_{k_j}} > 0$;

$v \leftarrow v + \left(\sum_{k_{j-1} < k \leq k_j} m_{\sigma_k} \right) \cdot \ell^{\sigma_{k_j}}$, where $k_0 = 1$, $k_{l+1} = m$, and m_σ denotes the multiplicity of σ .

Output v , and an objective vector in the line segment between $w + t_{k_j}e_i$ and $w + t_{k_{j+1}}e_i$, where $t_{k_{l+2}} := \infty$.

Algorithm 3.2: Walking: starting from an objective vector and corresponding vertex, compute the vertices obtained by changing the objective vector in negative coordinate directions.




Figure 3.2: Ray-shooting and walking algorithms combined. Starting from chamber C_0 we shoot and walk from chamber to chamber.

$(w_0 + \lambda e_i) \cap \sigma \neq \emptyset$ translates to the first eleven coordinates of the solution X of $X^t = (M_\sigma)^{-1} \cdot w_0$ being positive. Note that we can easily use the same systems to test $(w_0 - \lambda e_i) \cap \sigma \neq \emptyset$, just changing the sign condition for the first coordinate λ of X . This small modification allows us to walk in sixteen new directions (the negative coordinate axes), and find new adjacent vertices to vertex v_0 starting from objective vector w_0 . The step updating v in Algorithm 3.2 should be $v \leftarrow v - \left(\sum_{k_{j-1} < k \leq k_j} m_{\sigma_k} \right) \cdot \ell^{\sigma_{k_j}}$ instead of $v \leftarrow v + (\dots)$.

In Figure 3.2, the parameters λ associated to the intersection points in the positive directions are denoted by t_j^i , whereas we use \tilde{t}_j^i for the points obtained by shooting rays in the negative directions. The dashed arrows indicate all 32 shooting directions. The points in the codimension-one cones correspond to intersection points, whereas the points inside chambers are the objective vectors obtained for each vertex as described in Algorithm 3.2.

The dual walk in the Newton polytope is depicted in Figure 3.3. We start walking from vertex v_0 and via shooting we obtain the adjacent vertices v_3, v_5 and v_8 . If we walk in the negative directions, we obtain the adjacent vertices v_5, v_7, v_8 and v_{10} . Notice that by this




Figure 3.3: Walking from vertex to vertex in $\text{NP}(f) \subset \mathbb{R}^3$. In dash lines, we plot the tropical variety. The picture represents the local structure around v_0 .

procedure we miss vertices v_4, v_6 and v_9 , which are the remaining adjacent vertices to v_0 . However, we do get them if we start shooting from known adjacent chambers to \mathcal{C}_0 . For example, v_6 can be computed if we shoot rays and walk from chamber \mathcal{C}_8 , followed by a shoot from chamber \mathcal{C}_6 . Observe that this depends heavily on the choice of the objective vector w_0 .

3.4.3 Implementation

A few notes about the implementation of our algorithms are in order. As we started working on the problem, we used `Macaulay 2` [44] to do the ray-shooting (Algorithm 3.1). This script was fine for our first experiments, but it took three days to generate a single vertex of the polytope. It soon became evident that something faster was needed if we wanted to compute the entire polytope.

Our first step was to translate the `Macaulay 2` script for Algorithm 3.1 into `Python` [63]. We chose that language because of its fast speed of development and availability of arbitrary precision integers, which were needed by our program. We always scaled our objects (matrices and vectors) by positive integers so that our objects have integer coefficients. This step was crucial for numerical stability.

This new implementation brought the running time to about 10 hours. This was a remarkable improvement, but as the number of vertices of the polytope grew, we realized that something even faster was required. Therefore, we decided to resort to caching: instead of computing every inverse for each vector, we precomputed all the inverses and stored them using a binary format suitable for fast reading in `Python` (Pickles). This resulted in a file of a few tens of gigabytes, but dropped the time required for an individual ray-shooting

procedure down to under three hours.

Once the `Python` prototype was working at a reasonable speed, we translated it into `C++` [89], which brought the time required to do ray-shooting for a single vertex to 47 minutes on modest hardware. Moreover, ray-shooting for multiple objective vectors could be performed at the same time, thus amortizing the disk reads. Since we still needed large integers, we decided to use `GMP` [40] and its `C++` interface.

The procedure for walking is a more or less straightforward translation of the pseudocode presented in Algorithm 3.2. It is still implemented in `Python`, because it takes a short amount of time to walk from a few hundred vertices at a time, and the simplicity of the script far outweighs the time gains a `C++` translation would provide.

3.4.4 Certifying facets

We now discuss how to certify certain inequalities as facets of a polytope \mathcal{P} given by the dual tropical hypersurface $\mathcal{T}(f)$. By the duality between tropical hypersurfaces and Newton polytopes, each facet direction must be a ray in the tropical variety, equipped with the fan structure dual to \mathcal{P} . Lemma 3.4.3 provides a characterization for a vector in \mathbb{R}^n to be a ray of $\mathcal{T}(f)$ with the inherited Gröbner fan structure.

Lemma 3.4.3. *Let $w \in \mathbb{R}^n$ and $\mathcal{T}(f)$ be a tropical hypersurface given by a collection of cones, but with no prescribed fan structure. Let d be the dimension of its lineality space. Let $\mathcal{H} = \{\sigma_1, \dots, \sigma_l\}$ be the list of cones containing w . Let l^{σ_i} be the normal vector to cone σ_i for $i = 1, \dots, k$. Then, w is a ray of $\mathcal{T}(f)$ if and only if $\{l^{\sigma_1}, \dots, l^{\sigma_k}\}$ generates a $(n - d - 1)$ -dimensional vector space if and only if w is a facet direction of $\text{NP}(f)$.*

Proof. The vectors $\{l^{\sigma_1}, \dots, l^{\sigma_k}\}$ are precisely the directions of edges in the face \mathcal{P}^w of $\mathcal{P} := \text{NP}(f)$. Since the lineality space of $\mathcal{T}(f)$ has dimension d , the polytope \mathcal{P} has dimension $n - d$. The face \mathcal{P}^w is a facet of \mathcal{P} if and only if $l^{\sigma_1}, \dots, l^{\sigma_k}$ span a $(n - d - 1)$ -dimensional vector space. \square

For any objective vector $w \in \mathbb{R}^n$, we can compute a vertex in the face \mathcal{P}^w by applying ray-shooting (Algorithm 3.1) to a generic objective vector w' in a chamber of the normal fan of \mathcal{P} containing w . If we know that w is in fact a facet direction of \mathcal{P} , then any vertex in \mathcal{P}^w gives us the constant term a in the facet inequality $w \cdot x \geq a$. This is used in Algorithm 3.3 for checking if a given inequality is a facet inequality of \mathcal{P} . This step will be essential to certify that our partial list of vertices is indeed the complete list of vertices of the polytope \mathcal{P} . We discuss this approach in Section 3.4.5.

We now explain how to obtain a vector in the interior of a chamber containing a facet direction w . We start by applying a modified version of Algorithm 3.1 with input vector w and when we choose to shoot rays only in direction e_1 . Since w is a ray of the tropical variety given by the collection \mathcal{F} , it belongs to some cones $\{\tau_1, \dots, \tau_s\}$ in \mathcal{F} . Let $\sigma_1, \dots, \sigma_m$ be the cones we intersect along the e_1 direction (we allow intersections at boundary points

Input: An inequality $w \cdot x \geq a$ and a tropical hypersurface in \mathbb{R}^n (dual to polytope \mathcal{P}) given as a collection \mathcal{F} of weighted maximal cones, with $d = \text{dimension of the lineality space of the tropical hypersurface}$.

Output: *True* if the inequality is a valid facet inequality of \mathcal{P} ; *False* otherwise.

$\mathcal{N} \leftarrow \{\}$

for $\sigma \in \mathcal{F}$ **do**

if $w \in \sigma$ **then**
 $\mathcal{N} \leftarrow \mathcal{N} \cup \{\text{normal vector to } \sigma\};$

if $\dim\langle \mathcal{N} \rangle < n - d - 1$ **then**

Output *False*

else

$w' \leftarrow$ a vector in the interior of a chamber whose closure contains w
 Compute the vertex $\mathcal{P}^{w'}$ using ray-shooting (Algorithm 3.1).

if $w \cdot \mathcal{P}^{w'} = a$ **then**

Output *True*

else

Output *False*

Algorithm 3.3: Facet certificate: Check if a given inequality defines a facet of a polytope given by its normal fan.

of each cone). Note that we only pick those cones $\sigma \in \mathcal{F}$ whose corresponding edge direction satisfies $l_1^\sigma \neq 0$.

Now, we use Algorithm 3.2 with input vector w , the set \mathcal{S} corresponding to the cones $\sigma_1, \dots, \sigma_m$ and coordinate 1. We assume $(\sigma_k, 1, t_k)$ are ordered in an increasing order, with all $t_k \geq 0$. We have two possible scenarios: either \mathcal{S} is a subset of $\{0\}$ (that is, either the empty set or the set $\{0\}$) or it contains a positive real number. In the first case, we pick an objective vector $w_1 = w + te_1$ for a positive number t (for numerical stability, we choose t to be a big integer number). In the second case, pick a number t between zero and the first positive number from \mathcal{S} and let $w_1 = w + te_1$.

Third, we check if any cone in \mathcal{F} contains w_1 or not. If not, then we let $w' = w_1$. If yes, by the balancing condition for tropical varieties, this means that there exists a maximal cone in the tropical variety containing both w_1 and w . Note that this cone may be obtained by gluing and/or subdividing some cones in \mathcal{F} . In this case, we proceed as above, replacing the original input vector w by w_1 and shooting rays using coordinate 2 instead of coordinate 1. We repeat this process with all coordinates if necessary. Unless we have w_i not contained in any cone of \mathcal{F} , at step i we are guaranteed to have a cone containing $w, w_1, \dots, w_{i-1}, w_i$ by construction. By the dimensionality argument, at most in sixteen steps, we obtain a vector w_i not contained in any cone of \mathcal{F} . This vector will be the objective vector w' from Algorithm 3.3.

3.4.5 Completing the polytope

Once the ratio of new vertices computed with ray-shooting and walking decreases, the next natural question that arises is how to guarantee that we have found all vertices of our polytope. To answer this question, we construct the tangent cones at each vertex and try to certify their facets as facets of \mathcal{P} .

Definition 3.4.4. Let \mathcal{P} be a polytope in \mathbb{R}^n and v a vertex of \mathcal{P} . We define the *tangent cone of \mathcal{P} at v* to be the set:

$$\mathcal{T}_v^{\mathcal{P}} := v + \mathbb{R}_{\geq 0}\langle w - v : w \in \mathcal{P} \rangle = v + \mathbb{R}_{\geq 0}\langle e : e \text{ edge of } \mathcal{P} \text{ adjacent to } v \rangle.$$

By construction, $\mathcal{T}_v^{\mathcal{P}}$ is a polyhedron with only one vertex and $\mathcal{P} = \bigcap_{v \text{ vertex of } \mathcal{P}} \mathcal{T}_v^{\mathcal{P}}$. In particular, an inequality defines a facet of \mathcal{P} if and only if it defines a facet of one of the tangent cones.

Let \mathcal{Q} be the convex hull of the vertices of \mathcal{P} obtained via Algorithms 3.1 and 3.2. Our goal is to certify that $\mathcal{Q} = \mathcal{P}$. We proceed as follows. For each vertex v of \mathcal{Q} we wish to compare the tangent cones $\mathcal{T}_v^{\mathcal{Q}}$ and $\mathcal{T}_v^{\mathcal{P}}$. Since \mathcal{Q} has over seventeen million vertices and $\mathcal{T}_v^{\mathcal{Q}}$ has no symmetry, straightforward convex hull computations are infeasible. If $\mathcal{T}_v^{\mathcal{Q}} = \mathcal{T}_v^{\mathcal{P}}$ then the extreme rays of $\mathcal{T}_v^{\mathcal{Q}}$ would be edge directions of \mathcal{P} , which we have already computed as the normal directions to the list of maximal cones of the tropical hypersurface, and which are 15 788 in total. For a fixed vertex $v \in \mathcal{Q}$ we compute all differences $w - v$ for

all vertices w of \mathcal{Q} and test which of these vectors are parallel to the edges of \mathcal{P} . The number of such edge directions in $\mathcal{T}_v^{\mathcal{Q}}$ is expected to be very small (usually under 30 in practice). Let $C_v^{\mathcal{Q},\mathcal{P}}$ be the convex hull of all rays along the edge directions of \mathcal{P} in $\mathcal{T}_v^{\mathcal{Q}}$, translated by v . So we have $C_v^{\mathcal{Q},\mathcal{P}} \subseteq \mathcal{T}_v^{\mathcal{Q}}$ and we can test if $C_v^{\mathcal{Q},\mathcal{P}} \supseteq \mathcal{T}_v^{\mathcal{Q}}$ by computing facets of $C_v^{\mathcal{Q},\mathcal{P}}$ with `Polymake` [43]. If $C_v^{\mathcal{Q},\mathcal{P}} \supseteq \mathcal{T}_v^{\mathcal{Q}}$, we use Algorithm 3.3 to check whether each facet of $C_v^{\mathcal{Q},\mathcal{P}}$ is also a facet of \mathcal{P} . In this way, we can certify that $\mathcal{T}_v^{\mathcal{P}} \subseteq C_v^{\mathcal{Q},\mathcal{P}}$, hence $C_v^{\mathcal{Q},\mathcal{P}} = \mathcal{T}_v^{\mathcal{Q}} = \mathcal{T}_v^{\mathcal{P}}$. Certifying this for a vertex v of \mathcal{Q} in each symmetry class will give us $\mathcal{Q} = \bigcap_{v \text{ vertex of } \mathcal{Q}} \mathcal{T}_v^{\mathcal{Q}} \supseteq \bigcap_{v \text{ vertex of } \mathcal{P}} \mathcal{T}_v^{\mathcal{P}} = \mathcal{P}$, hence $\mathcal{Q} = \mathcal{P}$. We conclude:

Lemma 3.4.5. *Let \mathcal{P} be a polytope and $\mathcal{Q} \subset \mathcal{P}$ be the convex hull of a subset of the vertices in \mathcal{P} . If all facets of \mathcal{Q} are facets of \mathcal{P} , then $\mathcal{Q} \supset \mathcal{P}$, so $\mathcal{Q} = \mathcal{P}$.*

If we find that a facet $w \cdot x \geq a$ of $C_v^{\mathcal{Q},\mathcal{P}}$ is not a facet of \mathcal{P} from Algorithm 3.3, then we are missing vertices adjacent to v in \mathcal{P} in this “false facet direction” w , so we can perturb w with the techniques of Section 3.4.3 and Theorem 1.4.3, so that it lies in a chamber of the normal fan of \mathcal{P} and use ray-shooting (Algorithm 3.1) to find a new vertex in that direction. Using this method, we obtained the entire polytope in a finite number of steps. We describe the process of approximating \mathcal{P} by a subpolytope \mathcal{Q} in Algorithm 3.4. A schematic of complete tangent cones and incomplete tangent cones is depicted in Figure 3.4.




Figure 3.4: Approximation algorithm. We build a polytope $\mathcal{Q} \subset \mathcal{P}$, we compute the tangent cones at vertices of \mathcal{Q} and compare them with the corresponding tangent cones of \mathcal{P} . We certify if facets of \mathcal{Q} are also facets of \mathcal{P} by Algorithm 3.3. In the picture, we certify all facet directions in $\mathcal{T}_v^{\mathcal{Q}}$ containing vertex v but we cannot certify the facet direction w of the tangent cone $\mathcal{T}_v^{\mathcal{Q}}$. In addition, although we can certify the facet direction w' of \mathcal{Q} as a true facet direction of \mathcal{P} , we will not be able to certify the constant corresponding to this facet direction of \mathcal{P} since we are missing all its supporting vertices. The true constant will be obtained using Algorithm 3.3.

Input: A partial list V of vertices of \mathcal{P} , a collection of weighted cones \mathcal{F} whose union is the tropical hypersurface, $d = \text{dimension of the lineality space of the tropical hypersurface}$, and the group of symmetries of the tropical hypersurface.

Output: A complete list of vertices and facets of \mathcal{P} .

$\mathcal{S} \leftarrow \{\}$;

for *representatives v of orbits of V* **do**

$C_v^{\mathcal{Q}, \mathcal{P}} \leftarrow$ convex hull of v and all rays in directions $w - v$ where $w \in V$ and $w - v$ is normal to a cone in \mathcal{F} .

$A \leftarrow$ facets of $C_v^{\mathcal{Q}, \mathcal{P}}$ (using Polymake).

for $z \in A$ **do**

if z is a not facet of \mathcal{P} by Algorithm 3.3 **then**

$w' \leftarrow$ a vector in the interior of a chamber whose closure contains z .

Compute the vertex $\mathcal{P}^{w'}$ using ray-shooting (Algorithm 3.1).

$V \leftarrow V \cup \text{orbit of } \mathcal{P}^{w'}$

Break and restart the outermost for-loop with the new V .

else

$\mathcal{S} \leftarrow \mathcal{S} \cup \{z\}$

Output vertices V and facets \mathcal{S} .

Algorithm 3.4: Approximation of \mathcal{P} by a subpolytope \mathcal{Q} : Given a partial list of vertices of a polytope \mathcal{P} with unknown complete list of vertices, we construct the subpolytope \mathcal{Q} generated by this list. We certify when \mathcal{Q} equals \mathcal{P} .

In the final stages of the computation, if we find that $C_v^{\mathcal{Q},\mathcal{P}}$ is a strict subcone of the tangent cone $\mathcal{T}_v^{\mathcal{Q}}$, we enumerated the rays $w - v$ (with $w \in V$) that lie in the difference $\mathcal{T}_v^{\mathcal{Q}} \setminus C_v^{\mathcal{Q},\mathcal{P}}$. If the number of such rays is small (no more than a few hundred), we replace $C_v^{\mathcal{Q},\mathcal{P}}$ with the convex hull of $C_v^{\mathcal{Q},\mathcal{P}}$ and those rays (computed using `Polymake`) and proceed as in Algorithm 3.4. By executing Algorithm 3.4 in this way, we were able to compute and certify all vertices and facets of the polytope, solving Problem 3.1.1.

Chapter 4

Tropical secant graphs of monomial curves

This chapter is joint work with Shaowei Lin. It is on the arXiv with identifier 1005.3364v1, under the same title. The present version incorporates some minor changes, largely for consistency with other chapters. An extended abstract including part of this material appeared in the Proceedings of *FPSAC 2010*, edited by *Discrete Mathematics and Theoretical Computer Science*.¹

4.1 Introduction

In this chapter, we define and study four graphs that hide rich geometry: an abstract graph (the *abstract tropical secant surface graph*), a weighted graph in \mathbb{R}^{n+1} (the *tropical secant surface graph* or *master graph*), a weighted graph in the $(n - 2)$ -sphere (the *tropical secant graph*) and, finally, a weighted graph representing a simplicial complex embedded in the same sphere (the *Gröbner tropical secant graph*), obtained by an appropriate refinement of the third graph. All four graphs are parameterized by a sequence of n coprime distinct positive integers i_1, \dots, i_n , where $n \geq 4$. As their names suggest, these graphs are stepping stones to construct either a tropical surface or the tropicalization of a secant variety.

The *abstract tropical secant surface graph* is constructed by gluing two caterpillar trees and several star trees, according to the combinatorics of the given integer sequence. We pick a *realization* of this abstract graph in \mathbb{R}^{n+1} by assigning integer coordinates to each node, with no injectivity assumption. Our map has a key feature: we can endow this graph with weights on all edges in such a way that it becomes balanced (Theorem 4.2.5). We call this weighted graph the *tropical secant surface graph* or *master graph* for short (Definition 4.2.1). This balanced graph is closely related to a *tropical surface* and will be the main player in this story. More precisely, it is the building block for constructing the tropicalization of a

¹<http://www.dmtcs.org/dmtcs-ojs/index.php/proceedings/article/view/dmAN0147/3170>

classical threefold: the first secant variety of a monomial projective curve $(1 : t^{i_1} : \dots : t^{i_n})$. By definition, the secant variety of the curve is the closure of the union of all lines that meet the curve in two distinct points. These varieties have been studied extensively in the literature; see [16, 82] and references therein for more details. One of the main contributions of this chapter is a complete characterization of their tropical counterparts, which is carried out in full detail in Section 4.4.

As we explained in Section 1.3, computing the tropicalization of an algebraic variety without any information about its defining ideal is not an easy task. We will do so using techniques from geometric tropicalization (Theorem 1.3.1). Although this divisorial characterization is very explicit, computations can be quite difficult if the variety is non-generic, as it is the case of the binomial surface associated to the master graph. We can see this from the extensive number of pages we devote to computing this tropical surface (Section 4.3), and also from the small number of concrete examples available in the literature. As we said before, we should not expect them to be simple since they are the manifestation of the algebro-geometric process of *resolution of singularities*. In this secants example, our methods allow us to read off the tropical multiplicities directly from the master graph, which encodes the resolution diagrams of the surface at each singular point and the intersection numbers among boundary divisors (see Figure 4.1). This is carried out in Sections 4.2 and 4.3, in particular in Theorem 4.2.6. It is worth mentioning that this construction provides a compactification of the toric arrangement given by the $n + 1$ binomial curves $(w^{i_j} - \lambda = 0)$ in \mathbb{T}^2 , for $0 \leq j \leq n$. Such compactifications have been studied recently by L. Moci [71]. His construction, closely related to ours in spirit, realizes the wonderful compactification of binomial arrangements, following earlier work by De Concini and Procesi [19].

Equipped with the tropical interpretation of the master graph, we describe the tropicalization of the first secant variety of any monomial curve in three ways. First, as a collection of four-dimensional cones in \mathbb{R}^{n+1} with multiplicities (Theorem 4.4.6). Second, as a collection of three-dimensional weighted cones in the tropical projective torus \mathbb{TP}^n , with a one-dimensional lineality space, and third, as a weighted graph in the $(n - 2)$ -sphere. This graph is constructed by gluing together specific nodes in the master graph, and identifying edges accordingly. Weights for these edges are computed from the original ones by combinatorial methods. We call this graph the *tropical secant graph* (Definition 4.4.4). Strictly contained as a subgraph of this graph is the first tropical secant complex of our monomial curve (Propositions 4.6.2 and 4.6.4). This complex has recently been investigated by Develin and Draisma in [22, 25] in an attempt to study tropicalizations of secant varieties of affine toric varieties and it was briefly mentioned in Chapter 2 in connection with the tropical model of the restricted Boltzmann machine.

These tropical secant graphs are then exploited to recover geometric information about the original secant variety in \mathbb{P}^n . Using tropical implicitization techniques, we recover the multidegree of the classical secant threefold with respect to the rank-two lattice generated by the all-one's vector and the exponent vector of the curve. The degree of this variety was previously worked out in [82] and, unsurprisingly, our methods give similar combinatorial

formulas for this degree in terms of the used exponents. The main advantage of our approach is that, with the same effort, we can provide much more information about the secant variety, including its *Chow polytope*.

Our construction is particularly enlightening in the case when $n = 4$, where the secant variety becomes a hypersurface. In this special situation, we recover the Newton polytope of its defining equation. Although the lack of a fan structure in our description of this tropical variety is not an issue in our methods, as we discussed in Section 1.4, it would be desirable to have one to predict extra combinatorial information about the Newton polytope, such as the number of facets. For this reason, we devote the last part of Section 4.5 to refining the presentation of the tropical secant graph to turn it into a weighted simplicial complex. This structure is inherited from the Gröbner fan structure of the defining homogeneous ideal. The name *Gröbner tropical secant graph* highlights this property. We illustrate all our constructions and results with the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$, inspired by [82, Example 3.3], and with the rational normal curve in \mathbb{P}^n (Example 4.6.1).

Although secant varieties have been extensively studied in the past, we hope our work illustrates the power of *tropical implicitization* and how it can be used to go beyond implicitization methods even when looking at classical examples.

4.2 The master graph

In this section, we describe the main object of this chapter: the master graph. We start by defining an abstract graph parameterized by n coprime positive integers i_1, \dots, i_n , where $n \geq 4$. To simplify notation, we call $i_0 = 0$ and we assume $i_1 < i_2 < \dots < i_n$. This graph will be built from two types of graphs: two caterpillar graphs $G_{E,D}$, $G_{h,D}$ and a family of star graphs $\{G_{F_{\underline{a},D}}\}_{\underline{a}}$ parameterized by suitable subsets \underline{a} of the index set $\{0, i_1, \dots, i_n\}$. We construct the abstract graph by gluing all our graphs $G_{E,D}$, $G_{h,D}$ and the family $\{G_{F_{\underline{a},D}}\}_{\underline{a}}$ along the common labeled nodes D_{i_j} . We call this abstract graph the *abstract tropical secant surface graph*.

Our first building block is the caterpillar graph $G_{E,D}$, as illustrated on the top-left side of Figure 4.1. It consists of $2n - 1$ nodes and $2n - 2$ edges. Nodes are grouped in two levels, with labels $E_{i_1}, \dots, E_{i_{n-1}}$ and D_{i_1}, \dots, D_{i_n} . Similarly, our second graph, denoted by $G_{h,D}$ and depicted on the bottom-left side of Figure 4.1, consists of $2n$ nodes grouped in two levels with labels $h_{i_1}, \dots, h_{i_{n-1}}$ and $D_0, D_{i_1}, \dots, D_{i_n}$ respectively, and $2n - 1$ edges.

The third family of graphs consists of star trees and is denoted by $\{G_{F_{\underline{a},D}}\}$. These graphs are parameterized by subsets of size at least two, obtained by intersecting an arithmetic progression of integer numbers with the index set. They are illustrated in the rightmost picture in Figure 4.1. We allow the common difference of these progressions to be 1, so the set of all exponents $\{0, i_1, \dots, i_n\}$ is a valid subset. The size of each subset $\underline{a} \subseteq \{0, i_1, \dots, i_n\}$ associated to an arithmetic progression will coincide with the degree of the corresponding node $F_{\underline{a}}$ in the abstract graph. Note that several arithmetic progressions can give the same

subset of $\{0, i_1, \dots, i_n\}$, and hence the same node $F_{\underline{a}}$ in the graph $G_{F_{\underline{a}}, D}$. If $\underline{a} = \{i_{j_1}, \dots, i_{j_k}\}$, then the graph has $k + 1$ nodes and k edges: k nodes labeled $D_{i_{j_1}}, \dots, D_{i_{j_k}}$ and a central node $F_{\underline{a}}$, connected to the other k nodes in the graph. As Example 4.2.3 reveals, only nodes of degree at least three will be relevant for our constructions, so in principle we should only consider subsets of size at least three. However, to simplify the statements in the chapter, we will allow subsets of size two as well.




Figure 4.1: The graphs $G_{E,D}$, $G_{h,D}$ and $\{G_{F_{\underline{a}},D}\}_{\underline{a}}$ glue together to form the *abstract tropical secant surface graph*.

We now realize the abstract tropical secant surface graph in \mathbb{R}^{n+1} by mapping each node to an integer vector. Our chosen mapping has additional data, a weight for each edge in the graph. Note that the assignment need not be injective, which explains the choice of the word “realization” instead of “embedding.” We call this weighted graph the *tropical secant surface graph* or *master graph*. We explain this construction in full detail below. For a numerical example, see Figure 4.2.

Definition 4.2.1. The *master graph* is a weighted graph in \mathbb{R}^{n+1} parameterized by n distinct coprime numbers $\{i_1, \dots, i_n\}$ with nodes:

- (i) $D_{i_j} = e_j := (0, \dots, 0, 1, 0, \dots, 0) \quad (0 \leq j \leq n)$,
- (ii) $E_{i_j} = (0, i_1, \dots, i_{j-1}, i_j, \dots, i_j), h_{i_j} = (-i_j, \dots, -i_j, -i_{j+1}, \dots, -i_n) \quad (1 \leq j \leq n-1)$,
- (iii) $F_{\underline{a}} = \sum_{i_j \in \underline{a}} e_j$ where $\underline{a} \subseteq \{0, i_1, \dots, i_n\}$ has size at least two and is obtained by intersecting an arithmetic progression of integers with the index set $\{0, i_1, \dots, i_n\}$.

Its edges agree with the edges of the abstract tropical secant surface graph, and have weights:

- (i) $m_{D_{i_0}, h_{i_1}} = 1, m_{D_{i_n}, E_{i_{n-1}}} = \gcd(i_1, \dots, i_{n-1}), m_{D_{i_n}, h_{i_{n-1}}} = i_n$,

- (ii) $m_{D_{i_j, E_{i_j}}} = \gcd(i_1, \dots, i_j)$, $m_{D_{i_j, h_{i_j}}} = \gcd(i_j, \dots, i_n)$ $(1 \leq j \leq n - 1)$,
- (iii) $m_{E_{i_j, E_{i_{j+1}}}} = \gcd(i_1, \dots, i_j)$, $m_{h_{i_j, h_{i_{j+1}}}} = \gcd(i_{j+1}, \dots, i_n)$ $(1 \leq j \leq n - 2)$,
- (iv) $m_{F_{\underline{a}, D_{i_j}}} = \sum_r \varphi(r)$, where we sum over the common differences r of all arithmetic progressions containing i_j and giving the same subset \underline{a} . Here, φ denotes Euler's phi function.

Remark 4.2.2. As we mentioned earlier, if the subset \underline{a} coming from an arithmetic progression has two elements say i_j and i_k , then $F_{\underline{a}}$ is a bivalent node and we may eliminate it from the graph if desired, replacing its two adjacent edges by a single edge. Both edges $F_{i_j, i_k} D_{i_j}$ and $F_{i_j, i_k} D_{i_k}$ will have the same multiplicity, so we assign this number as the multiplicity of the new edge $D_{i_j} D_{i_k}$.

From Definition 4.2.1, it is immediate to check that the node E_{i_1} is always bivalent. However, we will always keep it in our graph, since it will greatly simplify our constructions in Section 4.4.

We illustrate the definition of the master graph with a numerical example. We should warn the reader that, unlike the case of this example, master graphs in general may have nodes $F_{\underline{a}}$ with $0 \notin \underline{a}$. This will be completely determined by the combinatorics of the set $\{i_1, \dots, i_n\}$.

Example 4.2.3. We compute the master graph associated to the set $\{30, 45, 55, 78\}$. For simplicity, we eliminate all nine bivalent nodes F_{i_j, i_k} from the graph, but we keep the bivalent grey node E_{i_1} . The resulting weighted graph has 16 vertices and 36 edges and it is depicted in Figure 4.2. There are five nodes of type $F_{\underline{a}}$, namely $F_{0,30,45,55,78} = (1, 1, 1, 1, 1)$, $F_{0,30,45,78} = (1, 1, 1, 0, 1)$, $F_{0,30,45,55} = (1, 1, 1, 1, 0)$, $F_{0,30,45} = (1, 1, 1, 0, 0)$ and $F_{0,30,78} = (1, 1, 0, 0, 1)$. They correspond to the five red unlabeled nodes in the picture. \diamond

Before stating the main result of this section, we recall the definition of a balanced graph.

Definition 4.2.4. Let $(G, m) \subset \mathbb{R}^N$ be a weighted graph where each node has integer coordinates. Let w be a node in G and let $\{w_1, \dots, w_r\}$ be the set of nodes adjacent to w . Consider the primitive lattices $\Lambda_w = \mathbb{R}\langle w \rangle \cap \mathbb{Z}^N$ and $\Lambda_{w, w_i} = \mathbb{R}\langle w, w_i \rangle \cap \mathbb{Z}^N$. Then $\Lambda_{w, w_i} / \Lambda_w$ is a rank one lattice, and it admits a unique generator which lifts to an element in the cone $\mathbb{R}_{>0}\langle w, w_i \rangle \subset \mathbb{R}^N$. Let $u_{w_i|w}$ be one such lifting. We say that the *node* w is *balanced* if $\sum_{i=1}^r m_{w_i, w} u_{w_i|w} \in \mathbb{R}\langle w \rangle$. If all nodes of G are balanced, then the weighted graph (G, m) satisfies the *balancing condition*.

