

NEWTON POLYTOPES AND NUMERICAL ALGEBRAIC GEOMETRY

A Dissertation
by
TAYLOR CHRISTIAN BRYSEWICZ

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Chair of Committee, Frank Sottile
Committee Members, Laura Matusevich
Andrea Bonito
Christopher Menzel
Head of Department, Sarah Witherspoon

May 2020

Major Subject: Mathematics

Copyright 2020 Taylor Christian Brysiewicz

ABSTRACT

We develop a collection of numerical algorithms which connect ideas from polyhedral geometry and algebraic geometry. The first algorithm we develop functions as a numerical oracle for the Newton polytope of a hypersurface and is based on ideas of Hauenstein and Sottile. Additionally, we construct a numerical tropical membership algorithm which uses the former algorithm as a subroutine. Based on recent results of Esterov, we give an algorithm which recursively solves a sparse polynomial system when the support of that system is either lacunary or triangular. Prior to explaining these results, we give necessary background on polytopes, algebraic geometry, monodromy groups of branched covers, and numerical algebraic geometry.

DEDICATION

To my father

ACKNOWLEDGMENTS

I offer my deepest gratitude to my family, friends, teachers, and mentors whom have helped me along my journey. Without their love and support, none of this work would have begun.

To my parents, Kathleen and Robert Brysiewicz, for their unconditional love. To Nicholas, Bobbi, Shelby, Alexandra, and Tobias, for constantly supporting me, always picking up the phone, and generally mapping out corners of the world before I need to.

To Alex, Hank, Fulvio, and all of the other friends and colleagues I have met throughout graduate school, for the countless conversations about mathematics. To my best friend, Jamie, for her unwavering love and reminders to breathe.

To my teachers, for believing in me, in particular, Alexandra Brysiewicz, Ginger Benning, Karen Yerly, Sue Samonds, Joan Kustak, Andrew Wang, Deepak Naidu, Michael Geline, and Seth Dutter. To Peter Howard and Monique Stewart for helping me navigate graduate school. To Laura Matusevich, Andrea Bonito, and Christopher Menzel for serving on my committee. To Jonathan Hauenstein, Michael Burr, Christopher O'Neill, Timo de Wolff, Laura Matusevich, Anton Leykin, Bernd Sturmfels, and Cynthia Vinzant, for their general mentorship.

Finally, to Frank (Sottile), whose commitment, thoughtfulness, and mathematical insight are more than I could ask for in an advisor. He has helped me grow in every facet of what it means to be a professional mathematician, and for that, I am forever grateful.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professor Frank Sottile [advisor], Professor Laura Matusevich, and Professor Andrea Bonito of the Department of Mathematics and Professor Christopher Menzel of the Department of Philosophy. The material in Section 8 is joint work with Jose Rodriguez, Frank Sottile, and Thomas Yahl.

All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by a graduate fellowship from Texas A&M University. The material in Section 7 was supported by NSF grant DMS-1501370 and completed during the ICERM-2018 semester on nonlinear algebra.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
DEDICATION	iii
ACKNOWLEDGMENTS	iv
CONTRIBUTORS AND FUNDING SOURCES	v
TABLE OF CONTENTS	vi
LIST OF FIGURES	ix
LIST OF TABLES.....	xiii
1. INTRODUCTION	1
2. POLYTOPES	3
2.1 Describing polytopes.....	3
2.2 Oracles	7
2.3 Mixed volume	11
2.4 Subdivisions	14
2.5 Monotonicity and positivity of mixed volume	17
3. ALGEBRAIC GEOMETRY	21
3.1 Affine varieties	21
3.2 Projective varieties	25
3.3 Charts on projective space	28
3.4 Homogenizing and dehomogenizing	28
3.5 Regular functions	29
3.6 Irreducibility and dimension.....	30
3.6.1 Irreducibility	30
3.6.2 Dimension	31
3.7 Function fields and rational functions	33
3.8 Products, graphs, and the degree of a variety	33
3.8.1 Segre maps.....	34
3.8.2 Products	34
3.8.3 Graphs.....	34
3.8.4 Dominant maps	35