Theorem 4.2.5. *The master graph satisfies the balancing condition.*




Figure 4.2: The master graph associated to the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$.

Proof. We proceed by analyzing the balance at each node, following Definition 4.2.4. The main difficulty will be to find the corresponding vector $u_{w_i|w}$ for each edge w_iw in the graph. We define $g_{i_j} := \gcd(i_1, \dots, i_j)$ and $g^{i_j} := \gcd(i_j, \dots, i_n)$. Note that these are the weights $m_{D_{i_j}, E_{i_j}}$ and $m_{D_{i_j}, h_{i_j}}$ of the master graph. To simplify notation, we set $E_{i_0} = E_{i_n} = h_{i_0} = h_{i_n} = \mathbf{0}$, add edges $E_{i_0}E_{i_1}$, $E_{i_{n-1}}E_{i_n}$, $h_{i_0}h_{i_1}$ and $h_{i_{n-1}}h_{i_n}$ to our graph and assign weight zero to these edges.

We start by checking the balance at all nodes E_{i_j} , for $1 \leq j \leq n-1$. In this case, we know that $\Lambda_{E_{i_j}} = \mathbb{Z}\langle E_{i_j}/g_{i_j} \rangle = \mathbb{Z}\langle (0, i_1/g_{i_j}, \dots, i_j/g_{i_j}, \dots, i_j/g_{i_j}) \rangle$ and $\Lambda_{E_{i_j}, D_{i_j}} = \mathbb{Z}\langle E_{i_j}/g_{i_j}, e_j \rangle$. So $u_{D_{i_j}|E_{i_j}} = e_j$. Similarly, we have $u_{D_{i_n}|E_{i_{n-1}}} = e_n$. On the other hand, $\Lambda_{E_{i_j}, E_{i_{j+1}}} = \mathbb{R}\langle E_{i_{j+1}}/g_{i_{j+1}}, E_{i_j}/g_{i_j} \rangle \cap \mathbb{Z}^{n+1}$. By definition, we need to extend the primitive vector E_{i_j}/g_{i_j} to a basis of $\Lambda_{E_{i_j}, E_{i_{j+1}}}$ by adding a single vector with appropriate sign. In this case, $u_{E_{i_{j+1}}|E_{i_j}} = \sum_{k=j+1}^n e_k$.

Next, we compute $u_{E_{i_{j-1}}|E_{i_j}}$. Here $\Lambda = \mathbb{Z}\langle E_{i_{j-1}}/g_{i_{j-1}}, E_{i_j}/g_{i_j} \rangle$ is not a primitive lattice, and we need to extend E_{i_j}/g_{i_j} to a basis of its saturation $\Lambda_{E_{i_{j-1}}, E_{i_j}}$. Our first candidate vector is $(E_{i_j} - E_{i_{j-1}})/(i_j - i_{j-1}) = -\sum_{k=j-1}^n e_k$. However, since $g_{i_{j-1}}$ need not equal g_{i_j} , we need

to slightly modify our choice. We propose $v = (0, a i_1/g_{i_{j-1}}, \dots, a i_{j-1}/g_{i_{j-1}}, -b, \dots, -b)$, for $a, b \in \mathbb{Z}$ such that $a i_j + b g_{i_{j-1}} = g_{i_j}$. We can check that all nonzero 2×2 -minors of the matrix with rows E_{i_j}/g_{i_j} and v are of the form: $-b i_k/g_{i_j} - (i_j i_k)/(g_{i_j} g_{i_{j-1}}) = -i_k/g_{j-1}$, for $j-1 \leq k \leq n$, so their gcd equals one. Therefore, E_{i_j}/g_{i_j} and $\pm v$ generate a primitive lattice, which contains $E_{i_{j-1}}$ by construction. To determine the correct choice of sign for $\pm v$, we write $E_{i_{j-1}}$ as a linear combination of E_{i_j}/g_{i_j} and v , and we require the coefficient of v to be positive. In this case,

$$E_{i_{j-1}} = g_{i_j}(1 - a(i_j - i_{j-1})/g_{i_j}) \cdot E_{i_j}/g_{i_j} + g_{i_{j-1}}(i_j - i_{j-1})/g_{i_j} \cdot v.$$

Thus, we conclude that $u_{E_{i_{j-1}}|E_{i_j}} = v$ for $j-1 \leq k \leq n$. With these weights, it is straightforward to check that the graph is balanced at E_{i_j} .

Working with g^{i_j} instead of g_{i_j} , a similar procedure to the one we just described proves that the graph is balanced at the nodes h_{i_j} for all $1 \leq j \leq n-1$. Balance at the nodes $F_{\underline{a}}$ follows by construction, so it remains to check the balance at the nodes D_{i_j} . In this case, $u_{E_{i_j}|D_{i_j}} = E_{i_j}/g_{i_j}$, $u_{h_{i_j}|D_{i_j}} = h_{i_j}/g^{i_j}$, $u_{F_{\underline{a}}|D_{i_j}} = F_{\underline{a}}$ ($i_j \in \underline{a}$), and $u_{E_{i_{n-1}}|D_{i_n}} = E_{i_{n-1}}$. The balancing equation at D_{i_j} gives

$$\sum_{\underline{a} \ni i_j} \left(\sum_r \varphi(r) \right) F_{\underline{a}} + E_{i_j} + h_{i_j} = \sum_{k=0}^n \left(\sum_{r \mid |i_k - i_j|} \varphi(r) - |i_k - i_j| \right) e_k.$$

Since $\sum_{l \mid k, l > 0} \varphi(l) = k$, we conclude that the graph is also balanced at D_{i_j} . \square

Knowing that the master graph is balanced, we can ask ourselves if this graph is related to the tropicalization of a surface in \mathbb{C}^{n+1} . In what follow we show that, indeed, this is the case. This explains the suggestive name of ‘‘tropical secant surface graph.’’

Theorem 4.2.6. *Fix a primitive strictly increasing sequence of n coprime positive integers $\{i_1, \dots, i_n\}$. Let Z be the surface in \mathbb{C}^{n+1} parameterized by $(\lambda, \omega) \mapsto (1 - \lambda, \omega^{i_1} - \lambda, \dots, \omega^{i_n} - \lambda)$. Then, the tropical surface $\mathcal{T}Z \subset \mathbb{R}^{n+1}$ coincides with the cone over the master graph as weighted polyhedral fans, with the convention that we assign the weight $i_1 + m_{F_{\underline{e}}, D_{i_1}}$ to the cone over the edge $D_{i_1} E_{i_1}$ if the ending sequence $\underline{e} = \{i_1, \dots, i_n\}$ gives a node $F_{\underline{e}}$ in the master graph.*

The proof of this statement involves techniques from geometric tropicalization and resolution of singularities of plane curves. Beautiful combinatorics are involved in its proof, as we show in Section 4.3.

Corollary 4.2.7. *With the notation of Theorem 4.2.6, the weighted graph obtained by identifying the nodes E_{i_1} and $F_{\underline{e}}$ in the master graph, and by assigning weight $i_1 + m_{F_{\underline{e}}, D_{i_1}}$ to the edge $D_{i_1} E_{i_1}$, agrees with the one-dimensional simplicial complex $\mathcal{T}Z \cap \mathbb{S}^n$.*

4.3 Combinatorics of monomial curves

In this section, we compute the tropical variety of the surface Z described in Theorem 4.2.6. Let $f_{i_j} := \omega^{i_j} - \lambda$ ($0 \leq j \leq n$) and consider the parameterization $\mathbf{f}: \mathbb{C}^2 \rightarrow Z$ given by these $n+1$ polynomials. Since geometric tropicalization involves subvarieties of tori, we restrict the domain of the function \mathbf{f} to the open set $X = \mathbb{C}^2 \setminus \bigcup_{j=0}^n (f_{i_j} = 0)$, which is the complement of a binomial arrangement. We give a compactification of X which, in turn, gives the desired compactification of $Z \cap \mathbb{T}^{n+1}$ with combinatorial normal crossing (CNC) boundary via the map \mathbf{f} . Recall that this CNC condition for surfaces means “no three irreducible (curve) components of the boundary intersect at a point.” This compactification will be precisely the *tropical compactification* of $Z \cap \mathbb{T}^{n+1}$ introduced in [94].

As we discussed in Section 1.3, most compactifications that one can construct for subvarieties of tori, such as the closure in suitable projective spaces, do not have the CNC boundary property. One method for producing such a compactification is to take the closure $\overline{Z \cap \mathbb{T}^{n+1}}$ of $Z \cap \mathbb{T}^{n+1}$ in \mathbb{P}^n and resolving the singularities of the boundary $\overline{Z} \setminus (Z \cap \mathbb{T}^{n+1})$ by blowing up points on curves to fulfill the CNC condition. We will follow this approach to construct the desired compactification. Along the way we will record intersection numbers among the boundary divisors. These numbers will allow us to compute tropical multiplicities, as stated in Theorems 5.2.4 and 1.3.3.

In the first part of this section, we explain in full detail the resolutions giving a compactification of X . Even though, in principle, we need a full resolution of our surface, in practice, this construction will give us several extra (exceptional) divisors that yield bivalent nodes in the intersection complex of the boundary of our surface. If we contract these curves with negative self-intersection, we obtain a singular surface whose boundary divisor has CNC. This weaker condition will suffice to construct our tropical surface. In the second part, we justify this contraction procedure from the tropical perspective, by means of Theorem 1.3.1. We end by computing the multiplicities of regular points in the tropical surface using Theorem 1.3.3. This will conclude the proof of Theorem 4.2.6.

We start by constructing a tropical compactification of X . First, we naively compactify X inside \mathbb{P}^2 . The components of the boundary divisor of X are $D_{i_j} = (f_{i_j}^h(\omega, \lambda, u) = 0)$ and $D_\infty = (u = 0)$, where $f_{i_j}^h$ is the homogenization of f_{i_j} with respect to the new variable u . Figures 4.3 and 4.4 illustrate this process in the case of the binomial arrangement X associated to the index set $\{0, 30, 45, 55, 78\}$, whose tropicalization was computed in Example 4.2.3. In these pictures, the black dots indicate the intersection of three or more of the corresponding binomial curves D_{i_j} or D_∞ , whereas grey dots indicates intersection of only two boundary components.

The boundary of X in \mathbb{P}^2 encounters three types of singularities: the origin $(0 : 0 : 1)$, the point $(0 : 1 : 0)$ at infinity, and isolated singularities in \mathbb{T}^2 . We resolve them all by blow-ups. After contracting appropriate exceptional curves, the resolutions diagrams from Figures 4.5 and 4.6 are *precisely* the graphs on the left side of Figure 4.1. The nodes E_{i_j} ($1 \leq j \leq n-1$)

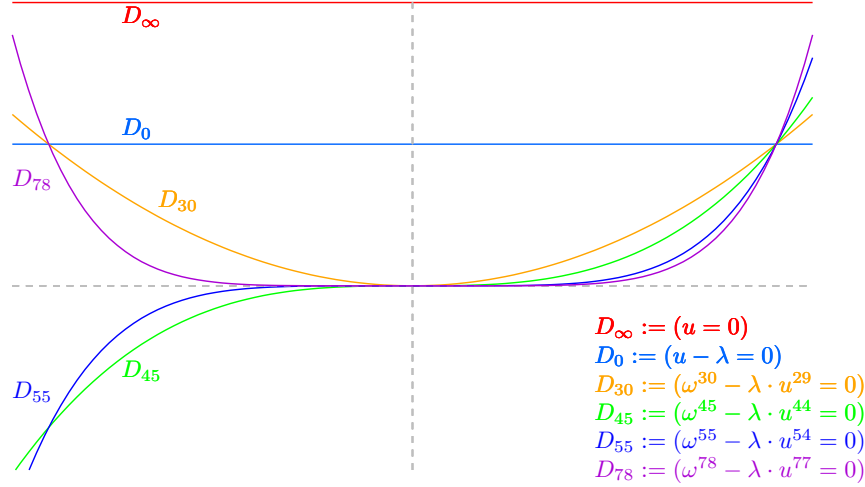


Figure 4.3: Compactification of Z in \mathbb{P}^n , with exponents $\{0, 30, 45, 55, 78\}$.

and h_{i_j} ($2 \leq j \leq n-1$) correspond to exceptional divisors, whereas h_{i_1} refers to the strict transform of the divisor D_∞ . All intersection multiplicities involving divisors E_{i_j} or h_{i_j} equal one.

We now describe the resolution process at the origin, whose affine chart is illustrated in the left of Figure 4.4. At this singular point, all n curves D_{i_1}, \dots, D_{i_n} intersect and they are tangential to each other. For any j , after a single blow-up, and by the change of coordinates $\lambda = w\lambda'$, we see that the strict transform of D_{i_j} is isomorphic to $D_{i_{j-1}}$ for all $1 \leq j \leq n$. This implies that we can resolve the singularity at the origin after i_{n-1} blow-ups. Moreover, we may use the previous isomorphism to compute the pull-back of each divisor D_{i_j} , one step at a time. For example, after the first blow-up, the proper transform of D_{i_j} will equal $\pi_1^*(D_{i_j}) = D'_{i_j} + E_1$, where $E_1 = (w = 0)$ is the exceptional divisor and D'_{i_j} is the strict transform of D_{i_j} . After a second blow-up, we get $\pi_2^*(E_1) = E'_1 + E_2$, and $\pi_2^*(D'_{i_j}) - E_2 = D''_{i_j} \simeq D_{i_{j-2}}$, so $(\pi_1 \circ \pi_2)^*(D_{i_j}) = D''_{i_j} + E'_1 + 2E_2$ with $D''_{i_j} \simeq D_{i_{j-2}}$ and $E'_1 \cdot D''_{i_j} = 0$.

To simplify notation, we label all exceptional divisors by E_l and we denote by D'_{i_j} the strict transform of D_{i_j} under the composition π of all blow-ups. Figure 4.5 illustrates the resolution diagram. All exceptional divisors satisfy:

$$E_l \cdot E_k = \begin{cases} 1 & \text{if } |l - k| = 1, \\ 0 & \text{otherwise,} \end{cases} \quad D'_{i_j} \cdot E_l = \begin{cases} 1 & \text{if } l = i_j, \\ 0 & \text{otherwise.} \end{cases}$$

Proceeding by induction, we conclude:

$$\pi^*(D_{i_j}) = D'_{i_j} + \sum_{l=1}^{i_j} l \cdot E_l + \sum_{l=i_j+1}^{i_{n-1}} i_j \cdot E_l \quad 1 \leq j \leq n. \quad (4.1)$$




Figure 4.4: From top to bottom: $(u = 1)$ and $(\lambda = 1)$ affine charts describing the singularities of the embedding of the surface Z from Theorem 4.2.6 in \mathbb{P}^n , with exponents $\{0, 30, 45, 55, 78\}$.

By convention, the sum over an empty set equals 0.

If we eliminate the bivalent nodes E_l from Figure 4.5 by contracting the corresponding curves with negative self-intersection, we obtain the graph $G_{D,E}$ depicted in the left side of Figure 4.1, where, by abuse of notation, the strict transform of D_{i_j} is also denoted by D_{i_j} . From (4.1), we see that the divisorial valuation for each exceptional divisor gives the integer vector E_{i_j} described in Theorem 4.2.5.

At infinity, the resolution process is more delicate. Here, the singular point $p = (0 : 1 : 0)$ corresponds to the intersection of D_∞ and all divisors D_{i_j} with $i_j \geq 2$. In this case, all divisors D_{i_j} are singular at p , as the rightmost picture in Figure 4.4 illustrates. Hence, we first need to perform a blow-up to smooth out these divisors. If π_0 denotes this blow-up, we obtain

$$\pi_0^*(D_{i_j}) = D'_{i_j} + (i_j - 1)H, \quad \pi_0^*(D_\infty) = D'_\infty + H,$$

where $H = (t = 0)$ is the exceptional divisor and $D'_{i_j} = (\omega - t^{i_j-1}) \simeq D_{i_j-1}$, $D'_\infty = (w = 0)$ are the strict transforms of the corresponding prime divisors.




Figure 4.5: Resolution by blow-ups at the origin, where E_{i_j} are exceptional divisors and D'_{i_j} are strict transforms of the boundary prime divisors D_{i_j} .




Figure 4.6: Resolution by blow-ups at infinity. Here, s is the minimum index with $i_s \geq 2$.

As the reader may have discovered already, the setting after applying π_0 is very similar to the one we described for the singularity at the origin, although there are some minor local differences between them that are worth pointing out. Firstly, there is a singularity coming from the intersection of the divisors $D_{i_s-1}, \dots, D_{i_n-1}$, where s is the minimum index satisfying $i_s \geq 2$. This singular point plays the role of the origin in the chart ($u = 1$). In addition, there are two extra divisors D''_{∞} and H , passing through this point. These curves had no counterpart at the origin in the chart ($u = 1$). Along the resolution, D''_{∞} will be separated from the other divisors after a single blow-up, whereas the strict transform of H will continue to be tangential to the strict transform of *all* divisors D'_{i_j} that it intersects.

The resolution diagram at infinity is shown in Figure 4.6. In that picture, all exceptional divisors are denoted by h_l ($2 \leq l \leq i_n$) and we label the strict transforms of D_{∞} , D_{i_j} and H by D''_{∞} , D''_{i_j} and H' respectively. In this case, the pull-backs under the composition π of the last $i_n - 1$ blow-ups give:

$$\begin{cases} \pi^*(D'_{\infty}) = D''_{\infty} + \sum_{l=2}^{i_n} h_l, & \pi^*(H) = H' + \sum_{l=2}^{i_n} (l-1) \cdot h_l, \\ \pi^*(D'_{i_j}) = D''_{i_j} + \sum_{l=2}^{i_j} (l-1) \cdot h_l + \sum_{l=i_j+1}^{i_n} (i_j-1) \cdot h_l & (i_j \geq 2). \end{cases}$$

Composing π with the initial blow-up π_0 at the point $(0 : 1 : 0)$, we get:

$$\begin{cases} (\pi \circ \pi_0)^*(D_{i_j}) = D''_{i_j} + (i_j - 1)H' + \sum_{l=2}^{i_j} i_j(l-1) \cdot h_l + \sum_{l=i_j+1}^{i_n} (i_j - 1)l \cdot h_l & (i_j \geq 2), \\ (\pi \circ \pi_0)^*(D_\infty) = H' + D''_\infty + \sum_{l=2}^{i_n} l \cdot h_l. \end{cases}$$

All intersection numbers $h_l \cdot h_{l+1}$, $h_{i_j} \cdot D''_{i_j}$, $D''_\infty \cdot h_2$ and $h_{i_n} \cdot H'$ equal one, whereas all other pairs have intersection number zero. In addition, we know that D_∞ intersects D_0 at a point, thus $D''_\infty \cdot D_0 = 1$. Finally, the divisor D_∞ will also intersect D_{i_1} at a point different from $(0 : 1 : 0)$, only if $i_1 = 1$. Thus, $D''_\infty \cdot D_1 = 1$ if $i_1 = 1$. In all other cases, $D''_\infty \cdot D_{i_1} = 0$.

We now explain the transition from the resolution diagram at infinity to the graph $G_{h,D}$, depicted at the bottom-left of Figure 4.1. As we did when blowing up the origin, we only keep the $n - 1$ exceptional divisors h_{i_j} giving non-bivalent nodes in the resolution diagram. We also contract the strict transform H' of the divisor H , since it has negative self-intersection number.

The degree of the node D''_∞ in the intersection complex is determined by the value of the index s . If $i_1 \geq 2$, then $s = 1$ and D''_∞ is a bivalent node adjacent to D_0 and h_{i_1} , so we remove it from the resolution diagram. On the contrary, if $i_1 = 1$, then $s = 2$ and D''_∞ has degree 3: it is adjacent to the nodes D_0 , D_1 and h_{i_2} . The node h_{i_1} in $G_{h,D}$ corresponds to the divisor D''_∞ . In both cases, and after removing all bivalent nodes and the node associated to the divisor H' , we obtain the graph $G_{h,D}$.

We now study multiple intersections of divisors inside \mathbb{T}^2 . If (ω, λ) satisfies $f_{i_j} = \omega^{i_j} - \lambda = 0$ and $f_{i_k} = \omega^{i_k} - \lambda = 0$, then $\omega^{i_j} = \lambda = \omega^{i_k}$, so ω is a primitive r^{th} root of unity for some $r \mid (i_k - i_j)$. Equivalently, $i_j \equiv i_k \equiv s \pmod{r}$, $\omega = e^{2\pi ip/r}$ and $\lambda = \omega^s$ for p coprime to r . All other curves ($f_{i_l} = 0$) with $i_l \equiv s \pmod{r}$ will also meet at (ω, λ) . We represent this crossing point by $x_{p,r,s}$ and the indices of curves meeting at $x_{p,r,s}$ by $\underline{a}_{r,s}$, or \underline{a} for short. That is,

$$x_{p,r,s} = (e^{2\pi ip/r}, e^{2\pi ips/r}), \quad \underline{a} = \underline{a}_{r,s} := \{i_j \mid i_j \equiv s \pmod{r}\}.$$

Furthermore, the gradients of the curves meeting at the point $x_{p,r,s}$ are pairwise independent, so the curves intersect transversally at $x_{p,r,s}$.

If three or more curves meet at a point $x_{p,r,s}$ in \mathbb{T}^2 , we blow up this point to separate the curves. After a single blow-up, we obtain a new exceptional divisor $F_{\underline{a},x_{p,r,s}}$ which intersects the strict transform of all D_{i_j} ($i_j \in \underline{a}$) with multiplicity one. The resolution diagram is the graph $G_{F_{\underline{a}},D}$ on the right-hand side of Figure 4.1, where we identify the node $F_{\underline{a}}$ with the corresponding divisor $F_{\underline{a},x_{p,r,s}}$. From these intersection numbers, we conclude that the divisorial valuation of the exceptional divisor $F_{\underline{a},x_{p,r,s}}$ gives $[F_{\underline{a},x_{p,r,s}}] = \sum_{i_j \in \underline{a}} e_j$ for all intersection points $x_{p,r,s}$ coming from the same subset \underline{a} . Thus, we get a single integer vector $F_{\underline{a}} = \sum_{i_j \in \underline{a}} e_j$ in the realization of the intersection complex, as desired. This explains the notation chosen for the graph $G_{F_{\underline{a}},D}$ in Figure 4.1, where we accounted only for the indices

of divisors intersecting at a point, rather than recording the point itself. To simplify the computation of multiplicities in the tropical variety $\mathcal{T}Z$, we also blow up crossings with $|\underline{a}| = 2$. Such blow-ups will give bivalent nodes F_{i_j, i_k} that we can easily discard in the end.

Equipped with the *smooth* set obtained by resolving the multiple intersection points in the boundary of $\bar{X} \subset \mathbb{P}^2$, we can focus on the tropical side of the construction. From this perspective, we explain why we can contract all divisors giving bivalent nodes and the strict transform of the divisor H , leading to a (possibly singular) surface, whose boundary divisor satisfies the CNC condition. This, in turn, will allow us to compute multiplicities at all regular points of $\mathcal{T}Z$.

First, we use the extended map $\mathbf{f}: \mathbb{P}^2 \dashrightarrow \mathbb{T}^{n+1} \subset \mathbb{P}^n$, together with the resolution $\pi: \tilde{X} \rightarrow X \subset \mathbb{P}^2$, to push-forward the intersection complex of the boundary of \tilde{X} to the intersection complex of $\bar{Z} \setminus (Z \cap \mathbb{T}^{n+1})$ for a suitable compactification \bar{Z} . The extended map \mathbf{f} is defined as

$$\mathbf{f}: X \subset \mathbb{P}^2 \rightarrow Z \cap \mathbb{T}^{n+1} \quad \mathbf{f}(\omega, \lambda, u) = \left(\frac{u - \lambda}{u}, \frac{\omega^{i_1} - \lambda u^{i_1-1}}{u^{i_1}}, \dots, \frac{\omega^{i_n} - \lambda u^{i_n-1}}{u^{i_n}} \right),$$

that is, $\mathbf{f}(\omega, \lambda, u) = (f_{i_j}^h(\omega, \lambda, u)/u^{\deg(f_{i_j}^h)})_{i=0}^n$. We compose \mathbf{f} with the resolution π to get the commutative diagram

$$\begin{array}{ccc} \tilde{X} & & \\ \pi \downarrow & \searrow \tilde{\mathbf{f}} & \\ X & \xrightarrow{\mathbf{f}} & Z \cap \mathbb{T}^{n+1}. \end{array}$$

Next, we pull-back a basis of characters $\{\chi_j : 0 \leq j \leq n\}$ of \mathbb{T}^{n+1} under $\tilde{\mathbf{f}}$. We know that the pullbacks of the characters under \mathbf{f} are units on X and rational functions on the closure of X in \mathbb{P}^2 . More precisely, they will have zeros and poles only along the boundary ∂X . By the universal property of the blow-up, the same holds for \tilde{X} and the pullback under $\tilde{\mathbf{f}}$ of the characters of \mathbb{T}^{n+1} . Therefore,

$$\tilde{\mathbf{f}}^*(\chi_0) = \pi^*(\mathbf{f}^*(\chi_0)) = \pi^*(D_{i_0} - D_\infty) \quad \text{and} \quad \tilde{\mathbf{f}}^*(\chi_j) = \pi^*(D_{i_j} - i_j D_\infty) \quad \text{for } j \geq 1.$$

For simplicity and to agree with the notation of the graphs in Figure 4.1, we denote strict transforms of all divisors with the label of the corresponding original divisors. With this convention, $\tilde{\mathbf{f}}^*(\chi_j)$ equals

$$\left\{ \begin{array}{ll} D_{i_j} + \sum_{l=1}^{i_j} l \cdot E_l + \sum_{l=i_j+1}^{i_n-1} i_j \cdot E_l - D_\infty - H - \sum_{l=2}^{i_j} l \cdot h_l - \sum_{l=i_j+1}^{i_n} l \cdot h_l + \sum_{\substack{\underline{a} \ni i_j \\ x_{p,r,s}}} F_{\underline{a}, x_{p,r,s}} & \text{if } i_j < 2, \\ D_{i_j} + \sum_{l=1}^{i_j} l \cdot E_l + \sum_{l=i_j+1}^{i_n-1} i_j \cdot E_l - i_j \cdot D_\infty - H - \sum_{l=2}^{i_j} i_j \cdot h_l - \sum_{l=i_j+1}^{i_n} l \cdot h_l + \sum_{\substack{\underline{a} \ni i_j \\ x_{p,r,s}}} F_{\underline{a}, x_{p,r,s}} & \text{else.} \end{array} \right.$$

The corresponding divisorial valuations will be read off from the columns of matrix of coefficients of $(\tilde{f}^*(\chi_j))_{j=1}^n$ with respect to the divisors $D_{i_j}, E_{i_j}, h_{i_j}, H$ and $F_{\underline{a}, x_p, r, s}$. Using Theorem 1.3.1, we get the following rays in the tropical variety TZ :

$$\left\{ \begin{array}{l} [D_{i_j}] = e_j, \quad [H] = -\mathbf{1}, \quad [F_{\underline{a}, x_p, r, s}] = \sum_{i_j \in \underline{a}} e_j, \quad [D_\infty] = -\sum_{i_j < 2} e_j - \sum_{i_j \geq 2} i_j \cdot e_j, \\ [E_l] = \sum_{l \leq i_j} l \cdot e_j + \sum_{l > i_j} i_j \cdot e_j \quad (1 \leq l \leq i_{n-1}), \\ [h_l] = -\sum_{i_j < l} l \cdot e_j - \sum_{l \leq i_j} i_j \cdot e_j \quad (2 \leq l \leq i_n). \end{array} \right. \quad (4.2)$$

We see that $[h_{i_n}] = i_n[H]$, so the cone over the edge $h_{i_n}H$ in the realized intersection complex is one dimensional. This explains why we do not see the divisor H in the graph $G_{h,D}$ from Figure 4.1. Likewise $i_1 \cdot [F_{i_1, \dots, i_n}] = [E_{i_1}]$ if $\gcd(i_n - i_1, \dots, i_2 - i_1) \neq 1$, so the cones over the edges $F_{i_1, \dots, i_n}D_{i_n}$ and $E_{i_1}D_{i_1}$ agree. In this case, we can replace these two cones by a single cone, adding the two weights. However, these are not the only identifications we can perform to simplify our construction. The next result implies that we can eliminate the bivalent nodes E_l, h_l ($l \neq i_j$) as well as the nodes h_{i_n} and D_∞ from this complex. Roughly speaking, it says that bivalent nodes E_{i_l} or h_{i_l} are contained in the two dimensional cone spanned by the corresponding nodes $E_{i_j}, E_{i_{j+1}}$ or $h_{i_j}, h_{i_{j+1}}$ with $i_j < l < i_{j+1}$, and similarly for D_∞ . It also asserts that there are no overlaps between cones over the edges other than the one we already mentioned. These two facts allow us to reduce our resolution graphs to $G_{E,D}, G_{h,D}$, and $G_{F_{\underline{a}}, D}$, thus proving the set theoretic equality in Theorem 4.2.6.

Lemma 4.3.1. *With the notation of (4.2) we have equalities*

- (i) $\mathbb{R}_{\geq 0}\langle [E_l], [E_{l+1}] \rangle \cap \mathbb{R}_{\geq 0}\langle [E_{l+1}], [E_{l+2}] \rangle = \mathbb{R}_{\geq 0}\langle [E_{l+1}] \rangle$ ($i_j \leq l \leq i_{j+1} - 2, 0 < j < n - 1$);
- (ii) $\mathbb{R}_{\geq 0}\langle [E_{i_j}], \dots, [E_{i_{j+1}}] \rangle = \mathbb{R}_{\geq 0}\langle [E_{i_j}], [E_{i_{j+1}}] \rangle$ ($1 \leq j \leq n - 2$);
- (iii) $\mathbb{R}_{\geq 0}\langle [h_l], [h_{l+1}] \rangle \cap \mathbb{R}_{\geq 0}\langle [h_{l+1}], [h_{l+2}] \rangle = \mathbb{R}_{\geq 0}\langle [h_{l+1}] \rangle$ ($2 \leq i_j \leq l \leq i_{j+1} - 2, 0 < j < n$);
- (iv) $\mathbb{R}_{\geq 0}\langle [h_{i_j}], \dots, [h_{i_{j+1}}] \rangle = \mathbb{R}_{\geq 0}\langle [h_{i_j}], [h_{i_{j+1}}] \rangle$ ($1 \leq j \leq n - 1$);
- (v) $[h_{i_n}] \in \mathbb{R}_{\geq 0}\langle [h_{i_{n-1}}], [D_{i_n}] \rangle$.
- (vi) $\mathbb{R}_{\geq 0}\langle [h_2], \dots, [h_{i_s}] \rangle = \mathbb{R}_{\geq 0}\langle [h_2], [h_{i_s}] \rangle$ and $\mathbb{R}_{\geq 0}\langle [E_1], \dots, [E_{i_1}] \rangle = \mathbb{R}_{\geq 0}\langle [E_{i_1}] \rangle$.
- (vii) If $s = 1$, then $[D_\infty] = [h_{i_1}]$. If $s = 2$, then $\mathbb{R}_{\geq 0}\langle [D_\infty], [h_{i_1}], \dots, [h_{i_2}] \rangle = \mathbb{R}_{\geq 0}\langle [h_{i_1}], [h_{i_2}] \rangle$.

Moreover, among maximal cones over the master graph, there are no two-dimensional intersections except when $F_{\underline{e}}$ is a node in the master graph where $\underline{e} = \{i_1, \dots, i_n\}$. In this case, $i_1[F_{\underline{e}}] = [E_{i_1}]$ and $\mathbb{R}_{\geq 0}\langle [F_{\underline{e}}], [D_{i_1}] \rangle = \mathbb{R}_{\geq 0}\langle [E_{i_1}], [D_{i_1}] \rangle$.

Proof. We prove the identities involving the rays $[E_l]$ ($1 \leq l \leq i_{n-1}$) in (i) and (ii). The claims for $[h_l]$ in (ii) and (iii) can be proven analogously. Assume $i_j \leq l \leq i_{j+1} - 2$. Then $[E_{l+1}] = [E_l] + \sum_{k \geq j+1} e_k$ and $[E_{l+2}] = [E_l] + 2 \sum_{k \geq j+1} e_k$ and the first identity follows by simple linear algebra arguments.

To prove the second claim, it suffices to show that $[E_l] \in \mathbb{R}_{\geq 0} \langle [E_{i_j}], [E_{i_{j+1}}] \rangle$ if $i_j < l < i_{j+1}$. In fact, by linear algebra calculations, we obtain $[E_l] = \frac{i_{j+1}-l}{i_{j+1}-i_j} \cdot [E_{i_j}] + \frac{l-i_j}{i_{j+1}-i_j} \cdot [E_{i_{j+1}}]$. The identities in (vi) are a direct consequence of the equalities $[h_l] = \frac{i_s-l}{i_s-2} [h_2] + \frac{l-2}{i_s-2} [h_{i_s}]$ for all $2 \leq l \leq i_s$, and $[E_l] = l \sum_{j \geq 1} e_j = \frac{l}{i_1} [E_{i_1}]$.

To prove (vii) we consider all pairs of maximal cones and compute their intersection. We get either the origin or the cone over a node in the master graph. \square

Next, we compute the weights of all edges in the \mathcal{TZ} using Theorem 1.3.3. From the resolution \tilde{X} , we know that the intersection number of any two boundary curves is zero or one. Using Lemma 4.3.1, we see that they are no 2-dimensional overlaps, except for the cones over the edges $D_{i_1}E_{i_1}$ and $F_{i_1, \dots, i_n}D_{i_1}$. The degree of the map \mathbf{f} equals one.

With the exception of the edge $D_{i_1}E_{i_1}$, the formula for computing weights on the edges containing D_{i_j} , h_{i_j} , E_{i_j} involves a single summand, namely the corresponding lattice index. This number is the gcd of the 2×2 -minors of a matrix whose rows are the two nodes of each edge, and it agrees with weights assigned to the master graph.

To end, we obtain the multiplicity of the cones over the edges $F_{\underline{a}}D_{i_j}$ in \mathcal{TZ} , with $\underline{a} \neq \{i_1, \dots, i_n\}$. In this case, all summands in the formula equal one and so the multiplicity equals the number of summands. Given \underline{a} of size at least two and $i_j \in \underline{a}$, we need to compute all possible common differences r of arithmetic sequences giving the set \underline{a} . Their number is $\sum_r \varphi(r)$. Finally, if $\underline{e} = \{i_1, \dots, i_n\}$ gives a node $F_{\underline{e}}$ in the master graph, the divisors E_{i_1} and $F_{\underline{e}}$ map to proportional rays $[E_{i_1}]$ and $[F_{\underline{e}}]$. The multiplicities of cones over edges $F_{\underline{e}}D_{i_j}$ with $j \geq 2$ equal $\sum_r \varphi(r)$ for all common differences r generating the set \underline{e} . The formula to compute the weight of the edge $F_{\underline{e}}D_{i_1}$ has an extra summand: the one involving the term $E_{i_1}D_{i_1}$. Hence, $F_{\underline{e}}D_{i_1}$ has weight $m_{D_{i_1}, E_{i_1}} + m_{F_{i_1, \dots, i_n}, D_{i_1}} = i_1 + \sum_r \varphi(r)$, as stated in Theorem 4.2.6. This concludes our proof.

4.4 The master graph under Hadamard products

In this section, we use the master graph to construct a new weighted graph: the *tropical secant graph*. This graph encodes with the tropicalization of the first secant variety of a monomial projective curve C parameterized by $(1 : t^{i_1} : \dots : t^{i_n})$, where $0 = i_0 \leq i_1 \leq \dots \leq i_n$ are integers. We define the first secant variety of the curve C as

$$\text{Sec}^1(C) = \overline{\{a \cdot p + b \cdot q : a, b \in \mathbb{C}, p, q \in C\}} \subset \mathbb{P}^n.$$

As discussed in Section 1.1.1, tropicalizations are toric in nature. Thus, for the rest of this section, instead of looking at the projective varieties C and $\text{Sec}^1(C)$, we will study the

corresponding very affine varieties obtained by intersecting their affine cones in \mathbb{C}^{n+1} with the torus \mathbb{T}^{n+1} . To simplify notation, we will also denote them by C and $\text{Sec}^1(C)$. The tropicalizations of the projective varieties and their corresponding very affine varieties are the same, but we will think of the projective one as living in \mathbb{TP}^n rather than in \mathbb{R}^{n+1} , reducing its dimension by one.

We parameterize this secant variety by the *secant map*

$$\phi: \mathbb{T}^4 \dashrightarrow \mathbb{T}^{n+1}, \quad \phi(a, b, s, t) = (as^{i_k} + bt^{i_k})_{0 \leq k \leq n}. \quad (4.3)$$

After a monomial change of coordinates $b = -\lambda a$ and $t = \omega s$, we rewrite ϕ as

$$\phi(a, s, \omega, \lambda) = (as^{i_k} (\omega^{i_k} - \lambda))_{0 \leq k \leq n}. \quad (4.4)$$

Therefore, the secant variety is the Hadamard product of the affine cone over the corresponding monomial curve C and the surface Z described in Theorem 4.2.6.

Proposition 4.4.1. *Let C be the monomial curve $(1 : t^{i_1} : \dots : t^{i_n})$ and let Z be the surface parameterized by $(\lambda, \omega) \mapsto (1 - \lambda, \omega^{i_1} - \lambda, \dots, \omega^{i_n} - \lambda)$. Then, the first secant variety $\text{Sec}^1(C) \subset \mathbb{T}^{n+1}$ is the Hadamard product $C \cdot Z$.*

We tropicalize the previous proposition using Corollaries 3.3.6 and 3.3.7 to obtain:

Proposition 4.4.2. *Given C, Z as in Proposition 4.4.1, then as sets*

$$\mathcal{T}\text{Sec}^1(C) = \mathcal{T}C + \mathcal{T}Z,$$

where the sum on the right-hand side denotes the Minkowski sum in \mathbb{R}^{n+1} . Moreover, since $\dim \text{Sec}^1(C) = \dim C + \dim Z$, we can effectively compute the multiplicities of regular points in $\mathcal{T}\text{Sec}^1(C)$ using formula (1).

Since the C is parameterized by monomials, its tropicalization $\mathcal{T}C$ equals the vector space $\mathbb{R} \otimes_{\mathbb{Z}} \Lambda$ with constant weight one, where Λ is the lattice $\mathbb{Z}\langle(1, \dots, 1), (0, i_1, \dots, i_n)\rangle \subset \mathbb{Z}^{n+1}$. In addition, $\mathcal{T}Z$ is a pointed polyhedral fan, and the lineality space of $\mathcal{T}\text{Sec}^1(C)$ is $\mathcal{T}C$ and the associated spherical complex $(\mathcal{T}\text{Sec}^1(C)/\mathcal{T}C) \cap \mathbb{S}^{n-2}$ is a graph. As we will show in this section this graph can be obtained by identifying nodes and edges in the master graph by their residue class modulo $\mathcal{T}C$. In the sequel, we explain this reduction process.

Before diving into the computation of $\mathcal{T}\text{Sec}^1(C)$ for any monomial curve C , we show that it suffices to treat the case where exponent vectors are primitive and have pairwise distinct coordinates. Recall that these were our assumptions for Sections 4.2 and 4.3. This will simplify the combinatorics and the computation of multiplicities. Here is the precise statement:

Lemma 4.4.3. *Via reparameterizations, we can assume that the exponent vector parameterizing the curve C is a primitive lattice vector $(0, i_1, \dots, i_n)$ with $0 = i_0 < i_1 < \dots < i_n$.*

The first claim follows by reparameterizing the curve C as $t \mapsto (1 : t^{\frac{i_1}{g}} : \dots : t^{\frac{i_n}{g}})$, where $g = \gcd(i_1, \dots, i_n)$. The second assertion is a direct consequence of Theorem 1.1.10:

Proof of Lemma 4.4.3. Let $\{0, i_1, \dots, i_r\}$ be the distinct values in the exponent vector defining the curve C in increasing order. We partition the set of indices $\{0, \dots, n\}$ into $S_0 \sqcup \dots \sqcup S_r$, where each S_j consists of all indices with the same value i_j . The monomial map

$$\alpha: \mathbb{T}^r \rightarrow \mathbb{T}^{n+1} \quad (t_1, \dots, t_r) \mapsto (\underbrace{1, \dots, 1}_{|S_0| \text{ times}}, \underbrace{t_1, \dots, t_1}_{|S_1| \text{ times}}, \dots, \underbrace{t_r, \dots, t_r}_{|S_r| \text{ times}})$$

is 1-1 and the linear map A induced by α is injective. Therefore,

$$\mathcal{T}Sec^1(\alpha(C)) = \mathcal{T}(\alpha(Sec^1(C))) = A(\mathcal{T}Sec^1(C)), \quad \mathcal{T}(\alpha(Z)) = A(\mathcal{T}Z), \quad \mathcal{T}(\alpha(C)) = A(\mathcal{T}C),$$

and any fan structure in $\mathcal{T}Sec^1(C)$ and $\mathcal{T}Z$ translates immediately to a fan structure in $\mathcal{T}Sec^1(\alpha(C))$ and $\mathcal{T}\alpha(Z)$ by the injectivity of A . Finally, using formula (1.4), we see that multiplicities match up, i.e. $m_v = m_{\alpha(v)}$ for any regular points $v, \alpha(v)$. This concludes our proof. \square


Our next goal is to explain the relationship between $\mathcal{T}Sec^1(C)$ and the *master graph* presented in Section 4.2. First, Proposition 4.4.2 expresses the tropical secant variety set-theoretically as the Minkowski sum of $\mathcal{T}C$ and $\mathcal{T}Z$. As we discussed in Section 1.2, we cannot hope to obtain a canonical fan structure from this description. Nonetheless, we can still use this characterization to describe $\mathcal{T}Sec^1(C)$ not just as a set, but as a collection of four-dimensional weighted cones $\{\mathcal{T}C + \sigma\}$ where σ varies over maximal cones of $\mathcal{T}Z$ whose sum with $\mathcal{T}C$ has dimension four. This presentation allows us to compute the multiplicity of any regular point ω in $\mathcal{T}Sec^1(C)$ as the sum of weights of cones containing ω , in agreement with the spirit of Theorem 5.2.4.

Intersections between cones in the collection $\mathcal{T}Sec^1(C)$ come in several flavors. If the intersection between two cones is four-dimensional, we call it an *overlap*. If these cones coincide, we say the overlap is *complete*; otherwise, it is *partial*. If their intersection is three-dimensional, we call it a *crossing*. If they intersect at a common face of each, we say that the crossing is *nodal*; otherwise, it is *internal*. We wish to summarize all complete overlaps and nodal crossings in our collection of weighted edges. This data is recorded in the *tropical secant graph* from Definition 4.4.4. For a numerical example, see Figure 4.7. Meanwhile, partial overlaps and internal crossings will be considered in Theorems 4.4.12, 4.4.13 and 4.5.3 when discussing fan structures. As we hinted in Theorem 4.2.6, a special role is played by the subsets $\underline{e} = \{i_1, \dots, i_n\}$ and $\underline{b} = \{0, i_1, \dots, i_{n-1}\}$, the “beginning” and “ending” subsets. We discuss this in Remark 4.4.5.

Definition 4.4.4. The *tropical secant graph* is a weighted subgraph of the master graph in \mathbb{R}^{n+1} , with nodes:

As a corollary we obtain a set theoretic description of $\mathcal{T}Sec^1(C)$:

Corollary 4.4.7. *The underlying graph of $\mathcal{T}Sec^1(C)$ is obtained by gluing the graphs*



along all nodes D_{i_j} , and gluing together the nodes $E_{i_1} \equiv F_{i_1, \dots, i_n}$, $E_{i_{n-1}} \equiv F_{0, \dots, i_{n-1}}$.

The remainder of this section is devoted to proving Theorem 4.4.6, and in particular, to explaining the mysterious formulas for the weights of the tropical secant graph. We obtain these numbers using formula (1). The next propositions and lemmas characterize each one of the quantities involved in (1).

Proposition 4.4.8. *Let $\alpha: \mathbb{T}^{2n+2} \rightarrow \mathbb{T}^{n+1}$ be the Hadamard monomial map associated to the matrix $(I_{n+1} \mid I_{n+1}) \in \mathbb{Z}^{(n+1) \times 2(n+1)}$ and C, Z as in Proposition 4.4.1. Then, the generic fiber of $\alpha|_{C \times Z}$ has size 2, giving $\delta = 2$ in formula (1).*

Proof. Generically, by equation (4.4), the elements of the fiber of α at a point p are in one-to-one correspondence with pairs of points in the curve C that are collinear with p . By switching the role of these two points in the secant map, we know that the generic fiber of $\alpha|_{C \times Z}$ has size at least two. Using Lemma 4.4.9 we conclude that it has exactly two points. \square

Lemma 4.4.9. *For almost all points p in the secant variety, p lies on a single one secant line which, in addition, intersects the curve C at exactly two points.*

Proof. We restrict the secant map ϕ to the open torus \mathbb{T}^3 mapping $(a, s, t) \mapsto (as^{i_k} + (1 - a)t^{i_k})_{0 \leq k \leq n}$. We claim that it suffices to prove the lemma for points in the image of ϕ , when $n = 4$ and the exponents are coprime.

Assume the statement is true for $n = 4$ and coprime exponents, and consider all maps ϕ_j obtained by composing the map ϕ with the projections π_j onto the five coordinates $0, 1, 2, 3, j$ ($4 \leq j \leq n$). Let $d_j = \gcd(i_1, i_2, i_3, i_j)$ and reparameterize ϕ_j using the identities $x := s^{d_j}$ and $y := t^{d_j}$, that is, define $\tilde{\phi}_j: \mathbb{T}^3 \rightarrow \mathbb{C}^{n+1}$ as

$$\tilde{\phi}_j(a, x, y) := (1, ax^{\frac{i_1}{d_j}} + (1 - a)y^{\frac{i_1}{d_j}}, ax^{\frac{i_2}{d_j}} + (1 - a)y^{\frac{i_2}{d_j}}, ax^{\frac{i_3}{d_j}} + (1 - a)y^{\frac{i_3}{d_j}}, ax^{\frac{i_j}{d_j}} + (1 - a)y^{\frac{i_j}{d_j}}).$$

Since the exponents of $\tilde{\phi}_j$ are coprime and the lemma holds for $n = 4$ by assumption, we know that the fiber of a point $\tilde{\phi}_j(a, x, y)$ contains only two points, namely the points (a, x, y)

and $(1 - a, y, x)$. Therefore, any two points (a, s, t) and (a', s', t') in the fiber of ϕ (up to symmetry) will satisfy $a = a'$, $(s')^{d_j} = s^{d_j}$ and $(t')^{d_j} = t^{d_j}$ for all $4 \leq j \leq n$. Since $\gcd(d_4, \dots, d_n) = 1$ we conclude $s = s'$ and $t = t'$.

We now treat the case where $n = 4$ and the exponents are coprime. To prove our result, it suffices to show that the Zariski closure of the set of points in $\text{Sec}^1(C) \setminus C$ which are intersections of two distinct secant lines of C has dimension at most two. We parameterize these points by tuples (s, t, u, v) of distinct complex numbers, corresponding to four coplanar, non-collinear points in C . It suffices to show that the set W of such tuples has dimension at most two.

The variety W is cut out by all 3×3 -minors of the 3×4 -matrix with rows $(t^{i_j} - s^{i_j})_{1 \leq j \leq 4}$, $(u^{i_j} - s^{i_j})_{1 \leq j \leq 4}$ and $(v^{i_j} - s^{i_j})_{1 \leq j \leq 4}$. Note that for all minors to vanish, it is enough to show that two of them do. We pick the ones corresponding to columns $\{1, 2, 3\}$ and $\{1, 2, 4\}$. These minors are precisely the determinants of the 4×4 generalized Vandermonde matrices

$$M_{i_1, i_2, i_3} := \begin{pmatrix} 1 & s^{i_1} & s^{i_2} & s^{i_3} \\ 1 & t^{i_1} & t^{i_2} & t^{i_3} \\ 1 & u^{i_1} & u^{i_2} & u^{i_3} \\ 1 & v^{i_1} & v^{i_2} & v^{i_3} \end{pmatrix}, \quad M_{i_1, i_2, i_4} := \begin{pmatrix} 1 & s^{i_1} & s^{i_2} & s^{i_4} \\ 1 & t^{i_1} & t^{i_2} & t^{i_4} \\ 1 & u^{i_1} & u^{i_2} & u^{i_4} \\ 1 & v^{i_1} & v^{i_2} & v^{i_4} \end{pmatrix}.$$

Because s, t, u and v are all distinct, we can divide the determinant of these two matrices by the product of pairwise differences among our four variables, that is, by the *Vandermonde determinant* $V(s, t, u, v)$. The resulting polynomials are the *Schur polynomials* S_{i_1, i_2, i_3} and $S_{i_1, i_2, i_4} \in \mathbb{Z}[s, t, u, v]$.

Let $g = \gcd(i_1, i_2, i_3)$ and $h = \gcd(i_1, i_2, i_4)$. Note that $\gcd(g, h) = 1$. By [29, Theorem 3.1] we can factorize the previous Schur polynomials over $\mathbb{Z}[s, t, u, v]$ as

$$\begin{aligned} S_{i_1, i_2, i_3} &= V(s^g, t^g, u^g, v^g)/V(s, t, u, v) \cdot T_{i_1, i_2, i_3}(s, t, u, v), \\ S_{i_1, i_2, i_4} &= V(s^h, t^h, u^h, v^h)/V(s, t, u, v) \cdot T_{i_1, i_2, i_4}(s, t, u, v), \end{aligned}$$

where T_{i_1, i_2, i_3} and T_{i_1, i_2, i_4} are either constants or irreducible over $\mathbb{C}[s, t, u, v]$. These two polynomials are homogeneous of total degree $g(i_3/g + i_2/g + i_1/g - 3)$ and $h(i_4/h + i_2/h + i_1/h - 3)$, and of degree $i_3 - 2g$ and $i_4 - 2h$ in each variable s, t, u and v . By comparing their multidegrees, we conclude that the polynomials T_{i_1, i_2, i_3} and T_{i_1, i_2, i_4} are coprime. Next, we claim that the polynomials $V(s^g, t^g, u^g, v^g)/V(s, t, u, v)$ and $V(s^h, t^h, u^h, v^h)/V(s, t, u, v)$ are coprime. This follows from the well-known identity $(x^g - y^g) = \prod_{\zeta^g=1} (x - \zeta y)$.

From the previous two observations, we see that if S_{i_1, i_2, i_3} and S_{i_1, i_2, i_4} have a common factor, then either T_{i_1, i_2, i_3} and $V(s^h, t^h, u^h, v^h)/V(s, t, u, v)$ or the polynomials T_{i_1, i_2, i_4} and $V(s^g, t^g, u^g, v^g)/V(s, t, u, v)$ would have a common factor over $\mathbb{C}[s, t, u, v]$. Without loss of generality, assume the first pair of polynomials is not coprime. By irreducibility of T_{i_1, i_2, i_3} and the factorization of $V(s^h, t^h, u^h, v^h)$ into linear factors involving only two of the variables, this forces T_{i_1, i_2, i_3} to be linear and to involve only two variables, contradicting the degree formulas provided above. Hence, we conclude that S_{i_1, i_2, i_3} and S_{i_1, i_2, i_4} are coprime. In particular, this

says that the hypersurfaces in \mathbb{C}^4 defined by S_{i_1, i_2, i_3} and S_{i_1, i_2, i_4} in have distinct reduced irreducible components of dimension 3. Therefore, $\dim W \leq 2$ and this ends our proof. \square

Next, we compute all cones of the form $\mathcal{TC} + \mathbb{R}_{\geq 0} \otimes \sigma$ of dimension at most three, where σ runs over edges of the master graph \mathcal{TZ} . We discard these cones from the $\mathcal{TSec}^1(C)$. In addition, we consider all possible pairs σ, σ' of such maximal cones in \mathcal{TZ} to find all intersections $(\mathcal{TC} + \mathbb{R}_{\geq 0} \otimes \sigma) \cap (\mathcal{TC} + \mathbb{R}_{\geq 0} \otimes \sigma')$ which are complete overlaps or nodal crossings. By an elementary exhaustive case by case analysis, we conclude:

Lemma 4.4.10. *After reducing the master graph by the linear space \mathcal{TC} , the only non-maximal cones, complete overlaps and nodal crossings are as follows.*

- (i) *The cones $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle D_0, h_{i_1} \rangle$, $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle D_{i_n}, E_{i_{n-1}} \rangle$, $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle D_{i_n}, h_{i_{n-1}} \rangle$ and $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{0, i_1, \dots, i_n}, D_{i_j} \rangle$ ($0 \leq j \leq n$) are not maximal, so we disregard them.*
- (ii) *The node $F_{\{0, i_1, \dots, i_n\}} = \mathbf{1} \in \mathcal{TC}$, so we eliminate it from the graph, together with all its $n + 1$ adjacent edges.*
- (iii) *For all $1 \leq j \leq n - 2$, we have equalities $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle E_{i_j}, D_{i_j} \rangle = \mathcal{TC} + \mathbb{R}_{\geq 0} \langle h_{i_j}, D_{i_j} \rangle$ and $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle E_{i_j}, E_{i_{j+1}} \rangle = \mathcal{TC} + \mathbb{R}_{\geq 0} \langle h_{i_j}, h_{i_{j+1}} \rangle$ because $E_{i_j} \equiv h_{i_j}$ modulo \mathcal{TC} . Hence, we disregard all nodes h_{i_j} and their adjacent edges.*
- (iv) *$i_1 \cdot F_{\underline{e}} = E_{i_1}$ and $(i_n - i_{n-1}) \cdot F_{\underline{b}} \equiv E_{i_{n-1}}$ modulo \mathcal{TC} , where $\underline{e} = \{i_1, \dots, i_n\}$ and $\underline{b} = \{0, i_1, \dots, i_{n-1}\}$. Thus, the maximal cones $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{\underline{e}}, D_{i_1} \rangle$ and $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle E_{i_1}, D_{i_1} \rangle$ coincide, as well as $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{\underline{b}}, D_{i_{n-1}} \rangle$ and $\mathcal{TC} + \mathbb{R}_{\geq 0} \langle E_{i_{n-1}}, D_{i_{n-1}} \rangle$.*

Proof of Theorem 4.4.6. Proposition 4.4.2 and Lemma 4.4.10 prove that the cone from \mathcal{TC} over the tropical secant graph coincides with the tropical variety $\mathcal{TSec}^1(C)$ as a collection of four-dimensional weighted cones. In particular, this shows that the tropical secant graph combines all nodes and edges coming from nodal crossings and complete overlaps. By formula (1.4), the multiplicity at a regular point ω of $\mathcal{TSec}^1(C)$ is the sum of all weights of four-dimensional cones $\mathcal{TC} + \sigma$ containing ω , where σ is a maximal two-dimensional cone of \mathcal{TZ} . Furthermore, if m_σ is the weight of σ in \mathcal{TZ} , the formula assigns to $\mathcal{TC} + \sigma$ the weight

$$\frac{1}{2} \cdot m_\sigma \cdot \text{index}((\mathbb{L}_\sigma + \mathcal{TC}) \cap \mathbb{Z}^{n+1}, (\mathbb{L}_\sigma \cap \mathbb{Z}^n) + (\mathcal{TC} \cap \mathbb{Z}^{n+1})), \quad (4.5)$$

by Proposition 4.4.8.

We now prove that this number yields the weight of the corresponding edge in the tropical secant graph, after combining weights in complete overlaps following Remark 4.4.5. Suppose the cone σ is generated by integer vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^{n+1}$. Let $\mathbf{l}_1 = \mathbf{1}$ and $\mathbf{l}_2 = (0, i_1, \dots, i_n)$ be the generators of the primitive lattice Λ in $\mathcal{TC} \cap \mathbb{Z}^{n+1}$. The lattice index in the above formula is the gcd of all 4×4 -minors of the matrix $(\mathbf{x} \mid \mathbf{y} \mid \mathbf{l}_1 \mid \mathbf{l}_2)$ divided by the gcd of all 2×2 -minors of the matrix $(\mathbf{x} \mid \mathbf{y})$. These gcd's are computed as the product of the nonzero diagonal elements of the Smith normal form of each matrix.

As an example, we show how to obtain the multiplicity $m_{D_{i_j}, E_{i_j}}$ ($2 \leq j \leq n-2$). The remaining multiplicities can be computed analogously. The edge $D_{i_j} E_{i_j}$ is associated to precisely two edges in the master graph giving two four-dimensional cones in $\mathcal{T}Sec^1(C)$ that overlap completely. These two edges are $\sigma = D_{i_j} E_{i_j}$ and $\sigma' = D_{i_j} h_{i_j}$. From Definition 2.1, the multiplicity m_σ equals $\gcd(i_1, \dots, i_j)$, which is also the gcd of the 2×2 -minors of the matrix $(D_{i_j} | E_{i_j})$. Likewise, $m_{\sigma'} = \gcd(i_j, \dots, i_n)$ is the gcd of the 2×2 -minors of $(D_{i_j} | h_{i_j})$. These numbers are precisely the denominators in the formulas for computing the indices associated to σ and σ' in (4.5). since $E_{i_j} \equiv h_{i_j} \pmod{\Lambda}$, we conclude

$$\begin{aligned} m_{D_{i_j}, E_{i_j}} &= \frac{1}{2} \left(\gcd(4 \times 4\text{-minors of } (D_{i_j} | E_{i_j} | \mathbf{l}_1 | \mathbf{l}_2)) + \gcd(4 \times 4\text{-minors of } (D_{i_j} | h_{i_j} | \mathbf{l}_1 | \mathbf{l}_2)) \right) \\ &= \gcd(4 \times 4\text{-minors of } (D_{i_j} | E_{i_j} | \mathbf{l}_1 | \mathbf{l}_2)). \end{aligned}$$

We now compute the gcd of all 4×4 -minors of the matrix $(D_{i_j} | E_{i_j} | \mathbf{l}_1 | \mathbf{l}_2)$. For simplicity, we work with the transpose of this matrix. By elementary operations between rows that do not change the minors, we alter the second, third and fourth rows, and expand the minors along the first row, reducing our problem to computing the 3×3 -minors of the matrix

$$\left(\begin{array}{cccc|ccc} 0 & i_1 & \dots & i_{j-1} & i_j & \dots & i_j \\ 1 & 1 & \dots & 1 & 1 & \dots & 1 \\ 0 & 0 & \dots & 0 & i_{j+1} - i_j & \dots & i_n - i_j \end{array} \right) \in \mathbb{Z}^{3 \times n},$$

where $i_0 = 0$. All non-vanishing 3×3 -minors must involve columns from the two constituent blocks. The gcd of the minors involving two columns of the left-hand side is $\gcd(i_1, \dots, i_{j-1}) \gcd_{j < l < n} (i_l - i_j)$, whereas the gcd of the minors involving two columns of the rightmost block equals the product $\gcd(i_1, \dots, i_j) \gcd_{j < l < n} (i_n - i_l)$. This justifies the formula for $m_{D_{i_j}, E_{i_j}}$ appearing in Definition 4.4.4. \square

We now study partial overlaps and internal crossings among cones from \mathcal{TC} over edges of the tropical secant graph. An exhaustive case by case analysis shows that if $n \geq 5$, there are no partial overlaps, and that if $n \geq 6$, there are no internal crossings either. When $n = 4$, Lemma 4.5.2 and Theorem 4.5.3 in the next section indicate that both partial overlaps and internal crossings are possible.

Partial overlaps and internal crossings prevent us from inferring a fan structure for $\mathcal{T}Sec^1(C)$ from the tropical secant graph. However, we may introduce new nodes at crossings, subdivide edges while preserving their weights or merge overlapping edges and their weights to create a new graph. If this surgery is performed appropriately, the new graph encodes the fan structure of our tropical variety as a subfan of the *Gröbner fan* of the homogeneous ideal defining the secant variety. This motivates the following definition:

Definition 4.4.11. A *Gröbner tropical secant graph* for a projective monomial curve C parameterized by n coprime distinct integers is a weighted graph in \mathbb{R}^{n+1} whose cone from \mathcal{TC} gives the Gröbner weighted fan structure on $\mathcal{T}Sec^1(C)$.

Unsurprisingly, the complexity of the surgery required to transform the tropical secant graph into a Gröbner tropical secant graph depends on the value of n . We present the results for $n > 4$ and leave the case $n = 4$ for the next section.

Theorem 4.4.12. *The tropical secant graph of a monomial curve in \mathbb{P}^n is a Gröbner tropical secant graph for $n \geq 6$.*

Proof. The proof of this result is elementary and it boils down to analyzing intersections between cones from \mathcal{TC} over pairs of edges in the tropical secant graph. If $n \geq 6$, we do not get any partial overlaps or internal crossings. \square

Theorem 4.4.13. *For a monomial curve in \mathbb{P}^5 , a Gröbner tropical secant graph may be constructed from the tropical secant graph by adding finitely many nodes and subdividing edges accordingly. More precisely, we need to add nodes $P_{\underline{a},j,\underline{a}',k} \in (\mathcal{TC} + \langle F_{\underline{a}}, D_{i_j} \rangle) \cap (\mathcal{TC} + \langle F_{\underline{a}'}, D_{i_k} \rangle)$ where $F_{\underline{a}}, F_{\underline{a}'}$ are nodes in the master graph, and $\underline{a}, j, \underline{a}', k$ together with the index set $\{i_1, \dots, i_5\}$ satisfy one of the following three conditions:*

- (i) $j = 5, k = 0, i_4 + i_1 = i_2 + i_3$ and the subsets $\underline{a}, \underline{a}'$ are either $\underline{a} = \{i_3, i_4, i_5\}, \underline{a}' = \{0, i_1, i_3\}$ or $\underline{a} = \{i_2, i_4, i_5\}, \underline{a}' = \{0, i_1, i_2\}$.
- (ii) $\underline{a} = \{i_j, i_r, i_k, i_u\}, \underline{a}' = \{i_u, i_k, i_t\}, i_l \notin \underline{a} \cup \underline{a}', i_r + i_t = i_l + i_u$ and either $j > r > l > t, r > u > t$ and $u > k$, or $j < r < l < t, r < u < t$ and $u < k$.
- (iii) $\underline{a} = \{i_j, i_u, i_k, i_r\}, \underline{a}' = \{i_j, i_u, i_k, i_t\}, i_l \notin \underline{a} \cup \underline{a}'$ and $i_r + i_t = i_l + i_u$, while $j > u > k, r > l > t$ and $r > u > t$.

For each new node $P_{\underline{a},j,\underline{a}',k}$, we subdivide the edges $F_{\underline{a}}D_{i_j}$ and $F_{\underline{a}'}D_{i_k}$ of the tropical secant graph to get edges $F_{\underline{a}}P_{\underline{a},j,\underline{a}',k}$, $P_{\underline{a},j,\underline{a}',k}D_{i_j}$, and $F_{\underline{a}'}P_{\underline{a},j,\underline{a}',k}$, $P_{\underline{a},j,\underline{a}',k}D_{i_k}$, preserving the original weights.