3.8.5	Degree of a variety	36
3.9	Singular points	37
4.	BRANCHED COVERS AND GROUPS	39
4.1	Branched covers	40
4.2	Permutation groups	40
4.3	Monodromy groups and Galois groups	41
4.4	Decomposable branched covers	41
4.5	Real branched covers	43
5.	NEWTON POLYTOPES, SUPPORT, TROPICAL GEOMETRY, AND SPARSE POLY- NOMIAL SYSTEMS	45
5.1	Newton polytopes	45
5.1.1	Basic observations about Newton polytopes	46
5.1.2	Integer linear algebra and coordinate changes	48
5.1.3	Smith normal form	49
5.1.4	Centroids and trace curves	49
5.2	Tropical geometry	53
5.3	Sparse polynomial systems	55
5.3.1	Geometry of sparse polynomial systems	55
6.	NUMERICAL ALGEBRAIC GEOMETRY	57
6.1	Core numerical methods	57
6.1.1	Approximate solutions	57
6.1.2	Euler's method	59
6.1.3	Predictor-corrector methods	60
6.1.4	Numerical errors	61
6.2	Homotopies	63
6.2.1	Homotopy continuation	65
6.2.2	Endgames	66
6.2.3	Cauchy endgame	67
6.3	Homotopy continuation methods	70
6.3.1	Parameter homotopies	70
6.3.2	The Bézout homotopy	71
6.3.3	The polyhedral homotopy	72
6.4	Witness sets	76
6.4.1	The witness cover	76
6.4.2	Witness sets for images of maps	77
6.5	Monodromy	80
6.5.1	Solving via monodromy	82
6.5.2	Monodromy solving for real branched covers	83
6.5.3	Expected success of monodromy solving	85

7. NEWTON POLYTOPES AND TROPICAL MEMBERSHIP VIA NUMERICAL ALGEBRAIC GEOMETRY	90
7.1 The HS-Algorithm	90
7.2 Tropical membership.....	93
7.3 Convergence rates of the HS-Algorithm	94
7.4 Implementation of Algorithms 7.1.2 and 7.2.2	99
7.5 A hypersurface from algebraic vision	104
7.6 The Lüroth invariant	107
7.6.1 The Lüroth invariant, hypersurface, and polytope.....	107
7.6.2 Computational setup.....	109
7.6.3 Vertices of the Lüroth polytope	111
8. SOLVING SPARSE DECOMPOSABLE SYSTEMS	123
8.1 Decompositions of sparse polynomial systems	124
8.1.1 Lacunary support	125
8.1.2 Triangular support	127
8.2 Computing the decompositions.....	131
8.3 Solving decomposable sparse systems	132
8.4 Start systems.....	135
8.5 A computational experiment.....	138
9. SUMMARY	143
REFERENCES	144

LIST OF FIGURES

FIGURE	Page
2.1 An integral polytope $Q \subset \mathbb{R}^2$	4
2.2 Left: Directions $\omega_1, \omega_2, \omega_3$, and $\omega_4 = (0, 0)$. Right: The polytope $Q \subset \mathbb{R}^2$ and three of its proper faces exposed by ω_1, ω_2 , and ω_3	5
2.3 A polytope and its corresponding normal fan.	6
2.4 Five halfspaces in \mathbb{R}^2 whose intersection is Q	7
2.5 A graphical interpretation of Algorithm 2.2.3 running on the polytope Q in Example 2.1.3.	10
2.6 A graphic expressing $\text{vol}(\lambda_1 A + \lambda_2 B)$ for two polygons $A, B \subset \mathbb{R}^2$	13
2.7 A lifting of a dilated square and the corresponding polyhedral subdivision.	17
2.8 A coherent fine mixed subdivision.	18
2.9 The projection of a coherent fine mixed subdivision.	18
2.10 Polytopes P_1, Q_1 , and Q_2 as in Example 2.5.3 along with a fine mixed subdivision displaying that $\text{MV}(P_1, Q_2) = \text{MV}(Q_1, Q_2)$	19
3.1 The set $\mathcal{V}(0)$ defines $\mathbb{C}^2 \subset \mathbb{C}^2$	21
3.2 The set $\mathcal{V}(1)$ defines $\emptyset \subset \mathbb{C}^2$	21
3.3 The set $\mathcal{V}(x^2 + y^2 - 1)$ defines the unit circle in \mathbb{C}^2	21
3.4 The set $\mathcal{V}(x - y)$ defines a line in \mathbb{C}^2	22
3.5 The set $\mathcal{V}(x - a, y - b)$ defines a single point $(a, b) \subset \mathbb{C}^2$	22
3.6 The set $\mathcal{V}(x^2 + y^2 - 1, x - y)$ defines two points.	24
3.7 The set $\mathcal{V}((x^2 + y^2 - 1) \cdot (x - y))$ defines the union of the unit circle and a line.	24
3.8 Left: An affine cone over a quartic curve in \mathbb{P}^2 . Right: An affine cone over a circle in \mathbb{P}^2	26
3.9 A twisted cubic.	29