Proof. The proof is tedious, yet elementary. As before, we consider all intersections between cones from \mathcal{TC} over edges of the tropical secant graph. While there are no partial overlaps, we have three types of internal crossings. We write down the intersection point for each case:

- (i) Suppose $|\underline{a}| = |\underline{a}'| = 3$, say $\underline{a} = \{i_j, i_r, i_u\}, \underline{a}' = \{i_u, i_k, i_t\}$. Let $i_l \notin \underline{a} \cup \underline{a}'$. Assume the cones over the edges $F_{\underline{a}}D_{i_j}$ and $F_{\underline{a}'}D_{i_k}$ intersect and $j > k$. Then, $j = 5, r = 4, t = 1, k = 0, u = 2$ or $3, i_4 + i_1 = i_2 + i_3$ and

$$(i_r - i_l) \cdot F_{\underline{a}} + (i_j - i_r) \cdot D_{i_j} = (i_r - i_u) \cdot F_{\underline{a}'} + (i_t - i_k) \cdot D_{i_k} + (-i_l)u \cdot \mathbf{1} + (0, i_1, \dots, i_5).$$

- (ii) Suppose $|\underline{a}| = 4, |\underline{a}'| = 3$, say $\underline{a} = \{i_j, i_r, i_k, i_u\}, \underline{a}' = \{i_k, i_u, i_t\}$. Let $i_l \notin \underline{a} \cup \underline{a}'$ and assume the cones over $F_{\underline{a}}D_{i_j}$ and $F_{\underline{a}'}D_{i_k}$ intersect. Then, $i_r + i_t = i_l + i_u$ and

$$(i_r - i_l) \cdot F_{\underline{a}} + (i_j - i_r) \cdot D_{i_j} = (i_l - i_t) \cdot F_{\underline{a}'} + (i_u - i_k) \cdot D_{i_k} + (-i_l) \cdot \mathbf{1} + (0, i_1, \dots, i_5).$$

In this case, all coefficients (except for $-i_l$ and 1) have the same sign, which can be negative. If the latter occurs, we multiply the previous identity by -1 , obtaining the internal crossing of the two cones. This expression gives us up to ten extra points, determined by the inequalities $j > r > l > t, r > u > t$ and $u > k$, or $j < r < l < t, r < u < t$ and $u < k$.

- (iii) Suppose $|\underline{a}| = |\underline{a}'| = 4$, say $\underline{a} = \{i_j, i_u, i_k, i_r\}$, $\underline{a}' = \{i_j, i_u, i_k, i_t\}$, and assume $j > k$. Let $i_l \notin \underline{a} \cup \underline{a}'$. If the cones over $F_{\underline{a}}D_{i_j}$ and $F_{\underline{a}'}D_{i_k}$ intersect, then $i_r + i_t = i_l + i_u$ and

$$(i_r - i_l) \cdot F_{\underline{a}} + (i_j - i_u) \cdot D_{i_j} = (i_l - i_t) \cdot F_{\underline{a}'} + (i_u - i_k) \cdot D_{i_k} + (-i_l) \cdot \mathbf{1} + (0, i_1, \dots, i_5).$$

By requiring all mandatory coefficients to be positive, we obtain $j > u > k, r > l > t$ and $r > u > t$. This gives twelve possibilities for such internal crossings. \square

4.5 The Newton polytope of the secant hypersurface in \mathbb{P}^4

In this section, we focus our attention on monomial curves in \mathbb{P}^4 . In this situation, the first secant variety becomes a hypersurface and we wish to obtain its definition homogeneous equation from the tropical secant graph. As we saw in Sections 1.4 and 3.4, a first step towards a complete solution would be to compute the *Newton polytope* of the defining equation f . In the same section, we described methods for performing such a task, based on two algorithms: ray-shooting (Algorithm 3.1) and walking (Algorithm 3.2). We illustrate these tropical implicitization techniques with an example appearing in the literature [82]. In addition, we construct the Gröbner tropical secant graph as promised in Section 4.4. As we already discussed, although a fan structure for the tropical secant graph is unnecessary for constructing the Newton polytope, knowing this structure would give important combinatorial data such as its f -vector.

Example 4.5.1. The first secant variety of the monomial curve $t \mapsto (1 : t^{30} : t^{45} : t^{55} : t^{78})$ in \mathbb{P}^4 is known to be a hypersurface of degree 1820 [82, Example 3.3]. We use geometric tropicalization to compute the tropicalization of this variety. Using this data as input for the ray-shooting and walking algorithms, we calculate the Newton polytope of this hypersurface.

By Theorem 4.4.6, we encode the tropical hypersurface $\mathcal{T}Sec^1(C) \subset \mathbb{R}^5$ by the 3-dimensional graph depicted in the left of Figure 4.7. The eleven nodes in the graph have coordinates $D_0 = e_0$, $D_{30} = e_1$, $D_{45} = e_2$, $D_{55} = e_3$, $D_{78} = e_4$, $E_{30} = (0, 30, 30, 30, 30)$, $E_{45} = (0, 30, 45, 45, 45)$, $F_{0,30,45,55} \equiv E_{55} = (0, 30, 45, 55, 55)$, $F_{0,30,45} = (1, 1, 1, 0, 0)$, $F_{0,30,78} = (1, 1, 0, 0, 1)$, and $F_{0,30,45,78} = (1, 1, 1, 0, 1)$. The unlabeled red nodes in the picture indicate nodes of type $F_{\underline{a}}$, where the subset \underline{a} consists of the indices of nodes D_{i_j} adjacent to it. Note that, in this example, the nodes E_{55} and $F_{0,30,45,55}$ are identified modulo the lineality

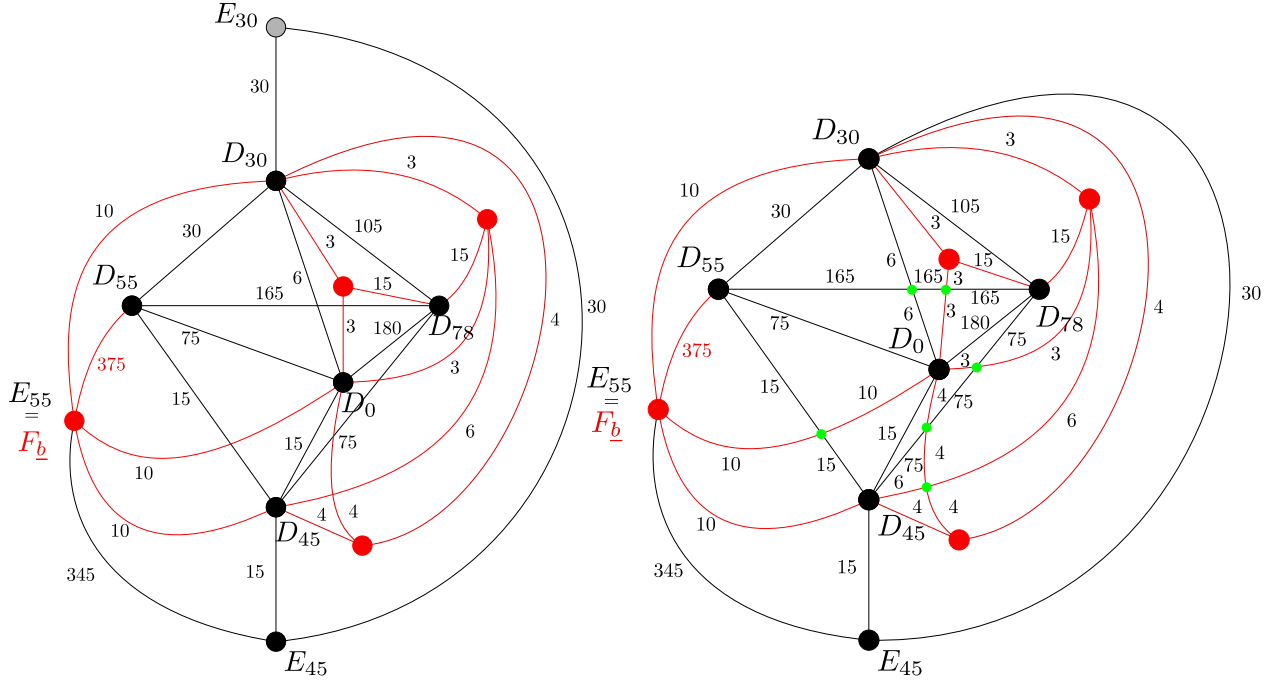


Figure 4.7: The tropical secant graph and the Gröbner tropical secant graph of the monomial curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$ in \mathbb{P}^4 .

space, as predicted by Definition 4.4.4. In particular, the edges $E_{55}D_{55}$ and $F_{0,30,45,55}D_{55}$ of the master graph coincide in the tropical secant graph and the old weights add up to $375 = 345 + 30$, as Figure 4.7 shows. After removing the bivalent gray node E_{30} , we have a graph with 10 nodes and 23 edges.

Finally, we apply the ray-shooting and walking algorithms to recover the Newton polytope of its defining equation. From our computation, we see that its multidegree with respect to the lattice $\Lambda = \mathbb{Z}\langle \mathbf{1}, (0, 30, 45, 55, 78) \rangle$ is $(1\ 820, 76\ 950)$, recovering the degree value from [82]. The polytope has 24 vertices and f -vector $(24, 38, 16)$. The difference between the number of facets of the polytope and the number of nodes in the tropical secant graph shows that this graph does not reflect the Gröbner fan structure of the tropical variety. In particular, we are missing six vertices which correspond to internal crossings of the graph. In the right of Figure 4.7, we indicate these six missing vertices with green nodes. After adding them to the picture we obtain the Gröbner tropical secant graph of the curve C , which is a planar graph in \mathbb{S}^2 with 16 nodes and 38 edges. Each edge emanating from a new green node corresponding to an internal crossing will inherit the weight of the original edge in the tropical secant graph. The complement of this graph has 24 connected components, which matches the number of vertices of our polytope. Using LaTTe, we see that the polytope contains 7 566 849 lattice points, which gives an upper bound for the number of monomials in its defining equation. \diamond

We conclude the section by building the Gröbner tropical secant graph of any monomial curve in \mathbb{P}^4 . From our previous example, we already know that the tropical secant graph can be non-planar. We provide a theorem that indicates which internal crossings need to be added to make it planar. However, these are not the only possible intersections: we can have partial overlaps between edges. Luckily, there are only three types of partial overlaps. We describe them in the next lemma, which follows notation from Definition 4.4.4:

Lemma 4.5.2. *The only partial overlaps among cones in the tropical secant graph of a monomial curve in \mathbb{P}^4 are:*

- (i) $F_{i_1, i_2, i_3} D_{i_2}$ and $D_{i_2} E_{i_2}$, where $(i_4 - i_2)i_1 = (i_4 - i_3)i_2$. In this case, E_{i_2} lies in the interior of the edge $F_{i_1, i_2, i_3} D_{i_2}$, and so we replace the edge $F_{i_1, i_2, i_3} D_{i_2}$ in the tropical secant graph by the two edges $F_{i_1, i_2, i_3} E_{i_2}$ (with weight $m_{F_{i_1, i_2, i_3} D_{i_2}}$), and $D_{i_2} E_{i_2}$ (with new weight $m_{D_{i_2} E_{i_2}} + m_{F_{i_1, i_2, i_3} D_{i_2}}$).
- (ii) $F_{\underline{a}} D_{i_0}$ and $F_{\underline{a}'} D_{i_0}$, where $\underline{a} = \{0, i_l, i_t\}$ and $\underline{a}' = \{0, i_u, i_t\}$, $l > t > u$. Furthermore, if m denotes the remaining index, then $i_m + i_t = i_l + i_u$ and $l > m > u$. Hence, $F_{\underline{a}} \in F_{\underline{a}'} D_{i_0}$, and we replace the edge $F_{\underline{a}'} D_{i_0}$ by the edges $F_{\underline{a}'} F_{\underline{a}}$ (with weight $m_{F_{\underline{a}'} D_{i_0}}$), and $F_{\underline{a}} D_{i_0}$ (endowed with the new weight $m_{F_{\underline{a}} D_{i_0}} + m_{F_{\underline{a}'} D_{i_0}}$).
- (iii) $F_{\underline{a}} D_{i_j}$ and $F_{\underline{a}'} D_{i_j}$, where $\underline{a} = \{0, i_j, i_t\}$ and $\underline{a}' = \{i_j, i_t, i_u\}$. Furthermore, if m is denotes the remaining index, then $i_u = i_t + i_m$. Assume $t < j$. Then, $F_{\underline{a}} \in F_{\underline{a}'} D_{i_j}$. We replace the edge $F_{\underline{a}'} D_{i_j}$ by the edges $F_{\underline{a}'} F_{\underline{a}}$ (with weight $m_{F_{\underline{a}'} D_{i_j}}$), and $F_{\underline{a}} D_{i_j}$ (changing its weight to $m_{F_{\underline{a}} D_{i_j}} + m_{F_{\underline{a}'} D_{i_j}}$). On the contrary, if $j < t$, then $F_{\underline{a}'} \in F_{\underline{a}} D_{i_j}$, and we replace the edge $F_{\underline{a}} D_{i_j}$ by the edges $F_{\underline{a}'} F_{\underline{a}}$ (with weight $m_{F_{\underline{a}} D_{i_j}}$) and $F_{\underline{a}'} D_{i_j}$ (with the new weight $m_{F_{\underline{a}} D_{i_j}} + m_{F_{\underline{a}'} D_{i_j}}$).

Proof. Let $\mathbf{l}_2 = (0, i_1, i_2, i_3, i_4)$. For any $\mu \geq 0$, each intersection point can be written as:

- (i) $\frac{i_2(i_4 - i_3)}{i_4} \cdot F_{\underline{a}} + \left(\frac{i_2(i_3 - i_2)}{i_4} + \mu\right) \cdot D_{i_2} = E_{i_2} + \mu \cdot D_{i_2} + \frac{-i_2}{i_4} \cdot \mathbf{l}_2$,
- (ii) $(i_l - i_m) \cdot F_{\underline{a}} + \mu \cdot D_{i_0} = (i_m - i_u) \cdot F_{\underline{a}'} + (i_t + \mu) \cdot D_{i_0} + (-i_m) \cdot \mathbf{1} + \mathbf{l}_2$,
- (iii) $i_m \cdot F_{\underline{a}} + (i_t + \mu) \cdot D_{i_j} = i_t \cdot F_{\underline{a}'} + (i_j + \mu) \cdot D_{i_j} + i_m \cdot \mathbf{1} - \mathbf{l}_2$.

□

From the previous lemma, we get a modification of the tropical secant graph possibly with internal crossings but without partial overlaps. Finally, using this new graph, we construct the Gröbner tropical secant graph as indicated by the following theorem, which we illustrate in Example 4.5.4:

Theorem 4.5.3. *The Gröbner tropical secant graph for the monomial curve $(1 : t^{i_1} : t^{i_2} : t^{i_3} : t^{i_4})$ in \mathbb{P}^4 may be obtained by adding finitely many internal crossings to the tropical secant graph (after modifications using Lemma 4.5.2). These internal crossings come in two types. The first one consists of points $P_{\underline{a}, j, \underline{a}', k} \in (\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{\underline{a}}, D_{i_j} \rangle) \cap (\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{\underline{a}'}, D_{i_k} \rangle)$, where:*

- (i) $\underline{a} = \{i_j, i_t\}, \underline{a}' = \{i_j, i_t, i_l, i_k\}$, where either $j > t > s > k$ and $u > s$ or $j < t < l < k$ and $u < l$;
- (ii) $\underline{a} = \{i_j, i_t\}, \underline{a}' = \{i_j, i_l, i_k\}$, where either $t > u > l > k, j > l$ and $i_j + i_u \geq i_l + i_t$ or $t < u < l < k, j < s$ and $i_j + i_u \leq i_l + i_t$;
- (iii) $\underline{a} = \{i_j, i_t\}, \underline{a}' = \{i_t, i_l, i_k\}$, where either $u > l > k, j > t > l$ and $i_j + i_l \geq i_u + i_t$ or $u < l < k, j < t < l$ and $i_j + i_l \leq i_u + i_t$;
- (iv) $\underline{a} = \{i_t, i_j, i_u, i_k\}, \underline{a}' = \{i_j, i_u, i_k\}$, where either $j > u > k, t > u > k$ and $t > l$ or $j < u < k, t < u < k$ and $t < l$;
- (v) $\underline{a} = \{i_j, i_u, i_k\}, \underline{a}' = \{i_u, i_t, i_k\}$, where either $j > s > t, j > u > k, u > t$ and $i_j + i_t \geq i_u + i_l$ or $j < l < t, j < u < k, u < t$ and $i_j + i_t \leq i_u + i_l$;
- (vi) $\underline{a} = \{i_u, i_j, i_k\}, \underline{a}' = \{i_j, i_k, i_t\}$, where $u > l > t, j > t, u > k, j > k, i_l + i_j \geq i_u + i_t \geq i_l + i_k$;
- (vii) $\underline{a} = \{i_j, i_u, i_k\}, \underline{a}' = \{i_k, i_t\}$, where either $j > u > l > t, u > k$ and $i_u + i_t \geq i_l + i_k$ or $j < u < l < t, u < k$ and $i_u + i_t \leq i_l + i_k$;
- (viii) $\underline{a} = \{i_j, i_l, i_u, i_k\}, \underline{a}' = \{i_u, i_t, i_k\}$, where either $j > l > u > k$ and $u > t$ or $j < l < u < k$ and $u < t$;
- (ix) $\underline{a} = \{i_l, i_j, i_u, i_k\}, \underline{a}' = \{i_j, i_u, i_t, i_k\}$, where $s > u > t$ and $j > u > k$;
- (x) $\underline{a} = \{i_4, i_3, i_2\}, \underline{a}' = \{i_2, i_1, 0\}$, $j = 4$ and $k = 0$;
- (xi) $\underline{a} = \{i_4, i_3\}, \underline{a}' = \{i_1, i_0\}$, $j = 4$ and $k = 0$.

The second class satisfies $P_{j, \underline{a}, k} \in (\mathcal{TC} + \mathbb{R}_{\geq 0} \langle E_{i_j}, E_{i_{j+1}} \rangle) \cap (\mathcal{TC} + \mathbb{R}_{\geq 0} \langle F_{\underline{a}}, D_{i_k} \rangle)$, where:

- (i) $\underline{a} = \{i_1, i_2, i_3\}$, $j = k = 2$ and $i_1(i_4 - i_2) \geq i_2(i_4 - i_3)$;
- (ii) $\underline{a} = \{i_1, i_2, i_4\}$, $j = k = 2$ and $i_3(i_2 - i_1) \geq i_2(i_4 - i_1)$;
- (iii) $\underline{a} = \{i_1, i_2, i_3\}$, $j = 1, k = 2$ and $i_4(i_2 - i_1) \geq i_2(i_3 - i_1)$;
- (iv) $\underline{a} = \{i_0, i_2, i_3\}$, $j = 1, k = 2$ and $i_3(i_2 - i_1) \geq i_2(i_4 - i_1)$.

Proof. The proof is very similar to the one we outlined for Theorem 4.4.13, so we only give the linear combination expressing each intersection point described in the statement. As usual, we let $\mathbf{l}_2 = (0, i_1, i_2, i_3, i_4) \in \Lambda \subset \mathbb{Z}^5$. The internal crossings of the first type are listed below:

- (i) For simplicity, assume $j > t > l > k$ and $u > l$ (if not, we multiply the expression by -1 to obtain the intersection point):

$$(i_t - i_l) \cdot F_{i_j, i_t} + (i_j - i_t) \cdot D_{i_j} = (i_u - i_l) \cdot F_{i_j, i_t, i_l, i_k} + (i_l - i_k) \cdot D_{i_k} - i_u \cdot \mathbf{1} + \mathbf{l}_2.$$

- (ii) Assume $i_j + i_u \geq i_l + i_t$. The intersection point is:

$$(i_t - i_u) \cdot F_{i_t, i_j} + (i_j + i_u - i_l - i_t) \cdot D_{i_j} = (i_u - i_l) \cdot F_{i_j, i_l, i_k} + (i_l - i_k) \cdot D_{i_k} - i_u \cdot \mathbf{1} + \mathbf{l}_2.$$

Note that if $i_j + i_u = i_l + i_t$, the intersection point is F_{i_j, i_t} .

- (iii) Assume $i_j + i_l \geq i_u + i_t$. The intersection point is:

$$(i_t - i_l) \cdot F_{i_t, i_j} + (i_j + i_l - i_u - i_t) \cdot D_{i_j} = (i_u - i_l) \cdot F_{i_t, i_l, i_k} + (i_l - i_k) \cdot D_{i_k} - i_u \cdot \mathbf{1} + \mathbf{l}_2.$$

Again, if $i_j + i_u = i_l + i_t$, the intersection point is F_{i_j, i_t} .

- (iv) Assume $u > k$. The intersection point is:

$$(i_t - i_l) \cdot F_{i_t, i_j, i_u, i_k} + (i_j - i_u) \cdot D_{i_j} = (i_t - i_u) \cdot F_{i_j, i_u, i_k} + (i_u - i_k) \cdot D_{i_k} - i_l \cdot \mathbf{1} + \mathbf{l}_2.$$

- (v) Assume $i_j + i_t \geq i_u + i_l$. The intersection point is:

$$(i_u - i_t) \cdot F_{i_j, i_u, i_k} + (i_j + i_t - i_u - i_l) \cdot D_{i_j} = (i_l - i_t) \cdot F_{i_u, i_t, i_k} + (i_u - i_k) \cdot D_{i_k} - i_l \cdot \mathbf{1} + \mathbf{l}_2.$$

If $i_j + i_t = i_u + i_l$, the intersection point is F_{i_j, i_u, i_k} .

- (vi) By symmetry, we can assume $j > k$. The intersection point is:

$$(i_r - i_l) \cdot F_{i_r, i_j, i_k} + (i_l + i_j - i_r - i_t) \cdot D_{i_j} = (i_l - i_t) \cdot F_{i_j, i_k, i_t} + (i_r + i_t - i_l - i_k) \cdot D_{i_k} - i_l \cdot \mathbf{1} + \mathbf{l}_2.$$

At most one of the equalities $i_l + i_j = i_i + i_t$, $i_r + i_t = i_l + i_k$ holds, and in this case, $F_{\underline{a}}$ or $F_{\underline{a}'}$ is the intersection point.

- (vii) Assume $i_r + i_t \geq i_l + i_k$. The intersection point is:

$$(i_r - i_l) \cdot F_{i_r, i_j, i_k} + (i_j - i_r) \cdot D_{i_j} = (i_l - i_t) \cdot F_{i_k, i_t} + (i_r + i_t - i_l - i_k) \cdot D_{i_k} - i_l \cdot \mathbf{1} + \mathbf{l}_2.$$

If $i_r + i_t = i_l + i_k$, the intersection point is F_{i_k, i_t} .

(viii) Assume $u > t$. The intersection point is:

$$(i_u - i_t) \cdot F_{i_j, i_l, i_u, i_k} + (i_j - i_l) \cdot D_{i_j} = (i_l - i_u) \cdot F_{i_u, i_t, i_k} + (i_u - i_k) \cdot D_{i_k} + (i_u - i_t - i_l) \cdot \mathbf{1} + \mathbf{l}_2.$$

(ix) The intersection point is:

$$(i_u - i_t) \cdot F_{i_j, i_l, i_u, i_k} + (i_j - i_u) \cdot D_{i_j} = (i_l - i_u) \cdot F_{i_j, i_u, i_t, i_k} + (i_u - i_k) \cdot D_{i_k} + (i_u - i_t - i_l) \cdot \mathbf{1} + \mathbf{l}_2.$$

(x) The intersection point is:

$$(i_2 - i_1) \cdot F_{i_4, i_3, i_2} + (i_4 - i_3) \cdot D_{i_4} = (i_3 - i_2) \cdot F_{i_2, i_1, 0} + i_1 \cdot D_0 + (i_2 - i_3 - i_1) \cdot \mathbf{1} + \mathbf{l}_2.$$

(xi) The intersection point is:

$$(i_3 - i_2) \cdot F_{i_4, i_3} + (i_4 - i_3) \cdot D_{i_4} = (i_2 - i_1) \cdot F_{i_1, i_0} + i_1 \cdot D_{i_0} - i_2 \cdot \mathbf{1} + \mathbf{l}_2.$$

We list the internal crossings of the second type:

(i) The intersection point is:

$$\frac{i_3 - i_1}{i_3 - i_2} \cdot E_{i_2} + \frac{i_1(i_4 - i_2) - i_2(i_4 - i_3)}{(i_4 - i_3)(i_3 - i_2)} \cdot E_{i_3} = i_1 \cdot F_{i_1, i_2, i_3} + (i_2 - i_1) \cdot D_{i_2} + \frac{i_1}{i_4 - i_3} \cdot \mathbf{l}_2.$$

The positivity of the coefficient of E_{i_3} gives the inequality constraint in the statement. Note that if $i_1(i_4 - i_2) = i_2(i_4 - i_3)$ then the crossing point is E_{i_2} , which is already a node in the graph, but in this case it lies in the interior of the edge $F_{i_1, i_2, i_3} D_{i_2}$.

(ii) The internal crossing is:

$$\frac{i_3}{i_3 - i_2} \cdot E_{i_2} + \frac{i_3(i_2 - i_1) - i_2(i_4 - i_1)}{(i_4 - i_3)(i_3 - i_2)} \cdot E_{i_3} = i_1 \cdot F_{i_1, i_2, i_4} + i_2 \cdot D_{i_2} - \frac{i_1}{i_4 - i_3} \cdot \mathbf{l}_2,$$

with the positivity constraint for the coefficient of E_{i_3} . If $i_3(i_2 - i_1) = i_2(i_4 - i_1)$ the crossing point is E_{i_2} and it lies in the interior of the edge $F_{i_1, i_2, i_4} D_{i_2}$.

(iii) The intersection point is:

$$\frac{i_4(i_2 - i_1) - i_2(i_3 - i_1)}{i_1(i_2 - i_1)} \cdot E_{i_1} + \frac{i_3 - i_1}{i_2 - i_1} \cdot E_{i_2} = (i_4 - i_3) \cdot F_{i_1, i_2, i_3} + (i_3 - i_2) \cdot D_{i_2} + \mathbf{l}_2,$$

with the positivity constraint for the coefficient of E_{i_1} . If $i_4(i_2 - i_1) = i_2(i_3 - i_1)$ the crossing point is E_{i_2} and it lies in the interior of the edge $F_{i_1, i_2, i_3} D_{i_2}$.

(iv) The intersection point is:

$$\frac{i_3(i_2 - i_1) - i_2(i_4 - i_1)}{i_1(i_2 - i_1)} \cdot E_{i_1} + \frac{i_4 - i_1}{i_2 - i_1} \cdot E_{i_2} = (i_4 - i_3) \cdot F_{i_0, i_2, i_3} + (i_3 - i_2) \cdot D_{i_2} - (i_4 - i_3) \cdot \mathbf{1} + \mathbf{l}_2,$$

with positivity constraint for the coefficient of E_{i_1} . If $i_3(i_2 - i_1) = i_2(i_4 - i_1)$ the crossing point is E_{i_2} and it lies in the interior of the edge $F_{i_0, i_2, i_3} D_{i_2}$. \square

Example 4.5.4 (Example 4.5.1, revisited). As we saw, there are no new overlaps of edges in the tropical secant graph of the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$, so Lemma 4.5.2 does not apply to this example. Using Theorem 4.5.3, we can explain the six new green nodes we added to build the Gröbner tropical secant graph (Figure 4.7). They come from the six internal crossings of the first type between the edges $F_{\underline{a}} D_{i_j}$ and $F_{\underline{a}'} D_{i_k}$, where: $\underline{a} = \{55, 78\}, j = 78, \underline{a}' = \{0, 30, 78\}, k = 0$ (case (ii)); $\underline{a} = \{55, 78\}, j = 78, \underline{a}' = \{0, 30\}, k = 0$ (case (xi)); $\underline{a} = \{45, 78\}, j = 78, \underline{a}' = \{0, 30, 45, 78\}, k = 0$ (case (i)); $\underline{a} = \{0, 30, 45, 78\}, j = 45, \underline{a}' = \{0, 30, 45\}, k = 0$ (case (iv)); $\underline{a} = \{45, 78\}, j = 78, \underline{a}' = \{0, 30, 45\}, k = 0$ (case (iii)); and $\underline{a} = \{45, 55\}, j = 55, \underline{a}' = \{0, 30, 45, 55\}, k = 0$ (case (i)). \diamond

4.6 Chow polytopes, tropical secants of lines, toric arrangements and beyond

The implicitization methods discussed in the previous section can be generalized to monomial curves in higher dimensional projective spaces, where the first secant no longer has codimension one. In this case, one can recover the *Chow polytope* of the secant variety by a generalization of the ray-shooting algorithm, known as the *orthant-shooting* algorithm [24, Theorem 2.2]. Instead of shooting rays, we shoot orthants (i.e. cones spanned by subsets of the canonical basis of \mathbb{R}^{n+1}) of dimension equal to $n - 3$ (the codimension of our variety). A similar formula to the one described in Theorem 1.4.1 gives us the vertex of the Chow polytope associated to any input objective vector. However, it is not as easy to given an analog to the walking algorithm. The difficulty comes from the fact that, a priori, there is no canonical way of walking along the complement of the tropical variety. Recently, Alex Fink has developed a method to reduce the computation of the Chow polytope to the codimension one setting, based on the *orthant shooting algorithm* [38, 37]. His approach allows us to use the techniques discussed for the secant hypersurface case [38]. Thanks to his results, existing software from [18] presented in Chapter 3 can be used in higher codimensions, for example, for rational normal curves in \mathbb{P}^n .

Before giving a numerical example, we explain the method presented in [38] for computing Chow polytopes. We define a map, called the *Chow map* ch , taking a tropical variety $\mathcal{T}X$ of dimension d in $\mathbb{T}\mathbb{P}^n$ to its *tropical Chow hypersurface*, $ch(\mathcal{T}X) = \mathcal{T}X \boxplus \mathcal{L}_{n-d-1}^{\text{refl}}$ [38,

Definition 5.2]. The set $ch(\mathcal{T}X)$ will be precisely the tropicalization of a hypersurface in \mathbb{P}^n whose Newton polytope equals the desired Chow polytope of X [38, Theorem 5.1]. Thus, we can use the machinery of Sections 1.4 and 3.4 to find this polytope.

We now describe the tropical hypersurface $ch(\mathcal{T}X)$ as the union of weighted $(n-1)$ -dimensional cones in $\mathbb{T}\mathbb{P}^n$. First, we pick all weighted maximal d -dimensional cones (σ, m_σ) in $\mathcal{T}X \subset \mathbb{T}\mathbb{P}^n$, and all $n-d-1$ cones C_J generated by subsets of $n-d-1$ vectors among $\{-e_0, -e_1, \dots, -e_n\}$, i.e. the negative of the elements in the canonical basis of \mathbb{R}^{n+1} . The subscript J indicates the indices of the vectors chosen from this basis. For simplicity, assume that each cone σ is simplicial and it is spanned by integer vectors $\{v_1^\sigma, \dots, v_d^\sigma\}$ in $\mathbb{T}\mathbb{P}^n$. Secondly, we take Minkowski sum of σ and C_J for every possible pair, and we check if the cone $\sigma + C_J$ in $\mathbb{T}\mathbb{P}^n$ has codimension 1. If so, this means that the matrix

$$A := \left(v_1^\sigma \mid \dots \mid v_d^\sigma \mid -e_{j_1} \mid \dots \mid -e_{j_{n-d-1}} \right),$$

is full dimensional, for $J = \{j_1, \dots, j_{n-d-1}\}$. We assign to the new cone $\sigma + C_J$ the weight

$$m_{\sigma+C_J} := m_\sigma \cdot \gcd(\text{maximal } (n-d-1) \times (n-d-1)\text{-minors of } A). \quad (4.6)$$

If the matrix A is not of full rank, we discard the cone $\sigma + C_J$ from the list of valid combinations and we move on to the next pair. The set $ch(\mathcal{T}X)$ will be the union of the valid combinations, with weights given by formula (4.6).

Example 4.6.1. The canonical example of a monomial curve in \mathbb{P}^n is the *rational normal curve* $t \mapsto (1 : t : t^2 : \dots : t^n)$. These curves and their secants have been extensively studied in the past. They are known to be determinantal varieties ([46, Proposition 9.7],[9, Proposition 2.2]) defined by the 3×3 -minors of the $j \times (n-j+2)$ *Hankel* matrix:

$$x_j := \begin{pmatrix} x_1 & x_2 & x_3 & \dots & x_{n-j+2} \\ x_2 & x_3 & \dots & \dots & \dots \\ x_3 & \dots & \dots & \dots & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ x_j & x_{j+1} & \dots & x_n & x_{n+1} \end{pmatrix}.$$

The ideal generated by the 3×3 -minors of this matrix is independent of the index j [13, Corollary 2.2] and it is a set-theoretic complete intersection [96] of degree $\binom{n-1}{2}$ [13].

Using Theorem 4.4.6, we compute the tropical secant graph of the rational normal curve in \mathbb{P}^n . It has $n+1$ nodes $D_j = e_j$ ($0 \leq j \leq n$), $n-3$ nodes $E_j = (0, 1, 2, \dots, j, \dots, j)$ ($2 \leq j \leq n-2$) and $(\lfloor n/2 \rfloor + 1)(\lfloor n/2 \rfloor - 2)/2$ nodes $F_{\underline{a}}$, where \underline{a} is the subset $\{s+k \cdot r : k \in \mathbb{N}\} \cap \{0, \dots, n\}$, with $0 \leq s < r$ and $1 < r < \lfloor n/2 \rfloor$. In addition, it has one or two nodes $F_{\underline{a}}$ where $r = \lfloor n/2 \rfloor$ and $s = 0$ (for n even), or $s = 0, 1$ (for n odd).

The graph has $2n-5$ edges labeled $E_j E_{j+1}$, ($2 \leq j \leq n-3$), $E_j D_j$ ($2 \leq j \leq n-2$), $D_1 E_2$, $E_{n-2} D_{n-1}$, with weight one. It also has edges $F_{\underline{a},s} D_j$ ($j \in \underline{a}, s$), with weight $\varphi(r)/2$

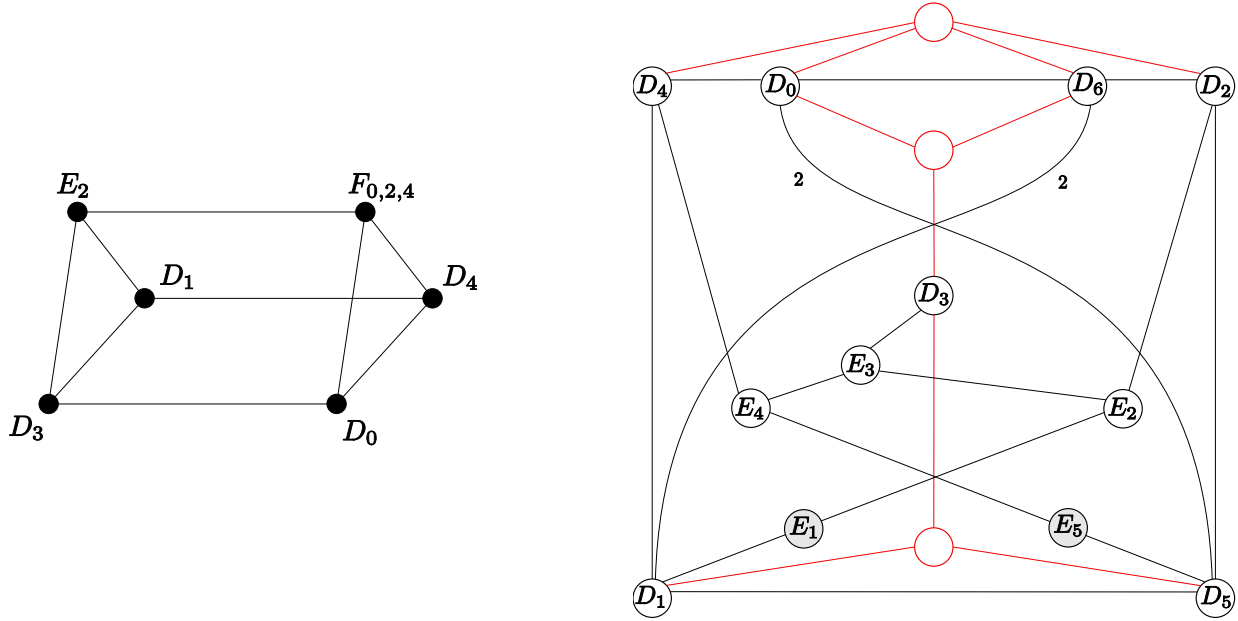


Figure 4.8: The Gröbner tropical secant graphs of the rational normal curves in \mathbb{P}^4 and \mathbb{P}^6 .

if $r > 2$, or weight 1 if $r = 2$. In addition, it has $\lfloor n/2 \rfloor (n + 3 + \lfloor n/2 \rfloor) / 2$ edges $D_j D_{j+r}$ with $0 \leq j < n - r$ and $r > \lfloor n/2 \rfloor$ with weight $\varphi(r)/2$, and $\lfloor n/2 \rfloor - 2$ edges $D_j D_{j+\lfloor n/2 \rfloor}$ ($2 \leq j \leq \lfloor n/2 \rfloor$) with weight $\varphi(\lfloor n/2 \rfloor) / 2$. Finally, if n is even, the edge $D_1 D_{1+n/2}$ has weight $\varphi(n/2) / 2$ if $n > 2$ and weight 1 if $n = 2$.

We illustrate the previous construction in the case $n = 4$. After removing the bivalent node $D_2 = e_2$, we are left with a graph with six nodes $D_0 = e_0$, $D_1 = e_1$, $D_3 = e_3$, $D_4 = e_4$, $E_2 = (0, 1, 2, 2, 2)$ and $F_{0,2,4} = (1, 0, 1, 0, 1)$ and nine edges, all with trivial weight 1. This graph is the 1-skeleton of a bipyramid (Figure 4.8).

Using ray-shooting and walking algorithms, we see that the equation has multidegree $(3, 6)$ with respect to the lattice $\Lambda = \mathbb{Z}\langle \mathbf{1}, (0, 1, 2, 3, 4) \rangle$ and its Newton polytope has vertices $(0, 0, 3, 0, 0)$, $(0, 1, 1, 1, 0)$, $(1, 0, 0, 2, 0)$, $(0, 2, 0, 0, 1)$, $(1, 0, 1, 0, 1)$, and f -vector $(5, 9, 6)$. This graph is precisely the tropical discriminant of the Veronese surface, regarded as the projectivization of the variety of all symmetric 3×3 matrices of rank at most one [24, Example 4.4]. Its defining equation is a dehomogenization of the Hankel 3×3 -determinant.

The case $n = 6$ was computed in [6, Example 20]. Its defining ideal is generated by the 3×3 -minors of the 4×4 -Hankel matrix. After changing the signs of the rays obtained by **Gfan** (to agree with our *min* convention) and reducing modulo the lattice Λ , we see that our construction (Figure 4.8) matches theirs for all rays except for the ones corresponding to three bivalent node E_1, E_5 and one that is absent in our graph (nodes 2, 3 and 15 in the notation of [6]). We follow our convention from Example 4.5.1 and keep the red nodes $F_{\underline{a}}$ unlabeled. All weights in our graph equal 1, except for two edges with weight 2 (indicated in Figure 4.8). From this computation, we confirm that the Gröbner tropical secant graph

coincides with the tropical secant graph.

We now use the method developed in [38] for computing the Chow polytope of the secant variety of the rational normal curve in \mathbb{P}^6 . In this case, the polytope has 289 vertices and f -vector $(289, 897, 981, 442, 71)$. All vertices have multi homogeneity $(30, 90)$, which matches the formula $\deg ch(X) = \text{codim}X \cdot \deg X$ from [38, Lemma 3.4], since $30 = 3 \cdot \binom{5}{2}$.

Changing the grading to reflect the torus action by the exponent vector $(0, 1, 2, \dots, n)$ rather than the action by the all-ones vector, will give us the weighted projective space $\mathbb{P}_{(0,1,2,\dots,n)}^n$ as the ambient space, rather than the usual \mathbb{P}^n . This new setting will result in the formula $\deg' ch(\text{Sec}^1(C)) = (n - 3) \cdot \deg' \text{Sec}^1(C)$ connecting the degree \deg' of $ch(\text{Sec}^1(C))$ with respect to this exponent vector to the codimension and the degree of the secant variety of C inside this new toric variety. We believe the exact same formula holds in this new setting. \diamond

We now switch gears to study the set of all tropical lines between points in the tropicalization of a monomial projective curve. We aim to highlight the differences between this set and the tropicalization of the first secant variety of the same curve. By definition, a tropical line segment between two points in the tropical curve \mathcal{TC} is the loci of all points obtained as the coordinatewise minima (i.e. tropical addition) of two fixed points in the classical plane spanned by the lattice $\Lambda = \langle \mathbf{1}, (0, i_1, \dots, i_n) \rangle$. We interpret this as the line spanned by the vector $(0, i_1, \dots, i_n)$ in \mathbb{TP}^n . The set of all tropical lines between points in \mathcal{TC} is often denoted by $S^1(\mathcal{TC})$ and it is called the *first tropical secant variety* of the line $\mathcal{TC} \subset \mathbb{TP}^n$. Since $S^1(\mathcal{TC})$ is the image of the tropicalization of the secant map ϕ from (4.3), we know it is contained in the tropicalization of the image of ϕ , hence $S^1(\mathcal{TC})$ is contained in $\mathcal{TSec}^1(C)$. These two tropical sets have been compared and their rich combinatorial structures studied by many, including Develin and Draisma [22, 25]. In particular, by [22, Corollary 2.2], we know that $S^1(\mathcal{TC})$ is a cone from \mathcal{TC} over a polytopal complex, called the *first tropical secant complex* of \mathcal{TC} .

Each point in $S^1(\mathcal{TC}) \subset \mathbb{R}^{n+1}$ may be thought of as a height vector for a configuration of points $\{0, i_1, \dots, i_n\}$ on \mathbb{R} which induces a regular subdivision of the convex hull defined by these $n + 1$ points. The faces of this polytopal complex correspond to regular subdivisions such that two facets cover all $n + 1$ points. These faces are ordered by refinements of the subdivisions. Since by assumption, our exponent vector has distinct coordinates, the classical line \mathcal{TC} is *generic* in the sense of Develin, and [22, Theorem 3.1] gives a very nice characterization of this complex. It is precisely the set of lower faces of the cyclic polytope $C(2, n - 1)$, defined as the convex hull of $n - 1$ generic points in the parabola $\{y = x^2\} \subset \mathbb{R}^2$. It is immediate to check that this complex is a chain graph with $n - 1$ vertices. Figure 4.9 illustrates this construction for a generic classical line in \mathbb{TP}^4 .

Proposition 4.6.2. *The first tropical secant complex of the tropicalization of the monomial curve $(1 : t^{i_1} : \dots : t^{i_n})$ in \mathbb{P}^n with $0 < i_1 < \dots < i_n$ is a chain graph in \mathbb{R}^{n-1} with $n - 1$ vertices*




Figure 4.9: The first tropical secant complex of the line $\mathbb{R}\langle(0, i_1, i_2, i_3, i_4)\rangle$ in \mathbb{TP}^4 .

$v^{(1)}, \dots, v^{(n-1)}$:

$$v^{(k)} = \sum_{j \leq k} i_j e_j + \sum_{k < j < n} \frac{i_k}{i_n - i_k} (i_n - i_j) e_j \quad k = 1, \dots, n - 1.$$

Moreover, $v^{(k)}$ corresponds to the regular subdivision of the configuration $\{0, i_1, \dots, i_n\}$ with exactly two facets $\{0, i_1, \dots, i_k\}$ and $\{i_k, \dots, i_n\}$. It is embedded as a height vector, where the points 0 and i_n have height zero, the point i_k has height i_k and the remaining points lie in the interior of the two facets.

Example 4.6.3. We describe the first secant complex of the curve $(1 : t^{30} : t^{45} : t^{55} : t^{78})$, as shown in Figure 4.9. It consists of three nodes $v^{(1)} = (0, 30, \frac{165}{8}, \frac{115}{8}, 0)$, $v^{(2)} = (0, 30, 45, \frac{345}{11}, 0)$ and $v^{(3)} = (0, 30, 45, 55, 0)$ and two edges $v^{(1)}v^{(2)}$ and $v^{(2)}v^{(3)}$. By taking the cone from the linear space \mathcal{TC} over this complex, we get the first tropical secant variety of the line $\mathbb{R}\langle(0, 30, 45, 55, 78)\rangle$. \diamond

We now describe $S^1(\mathcal{TC})$ as a subgraph of the tropical secant graph, and we show the containment of $S^1(\mathcal{TC})$ in $\mathcal{T}Sec^1(C)$ for a monomial curve C is strict in the generic case.

Proposition 4.6.4. *In the notation of Definition 4.4.4, the first tropical secant complex of the tropicalization of the monomial curve $(1 : t^{i_1} : \dots : t^{i_n})$ is the chain subgraph of the tropical secant graph with nodes $E_{i_1}, \dots, E_{i_{n-1}}$.*

Proof. In the notation of Proposition 4.6.2, $v^{(k)}$ and E_{i_k} generate the same ray in $\mathbb{R} \otimes \Lambda$ because

$$v^{(k)} = \frac{-i_k}{i_n - i_k} \cdot (0, i_1, \dots, i_n) + \frac{i_n}{i_n - i_k} \cdot E_{i_k} \quad \text{for } k = 1, \dots, n - 1.$$

The result follows immediately. \square

We finish this section by discussing briefly the relationship between our tropical secant surface graph and compactifications of toric arrangements. As we saw in Section 4.3, the process of geometric tropicalization of surfaces gives a recipe for finding a suitable compactification (a *tropical compactification*) of a parametric surface inside the torus \mathbb{T}^{n+1} , or, equivalently, of the complement of $n+1$ divisors in the torus \mathbb{T}^2 . In our setting, these divisors were the $n+1$ binomial curves $\{w^{i_j} - \lambda = 0\}$ ($0 \leq j \leq n$). It is well-known that this toric arrangement can be compactified by a wonderful model in the sense of De Concini-Procesi [19]. Recently, L. Moci [71] has given an explicit compactification for binomial arrangements and, at first glance, his techniques are very similar to the ones we discussed in Section 4.3. We hope to make this connection more explicit in the near future.

Chapter 5

Implicitization of surfaces via geometric tropicalization

In this chapter we describe tropical methods for implicitization of surfaces in three-space. We construct the corresponding tropical surfaces via the theory of geometric tropicalization, which we enrich with a formula for computing tropical multiplicities of regular points. This formula was used in Chapter 4 for obtaining a tropical surface in \mathbb{T}^n parameterized by binomials.

5.1 Introduction

In this chapter we further develop on the techniques described in Chapter 4 for calculating tropical surfaces given a polynomial parameterization of rational surfaces in \mathbb{T}^n . Such methods are known as *tropical elimination* and *tropical implicitization* and have been used recently for computations going beyond the power of classical elimination tools, including multidimensional resultants and Gröbner bases. Successful applications of tropical elimination techniques were presented in Chapters 3 and 4.

The foundations of tropical elimination theory were laid in the work of Sturmfels, Tevelev and Yu [91, 92], including the software implementation **TrIm** for **tropical implicitization** [93]. However, their techniques apply only for generic varieties, that is surfaces parameterized by Laurent polynomials which are generic relative to their support. The main goal of this chapter is to make this genericity condition precise and to introduce tropical implicitization methods for non-generic varieties. Even though these techniques were outlined in [91, Section 5], they rely on resolution of singularities, known to be hard to use in practice. Moreover, the question of computing tropical multiplicities in the non-generic case was not treated in [91]. These numbers are essential for recovering information about algebraic varieties from their tropical counterparts. Our main contribution to the theory is an explicit combinatorial formula for computing tropical multiplicities from the geometry of algebraic varieties (The-

orem 5.2.4 and Corollary 5.2.5), complementing the set-theoretic results of Hacking, Keel and Tevelev [45, §10]. The surface case is further discussed in Proposition 5.2.10.

In Section 5.2 we describe the basics on tropical elimination and tropical implicitization theory and show how to derive tropical implicitization methods from *geometric tropicalization*. Particular attention is devoted to the foundational paper on due to Hacking, Keel and Tevelev [45] and on ways of weakening the required hypothesis to apply their techniques in practical calculations.

As mentioned in Section 1.3, the crux of geometric tropicalization is to read off the tropicalization of an algebraic variety directly from the combinatorics of its boundary in a tropical compactification. To do so, this boundary is required to have “simple normal crossings,” that is, to behave locally like an arrangement of coordinate hyperplanes. These ideal conditions rarely hold in concrete examples. Thus, we need to overcome this difficulty by resolving these “bad crossings” and singularities with other tools, such as blow-ups. This is not an easy task. Nonetheless, as examples suggest and [91, Proposition 5.3] confirms, the full power of resolution of singularities is not required to obtain a set-theoretic description of tropicalized *surfaces*. More precisely, by looking at things from the tropical side, the right condition is a boundary with “combinatorial normal crossings:” k irreducible divisors intersect in codimension k (Conjecture 5.2.8). Proposition 5.2.10 addresses this question, which was previously stated in [91, Proposition 5.3] and [94, Theorem 1.2].

In Section 5.3 we focus our attention on tropical implicitization of generic surfaces parameterized by polynomial maps. First, we use the given map to embed our rational surface in a smooth projective toric variety. The genericity condition is chosen to certify the correctness of the method as well as to simplify the construction of the smooth ambient toric variety and compute tropical multiplicities via mixed volumes. Our approach allows to weaken this genericity condition and still be able to compute tropical multiplicities from the input map. As a byproduct, in Theorem 5.3.1 we show that the smoothness condition on the ambient space is unnecessary to obtain the tropical surface as a weighted graph. We call this graph the *tropical surface graph* of the rational surface. We illustrate our approach with several numerical examples of surfaces in \mathbb{C}^3 . These examples will be revisited in Section 5.4 to highlight the differences between generic and non-generic surfaces.

In Section 5.4 we discuss on tropical implicitization of non-generic surfaces. We start by clarifying what we mean by *special* surfaces. Then, we describe a general procedure to obtain the tropical surface graphs of special surfaces. Singularities coming from excessive intersections are the main obstruction to apply the methods of Section 5.3 in the special case. To fix this bad behavior, there are several ways to proceed. The first approach is to embed our algebraic object in a smooth ambient projective variety, as in the generic case, and then modify the latter by blow-ups of points on curves. Unfortunately, these singular points need not be torus invariant, so toric blow-ups cannot be used to resolve them. For this reason, we will not follow this approach in practice.

Inspired by Chapter 4, we translate our task to the one of finding a tropical compactification of an arrangement of n plane curves in \mathbb{T}^2 . This arrangement coincides with the

preimage of the torus \mathbb{T}^n under the given polynomial parameterization. We compactify this open subset of \mathbb{T}^2 inside \mathbb{P}^2 and then resolve the singular points in its boundary, by ordinary blow-ups. By applying geometric tropicalization methods to this compact set we obtain an abstract weighted graph which reflects the combinatorics of the tropical surfaces we wish to compute. Our given parameterization allows us to realize this abstract graph in \mathbb{R}^n by assigning coordinates to all of its nodes. The resulting graph is precisely the tropical surface graph of our rational surface.

We end this chapter with some remarks and open questions, discussed in Section 5.5. As our running examples illustrate, rational surfaces in \mathbb{T}^3 serve as a nice test case to explore tropical implicitization techniques. In this setting, these methods require to analyze the combinatorics of the intersection points of three curves in \mathbb{T}^2 , and the local behavior of each curve at these points. Using this data we discuss methods to classify tropical surface graphs in three-space, and to extend the theory to fields with non-trivial valuation. This leads us to Berkovich spaces. Even though the theory of tropical implicitization is at an early stage and it is still evolving, we expect Theorem 5.2.4 to become a valuable tool for future applications.

5.2 Tropical elimination and tropical implicitization

In this section we discuss tropical elimination and implicitization theory from the perspective of geometric tropicalization. Particular attention is devoted to the surface case, which is treated in more detail in subsequent sections. The exposition is based on [45], [91, Section 5] and [92]. Our main result in this section is a formula for computing tropical multiplicities, complementing the set-theoretic description of tropical varieties using geometric tropicalization (Theorem 5.2.4). This formula was announced without proof in [17]. It was used in Chapter 4 to compute the tropical secant surface graph of a monomial projective curve with arbitrary set of exponents.

Notation 5.2.1. *Given a variety $X \subset \mathbb{T}^r$ and a projective toric variety $\mathbb{P}(\Sigma)$ with torus \mathbb{T}^r associated to a fan Σ , we denote by $X(\Sigma)$ the closure of X in $\mathbb{P}(\Sigma)$ as in (1.8).*

Classical elimination is the art of computing the defining ideal of the projection α of a variety in \mathbb{T}^r to a coordinate subspace \mathbb{T}^n . Here, we assume that α is a monomial map associated to a matrix $A \in \mathbb{Z}^{n \times r}$ as in Section 1.1.1:

$$\begin{array}{ccc} X & & \\ \downarrow & \searrow & \\ \mathbb{T}^r & \xrightarrow{\alpha} & \overline{\alpha(X)} \subset \mathbb{T}^n. \end{array} \tag{5.1}$$

By definition, the ideal $I(\overline{\alpha(X)})$ is the inverse image of the ideal $I(X)$ under the monomial map $\alpha^*: \mathbb{C}[y_1^{\pm 1}, \dots, y_n^{\pm 1}] \rightarrow \mathbb{C}[x_1^{\pm 1}, \dots, x_r^{\pm 1}]$ where $y_i \mapsto \prod_{j=1}^r x_j^{a_{ij}}$ for all $i = 1, \dots, n$.

As we discussed in Section 1.1.1, we can tropicalize the diagram (5.1) and use the functoriality of tropicalization on monomial maps and subvarieties of tori to obtain a diagram for *tropical elimination*:

$$\begin{array}{ccc} \mathcal{T}X & & \\ \downarrow & \searrow & \\ \mathbb{R}^r & \xrightarrow{A} & \mathcal{T}(\alpha(X)) = A(\mathcal{T}X) \subset \mathbb{R}^n. \end{array}$$

For simplicity, assume that $\alpha|_X$ has finite generic fibers of size δ . Under this condition, Theorem 1.1.10 gives a way of computing multiplicities on $\mathcal{T}(\alpha(X))$ from the multiplicities on $\mathcal{T}X$, the degree δ and the fibers of A , known as the push-forward formula for multiplicities.

Classical implicitization is the art of transforming a polynomial parametric representation of a rational algebraic variety into its implicit representation as the zero set of finitely many polynomials. It is a special instance of elimination, where our variety is the graph of a parameterization given by n Laurent polynomials $\mathbf{f} = (f_1, \dots, f_n): X \subset \mathbb{T}^d \rightarrow \mathbb{T}^n$, i.e.

$$X' := \text{Graph}(\mathbf{f}) := \{(x, \mathbf{f}(x)) : x \in X\} \subset \mathbb{T}^{d+n}, \quad (5.2)$$

and the monomial map α is the projection to the last n coordinates of \mathbb{T}^{d+n} (the image coordinates):

$$\begin{array}{ccc} \mathbb{T}^d \supseteq X & \xrightarrow{\mathbf{f}} & \overline{\mathbf{f}(X)} \subset \mathbb{T}^n \\ \downarrow (id, \mathbf{f}) & & \downarrow id \\ \mathbb{T}^{d+n} \supseteq X' & \xrightarrow{\alpha} & \mathbb{T}^n. \end{array} \quad (5.3)$$

Tropical implicitization will aim to compute the tropical variety $\mathcal{T}\overline{\mathbf{f}(X)}$ from the geometry of X and the map \mathbf{f} . Our main tool is the theory of geometric tropicalization due to Hacking, Keel and Tevelev, which we introduced in Section 1.3. In what follows, we explain how to compute $\mathcal{T}\overline{\mathbf{f}(X)}$ from $\mathcal{T}X$ and the projection α . For simplicity, we assume that X is a dense open set in \mathbb{T}^d and that \mathbf{f} is a generically finite *polynomial* map on X of degree δ . Similar results will hold for finite rational maps.

From now on, we fix $Y = \overline{\mathbf{f}(X)} \subset \mathbb{T}^n$. The variety $X' \subset \mathbb{T}^{d+n}$ from (5.2) is a *complete intersection* defined by the ideal

$$I = (y_1 - f_1(x), \dots, y_n - f_n(x)) \subset \mathbb{C}[x_1^{\pm 1}, \dots, x_d^{\pm 1}, y_1^{\pm 1}, \dots, y_n^{\pm 1}]$$

and it is isomorphic to $X \subset \mathbb{T}^d$. Following Theorem 1.1.10, the tropical variety $\mathcal{T}Y$ equals the image of $\mathcal{T}X'$ under the linear projection $(0 \mid I_d) \subset \mathbb{Z}^{d \times (d+n)}$ and its multiplicities can be obtained using the push-forward formula (1.4) of Sturmfels and Tevelev. Thus, tropical implicitization reduces to the task of computing $\mathcal{T}X'$, which we do by geometric

tropicalization. Since $X \subset \mathbb{T}^d$ and $X' \subset \mathbb{T}^{n+d}$ are isomorphic and the construction of their tropicalization is independent of the auxiliary tropical compactification, we can choose to compactify X or X' to build the abstract simplicial complex that encodes the tropical variety. The realization of the boundary intersection complex of the compactification in either \mathbb{R}^d or \mathbb{R}^{n+d} will reflect our choice. The next result justifies our strategy:

Theorem 5.2.2 ([91, Corollary 2.9]). *Let $X \subset \mathbb{T}^d$ be a dense open subset and $\mathbf{f}: X \rightarrow \mathbb{T}^n$ a generically finite Laurent polynomial map. Let Y be the Zariski closure of the image of \mathbf{f} . Let $X \subset \overline{X}$ be a smooth tropical compactification whose boundary $D = D_1 \cup \dots \cup D_k$ has simple normal crossings. Let $\Delta_{\overline{X}, D}$ be the intersection complex of the boundary D and consider the commutative diagram (5.3). For each component D_k of D , define*

$$[\tilde{D}_k] := \text{val}_{D_k}(\chi \circ \mathbf{f}) = (\text{val}_{D_k}(\chi_1 \circ \mathbf{f}), \dots, \text{val}_{D_k}(\chi_n \circ \mathbf{f})) \in \mathbb{Z}^n.$$

For each cell $\sigma \in \Delta_{\overline{X}, D}$ define its realization in \mathbb{R}^n as the semigroup $[\tilde{\sigma}]$ spanned by $\{[\tilde{D}_k] : k \in \sigma\}$. Then, the tropical variety $\mathcal{T}Y$ equals the set

$$\mathcal{T}Y = \bigcup_{\sigma \in \Delta_{\overline{X}, D}^{\text{top}}} \mathbb{R}_{\geq 0}[\tilde{\sigma}],$$

where the superscript top indicates that we only consider top-dimensional cells.

Remark 5.2.3. It is in this sense that the simplicial boundary intersection complex of \overline{X} is “push-forward” via the map $\mathbf{f}: X \rightarrow Y$ to give a simplicial complex associated to Y (Section 1.3). The key fact in the proof of this result is that \mathbf{f} induces a map on function fields $\mathbf{f}^\#: \mathbb{C}(Y) \hookrightarrow \mathbb{C}(X)$. Since the field $\mathbb{C}(X)$ is a finite extension of $\mathbb{C}(Y)$ of degree δ , we can always extend any discrete valuation on $\mathbb{C}(Y)$ to a discrete valuation on $\mathbb{C}(X)$ via the map $\mathbf{f}^\#$. Likewise, valuations on $\mathbb{C}(X)$ can be restricted to $\mathbb{C}(Y)$. The realization of each node D_k by the lattice point $[\tilde{D}_k]$ corresponds to the image of the realization of D_k in \mathbb{R}^{n+d} under the linear map associated to the monomial map α from (5.3). This highlights the deep connections between tropical implicitization and homomorphisms of tori.

Before continuing our discussion on tropical implicitization, we provide a combinatorial formula for computing multiplicities from the geometric perspective, complementing the set-theoretic result of geometric tropicalization due to Hacking, Keel and Tevelev. In the complete intersection case, our theorem is equivalent to [91, Theorem 4.6]. The lattice index factor in our formula accounts for the change in the lattice structure from \mathbb{Z}^d to $\mathbb{R}[\sigma] \cap \mathbb{Z}^d$.

Theorem 5.2.4. *Let \overline{X} be a compactification of an s -dimensional variety $X \subset \mathbb{T}^d$ whose boundary has simple normal crossings. In the notation of Theorem 1.3.1, the multiplicity of a regular point w in the tropical variety $\mathcal{T}X$ equals*

$$m_w = \sum_{\substack{\sigma \in \Delta_{\overline{X}, D}^{\text{top}} \\ w \in \mathbb{R}_{\geq 0}[\sigma]}} (D_{k_1} \cdot \dots \cdot D_{k_s}) \text{index}(\mathbb{R}[\sigma] \cap \mathbb{Z}^d, \mathbb{Z}[\sigma]), \quad (5.4)$$

where $D_{k_1} \cdots D_{k_s}$ denotes the intersection number of these s divisors and we sum over all $(s - 1)$ -dimensional cells σ in $\Delta_{\bar{X}, D}$ whose associated rational cone $\mathbb{R}_{\geq 0}[\sigma]$ contains w .

Proof. Since our question is local, it suffices to show the result holds for the choice of a compactification whose underlying boundary intersection complex gives a rational polyhedral fan in \mathbb{R}^d rather than just a collection of cones that support a fan. In this setting, each regular point of $\mathcal{T}X$ will come from a single top-dimensional cell σ from $\Delta_{\bar{X}, D}$. By the additivity of tropical multiplicities, the result for general choices of tropical compactifications will be obtained by adding the multiplicity associated to the realization of each cell in the complex.

We fix a tropical fan \mathcal{F} of \mathbb{R}^d and the corresponding tropical compactification \bar{X} . Following the notation of Sections 1.1.1 and 1.1.2, we let $Z \subset \mathbb{P}(\mathcal{F})$ be the torus orbit of a maximal cone $\Sigma := \mathbb{R}_{\geq 0}[\sigma] \subset \mathcal{F}$, for $\sigma \in \Delta_{\bar{X}, D}^{\text{top}}$. Then, $X(\mathcal{F}) \cap Z$ is a zero-dimensional scheme, whose length equals the tropical multiplicity of the cone Σ . Moreover, by [91, Lemma 3.2], we know that this multiplicity equals $m_\Sigma = \deg([Z] \cdot [X(\mathcal{F})])$, where the intersection product of cycles is taken in $\mathbb{P}(\mathcal{F})$. Since Gröbner degenerations are flat, we have

$$[X(\mathcal{F})] \cdot [Z] = [\text{in}_w(X(\mathcal{F}))] \cdot [Z]$$

for any vector w in the relative interior of Σ . By [94, Lemma 2.2] we know that regular points of $\mathcal{T}X$ lie in the orbit closure of \bar{X} in $\mathbb{P}(\mathcal{F})$. If we restrict X to the orbit of the cone Σ in $\mathbb{P}(\mathcal{F})$ and the underlying torus \mathbb{T}_Σ , the same is true for $X \cap \mathbb{T}_\Sigma$ and the toric variety $\mathbb{P}(\Sigma)$. The boundary components of the closure of $X \cap \mathbb{T}_\Sigma$ correspond to the s divisors D_k with $k \in \sigma$. Thus, $X \cap \mathbb{T}_\Sigma$ is a zero-dimensional scheme inside the s -dimensional torus \mathbb{T}_Σ hence a complete intersection. By [91, Theorem 4.6] we get $X \cap \mathbb{T}_\Sigma = \cap_{k \in \sigma} D_k$, i.e. the intersection number of these divisors inside \mathbb{T}_Σ .