3.10	The affine varieties $\mathcal{V}(xy - 1)$ and $\mathcal{V}(x)$	32
3.11	Three points on a quartic curve in $\mathcal{V}(f) \subset \mathbb{C}_t \times \mathbb{C}_x$ such that the matrix $D_x f$ has rank zero when evaluated at these points.....	38
4.1	The parameter space $\mathbb{C}_{b,c}^2$ along with the discriminant $\mathcal{V}(b^2 - 4c)$	39
4.2	The Barth sextic.	44
4.3	The discriminant of the projection of the Barth sextic onto $\mathbb{C}_{x,y}^2$ with the number of real points in the fiber of any point in each discriminant chamber indicated.	44
4.4	One of the small discriminant chambers not easily noticeable in Figure 4.3.	44
5.1	A plane cubic $\mathcal{V}(f)$ (blue), the trace line $\mu(\mathcal{V}(f, x - t))$ (red), and the specific centroids $\mu(\mathcal{V}(f, x + 3)), \mu(\mathcal{V}(f, x - 1)), \mu(\mathcal{V}(f, x - 2))$	51
5.2	Left: The Newton polytope, support, and coefficients of f . Right: The curve $\mathcal{V}(f)$ (blue), the trace curve of $\mathcal{V}(f)$ with respect to $\mathcal{V}(x - t)$ (red), and five lines in the family L_t along with the centroids of their intersections with $\mathcal{V}(f)$	52
5.3	(Reprinted from [1]) An example of two tropical curves which cannot be distinguished from their coordinate projections	54
6.1	Algorithm 6.1.4 with $h = 0.1, x(1) = 0.75$, and $m = 10$	60
6.2	Algorithm 6.1.6 applied to the same differential equation, initial value, and step size as in Example 6.1.5	62
6.3	A visual display of path-jumping.	62
6.4	A homotopy displaying possible behaviors at $t = 0$	67
6.5	A depiction of the local behavior of a path $x(t)$ of a homotopy near a branch point with winding number 2.....	68
6.6	A visual depiction of the full path tracking algorithm.	69
6.7	A cartoon describing the Polyhedral homotopy.	74
6.8	(Reprinted from [1]) Constructing a pseudo-witness set for a projection of a twisted cubic	79
6.9	Constructing a pseudo-witness set for a projection of a cylinder via slicing.	80
6.10	A schematic of a single monodromy loop tracked numerically.	82
6.11	An example of two monodromy loops generating a transitive subgroup of a monodromy group.....	84