From the previous discussion we know that the multiplicity of the point w in $\mathcal{T}(X \cap \mathbb{T}_\Sigma)$ equals $D_{k_1} \cdots D_{k_s}$. We wish to view this point in $\mathcal{T}X$.

The inclusion $\Sigma \subset \mathcal{F}$ gives a monomial inclusion $\mathbb{T}_\Sigma \hookrightarrow \mathbb{T}^n \subset \mathbb{P}(\mathcal{F})$ to the open torus of the smooth toric variety $\mathbb{P}(\mathcal{F})$. Using the push-forward formula (1.4) of tropical multiplicities, the multiplicity of w in $\mathcal{T}X$ equals the intersection number $D_{k_1} \cdots D_{k_s}$ times the index of the lattice $\mathbb{Z}[\sigma]$ in its saturation, for $\sigma = \{k_1, \dots, k_s\}$. This concludes our proof. \square

Combining this result with Theorems 5.2.2 and 1.1.10 we obtain a formula for multiplicities of $\mathcal{T}Y$.

Corollary 5.2.5. *With the notation of Theorem 5.2.2, the multiplicity of a regular point w in the tropical variety $\mathcal{T}Y$ equals*

$$m_w = \frac{1}{\delta} \sum_{\substack{\sigma \in \Delta_{\bar{X}, D}^{\text{top}} \\ w \in \mathbb{R}_{\geq 0}[\sigma]}} (D_{k_1} \cdots D_{k_d}) \text{index}(\mathbb{R}[\tilde{\sigma}] \cap \mathbb{Z}^n, \mathbb{Z}[\tilde{\sigma}]), \tag{5.5}$$

where δ is the degree of the map \mathbf{f} .

Proof. By construction, δ equals the degree of the monomial map α restricted to the variety X' . The push-forward formula of multiplicities implies the transition from (5.4) to (5.5) and in particular, the addition of the factor $1/\delta$ and the replacement of the lattice index factor in \mathbb{Z}^d by the corresponding lattice index factor in \mathbb{Z}^n . \square

Equipped with these tools, we now discuss tropical implicitization in the generic case, that is, when the polynomials f_i defining our map $\mathbf{f}: X \rightarrow Y$ have fixed support and we let the coefficients take non-zero values outside a codimension one set. Our approach follows [91, Section 5]. The key fact in this construction is the complete intersection nature of X' . We will come back to this in Section 5.3 when presenting the surface case.

Let $\mathcal{P}_1, \dots, \mathcal{P}_n$ be the Newton polytopes of the polynomials $f_i \in \mathbb{C}[x_1^{\pm 1}, \dots, x_d^{\pm 1}]$ and let \mathcal{P} be their Minkowski sum. Fix C any cone in the inner normal fan of \mathcal{P} , and J any subset of $\{1, \dots, n\}$ of size at most d . We consider the lattice spanned by $\text{trop}(\mathbf{f})(C \cap \mathbb{Z}^d) + \mathbb{Z}^J$, where $\mathbb{Z}^J = \sum_{j \in J} \mathbb{Z}e_j$. If the rank of this lattice is d , denote by $\text{index}(C, J)$ its index in its saturation $(\text{trop}(\mathbf{f})(C) + \mathbb{R}^J) \cap \mathbb{Z}^n$. In all other cases, we take $\text{index}(C, J) = 0$. For every cone C , pick a vector w_C in its relative interior. Given $j \in J$, we define $\text{face}_{w_C}(\mathcal{P}_j)$ to be the face of the polytope \mathcal{P}_j defined by the objective vector w_C . From [91, (5.22)] we know that the $|J|$ -dimensional normalized mixed volume

$$\text{MV}(\text{face}_{w_C}(\mathcal{P}_j) : j \in J) \tag{5.6}$$

is positive if and only if $\dim(\sum_{i \in K} \text{face}_{w_C}(\mathcal{P}_i)) \geq |K|$ for all subsets $K \subseteq J$.

Theorem 5.2.6 ([91, Theorem 5.1]). *Under generic conditions on \mathbf{f} , the tropical variety $\mathcal{T}Y$ is the union of the cones $\text{trop}(\mathbf{f})(C) + \mathbb{R}^J$ over all pairs (C, J) such that (5.6) is positive. The multiplicity m_w at any regular point $w \in \mathcal{T}Y$ is the sum of the quantities*

$$\frac{1}{\delta} \text{index}(C, J) \text{MV}(\text{face}_{w_C}(\mathcal{P}_j) : j \in J),$$

where (C, J) runs over all pairs such that $w \in \text{trop}(\mathbf{f})(C) + \mathbb{R}^J$ and δ is the degree of the map \mathbf{f} .

Remark 5.2.7. The previous result follows from Theorem 5.2.2 and Corollary 5.2.5, after showing that X can be compactified as a projective toric variety $\mathbb{P}(\Sigma)$, whose fan is a strictly simplicial refinement of the inner normal fan of \mathcal{P} . The genericity condition ensures that this compact variety has simple normal crossing boundary. The cells of the boundary complex have two types of vertices, and are encoded by the pairs (C, J) . The realization of these cells give the semigroups $\text{trop}(\mathbf{f})(C \cap \mathbb{Z}^n) + \mathbb{Z}^J$. The index of the associated lattice in its saturation equals $\text{index}(C, J)$.

By Bernstein's theorem, the mixed volume $\text{MV}(\text{face}_{w_C}(\mathcal{P}_j) : j \in J)$ equals the number of non-zero solutions of initial forms associated to the $|J|$ polynomials $\{\text{in}_{w_C}(f_j) : j \in J\}$ in the generic case. This quantity is precisely the intersection number of the boundary divisors of $\mathcal{P}(\mathcal{F})$ index by the pair (C, J) . We refer to [91, Section 4] for further details on this connection.

In our previous discussion, the genericity condition of \mathbf{f} ensures a good choice for a tropical compactification of X by a projective toric variety. As we mentioned in Remark 1.3.2, the Bieri-Groves theorem imposes a necessary condition on the boundary of a compactification to apply geometric tropicalization: it must satisfy the combinatorial normal crossing condition. Following [91], we conjecture that this condition is also sufficient for computing tropical varieties set-theoretically from this compactification.

Conjecture 5.2.8. *Geometric tropicalization holds for any compactification \overline{X} that is normal and whose boundary has combinatorial normal crossings.*

From the previous conjecture it is natural to ask the following question:

Question 5.2.9. *Does the combinatorial normal crossing condition suffice to compute tropical multiplicities with formula (5.5)?*

At first glance, this combinatorial condition suffices in the surface case if we add an extra tangency condition on the branching of the curves. By definition, we have that $D_1 \cdot D_2 \geq k_1 k_2$, and equality holds if and only if the branches of D_1 and D_2 share no tangent directions. Without this extra hypothesis, after a single blow-up, the right-hand side of (5.5) will be strictly greater than the multiplicity of the corresponding regular point in $\mathcal{T}X$. The reason for this gap is that the combinatorial normal crossing condition fails to hold after this blow-up. It would be interesting to compare these two numbers and predict their difference from the geometry of X and its compactification. From simple experiments, it seems that, in the end of the resolution, the formula gives the correct number, suggesting that the tangency hypothesis is not necessary to apply (5.5).

The next proposition proves Conjecture 5.2.8 for surfaces and show that the multiplicity formula holds (5.5) where the intersection points among pairs of divisors come from curves whose branches have no common tangents. The set-theoretic statement is the content of [91, Proposition 5.4]: what is new here is the formula for multiplicities.

Proposition 5.2.10. *Let \overline{X} be a compactification of a smooth surface X whose boundary satisfies the combinatorial normal crossing condition and such that no pair of boundary components have branches with the same tangent directions. Then, geometric tropicalization holds for \overline{X} .*

Proof. The set-theoretic statement is essentially contained in [91, Proposition 5.3]. Without loss of generality, we can assume the tropical fan Σ is strictly simplicial. If this is not the case, we can replace \overline{X} by the fiber product $\overline{X} \times_X \pi^{-1}(X)$, where $\pi: \mathbb{P}(\Sigma') \rightarrow \mathbb{P}(\Sigma)$ is the toric resolution associated to the refinement of the fan Σ . The surface \overline{X} can be resolved by blow-ups of points and the final result is independent of the ordering of these points. The key observation is that new exceptional divisors added by a resolution will not change the tropicalization. To view this, we study two possible scenarios: either the point p we blow up belongs to a unique component of \overline{X} or to exactly two components.

Denote by π this blow-up. In the first situation, we replace the component D_i by two divisors, namely, its strict transform D'_i and a new exceptional divisor E . Notice that the divisor E intersects the boundary of $\pi^{-1}(\overline{X})$ only at the divisor D'_i . The boundary intersection complex of this blow-up is obtained by adding one extra node to $\Delta_{\overline{X}, \overline{D}}$ corresponding to E , identifying the old node D_i with the new node D'_i and adding one edge connecting the node D_i to E . The formula $\pi^*(D_i) = D'_i + mE$, where $m = \text{mult}_p(D_i)$, tells us that the divisorial valuations satisfy the relation $\text{val}_{D'_i} = \text{val}_{D_i}$ and $\text{val}_E = m \text{val}_{D_i}$. All other divisorial valuations stay the same because the point p does not belong to other boundary components. Therefore, the two complexes have the same realization. Likewise, multiplicities are also preserved because the edge $D'_i E$ maps to a ray in its realization. If the divisor D'_i is not smooth, we keep blowing-up the intersection points of D'_i and E until we obtain a resolution. The resolution diagram will be realized by a ray in the tropical variety $\mathcal{T}X$, so we do not need to analyze these points.

Assume now that the point p belongs to exactly two boundary components. From the previous paragraph we know that the question is local, so we can assume the boundary of \overline{X} has only two components, say D_1 and D_2 and that they only intersect at p . The general result will follow by additivity of the intersection multiplicity at each common point. If we blow-up the point p we obtain the exceptional divisor E , the two strict transforms D'_1 and D'_2 , and the following known relations

$$\pi^*(D_1) = D'_1 + k_1 E, \quad \pi^*(D_2) = D'_2 + k_2 E,$$

where $k_i = \text{mult}_p(D_i)$ for $i = 1, 2$. In addition, the intersection numbers are $D'_i \cdot E = k_i$ for $i = 1, 2$ and $D'_1 \cdot D'_2 = D_1 \cdot D_2 - k_1 k_2$ [85, Chapter IV §3.2].

The main difficulty that arises in this construction is that after a single blow-up we could destroy the combinatorial normal crossing condition, since we could have a triple intersection point of the divisors D'_1, D'_2 and E . If so, we would have a triangle with vertices D'_1, D'_2 and E as the boundary intersection complex of the blow-up. However, since the divisorial valuations satisfy $\text{val}_{D'_i} = \text{val}_{D_i}$ for $i = 1, 2$ and $\text{val}_E = k_1 \text{val}_{D_1} + k_2 \text{val}_{D_2}$, the realization of this complex will be two-dimensional, generated by the same two generators $[D_1]$ and $[D_2]$ as the realization of $\Delta_{\overline{X}, \overline{D}}$. If we perform a full resolution with a chain of blow-ups, we end up with a boundary intersection complex that is also realized by the original cone $\mathbb{R}_{\geq 0}\langle [D_1], [D_2] \rangle$. This proves the set-theoretic statement under no extra extra conditions on the boundary intersection points other than disallowing triple intersections.

The question regarding multiplicities is a bit more subtle. The realization of the abstract triangle with vertices D'_1, D'_2, E has the same multiplicities as the original cone spanned by $[D_1]$ and $[D_2]$. More precisely, the formula for any regular point w in the cone over the edge $[D'_i][E]$ gives

$$\begin{aligned} m_w &= (D'_i \cdot E) \gcd(2 \times 2 - \text{minors}([D'_i] \mid [E])) \\ &= k_i \gcd(2 \times 2 - \text{minors}([D_i] \mid k_i [D_i] + k_j [D_j])) \\ &= k_1 k_2 \gcd(2 \times 2 - \text{minors}([D_1] \mid [D_2])), \end{aligned}$$

for $i = 1, 2$. Since the divisors D_1 and D_2 had branches with no common tangents, the number $k_1 k_2$ is the intersection number $D_1 \cdot D_2$. Therefore, the previous expression agrees with the multiplicity of w in the realization of $\Delta_{\overline{X}, D}$. Notice that the hypothesis on the boundary ensure that $D'_1 \cap D'_2 \cap E = \emptyset$. In addition, E and D'_i share no tangent directions, and we can repeat the process. The standard invariants of resolution of singularities ensure the finiteness of this procedure. \square

Even though the hypothesis of the previous result are hard to verify in concrete examples, we know that they hold, for example, for contraction of exceptional divisors giving bivalent nodes in the intersection complex of a full resolution. In practice, this complex is obtained by constructing a resolution diagram, contracting bivalent nodes and then tracing back the corresponding intersection numbers. We refer to Section 4.3 for a concrete application of the previous result.

5.3 Tropical implicitization for generic surfaces

In this section, we specialize the constructions of Section 5.2 to the case of generic rational surfaces. In agreement with the discussion ending that section, we show that the *genericity* condition for surfaces is precisely the combinatorial normal crossing boundary condition plus a tangency condition at crossing points. As stated in Proposition 5.2.10, we show that no smoothness condition on the ambient space is required.

We start by describing our input, a polynomial parameterization $\mathbf{f}: \mathbb{T}^2 \rightarrow \mathbb{T}^n$ where $n \geq 3$, $\mathbf{f} = (f_1, \dots, f_n)$ and $f_i \in \mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}]$ for all $i = 1, \dots, n$. We let Y be the Zariski closure of the image of \mathbf{f} in \mathbb{T}^n . We assume our map \mathbf{f} is generically finite, so Y is a surface in \mathbb{T}^n . Our goal is to compute the tropical surface graph associated to Y .

Following Khovanskii's philosophy [35, 57], we assume each f_i is generic *relative to its support*. We now state the main result in this section. The remainder of the section will be devoted to its proof and to give several numerical examples. For simplicity, we assume that our choices of coefficients give *irreducible* polynomials. The genericity condition says that the plane curves defined by the polynomials f_i satisfy the combinatorial normal crossing condition in \mathbb{T}^n . Algorithm 5.1 gives the pseudo-code implementation of this result.

Theorem 5.3.1. *Let $\mathbf{f} = (f_1, \dots, f_n): \mathbb{T}^2 \rightarrow \mathbb{T}^n$ be a generically finite map. Assume that the polynomials f_i are irreducible and generic relative to their Newton polytopes \mathcal{P}_i , i.e. no three plane curves in \mathbb{T}^2 defined by these n polynomials intersect at a point and the tangent directions at pairwise crossing points of the branches of two curves are distinct. Then, the tropical surface graph associated to the variety $\overline{\text{im } \mathbf{f}} \subset \mathbb{T}^n$ can be described as follows. Let \mathcal{N} the common refinement of the n inner normal fans of the polytopes \mathcal{P}_i ($i = 1, \dots, n$), and let ρ_1, \dots, ρ_l be the rays of \mathcal{N} , oriented counterclockwise, with primitive generators n_{ρ_j} for all $j = 1, \dots, l$. With this notation, the nodes of the graph are*

$$\{e_i : \dim \mathcal{P}_i \neq 0, 1 \leq i \leq n\} \cup \{[D_\rho] := \text{trop}(\mathbf{f})(n_\rho) : [D_\rho] \neq 0, \rho \in \mathcal{N}^{[1]}\}.$$

The list of edges and their weights is:

- (i) $m_{([D_{\rho_j}], [D_{\rho_k}])} = \delta^{-1} \gcd(2 \times 2 - \text{minors}([D_{\rho_j}] \mid [D_{\rho_k}])) / \gcd(2 \times 2 - \text{minors}(n_{\rho_k} \mid n_{\rho_j}))$,
if $|j - k| = 1 \pmod{l}$ or 0 if not.
- (ii) $m_{(e_i, [D_{\rho}])} = \delta^{-1} (|\text{face}_{n_{\rho}}(\mathcal{P}_i) \cap \mathbb{Z}^2| - 1) \gcd([D_{\rho}]_j : j \neq i)$, if $n_{\rho} \in \mathcal{T}(f_i)$, or 0 if not.
- (iii) $m_{(e_i, e_j)} = \delta^{-1} \text{length}((f_i = f_j = 0) \cap \mathbb{T}^2)$ if $\dim(\mathcal{P}_i \oplus \mathcal{P}_j) = 2$, and 0 if not. If the coefficients are generic enough, this number equals $1/\delta$ times the mixed volume of the polytopes \mathcal{P}_i and \mathcal{P}_j .

It is important to point out that the previous algorithm was already presented in [92] and further studied in [91]. We contribute to the subject by elucidating the right genericity condition to impose. The proof of [92, Theorem 2.1] requires the genericity of both the coefficients and the Newton polytopes, to ensure that the Minkowski sum of the n polytopes $\mathcal{P}_1, \dots, \mathcal{P}_n$ is a simple *smooth* polytope. Our proof discards this extra assumption on the polytopes, unraveling the key aspects in their argumentation. In addition, we correct the missing factor of $1/\delta$ in [92], following [91, Theorem 5.1].

Proof. As we stated in the previous section, our task is to compactify $Y \subset \mathbb{T}^n$. Instead, we work with $X := \mathbf{f}^{-1}(Y) = \mathbb{T}^2 \setminus (f_i = 0)$. Using the knowledge of the Newton polytopes $\mathcal{P}_1, \dots, \mathcal{P}_n$, we construct a projective toric variety $\mathbb{P}(\Sigma)$ and compactify X inside this space [92, Theorem 2.1]. We now explain how to construct the fan Σ from these n input polynomials. Consider the Minkowski sum of the n polytopes $\mathcal{P}_1, \dots, \mathcal{P}_n$. Since the polynomials are generic, its inner normal fan is the common refinement of the n normal fans of these polytopes. We let \mathcal{N} be this fan. Notice that the fan need not be strictly simplicial, thus $\mathbb{P}(\mathcal{N})$ need not be a smooth toric surface. To fix this, we perform a refinement of the fan \mathcal{N} by a strictly simplicial fan \mathcal{N}' . This is done by subdividing the two-dimensional cones that are not unimodular [15, Section 5]. On the geometric side, this operation corresponds to performing toric blow-ups on the surface $\mathbb{P}(\mathcal{N})$: we blow-up torus invariant points associated to these two-dimensional cones. The output is a smooth projective toric surface $\mathbb{P}(\mathcal{N}')$ that compactifies X . We set Σ to be the fan \mathcal{N}' .

The boundary of $\mathbb{P}(\mathcal{N}')$ consists of two types of divisors. The first class of divisors are the toric divisors D_{ρ} indexed by the rays ρ in \mathcal{N}' . They correspond to facets of the Minkowski sum $\bigoplus_{i=1}^n \mathcal{P}_i$ and, thus, to the rays of \mathcal{N}' . The toric boundary $\bigcup_{\rho \in \mathcal{N}'[1]} D_{\rho}$ has simple normal crossings because \mathcal{N}' is a strictly simplicial fan. Similarly, the toric divisor $\bigcup_{\rho \in \mathcal{N}'[1]} D_{\rho}$ has combinatorial normal crossings because \mathcal{N} is a simplicial fan.

The second type of components consist of n divisors $\overline{E}_1, \dots, \overline{E}_n$, which are the closures in $\mathbb{P}(\mathcal{N}')$ of the divisors $E_i = (f_i = 0) \subset \mathbb{T}^2$ in X . The irreducibility of the polynomials f_i and the proof of [92, Theorem 2.1] show that these divisors are smooth and irreducible and that the union of all D_{ρ} 's and all \overline{E}_j 's has simple normal crossings. Notice that if f_j consists of a single monomial, then E_j is the empty set. Such indices will not induce a node

in the boundary intersection complex of X , so from now on we may assume $\dim \mathcal{P}_i > 0$ for all $i = 1, \dots, n$.

Thus, the boundary divisor of X and its smooth compactification $\mathbb{P}(\mathcal{N}')$ can be decomposed as

$$\partial := \mathbb{P}(\mathcal{N}') \setminus X = D_1 \cup \dots \cup D_l \cup D_{l+1} \cup \dots \cup D_m \cup \overline{E}_1 \cup \dots \cup \overline{E}_n,$$

where ρ_1, \dots, ρ_l are the rays of \mathcal{N} , $\rho_{l+1}, \dots, \rho_m$ are the rays in $\mathcal{N}'^{[1]} \setminus \mathcal{N}^{[1]}$ and D_i denotes the toric divisor D_{ρ_i} ($i = 1, \dots, m$).

The simplicial complex $\Delta_{\mathbb{P}(\mathcal{N}'), \partial}$ is a graph with $m + n$ vertices. The edges of this graph consist of pairs of vertices $C \cup J$ where $C \subset \{1, \dots, m\}$ and $J \subset \{1, \dots, n\}$ and $|C| + |J| = 2$. Following [92], we denote by $\Delta_{\mathbb{P}(\mathcal{N}'), \partial}[J]$ the subset of edges in $\Delta_{\mathbb{P}(\mathcal{N}'), \partial}$ with fixed J . By construction, we have three possibilities: $|J| = 0, 1$ or 2 . For $|J| = 0$, the edges $\Delta_{\mathbb{P}(\mathcal{N}'), \partial}[J]$ are of the form $(D_\rho, D_{\rho'})$ for ρ and ρ' rays in the fan \mathcal{N}' . By standard intersection theory on smooth toric varieties, we know that the intersection numbers among the torus-invariant divisors D_ρ are:

$$D_\rho \cdot D_{\rho'} = \begin{cases} 1 & \text{if } \rho \text{ and } \rho' \text{ generate a two-dimensional cone in } \mathcal{N}', \\ 0 & \text{else.} \end{cases} \quad (5.7)$$

In particular, this says that we only have edges among consecutive rays of \mathcal{N}' , if counter-clockwise oriented.

When $|J| = 1$, we seek to identify edges of the form (\overline{E}_j, D_ρ) , for $\rho \in \mathcal{N}'$ and $j = 1, \dots, n$. Again, this is done by toric methods. Since \overline{E}_j represents a Cartier divisor with local equation f_j and D_ρ is a torus invariant divisor, the intersection number will correspond to the intersection number of the initial form in $n_\rho(f_j)$ and D_ρ . This quantity coincides with the number of nonzero solutions of a univariate polynomial. The Newton polytope of $\text{in}_{n_\rho}(f_j)$ agrees with $\text{face}_{n_\rho}(\mathcal{P}_j)$. Since $n_\rho \neq 0$, this polytope has dimension zero or one. If this number is zero, the initial form is a monomial, and the intersection number is zero. If the dimension is one, Newton-Puiseux's theorem implies that the intersection number is the lattice length of the edge $\text{face}_{n_\rho}(\mathcal{P}_j)$. Thus, we see that \overline{E}_j is adjacent to a node D_ρ if and only if n_ρ is a ray in the normal cone of the polytope \mathcal{P}_j , or equivalently, if it belongs to the tropical hypersurface $\mathcal{T}(f_j)$. In addition:

$$\overline{E}_j \cdot D_\rho = \text{lattice length of } \text{face}_{n_\rho}(\mathcal{P}_j) = |\text{face}_{n_\rho}(\mathcal{P}_j) \cap \mathbb{Z}^2| - 1. \quad (5.8)$$

Finally, if $|J| = 2$, we want to certify which edges $(\overline{E}_i, \overline{E}_j)$ belong to the boundary complex. We claim it suffices to check if the equations f_i and f_j have a common root in \mathbb{T}^2 . In fact, the remaining intersection points outside the big open torus will lie in the toric boundary of $\mathbb{P}(\mathcal{N}')$ and, thus, will yield a triple intersection among the divisors $\overline{E}_i, \overline{E}_j$ and some toric divisor D_ρ . Therefore, the intersection number is the length of the zero-dimensional scheme $(f_i = f_j = 0) \cap \mathbb{T}^2$. If the coefficients of these polynomials are generic,

Bernstein’s theorem show that this length equals the mixed volume of the polytopes \mathcal{P}_i and \mathcal{P}_j , whereas for special choices of coefficients, this number is an upper bound for the true intersection number [3, Theorem A]. The mixed volume is nonzero if and only if the Minkowski sum of the corresponding polytopes is two-dimensional. Therefore, the edges of the form $(\overline{E}_i, \overline{E}_j)$ in the boundary complex must satisfy $\dim(\mathcal{P}_i \oplus \mathcal{P}_j) = 2$.

Notice that since we are interested in the realization of the boundary complex, we can safely assume that the dimension restriction characterizes the edges $(\overline{E}_i, \overline{E}_j)$. “Artificial” edges in the abstract graph will have weight zero in its realization. For this reason, we safely add these artificial edges to the boundary complex and still obtain the correct tropical surface graph. Summarizing:

$$(\overline{E}_i, \overline{E}_j) \text{ is an edge of } \Delta_{\mathbb{P}(\mathcal{N}'), \partial} \iff \dim(\mathcal{P}_i \oplus \mathcal{P}_j) = 2.$$

We now discuss the realization of the boundary complex. As we know, this is done by describing the divisorial valuation of each component of the boundary ∂ . By construction, $val_{\overline{E}_j}(f_i) = \delta_{i,j}$, so the node corresponding to \overline{E}_j maps to e_j , the j^{th} element of the canonical basis. We compute the divisorial valuation of all D_ρ ’s with the tools of toric geometry [39, Section 5.2]. Without loss of generality, assume $\rho = \rho_1$. We picking the primitive generator n_ρ of ρ and we extend this vector to a \mathbb{Z} -basis of \mathbb{R}^2 . Thus, we can assume $\rho = e_1$. By definition, the divisorial valuation val_{D_ρ} is obtained from the order of vanishing of all the polynomials f_j at D_ρ , that is, by the maximal exponent of the variable x_1 dividing f_j in the polynomial ring $\mathbb{C}[x_2^{\pm 1}][x_1]$. Notice that this number can be negative. The maximum exponent is precisely the minimum value among the inner products $\langle e_1, \nu \rangle$, for all possible monomials with exponent $\nu \in \mathbb{Z}^2$ that appear in f_j . Hence, $val_{D_\rho}(\chi^*(t_j)) = val_{D_\rho}(f_j) = trop(f_j)(e_1)$. With the same reasoning, we obtain

$$[D_i] := (val_{D_i}(f_j))_{j=1}^n = (trop(f_1)(n_{\rho_i}), \dots, trop(f_n)(n_{\rho_i})) = trop(\mathbf{f})(n_{\rho_i}) \quad \forall i = 1, \dots, m.$$

Using Corollary 5.2.5 and expressions (5.7) and (5.8), we compute the weights of the edges in the realization of $\Delta_{\mathbb{P}(\mathcal{N}'), \partial}$:

- (i) $m_{([D_\rho], [D_{\rho'}])} = \delta^{-1} \gcd(2 \times 2 - \text{minors}([D_\rho] \mid [D_{\rho'}])) / \gcd(2 \times 2 - \text{minors}(n_\rho \mid n_{\rho'}))$, if $n_\rho, n_{\rho'}$ span a two-dimensional cone in \mathcal{N}' . In other cases, this number is 0.
- (ii) $m_{([\overline{E}_i], [D_\rho])} = \delta^{-1} (|\text{face}_{n_\rho}(\mathcal{P}_i) \cap \mathbb{Z}^2| - 1) \gcd(trop(f_j)(n_\rho) : j \neq i)$, if $n_\rho \in \mathcal{T}(f_i)$, or 0 if not.
- (iii) $m_{([\overline{E}_i], [\overline{E}_j])} = \delta^{-1} \text{length}((f_i = f_j = 0) \cap \mathbb{T}^2)$ if $\dim(\mathcal{P}_i \oplus \mathcal{P}_j) = 2$ or zero otherwise. If the coefficients of f_i , and f_j are generic enough, this number is precisely the mixed volume of the polytopes \mathcal{P}_i and \mathcal{P}_j , multiplied by a factor of $1/\delta$.

Notice that the statement of our theorem constructs the tropical surface graph by means of the fan \mathcal{N} rather than our choice $\Sigma = \mathcal{N}'$. We now explain how can we avoid performing the refinement \mathcal{N}' of the fan \mathcal{N} in the previous argument, in agreement with the

combinatorial normal crossing genericity condition required for the polynomials f_1, \dots, f_n and the spirit of Proposition 5.2.10. First we provide set-theoretic reasons and then focus on the weights of the tropical surface graph. By induction on the number of elements in $\mathcal{N}'^{[1]} \setminus \mathcal{N}^{[1]}$, it suffices to show that subdividing a two-dimensional cone on \mathcal{N} by adding a single ray gives the same tropical surface graph.

Our first goal is to show that the rays added to \mathcal{N} give nodes and edges in the realization of the complex that are contained in cones over edges of the form $([D_\rho], [D_{\rho'}])$ for consecutive edges ρ, ρ' in $\mathcal{N}^{[1]}$ with counterclockwise orientation. Each ray that we add to \mathcal{N} corresponds to a blow-up at a torus-invariant point p of $\mathbb{P}(\mathcal{N})$. This point is the torus orbit associated to the two-dimensional cone spanned by ρ and ρ' and supports the intersection of the two torus-invariant divisors, thought of as \mathbb{Q} -Cartier divisors:

$$D_\rho \cap D_{\rho'} = (\gcd(2 \times 2 - \text{minors}(n_\rho \mid n_{\rho'})))^{-1} p.$$

This implies $D_\rho \cdot D_{\rho'} = (\gcd(2 \times 2 - \text{minors}(n_\rho \mid n_{\rho'})))^{-1}$, where the intersection product is considered in the (possibly singular) toric variety $\mathbb{P}(\mathcal{N})$ [39, Section 5.1].

At each step of the toric resolution, we subdivide the cone spanned by consecutive rays ρ and ρ' of \mathcal{N} by adding a ray τ in this cone. Call \mathcal{N}' the corresponding refinement. Note that this extra ray modifies the boundary intersection complex associated to $\mathbb{P}(\mathcal{N})$ as follows. It adds a node D_τ and replaces the edge $(D_\rho, D_{\rho'})$ in the boundary complex associated to $\mathbb{P}(\mathcal{N})$ with two edges: (D_ρ, D_τ) and $(D_\tau, D_{\rho'})$. The key fact is the combinatorial normal crossing condition on the boundary of $\mathbb{P}(\mathcal{N})$, satisfied by the genericity of the polynomials f_i . This property is preserved under refinements of the fan \mathcal{N} . Notice that every ray that we add to \mathcal{N} optimizes a single vertex of all Newton polytopes \mathcal{P}_i , so it does not intersect any of the divisors \overline{E}_i .

By construction, the valuations of the torus-invariant divisors associated to rays in $\mathcal{N}^{[1]}$ viewed as divisors in $\mathbb{P}(\mathcal{N})$ or $\mathbb{P}(\mathcal{N}')$ are the same. If we write $n_\tau = a n_\rho + b n_{\rho'}$, for $a, b \in \mathbb{Q}_{>0}$, we see that the divisorial valuation of D_τ equals $a \text{val}_{D_\rho} + b \text{val}_{D_{\rho'}}$. Thus, the cone over the realization of the edge $([D_\rho], [D_{\rho'}])$ coincides with the union of the cones over the edges $([D_\rho], [D_\tau])$ and $([D_\tau], [D_{\rho'}])$, with no partial overlap between the last two cones. If we write $a = c/q$ and $b = d/q$ with $c, d, q \in \mathbb{Z}$, we see that

$$\begin{cases} \gcd(2 \times 2 - \text{minors}([D_\tau] \mid [D_\rho])) &= b \gcd(2 \times 2 - \text{minors}([D_\rho] \mid [D_{\rho'}])), \\ \gcd(2 \times 2 - \text{minors}([D_\tau] \mid [D_{\rho'}])) &= a \gcd(2 \times 2 - \text{minors}([D_\rho] \mid [D_{\rho'}])), \\ \gcd(2 \times 2 - \text{minors}([n_\tau] \mid [n_\rho])) &= b \gcd(2 \times 2 - \text{minors}([n_\rho] \mid [n_{\rho'}])), \\ \gcd(2 \times 2 - \text{minors}([n_\tau] \mid [n_{\rho'}])) &= a \gcd(2 \times 2 - \text{minors}([n_\rho] \mid [n_{\rho'}])). \end{cases} \quad (5.9)$$

When analyzing the intersection numbers among pairs of boundary divisors we only need to look at the three divisors $D_\rho, D_{\rho'}$ and D_τ . As we mentioned, the intersection of the two divisors D_ρ and $D_{\rho'}$ in $\mathbb{P}(\mathcal{N})$ is replaced by the chain of intersections D_ρ, D_τ , and $D_\tau, D_{\rho'}$ in $\mathbb{P}(\mathcal{N}')$. Using (5.9) these two intersection numbers in $\mathbb{P}(\mathcal{N}')$ are related to the intersection

number $D_\rho \cdot D_{\rho'}$ in $\mathbb{P}(\mathcal{N})$ by the expressions

$$D_\rho \cdot D_\tau = \frac{1}{b} D_\rho \cdot D_{\rho'}, \quad D_\tau \cdot D_{\rho'} = \frac{1}{a} D_\rho \cdot D_{\rho'}.$$

Combining these identities with Theorem 5.2.10, we see that the multiplicities of the edges $([D_\rho], [D_\tau])$ and $([D_\tau], [D_{\rho'}])$ in the realization of the boundary complex of $\mathbb{P}(\mathcal{N}')$ equal the multiplicity of the edge $([D_\rho], [D_{\rho'}])$ in the realization of the boundary complex of $\mathbb{P}(\mathcal{N})$. This concludes our proof. \square

Remark 5.3.2. The last part of the previous proof shows the effect of the blow-ups on the realization of the boundary intersection complex for a given compactification: both weighted graphs in \mathbb{R}^n yield the same weighted set $\mathcal{T}Y$. This effect was discussed in the proof of Proposition 5.2.10 and in the construction of the tropical secant surface graph of Section 4.3. It also remarks the key role of the combinatorial normal crossing boundary as the generic condition to ensure the validity of the conclusion in Theorem 5.3.1.

We illustrate the previous methods for computing tropicalizations of generic surfaces with three examples. In the next section, we revisit the last two examples for special choices of coefficients that violate the genericity condition. In particular, we show how the resulting tropical surface graphs need to be modified to reflect the non-genericity of these surfaces.

Example 5.3.3. Our first example is a modification of [92, Example 3.4], where we remove a monomial factor from each polynomial to make them irreducible over the polynomial ring $\mathbb{C}[x, y, z]$. This change will have no effect on the combinatorics of the tropical surface graph, but will change its coordinates. Consider the general surface in $Y \subset \mathbb{T}^3$ parameterized by three bivariate polynomials

$$\begin{cases} f_1(s, t) = a_1 + a_2 s^2 t + a_3 s t^2, \\ f_2(s, t) = b_1 s t + b_2 s + b_3 t, \\ f_3(s, t) = c_1 t + c_2 s^2 + c_3 s t^2, \end{cases}$$

where $a_1, a_2, a_3, b_1, b_2, b_3, c_1, c_2, c_3 \in \mathbb{C}$ are generic nonzero coefficients. By construction, the map \mathbf{f} has degree $\delta = 1$. Their Newton polytopes and corresponding inner normal fans are depicted in Figures 5.1 and 5.2. All rays in these fans have multiplicity one, since the lattice lengths of the corresponding edges in the polytopes equal one.

Next, we compute the Minkowski sum $\mathcal{P} := \mathcal{P}_1 \oplus \mathcal{P}_2 \oplus \mathcal{P}_3$ of the three Newton polytopes, and its associated inner normal fan. This fan corresponds to the common refinement of the normal fans of the three constituent polytopes. Notice that the resulting fan $\mathcal{N}(\mathcal{P})$ is not strictly simplicial, since the maximal cones corresponding to the vertices v_1, v_7 and v_9 of \mathcal{P} are not unimodular. Thus, we subdivide these three cones by adding the rays r_{19} , r_{89} and r_{67} . The resulting fan \mathcal{N}' has f -vector $(1, 12, 12)$. The polytope \mathcal{P} and the fan \mathcal{N}' are depicted in Figure 5.3.




Figure 5.1: From left to right: Newton polytopes of the polynomials f_1, f_2 and f_3 .




Figure 5.2: From left to right: inner normal fans of the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 .

Following Algorithm 5.1, we construct the nodes of the graph encoding $\mathcal{T}Y$. To simplify notation, we denote by D_i the toric divisor associated to the ray r_i . The corresponding primitive vectors n_{r_i} are indicated in Figure 5.3. The nodes associated to the nine toric divisors are:

$$\begin{cases} [D_1] := (-2, -1, -2), & [D_4] := (-2, -1, -2), & [D_7] := (0, 1, 2), \\ [D_2] := (-5, -3, -4), & [D_5] := (-1, -1, -1), & [D_8] := (0, -1, -2), \\ [D_3] := (-3, -2, -3), & [D_6] := (0, -1, -1), & [D_9] := (-1, -1, -1). \end{cases} \quad (5.10)$$

Notice that $[D_1] = [D_4]$, so the tropical surface graph in \mathbb{R}^3 has fewer nodes than expected: there are eleven nodes in this realization. Likewise, some edges $([D_\rho], e_i)$ or $([D_\rho], [D_\eta])$ may give one-dimensional cones in the tropical variety $\mathcal{T}Y$. For this reason, in Algorithm 5.1, we tested the dimension of the corresponding cones before adding any of these pairs to the list of edges of our tropical surface graph. After computing these numbers, we obtain 19 edges: the three edges (e_i, e_j) , the eight edges $([D_9], e_1)$, $([D_3], e_1)$, $([D_4], e_1)$, $([D_7], e_2)$, $([D_1], e_2)$, $([D_8], e_3)$, $([D_2], e_3)$ and $([D_5], e_3)$, the seven edges $([D_i], [D_{i+1}])$ ($i = 1, \dots, 8, i \neq 6$) and $([D_9], [D_1])$.

The weights of these edges are computed using mixed volumes, and are indicated in the left-most picture in Figure 5.4. We start with the edges corresponding to $|J| = 2$ and $C = \{(0, 0)\}$, that is, to the three edges (e_i, e_j) . These mixed volumes are obtained using the formula $MV(P_i, P_j) = \text{Vol}(P_i \oplus P_j) - \text{Vol}(P_i) - \text{Vol}(P_j)$, where $\text{Vol}(-)$ denotes the Euclidean




Figure 5.3: From left to right: Minkowski sum \mathcal{P} of $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 and a strictly simplicial fan \mathcal{N}' refining the normal fan of \mathcal{P} . The different colors indicate the corresponding normal fans of each \mathcal{P}_i ($i = 1, 2, 3$). The dashed rays r_{19}, r_{89} and r_{67} are introduced to refine the singular cones of $\mathcal{N}(\mathcal{P})$. Ignoring these three rays gives the nine chambers in the normal fan of \mathcal{P} , which are dual to the nine vertices of \mathcal{P} .

volume in \mathbb{R}^2 . Figure 5.4 shows the pairwise Minkowski sums of the polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$. This gives $m_{12} = \frac{1}{2}(12 - 3 - 1) = 4$, $m_{23} = \frac{1}{2}(10 - 1 - 3) = 3$, $m_{13} = \frac{1}{2}(18 - 3 - 3) = 6$.

We now consider the case where $|J| = 1$ and C is a ray in the fan \mathcal{N} . These are edges of the form $([D_\rho], e_i)$, where ρ is a ray in the normal fan $\mathcal{N}(\mathcal{P}_i)$ of the polytope \mathcal{P}_i . Since the edges of all these polytopes have lattice length one and the generating rays of \mathcal{N} are primitive, our task reduces to computing the gcd of the maximal minors of the 3×2 -matrices $([D_\rho] \mid e_i)$, that is, the gcd of all coordinates $[D_\rho]$ except for the i^{th} one. All such numbers equal one, except for the edges with $\rho = r_1 = r_4$ and $i = 2$, whose value is two. Since the nodes D_1 and D_4 are proportional, we need to add these numbers to get the true multiplicity of the edge $([D_1], e_2)$, which equals four. This explains the transition from the left to the right of Figure 5.5.

Finally we compute the values of the edges with $|J| = 0$, that is, edges of the form $([D_\rho], [D_\eta])$ associated to consecutive rays in the normal fan \mathcal{N} from the left-most picture in Figure 5.3. Such numbers are computed as the quotient of the gcd of maximal minors of the matrix $([D_\rho] \mid [D_\eta])$ by the determinant of the 2×2 -matrix $(\rho \mid \eta)$. In our example, all such numbers equal one. In particular, this shows why we need not consider the fan \mathcal{N}' and, instead, work safely with the fan \mathcal{N} .

The resulting weighted graph has four bivalent nodes (in gray) and it is depicted on the right of Figure 5.5. After removing these gray nodes, we obtain a graph with f -vector $(7, 13)$. The complement of the graph has eight connected components. Notice that the nodes $e_2, [D_1] = [D_4], [D_3]$ and $[D_5]$ are aligned in the picture. This reflect the fact that




Figure 5.4: From left to right: Minkowski sum of the Newton polytopes $\mathcal{P}_1 \oplus \mathcal{P}_2$, $\mathcal{P}_2 \oplus \mathcal{P}_3$ and $\mathcal{P}_1 \oplus \mathcal{P}_3$.

these four vectors generate a two-dimensional cone in \mathbb{R}^3 . In addition to the four bivalent nodes, this also explains the difference between the number of edges in the tropical surface graph and the number of edges in the abstract graph, seen on the left-side of Figure 5.5. In particular, the predicted edge $([D_4], [D_5])$ can be seen in the graph as the line segment containing the points $[D_4], [D_3]$ and $[D_5]$.

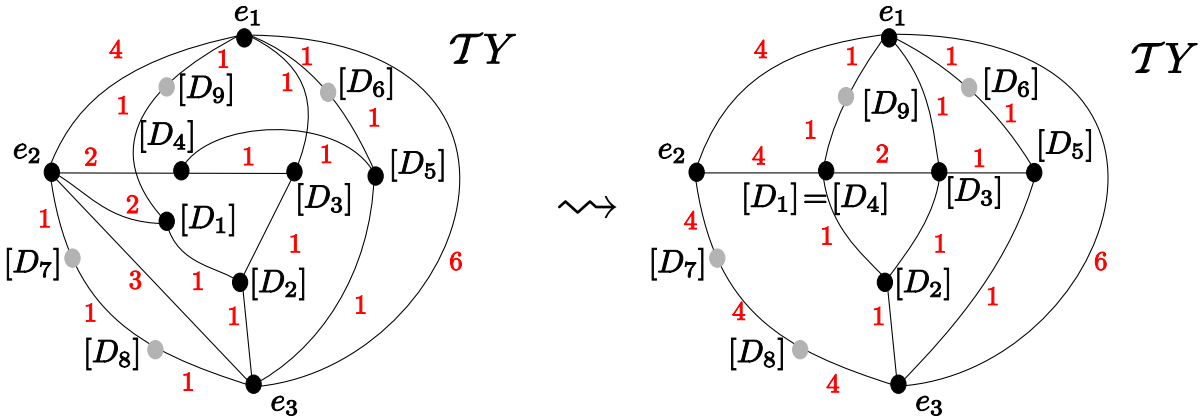


Figure 5.5: From left to right: weighted simplicial complexes representing TY . The left one corresponds to the abstract graph and the right one is the planar graph obtained by realizing the abstract graph and combining weights of overlapping edges.

We certify our calculations by computing the Newton polytope of the generator of the principal ideal $(x - f_1(s, t), y - f_2(s, t), z - f_3(s, t)) \cap \mathbb{C}[x, y, z]$ using Singular [21]. For generic choices of coefficients a_1, \dots, c_3 , this polynomial has degree 14. Its Newton polytope has f -vector $(8, 13, 7)$. Using Gfan, we compute its tropicalization:

RAYS		MAXIMAL_CONES	MULTIPLICITIES
-1 0 0	# 0	{0 1} # Dimension 3	4 # Dimension 3
0 -1 0	# 1	{0 2}	1
1 1 1	# 2	{0 3}	6
0 0 -1	# 3	{0 4}	1
2 1 2	# 4	{0 5}	1
3 2 3	# 5	{5 6}	1
5 3 4	# 6	{2 3}	1
		{4 5}	2
N_RAYS		{1 4}	4
7		{2 5}	1
		{1 3}	4
F_VECTOR		{4 6}	1
1 7 13		{3 6}	1

After changing the signs of the seven rays in the left-most column to overcome the *max* convention of *Gfan*, we recover the nodes $e_1, e_2, [D_5], e_3, [D_1], [D_3]$ and $[D_2]$, as expected. \diamond

Example 5.3.4. We consider the morphism $f = (f_1, f_2, f_3): \mathbb{C}^2 \rightarrow Y \subset \mathbb{C}^3$ given by

$$\begin{cases} f_1(s, t) = a_1 s^2 + a_2 s^3 + a_3 t^2, \\ f_2(s, t) = b_1 t^2 + b_2 t^3 + b_3 s^2, \\ f_3(s, t) = (s + t)^2 - (s + t)^3 - (s - t)^2 = c_1 st + c_2 s^3 + c_3 t^3 + c_4 st^2 + c_5 s^2t, \end{cases}$$

with generic coefficients $a_1, \dots, c_5 \in \mathbb{C}^*$. By construction, the degree of the map \mathbf{f} equals one. As in the previous example, we draw the Newton polytopes of each f_i (Figure 5.6), the normal fans (Figure 5.7), the Minkowski sum of the three polytopes and its normal fan (Figure 5.8). In this case, ray r_7 is a common ray of two of the constituent fans. We indicate this phenomenon on the right-most picture in Figure 5.8 by drawing a purple edge, which combines the red and blue edges from the polytopes \mathcal{P}_1 and \mathcal{P}_2 . Similarly, the purple ray in the normal fan on the left-most picture in Figure 5.8 combines the red and blue rays of $\mathcal{N}(\mathcal{P}_1)$ and $\mathcal{N}(\mathcal{P}_2)$. In addition, three rays out of the nine rays in $\bigvee_{i=1}^3 \mathcal{N}(\mathcal{P}_i)$ have non-trivial weights as indicated in the pictures on Figure 5.7.




Figure 5.6: From left to right: Newton polytopes of the polynomials f_1, f_2 and f_3 .




Figure 5.7: From left to right: inner normal fans of the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 .

Following the notation of Figure 5.8, the nodes of the tropical surface graph have coordinates: e_1, e_2, e_3 , $[D_1] = (0, 0, 0)$, $[D_2] = (-9, -6, -9)$, $[D_3] = (-3, -3, -3)$, $[D_4] = (-6, -9, -9)$, $[D_5] = (0, 0, 0)$, $[D_6] = (2, 2, 3)$, $[D_7] = (2, 2, 2)$ and $[D_8] = (2, 2, 3)$. After going through dimension testings, we obtain a list of fourteen edges: three with $|J| = 2$, (e_i, e_j) , seven with $|J| = 1$, $([D_4], e_1)$, $([D_7], e_1)$, $([D_2], e_2)$, $([D_7], e_2)$, $([D_3], e_3)$, $([D_6], e_3)$, $([D_8], e_3)$ and four with $|J| = 0$, $([D_2], [D_3])$, $([D_3], [D_4])$, $([D_6], [D_7])$, $([D_7], [D_8])$. Finally, following the approach of the previous example, we compute tropical multiplicities as mixed volumes. The transition from the weighted abstract graph to its realization is indicated in Figure 5.9. \diamond

Example 5.3.5. As our third example we consider the surface parameterized by the morphism $f = (f_1, f_2, f_3): \mathbb{C}^2 \rightarrow Y \subset \mathbb{C}^3$, where

$$\begin{cases} f_1(s, t) = a_1 + a_2 s + a_3 t, \\ f_2(s, t) = b_1 + b_2 t + b_3 s^2, \\ f_3(s, t) = c_1 + c_2 st. \end{cases} \quad (5.11)$$

The degree of this map is $\delta = 1$. Figure 5.10 depicts the Newton polytopes of our three polynomials and the common refinement of their inner normal fans. Using the methods described in this section we obtain a weighted graph with seven nodes (Figure 5.11): e_1, e_2, e_3 , $[D_2] = (-1, -2, 0)$, $[D_3] = (-1, -2, -2)$, $[D_4] = (-2, -2, -3)$ and $[D_5] = (-1, -1, 0)$. If we remove the two bivalent nodes $[D_2]$ and $[D_5]$, we obtain a graph with five nodes and eight edges, whose complement has five connected components. This modified graph is depicted in Figure 5.11. The edges with no label have weight one.

On the other hand, if we use standard elimination techniques implemented in **Singular** we obtain an equation of degree 3 in x, y, z whose coefficients are polynomials in the indeterminates a_1 through c_2 :

$$b_2 \cdot b_3 \cdot c_2^2 \cdot x^3 - a_3 \cdot b_3 \cdot c_2^2 \cdot x^2 \cdot y + (a_3^3 \cdot b_3^2) \cdot z^2 + a_2^2 \cdot a_3 \cdot c_2^2 \cdot y^2 +$$




Figure 5.8: From left to right: Minkowski sum \mathcal{P} of $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 and a strictly simplicial fan \mathcal{N}' refining the normal fan of \mathcal{P} . In this case, we need to add three rays: two between r_1 and r_2 (r_{012}, r_{112}), and one between r_4 and r_5 (r_{45}).

$$\begin{aligned}
 & b_3 \cdot c_2^2 \cdot (a_3 \cdot b_1 - 3 \cdot a_1 \cdot b_2) \cdot x^2 + c_2^2 \cdot (2 \cdot a_1 \cdot a_3 \cdot b_3 - a_2^2 \cdot b_2) \cdot x \cdot y + \\
 & 2 \cdot a_2 \cdot a_3^2 \cdot b_3 \cdot c_2 \cdot y \cdot z - 3 \cdot a_2 \cdot a_3 \cdot b_2 \cdot b_3 \cdot c_2 \cdot x \cdot z - (2 \cdot a_3^3 \cdot b_3^2 \cdot c_1 + a_2^3 \cdot b_2^2 \cdot c_2 - \\
 & 2 \cdot a_2 \cdot a_3^2 \cdot b_1 \cdot b_3 \cdot c_2 + 3 \cdot a_1 \cdot a_2 \cdot a_3 \cdot b_2 \cdot b_3 \cdot c_2) \cdot z + c_2 \cdot (-2 \cdot a_2 \cdot a_3^2 \cdot b_3 \cdot c_1 - \\
 & 2 \cdot a_2^2 \cdot a_3 \cdot b_1 \cdot c_2 + a_1 \cdot a_2^2 \cdot b_2 \cdot c_2 - a_1^2 \cdot a_3 \cdot b_3 \cdot c_2) \cdot y + c_2 \cdot (3 \cdot a_1^2 \cdot b_2 \cdot b_3 \cdot c_2 - \\
 & 2 \cdot a_1 \cdot a_3 \cdot b_1 \cdot b_3 \cdot c_2 + 3 \cdot a_2 \cdot a_3 \cdot b_2 \cdot b_3 \cdot c_1 + a_2^2 \cdot b_1 \cdot b_2 \cdot c_2) \cdot x + (-a_2^3 \cdot b_2^2 \cdot c_1 \cdot c_2 + \\
 & 2 \cdot a_2 \cdot a_3^2 \cdot b_1 \cdot b_3 \cdot c_1 \cdot c_2 - 3 \cdot a \cdot a_2 \cdot a_3 \cdot b_2 \cdot b_3 \cdot c_1 \cdot c_2 + a_2^2 \cdot a_3 \cdot b_1^2 \cdot c_2^2 - \\
 & a_1 \cdot a_2^2 \cdot b_1 \cdot b_2 \cdot c_2^2 + a_1^2 \cdot a_3 \cdot b_1 \cdot b_3 \cdot c_2^2 - a_1^3 \cdot b_2 \cdot b_3 \cdot c_2^2 + a_3^3 \cdot b_3^2 \cdot c_1^2).
 \end{aligned}$$

The Newton polytope of the equation and its dual graph are illustrated in Figure 5.12. As expected, this dual graph agrees with the constructed tropical surface graph. In Section 5.4 we revisit this example and explain how certain specializations of the coefficients a_1 through c_2 give a new facet of the polytope, by setting the constant coefficient of the defining equation to zero, and destroy the genericity conditions on the polynomial map \mathbf{f} . \diamond

5.4 Tropical implicitization for non-generic surfaces

In this section, we discuss methods for computing the tropicalization of a parametric surface in the non-generic case. The obstruction to apply generic methods lies in the failure of the




Figure 5.9: Weighted simplicial complex representing TY .




Figure 5.10: From left to right: Newton polytopes \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}_3 , together with the common refinement of their inner normal fans. The weight 2 (in blue) on r_1 indicates the weight of this ray in $\mathcal{N}(\mathcal{P}_2)$.

combinatorial normal crossing condition. We explain how to solve this issue and present several numerical examples.

Our setting is as follows: we are given n plane curves defined by Laurent polynomials f_1, \dots, f_n with fixed support and we allow special choices of coefficients that preserve these supports. Assume that the map $\mathbf{f} = (f_1, \dots, f_n): \mathbb{T}^2 \dashrightarrow \mathbb{T}^n$ is generically finite. As in the previous section, our goal is to tropicalize the open set $X = \mathbb{T}^2 \setminus \bigcup_{i=1}^n (f_i = 0)$ using geometric tropicalization, by finding a tropical compactification of X . The following lemma implies that we can assume all f_i 's are *irreducible*.

Lemma 5.4.1. *Assume \mathbf{f} is a finite map and that f_1 factors as $f_1 = gh$ with $\deg g, \deg h < \deg f_1$. Then, the map $\mathbf{f}' = (g, h, f_2, \dots, f_n): X \rightarrow \mathbb{T}^{n+1}$ is generically finite and $\mathbf{f} = \alpha \circ \mathbf{f}'$, where $\alpha: \mathbb{T}^{n+1} \rightarrow \mathbb{T}^n$ sends (t_0, t_1, \dots, t_n) to $(t_0 t_1, t_2, \dots, t_n)$. In addition, α restricted to the image of \mathbf{f}' is generically finite.*

the ambient space $\mathbb{P}(\Sigma)$ by toric blow-ups, refining the fan to a strictly simplicial fan $\Sigma' \subset \mathbb{R}^2$, performing classical point blow-ups on the smooth surface $\mathbb{P}(\Sigma')$ and finally pull-back the variety $X(\Sigma')$ along this resolution. This procedure will be tedious to do in practice, so we will not follow this approach in concrete examples.

Our alternative strategy in the non-generic case does not take advantage of the combinatorial data provided by the support of our polynomial map. However, it will have the advantage of being simpler to carry out in concrete examples. Here is our main result:

Theorem 5.4.2. *Let $\mathbf{f} = (f_1, \dots, f_n) : \mathbb{T}^2 \dashrightarrow \mathbb{T}^n$ by a finite map of degree δ , where $f_i \in \mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}]$ are irreducible Laurent polynomials. Let X' be the compactification of $X = \mathbb{T}^2 \setminus \bigcup_{i=1}^n (f_i = 0)$ inside \mathbb{P}^2 . The compact space X' has $n+1$ boundary divisors: $D_i = (f_i = 0)$ and $D_\infty = (x_3 = 0)$. Assume X' does not satisfy the combinatorial normal crossing boundary condition. Let $\phi: \tilde{X} \rightarrow X'$ be any resolution of X' obtained by blowing up all intersection points of three or more boundary components, so that \tilde{X} satisfies the combinatorial normal crossing boundary condition and the tangency directions of the branches of intersecting curves at crossing points are different. Let E_1, \dots, E_s be the corresponding exceptional divisors and D'_i, D'_∞ be the strict transforms of the divisors D_i, D_∞ , $i = 1, \dots, n$. Write:*

$$\phi^*(D_i) = D'_i + \sum_{j=1}^s b_{ij} \cdot E_j, \quad i = 1, \dots, n, \quad \phi^*(D_\infty) = D'_\infty + \sum_{j=1}^s b_j E_j,$$

for suitable $b_{ij}, b_j \in \mathbb{Z}$. Then, the tropical surface graph $\overline{\text{Tim } \mathbf{f}}$ has $n + 1 + s$ nodes

$$[D'_i] = e_i, \quad i = 1, \dots, n, \quad [D'_\infty] = (-\deg f_1, \dots, -\deg f_n),$$

$$[E_r] = (b_{1r} - b_r \deg f_1, \dots, b_{nr} - b_r \deg f_n), \quad r = 1, \dots, s.$$

The edges of this graph are given by all pairs of nodes whose associated divisors intersect in \tilde{X} . The weight of an edge (v, w) equals

$$m_{(v,w)} = \frac{1}{\delta} i(v, w) \gcd(2 \times 2 - \text{minors}(v | w)),$$

where $i(v, w)$ is the intersection number of the corresponding boundary divisors in \tilde{X} .

Proof. We start by constructing a naive compactification of the set X , by taking its closure inside \mathbb{P}^2 , which we denote by X' . In addition, we extend the map \mathbf{f} from X to X' , indicated in the following commutative diagram:

$$\begin{array}{ccc} X' & & \\ \pi \downarrow & \searrow \tilde{\mathbf{f}} & \\ X & \xrightarrow{\mathbf{f}} & Y \subset \mathbb{T}^n. \end{array}$$

The i^{th} coordinate of $\tilde{\mathbf{f}}$ is a rational function whose numerator is the homogenization f_i^h of the polynomial f_i with respect to a new variable x_0 and whose denominator equals $x_0^{\deg f_i}$. In particular, \tilde{f}_i is a degree zero rational function in three variables, hence defined over \mathbb{P}^2 .

The variety $X' \subset \mathbb{P}^2$ has a divisorial boundary with $n + 1$ irreducible components: the n divisors $D_i = (f_i^h = 0) \subset \mathbb{P}^2$, $i = 1, \dots, n$ and the divisor at infinity $D_\infty = (x_0 = 0)$. By construction, the pull-back along $\tilde{\mathbf{f}}$ of the basis of characters $\{\chi_1, \dots, \chi_n\}$ is

$$\tilde{\mathbf{f}}^*(\chi_j) = D_j + (-\deg f_i) \cdot D_\infty \quad j = 1, \dots, n.$$

Finally, we take a resolution $\phi: \tilde{X} \rightarrow X'$ by blowing up the excessive intersection points on the boundary of X' , where at least three boundary components meet. The set \tilde{X} together with the map $g = \tilde{\mathbf{f}} \circ \phi$ constructs the desired tropical compactification. If E_1, \dots, E_s denote the exceptional divisors on \tilde{X} , we obtain

$$g^*(\chi_j) = D_j + (-\deg f_i) \cdot D_\infty + \sum_{r=1}^s a_{jr} \cdot E_r \quad j = 1, \dots, n.$$

As a result, the divisorial valuations of the $n + 1 + r$ boundary divisors on \tilde{X} give the following nodes in the tropical surface graph $\mathcal{T}Y$:

$$[D_j] = e_j \quad j = 1, \dots, n, \quad [D_\infty] = (-\deg f_1, \dots, -\deg f_n), \quad [E_r] = (a_{jr})_{j=1}^n \quad r = 1, \dots, s.$$

Note that some of these lattice points in \mathbb{Z}^n could coincide. Thus, the tropical surface graph has at most $n + 1 + r$ nodes. The identity $a_{jr} = b_{jr} - b_r \deg f_j$, $j = 1, \dots, s$ follows by construction.

We now explain the transition from $\Delta_{X', D'}$ to $\Delta_{\tilde{X}, \tilde{D}}$. Assume a maximal subfamily of $k \geq 3$ boundary prime divisors of X' indexed by a subset I of $\{1, \dots, n\} \cup \{\infty\}$ meets at a point. In particular, the abstract complex $\Delta_{X', D}$ would contain a $(k-1)$ -dimensional cell with vertices indexed by I . Each blow-up produces a subdivision of this cell, ultimately leading to a graph embedded in this k -cell. At each step of the resolution, excessive intersection points give an exceptional divisor and the remaining bad crossing points have lower multiplicity. The intersection complex $\Delta_{\tilde{X}, \tilde{D}}$ is obtained by gluing all these resolution diagrams along common labeled nodes and also adding edges corresponding to pairwise intersections of boundary components. The realization of this abstract graph in \mathbb{R}^n is obtained by keeping track of the proper transforms of the boundary components of D' and all exceptional divisors along the resolution.

The formula for the weights on the tropical surface graph follows from Corollary 5.2.5 and Proposition 5.2.10. \square

We conclude this section with two numerical examples illustrating the previous construction. Further examples were given in Section 4.3, when studying a family of surfaces in \mathbb{T}^n parameterized by binomials. To simplify notation, we let s, t be our domain parameters and we call u the homogenizing variable.

Example 5.4.3. [Example 5.3.4 revisited] We consider a particular choice of coefficients for the parameterized surface in Example 5.3.4. In this case, the coordinates of our map are the following three bivariate polynomials:

$$\begin{cases} f_1(s, t) = s^2 - s^3 - t^2, \\ f_2(s, t) = t^2 - t^3 - s^2, \\ f_3(s, t) = (s + t)^2 - (s + t)^3 - (s - t)^2 = 4st - s^3 - t^3 - 3st^2 - 3s^2t. \end{cases} \tag{5.12}$$

This map has degree $\delta = 1$. As we see from the left-most picture in Figure 5.13, the curves $(f_i = 0)$, $i = 1, 2, 3$, intersect at the origin. Therefore, our surface is non-generic. We remove this triple intersection by means of resolution of singularities.

Since our polynomials f_1, f_2, f_3 have nonnegative exponents, we consider $X = \mathbb{C}^2 \setminus \bigcup_{i=1}^3 (f_i = 0)$ and its compactification in \mathbb{P}^2 . From the figure we see that the divisors $D_i = (f_i = 0)$ intersect at the origin. The intersection complex of the boundary divisors of X' and its resolution by four blow-ups is illustrated in the right of Figure 5.13.




Figure 5.13: Boundary of the closure of X in \mathbb{P}^2 and its resolution after four blow-ups.

Let $g = f \circ \pi: \tilde{X} \rightarrow X' \rightarrow Y$. Then, $g^*(\chi_1) = D_1 + 2E_1 + 3E_2 + 3E_3 + 4E_4 - 3D_\infty$, $g^*(\chi_2) = D_2 + 2E_1 + 3E_2 + 3E_3 + 4E_4 - 3D_\infty$, $g^*(\chi_3) = D_3 + 2E_1 + 2E_2 + 2E_3 + 2E_4 - 3D_\infty$. Thus, $[D_i] = e_i$, $[D_\infty] = (-3, -3, -3)$, $[E_1] = (2, 2, 2)$, $[E_2] = [E_3] = (3, 3, 2)$, and $[E_4] = (4, 4, 2)$. The tropical surface graph of Y has six nodes and twelve edges and is illustrated in Figure 5.14. Notice that the abstract graph of the intersection complex has one bivalent node and two nodes E_2 and E_3 that map to the same integer point. \diamond




Figure 5.14: Weighted simplicial complex representing $\mathcal{T}Y$.

Example 5.4.4. [Example 5.3.5 revisited] We choose parameter values for the parameterization (5.11) to obtain a non-generic surface Y in \mathbb{T}^2 and its defining degree one map:

$$\begin{cases} f_1(s, t) = s - t, \\ f_2(s, t) = t - s^2, \\ f_3(s, t) = -1 + st. \end{cases} \quad (5.13)$$

This choice of specialization of the coefficients in (5.11) forces the constant term of the implicit equation of Y to equal one. In particular, this implies that the tropical surface graph has one extra node: the one corresponding to the extra facet that appears in the Newton polytope, shown in Figure 5.17).

The left-most picture in Figure 5.15 shows the plane curves defined by each coordinate of the parameterization (5.13). As we see, these three curves intersect at the point $(1, 1)$ and, thus, the surface Y is non-generic.




Figure 5.15: From left to right: affine charts and boundary divisors at $u = 1$ and $t = 1$.

We now compute $\mathcal{T}Y$. We start by embedding $X = \mathbb{T}^2 \setminus \bigcup_{i=1}^3 (f_i = 0)$ in \mathbb{P}^2 and considering its compactification X' . This adds the new component $D_\infty = (u = 0)$ to the

boundary of $X' \subset \mathbb{P}^2$ and an extra singularity, coming from the intersection of $(f_2^h = 0)$, $(f_3^h = 0)$ and D_∞ . The resolution diagrams of these two singularities are depicted in Figure 5.16. The abstract boundary intersection complex is obtained by gluing these two resolution diagrams and adding edges between nodes representing prime divisors that intersect away from these two singularities. The resolution graph of the arrangement X is depicted in the right of Figure 5.17.



Figure 5.16: From left to right: resolution diagrams at $(1 : 1 : 1)$ and $(0 : 1 : 0)$.

Finally, we realize this abstract graph in \mathbb{R}^3 by pulling back the basis of characters $\{\chi_1, \chi_2, \chi_3\}$ of the torus \mathbb{T}^3 along the composition of the resolution $\pi: \tilde{X} \rightarrow X'$ and the map $\tilde{\mathbf{f}}: X' \rightarrow \mathbb{T}^3$, and computing the intersection numbers of the boundary components. If we set $D_i = (f_i(s, t) = 0) \subset \mathbb{P}^2$, $i = 1, 2, 3$, $D_\infty = (u = 0) \subset \mathbb{P}^2$ and follow the notation of Figure 5.16, we obtain

$$\begin{cases} (\tilde{\mathbf{f}} \circ \pi)^*(\chi_1) = D_1 - D_\infty - E_1 - 2E_2 + E_3, \\ (\tilde{\mathbf{f}} \circ \pi)^*(\chi_2) = D_2 - 2D_\infty - E_1 - 2E_2 + E_3, \\ (\tilde{\mathbf{f}} \circ \pi)^*(\chi_3) = D_2 - 2D_\infty - E_1 - 3E_2 + E_3. \end{cases}$$

Therefore, $[D_i] = e_i$ ($i = 1, 2, 3$), $[D_\infty] = (-1, -2, -2)$, $[E_1] = (-1, -1, -1)$, $[E_2] = (-2, -2, -3)$ and $[E_3] = (1, 1, 1)$. In addition, the nonzero intersection multiplicities are $D_1 \cdot D_2 = D_1 \cdot D_3 = E_1 \cdot D_3 = E_2 \cdot D_2 = E_2 \cdot D_\infty = E_2 \cdot E_3 = E_3 \cdot D_i = 1$ ($i = 1, 2, 3$) and $D_2 \cdot D_3 = 2$. By applying Theorem 5.2.4 and Proposition 5.2.10 we conclude that all edges have weight one, except for the edges (e_2, e_3) and $(e_1, [D_\infty])$, whose weight equals two. The resulting graph and the Newton polytope of the defining equation are shown in Figure 5.17. By comparing the latter to Figure 5.12 we see the transition from a generic to a non-generic surface obtained by specializing coordinates of a parameterization. \diamond

5.5 Further remarks

We conclude the chapter with some remarks and open questions that deserve to be studied in more detail, hoping to continue this line of research in the near future.




Figure 5.17: Newton polytope and dual graph of a non-generic surface in \mathbb{C}^3 parameterized by (5.13).

In Sections 5.3 and 5.4 we discussed the problem of computing tropical surface graphs associated to rational surfaces in \mathbb{T}^n . We assumed the given polynomial parameterization had coordinates with fixed support and we explored two solutions, based on the genericity of the complex coefficients of our map. This genericity condition was stated as a the combinatorial normal crossing boundary and a tangency requirement, i.e. the lack of triple intersections among boundary divisors in a compactification of an arrangement of plane curves plus the condition that tangent directions at the crossing points of the branchings curves were distinct.

In the non-generic case, we explained a reduction to this generic setting by blowing up points on plane curves and carrying the intersection numbers and divisorial valuations along the resolution. The examples presented show how special these resolutions are, i.e. how hard is to predict the combinatorics of the resolution by looking at the initial curve arrangement. The final divisorial valuations of the exceptional divisors heavily depend on the topology of the plane curves containing the points we blow-up.

The standard approach to obtain such valuations was introduced in work of Enriques and Chisini [33] and further developed with the notions of Enriques and dual diagrams [98]. Such methods are based on the topological type of the branches of the curves we resolve. The pairwise intersection numbers are hard to compute in concrete examples. The main difficulty lies in the construction of clusters of infinitely near points of a singularity [8]. These clusters are precisely the point configurations emanating from successive blow-ups.

From our discussion in Section 5.4 we know that the construction of the intersection complex $\Delta_{\tilde{X}, \tilde{D}}$ corresponds to the gluing of resolution diagrams (cf. Figure 5.16). If we take a careful look at the local picture near each singularity and its resolution, we can interpret these graphs as *phylogenetic trees*. The original prime divisors D_j , $j \in \{1, \dots, n\} \cup \{\infty\}$, correspond to the leaves of this tree, whereas the exceptional divisors are internal nodes in the tree. In addition, we can place the leaves of the tree at infinity and interpret these graphs as tropical trees. Since the tropical surface graph is obtained by gluing these graphs,

we expect the tropicalization of a rational surface to be a gluing of phylogenetic trees.

We propose a method to construct these phylogenetic trees. Fix an intersection point p of at least three boundary divisors of $X' \subset \mathbb{P}^2$ and perform a base change to the field of Puiseux series, or any other non-archimedean valued field whose valued group is \mathbb{Q} . Since blow-ups of points on curves locally correspond to monomial changes of variables, we can pick a value u for a slope and replace the bivariate equation f_k defining a divisor D_k containing the point p in an affine chart by the equation $g_k(t, q) = f_k(t, q^u t)$ for a fixed parameter q . We repeat this process for all divisors containing p . The polynomial g_k is interpreted as a univariate polynomial in q with coefficients in $\mathbb{C}[[t]] \subset \mathbb{C}\{\{t\}\}$. In particular, the surface parameterized by \mathbf{f} can be locally viewed as a curve in n -space over $\mathbb{C}\{\{t\}\}$. By tropicalizing the image of the map $\mathbf{g}: \mathbb{C}\{\{t\}\} \rightarrow \mathbb{C}\{\{t\}\}^n$ we obtain a tropical tree. This tropical tree will be precisely the phylogenetic tree associated to the point p .

In the last years, a new object combining both Enriques and dual graphs was introduced by Popescu-Pampu under the name of *kite* [80]. In his language, clusters of infinitely near points are called constellations. This kite has a natural interpretation in the valuative tree of Favre and Jonsson [36] and it seems to provide the best framework to study arrangements of plane curves [42]. It also shows us how to glue local resolutions of intersection points. Combining this new setting with the tools of *combinatorial resolutions* derived from Max Noether's fundamental theorem [11] we hope to obtain a classification of tropical surface graphs in three-space.

As explained in [92, Section 5.3], even if we want to tropicalize a non-generic rational surface in three-space we can still apply generic methods to obtain a tropical surface graph. As expected, this graph does not correspond to our rational surface but rather to a deformation of it, whose generic fiber is a generic surface parameterized by a polynomial map with the same support. Using this tropical surface graph as input for the implicitization algorithms described in Section 1.4, we obtain a polynomial multiple of the equation defining the original surface. More precisely:

Proposition 5.5.1. [92, Proposition 5.3] *Let f_1, f_2, f_3 be Laurent bivariate polynomials defining a surface Y in \mathbb{T}^3 by the polynomial map $\mathbf{f} = (f_1, f_2, f_3)$. Let \mathcal{Q} be the polytope constructed from the Newton polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ as in Section 5.3. Then, the polytope \mathcal{Q} contains a translate of the Newton polytope of the equation defining Y .*

Example 5.5.2. Figures 5.12 and 5.17 illustrate the previous result. Namely, the polytope \mathcal{Q} corresponds to the polytope with five vertices $(0, 0, 0)$, $(0, 0, 2)$, $(0, 2, 0)$, $(2, 1, 0)$ and $(3, 0, 0)$, which contains the Newton polytope \mathcal{P} corresponding to a non-generic surface. This polytope has seven vertices $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$, $(0, 0, 2)$, $(0, 2, 0)$, $(2, 1, 0)$ and $(3, 0, 0)$. In this case, no translation is required. \diamond

As we discussed throughout this dissertation, the tropicalization of algebraic varieties heavily depends on the field of definition. As it stands, tropical implicitization was presented

in the constant coefficient case, i.e. of algebraic varieties over the complex numbers. Since arbitrary tropical complexes behave as tropical fans from the degenerations perspective [87, Prop 2.2.3], our results for complex varieties should have analogs in the non-trivial valued case. This philosophical principle needs to be carefully verified in our constructions.

One particular instance of this extension is the case of generic rational surfaces parameterized by polynomial maps with fixed support. Instead of allowing complex coefficients, we may allow coefficients on any non-archimedean valued field \mathbb{K} with valued group \mathbb{Q} as in Section 1.1.2. In that setting, the tropical variety is a polyhedral complex in \mathbb{Q}^n and we need to extend the basics on geometric tropicalization to the arbitrary coefficient case, where toric varieties are defined over discrete valuation rings. The notion of tropical compactification of varieties defined over \mathbb{K} was recently explored by Luxton and Qu [64]. Note that the realization of the boundary intersection complex of such compactification gives a complex in \mathbb{Q}^n whose dimension is off by one from the dimension of the tropical complex. In the constant coefficient case, this is fixed by taking the cone over such complex, but we are not able to do that in the arbitrary coefficient case. Likewise, the push-forward formula for tropical multiplicities needs to be extended to the arbitrary coefficient case to obtain an analog of Theorem 5.2.4 for varieties defined over \mathbb{K} .

In recent articles, a connection between tropical geometry and Berkovich spaces was established, endowing tropical geometry with powerful analytic tools [77]. In particular, questions regarding tropical compactifications should have a natural interpretation in this new setting. Thuillier has successfully applied this principle to study homotopy types of the intersection complex of a regular compactification with simple normal crossing boundary divisor, showing that this type depends only on the original variety [95, Section 4.2]. We expect his analysis can provide insight to the question of tropical compactifications discussed in this chapter and give relevant information to construct tropical surface graphs. It would be very interesting to compare tropical implicitization techniques with similar notions on Berkovich spaces.

Input: The supports $\{\mathcal{A}_i : i = 1, \dots, n\}$ of n bivariate polynomials f_1, \dots, f_n with generic coefficients and the degree δ of the map $\mathbf{f} = (f_1, \dots, f_n)$.

Assumption: The coefficients of the polynomials are generic enough to guarantee that the intersection of any triple of distinct curves $(f_i = 0)$, $(f_j = 0)$ and $(f_k = 0)$ ($i, j, k = 1, \dots, n$) is empty.

Output: The tropical surface graph of the image of $\mathbf{f} : \mathbb{T}^2 \dashrightarrow \mathbb{T}^n$.

for each $i = 1, \dots, n$ **do**

$\mathcal{P}_i \leftarrow$ convex hull of $\mathcal{A}_i \subset \mathbb{Z}^2$.

$\mathcal{N}_i \leftarrow$ inner normal fan of \mathcal{P}_i .

$\mathcal{P} \leftarrow \bigoplus_{i=1}^n \mathcal{P}_i$.

$\mathcal{N} \leftarrow$ normal fan of \mathcal{P} .

$V \leftarrow \emptyset$.

while $\dim \mathcal{P}_i > 0$ **do**

$V \leftarrow V \cup \{e_i\}$.

for all ρ ray in $\mathcal{N}^{[1]}$ **do**

$n_\rho \leftarrow$ primitive integer generator of ρ .

$[D_\rho] \leftarrow \text{trop}(\mathbf{f})(n_\rho) \in \mathbb{Z}^n$.

if $[D_\rho] \neq \mathbf{0}$ **then**

$V \leftarrow V \cup \{[D_\rho]\}$.

$E \leftarrow \emptyset$.

for each $i = 1, \dots, n-1$ **do**

for each $j = i+1, \dots, n$ **do**

if $\dim(\mathcal{P}_i \oplus \mathcal{P}_j) = 2$ **then**

$E \leftarrow E \cup \{(e_i, e_j)\}$.

$m_{(e_i, e_j)} = \frac{1}{\delta} \text{MV}(\mathcal{P}_i, \mathcal{P}_j)$.

for each $\rho \in \mathcal{N}^{[1]}$ **do**

while $\dim \mathcal{P}_i > 0$ **do**

if $\rho \in \mathcal{N}_i^{[1]}$ and $\dim \mathbb{R}\langle [D_\rho], e_i \rangle = 2$ **then**

$E \leftarrow E \cup \{([D_\rho], e_i)\}$.

$m_{([D_\rho], e_i)} = \frac{1}{\delta} (|\text{face}_{n_\rho}(\mathcal{P}_i) \cap \mathbb{Z}^2| - 1) \text{gcd}([D_\rho]_j : j \neq i)$.

for each $\eta \in \mathcal{N}^{[1]}$, $\eta \neq \rho$ **do**

if ρ and η generate a cone in $\mathcal{N}^{[2]}$ and $\dim \mathbb{R}\langle [D_\rho], [D_\eta] \rangle = 2$ **then**

$E \leftarrow E \cup \{([D_\rho], [D_\eta])\}$.

$m_{([D_\rho], [D_\eta])} = \frac{1}{\delta} \frac{\text{gcd}(2 \times 2\text{-minors}([D_\rho][D_\eta])}{\text{gcd}(2 \times 2\text{-minors}(n_\rho | n_\eta))}$.

return The graph (V, E) .

Algorithm 5.1: Tropical Implicitization of generic surfaces: From the support of n polynomials with generic coefficients, we compute a weighted graph encoding the tropicalization of a rational variety in \mathbb{T}^n parameterized by these n polynomials.

Bibliography

- [1] Oswin Aichholzer and Franz Aurenhammer. Classifying hyperplanes in hypercubes. *SIAM J. Discrete Math.*, 9(2):225–232, 1996.
- [2] George M. Bergman. The logarithmic limit-set of an algebraic variety. *Trans. Amer. Math. Soc.*, 157:459–469, 1971.
- [3] D. N. Bernstein. The number of roots of a system of equations. *Funkcional. Anal. i Priložen.*, 9(3):1–4, 1975.
- [4] Mark R. Best and Andries E. Brouwer. The triply shortened binary Hamming code is optimal. *Discrete Math.*, 17(3):235–245, 1977.
- [5] Robert Bieri and John Richard James Groves. The geometry of the set of characters induced by valuations. *J. Reine Angew. Math.*, 347:168–195, 1984.
- [6] Tristram Bogart, Anders N. Jensen, David Speyer, Bernd Sturmfels, and Rekha R. Thomas. Computing tropical varieties. *J. Symbolic Comput.*, 42(1-2):54–73, 2007.
- [7] Andries E. Brower and Jan Draisma. Equivariant gröbner bases and the gaussian two-factor model. To appear in *Mathematics of Computation*. [arXiv:0908.1530](https://arxiv.org/abs/0908.1530).
- [8] Antonio Campillo, Gérard Gonzalez-Sprinberg, and Monique Lejeune-Jalabert. Clusters of infinitely near points. *Math. Ann.*, 306(1):169–194, 1996.
- [9] Michael L. Catalano-Johnson. The possible dimensions of the higher secant varieties. *Amer. J. Math.*, 118(2):355–361, 1996.
- [10] Maria Virginia Catalisano, Anthony V. Geramita, and Alessandro Gimigliano. Secant varieties of $\mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ (n -times) are not defective for $n \geq 5$. [arXiv:0809.1701](https://arxiv.org/abs/0809.1701).
- [11] José Ignacio Cogolludo Agustín and Miguel Angel Marco-Buzunariz. The Max Noether fundamental theorem is combinatorial. [arXiv:1002.2325](https://arxiv.org/abs/1002.2325), 2010.
- [12] Gérard Cohen, Iiro Honkala, Simon Litsyn, and Antoine Lobstein. *Covering codes*, volume 54 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam, 1997.

- [13] Aldo Conca. Straightening law and powers of determinantal ideals of Hankel matrices. *Adv. Math.*, 138(2):263–292, 1998.
- [14] Tom M. Cover and Joy A. Thomas. *Elements of information theory*. John Wiley and sons, second edition, 2006.
- [15] David Cox. Toric varieties and toric resolutions. In *Resolution of singularities (Obergurgl, 1997)*, volume 181 of *Progr. Math.*, pages 259–284. Birkhäuser, Basel, 2000.
- [16] David Cox and Jessica Sidman. Secant varieties of toric varieties. *J. Pure Appl. Algebra*, 209(3):651–669, 2007.
- [17] María Angélica Cueto and Shaowei Lin. Tropical secant graphs of monomial curves. Contribution FPSAC 2010 (San Francisco, CA, USA), [arXiv:1005.3364v1](https://arxiv.org/abs/1005.3364v1), 2010.
- [18] Mara Anglica Cueto, Enrique A. Tobis, and Josephine Yu. An implicitization challenge for binary factor analysis. *J. Symbolic Comput.*, 45(12):1296–1315, 2010.
- [19] Corrado De Concini and Claudio Procesi. Wonderful models of subspace arrangements. *Selecta Math. (N.S.)*, 1(3):459–494, 1995.
- [20] Jesús A. De Loera, David Haws, Rraymond Hemmecke, Peter Huggins, Jeremy Tauzer, and Ruriko Yoshida. A user’s guide for latte v1.1. Available at <http://www.math.ucdavis.edu/~latte>, 2003.
- [21] Wolfram Decker, Gert-Martin Greuel, Gerhard Pfister, and Hans Schönemann. SINGULAR 3-1-1 – A computer algebra system for polynomial computations, 2010.
- [22] Mike Develin. Tropical secant varieties of linear spaces. *Discrete Comput. Geom.*, 35(1):117–129, 2006.
- [23] Mike Develin, Francisco Santos, and Bernd Sturmfels. On the rank of a tropical matrix. In *Combinatorial and computational geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 213–242. Cambridge Univ. Press, Cambridge, 2005.
- [24] Alicia Dickenstein, Eva Maria Feichtner, and Bernd Sturmfels. Tropical discriminants. *J. Amer. Math. Soc.*, 20(4):1111–1133 (electronic), 2007.
- [25] Jan Draisma. A tropical approach to secant dimensions. *J. Pure Appl. Algebra*, 212(2):349–363, 2008.
- [26] Mathias Drton, Bernd Sturmfels, and Bernd Sullivant. Algebraic factor analysis: Tetrads, pentads and beyond. *Probab. Theory Related Fields*, 138:463–493, 2007.
- [27] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on Algebraic Statistics*, volume 39 of *Oberwolfach Seminars*. Birkhäuser, 2009.

- [28] Mathias Drton and Seth Sullivant. Algebraic statistical models. *Statist. Sinica*, 17(4):1273–1297, 2007.
- [29] Roberto Dvornicich and Umberto Zannier. Newton functions generating symmetric fields and irreducibility of Schur polynomials. *Adv. Math.*, 222(6):1982–2003, 2009.
- [30] Manfred Einsiedler, Mikhail M. Kapranov, and Douglas Lind. Non-Archimedean amoebas and tropical varieties. *J. Reine Angew. Math.*, 601:139–157, 2006.
- [31] David Eisenbud. *Commutative algebra with a view toward algebraic geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [32] Sergi Elizalde and Kevin Woods. Bounds on the number of inference functions of a graphical model. *Statist. Sinica*, 17(4):1395–1415, 2007.
- [33] Federigo Enriques and Oscar Chisini. *Lezioni sulla teoria geometrica delle equazioni e delle funzioni algebriche. 2. Vol. III, IV*, volume 5 of *Collana di Matematica [Mathematics Collection]*. Nicola Zanichelli Editore S.p.A., Bologna, 1985. Reprint of the 1924 and 1934 editions.
- [34] Nicholas Eriksson, Kristian Ranestad, Bernd Sturmfels, and Seth Sullivant. Phylogenetic algebraic geometry. In *Projective varieties with unexpected properties*, pages 237–255. Walter de Gruyter GmbH & Co. KG, Berlin, 2005.
- [35] A. Esterov and Askold G. Khovanskiĭ. Elimination theory and Newton polytopes. *Funct. Anal. Other Math.*, 2(1):45–71, 2008.
- [36] Charles Favre and Mattias Jonsson. *The valuative tree*, volume 1853 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2004.
- [37] Alex Fink. *Matroid polytope subdivisions and valuations*. PhD thesis, University of California - Berkeley, 2010.
- [38] Alex Fink. Tropical cycles and Chow polytopes. [arXiv:1001.4784v2](https://arxiv.org/abs/1001.4784v2), 2010.
- [39] William Fulton. *Introduction to toric varieties*, volume 131 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1993. The William H. Roever Lectures in Geometry.
- [40] Mark Galassi, Jim Davies, James Theiler, Brian Gough, Gerard Jungman, Michael Booth, and Fabrice Rossi. *GNU Scientific Library Reference Manual - Third Edition*. Network Theory Ltd., second edition, 2006. <http://www.gnu.org/software/gsl/>.
- [41] Luis David García, Michael E. Stillman, and Bernd Sturmfels. Algebraic geometry of Bayesian networks. *J. Symbolic Comput.*, 39(3-4):331–355, 2005.

- [42] Evelia García Barroso and Patrick Popescu-Pampu. The kite of a plane curve singularity. In preparation, 2010.
- [43] Ewgenij Gawrilow and Michael Joswig. polymake: a framework for analyzing convex polytopes. In Gil Kalai and Günter M. Ziegler, editors, *Polytopes — Combinatorics and Computation*, pages 43–74. Birkhäuser, 2000.
- [44] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www.math.uiuc.edu/Macaulay2/>, 2009.
- [45] Paul Hacking, Sean Keel, and Jenia Tevelev. Stable pair, tropical, and log canonical compactifications of moduli spaces of del Pezzo surfaces. *Invent. Math.*, 178(1):173–227, 2009.
- [46] Joe Harris. *Algebraic geometry, A first course*, volume 133 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. Corrected reprint of the 1992 original.
- [47] Robin Hartshorne. *Algebraic geometry*, volume 52 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1977.
- [48] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning*. Springer Series in Statistics. Springer-Verlag, New York, 2nd edition, 2009. Data mining, inference, and prediction.
- [49] Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Comput.*, 18(7):1527–1554, 2006.
- [50] Geoffrey E. Hinton and Ruslan R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [51] W. Cary Huffman and Vera Pless. *Fundamentals of error-correcting codes*. Cambridge Univ Pr, 2003.
- [52] Ilia Itenberg and Oleg Viro. Patchworking algebraic curves disproves the Ragsdale conjecture. *Math. Intelligencer*, 18(4):19–28, 1996.
- [53] Anders Jensen. *Algorithmic Aspects of Gröbner Fans and Tropical Varieties*. PhD thesis, Aarhus Universitet, 2007.
- [54] Anders N. Jensen. Gfan, a software system for Gröbner fans and tropical varieties. Available at <http://www.math.tu-berlin.de/~jensen/software/gfan/gfan.html>, 2009.
- [55] Mikhail M. Kapranov, Bernd Sturmfels, and Andrei V. Zelevinsky. Chow polytopes and general resultants. *Duke Math. J.*, 67(1):189–218, 1992.

- [56] Eric Katz. A tropical toolkit. *Expo. Math.*, 27(1):1–36, 2009.
- [57] Askold G. Khovanskii. Newton polyhedra, a new formula for mixed volume, product of roots of a system of equations. In *The Arnoldfest (Toronto, ON, 1997)*, volume 24 of *Fields Inst. Commun.*, pages 325–364. Amer. Math. Soc., Providence, RI, 1999.
- [58] Joseph M. Landsberg and Laurent Manivel. On the ideals of secant varieties of Segre varieties. *Found. Comput. Math.*, 4(4):397–422, 2004.
- [59] Joseph M. Landsberg and Jerzy Weyman. On the ideals and singularities of secant varieties of Segre varieties. *Bull. Lond. Math. Soc.*, 39(4):685–697, 2007.
- [60] Serge Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 2002.
- [61] Nicolas Le Roux and Yoshua Bengio. Representational power of restricted Boltzmann machines and deep belief networks. *Neural Comput.*, 20(6):1631–1649, 2008.
- [62] Simon Litsyn, E. M. Rains, and Neil J. A. Sloane. Table of nonlinear binary codes. <http://www.eng.tau.ac.il/~litsyn/tableand/index.html>, September 1999.
- [63] Mark Lutz and David Ascher. *Learning Python*. O’Reilly & Associates, Inc. Sebastopol, CA, USA, first edition, 1999.
- [64] Mark Luxton and Zhenhua Qu. On tropical compactifications. [arXiv:0902.2009v2](https://arxiv.org/abs/0902.2009v2), 2009.
- [65] Diane Maclagan. notes from the AARMS Tropical Geometry summer school. <http://www.warwick.ac.uk/staff/D.Maclagan/AARMS/AARMSTropical.pdf>, 2008.
- [66] Hannah Markwig and Josephine Yu. The space of tropically collinear points is shellable. *Collect. Math.*, 60(1):63–77, 2009.
- [67] William M. McEneaney. *Max-plus methods for nonlinear control and estimation*. Systems & Control: Foundations & Applications. Birkhäuser Boston Inc., Boston, MA, 2006.
- [68] Grigory Mikhalkin. Counting curves via lattice paths in polygons. *C. R. Math. Acad. Sci. Paris*, 336(8):629–634, 2003.
- [69] Grigory Mikhalkin. Tropical geometry and its applications. In *International Congress of Mathematicians. Vol. II*, pages 827–852. Eur. Math. Soc., Zürich, 2006.
- [70] Marvin L. Minsky and Seymour Papert. *Perceptrons: An introduction to computational geometry*. MIT Press, Cambridge, MA, 1969.

- [71] Luca Moci. Wonderful models for toric arrangements. [arXiv:0912.5461](#), accepted for publication in *Int. Math. Res. Not.*, 2010.
- [72] Michael B. Monagan, Keith O. Geddes, K. Michael Heal, George Labahn, Stefan M. Vorkoetter, James McCarron, and Paul DeMarco. *Maple 10 Programming Guide*. Maplesoft, Waterloo ON, Canada, 2005.
- [73] Piyush C. Ojha. Enumeration of linear threshold functions from the lattice of hyperplane intersections. *IEEE Trans. Neural Netw.*, 11(4):839–850, 2000.
- [74] Brian Osserman and Sam Payne. Lifting tropical intersections. [arXiv:1007.1314](#), 2010.
- [75] Lior Pachter and Bernd Sturmfels. Tropical geometry of statistical models. *Proc. Natl. Acad. Sci. USA*, 101(46):16132–16137 (electronic), 2004.
- [76] Lior Pachter and Bernd Sturmfels, editors. *Algebraic statistics for computational biology*. Cambridge University Press, New York, 2005.
- [77] Sam Payne. Analytification is the limit of all tropicalizations. *Math. Res. Lett.*, 16(3):543–556, 2009.
- [78] Sam Payne. Fibers of tropicalization. *Math. Z.*, 262(2):301–311, 2009.
- [79] Jean-Eric Pin. Tropical semirings. In *Idempotency (Bristol, 1994)*, volume 11 of *Publ. Newton Inst.*, pages 50–69. Cambridge Univ. Press, Cambridge, 1998.
- [80] Patrick Popescu-Pampu. Le cerf-volant d’une constellation. [arXiv:0906.2932](#), 2009.
- [81] Claudiu Raicu. The GSS conjecture. [arXiv:1011.5867v1](#), 2010.
- [82] Kristian Ranestad. The degree of the secant variety and the join of monomial curves. *Collect. Math.*, 57(1):27–41, 2006.
- [83] Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald. First steps in tropical geometry. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 2005.
- [84] Frank Rosenblatt. *Principles of neurodynamics: Perceptrons and the theory of brain mechanisms*. Spartan Books, Washington, DC, 1962.
- [85] Igor Shafarevich. *Basic Algebraic Geometry I - Varieties in Projective Space*. Springer-Verlag, second edition, 1994.
- [86] Neil J.A. Sloane. The on-line encyclopedia of integer sequences. [www.research.att.com/~njas/sequences/](#), 2008.

- [87] David Speyer. *Tropical geometry*. PhD thesis, University of California - Berkeley, 2005.
- [88] David Speyer and Bernd Sturmfels. Tropical mathematics. *Math. Mag.*, 82(3):163–173, 2009.
- [89] Bjarne Stroustrup. *The C++ programming language*. Addison-Wesley Reading, MA, 1997.
- [90] Bernd Sturmfels. *Solving systems of polynomial equations*, volume 97 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2002.
- [91] Bernd Sturmfels and Jenia Tevelev. Elimination theory for tropical varieties. *Math. Res. Lett.*, 15(3):543–562, 2008.
- [92] Bernd Sturmfels, Jenia Tevelev, and Josephine Yu. The Newton polytope of the implicit equation. *Mosc. Math. J.*, 7(2):327–346, 351, 2007.
- [93] Bernd Sturmfels and Josephine Yu. Tropical implicitization and mixed fiber polytopes. In Michael E. Stillman, Nobuki Takayama, and Jan Verschelde, editors, *Software for Algebraic Geometry*, volume 148 of *I.M.A. Volumes in Mathematics and its Applications*, pages 111–132, New York, 2008. Springer.
- [94] Jenia Tevelev. Compactifications of subvarieties of tori. *Amer. J. Math.*, 129(4):1087–1104, 2007.
- [95] Amaury Thuillier. Géométrie toroïdale et géométrie analytique non archimédienne. Application au type d’homotopie de certains schémas formels. *Manuscripta Math.*, 123(4):381–451, 2007.
- [96] Giuseppe Valla. On determinantal ideals which are set-theoretic complete intersections. *Compositio Math.*, 42(1):3–11, 1980/81.
- [97] Oleg Viro. From the sixteenth Hilbert problem to tropical geometry. *Jpn. J. Math.*, 3(2):185–214, 2008.
- [98] C. T. C. Wall. *Singular points of plane curves*, volume 63 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2004.
- [99] Roland Wunderling. Paralleler und objektorientierter simplex-algorithmus. *ZIB Technical Report TR 96-09, TU Berlin*, 1996.
- [100] Piotr Zwiernik and Jim Q. Smith. The geometry of conditional independence tree models with hidden variables. [arXiv:0904.1980](https://arxiv.org/abs/0904.1980).