6.12	A schematic describing step (2) of Algorithm 6.5.6.....	85
7.1	An example of locations of $\{\rho_i\}_{i=1}^4$ and $\{\tau_j\}_{j=1}^3$ in \mathbb{C}_s . The smaller circle has radius γ_{τ_1} and the larger circle has radius Γ_{τ_1}	95
7.2	A unit vector ω and a geometric description of d_ω	96
7.3	The Newton polytope and tropicalization of a hypersurface.	98
7.4	For directions $\omega \in S^1$ we draw the ray in direction ω with length $\min(1, d^{-\omega})$ describing the convergence rate proven in Theorem 7.3.1.	99
7.5	(Reprinted from [1]) Left: Values of t (magnitude of rays) such that queryOracle finishes for different ω (direction of rays) on a hypersurface with Newton polytope (center) and normal fan (right).	101
7.6	Three snapshots of Sage animation from example with viewing window $[-4, 4]^2$	103
7.7	(Reprinted from [1]) Two tropical space curves with the same tropical coordinate projections	104
7.8	A Lüroth quartic.	108
7.9	A fiber product construction to compute witness points.....	110
7.10	(Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 1$ and $t = 4$ resp.) of Algorithm 7.1.2.....	112
7.11	(Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 8$ and $t = 20$ resp.) of Algorithm 7.1.2.	112
7.12	(Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 30$ and $t = 75$ resp.) of Algorithm 7.1.2.	113
8.1	The lacunary image $(\mathcal{A}_1, \mathcal{A}_2)$ of the support $(\mathcal{B}_1, \mathcal{B}_2)$ under the map φ	126
8.2	An example of triangular support.....	128
8.3	A support \mathcal{A} such that $(\mathcal{A}, \mathcal{A})$ is neither triangular nor lacunary.....	136
8.4	Dense support $(\mathcal{A}_1, \mathcal{A}_2)$, the support $(\text{vert}(\mathcal{A}_1), \text{vert}(\mathcal{A}_2))$, and the support of the Bézout start system.	137
8.5	The four supports involved in a computational experiment.	138
8.6	A schematic of the process in Example 8.5.1.	140
8.7	Scatter plot of timings.....	141

8.8	Box plot of timings.....	142
-----	--------------------------	-----

LIST OF TABLES

TABLE	Page
6.1 Numerical data for Algorithm 6.1.6 on input from Example 6.1.5.	61
6.2 Ingredients for homotopy methods.	78
6.3 Some probabilities of generating a transitive action by uniformly choosing from $S_d \times R_d$	88
7.1 Vertices of the Lüroth polytope found.....	115

1. INTRODUCTION

Understanding the solution sets of polynomial systems,

$$f_1(x_1, \dots, x_n) = f_2(x_1, \dots, x_n) = \dots = f_k(x_1, \dots, x_n) = 0, \tag{1.1}$$

is a ubiquitous need throughout mathematics, as well as the primary goal of algebraic geometry. Such solution sets,

$$\mathcal{V}(f_1, \dots, f_k) = \{(a_1, \dots, a_n) \in \mathbb{C}^n \mid f_i(a_1, \dots, a_n) = 0 \text{ for } i = 1, \dots, k\},$$

are called varieties. One way to study varieties is to partition them into families with respect to some structure so that most varieties in the same family have the same properties. Those that do not exhibit these generic properties may still be understood through the role they play in their family. In this dissertation, we study families of varieties delineated via the monomials appearing in their defining polynomials.

The support of a polynomial,

$$f(x_1, \dots, x_n) = \sum_{\alpha=(\alpha_1, \dots, \alpha_n) \in \mathbb{Z}^n} c_\alpha x_1^{\alpha_1} \cdots x_n^{\alpha_n}, \quad c_\alpha \in \mathbb{C},$$

is the set $\text{supp}(f) = \{\alpha \in \mathbb{Z}^n \mid c_\alpha \neq 0\}$. Studying a polynomial system $F = (f_1, \dots, f_k)$ through its support $\mathcal{A}_\bullet = (\text{supp}(f_1), \dots, \text{supp}(f_k))$ endows it with the structure of a sparse polynomial system and identifies F as a point in the coefficient space $\mathbb{C}^{\mathcal{A}_\bullet}$. Sparse polynomial systems belonging to the same family share a striking number of properties, many depending only on the collection P_\bullet of convex hulls of the supports in \mathcal{A}_\bullet , called Newton polytopes.

The polyhedral geometry of the Newton polytopes P_\bullet encodes much information about $\mathcal{V}(F)$. For example, the famous Bernstein-Kushnirenko Theorem (Proposition 5.3.1) states that when F is a square system ($k = n$) the number of isolated points of $\mathcal{V}(F)$ in $(\mathbb{C}^\times)^n$ is bounded by a numerical value called the mixed volume of P_\bullet . It also states that this bound is almost always attained, inducing a branched cover

$$\begin{aligned} \pi_{\mathcal{A}_\bullet}: X_{\mathcal{A}_\bullet} &\rightarrow \mathbb{C}^{\mathcal{A}_\bullet} \\ (x, F) &\mapsto F \end{aligned} \tag{1.2}$$

from the incidence variety $X_{\mathcal{A}_\bullet} = \{(x, F) \mid x \in (\mathbb{C}^\times)^n, F(x) = 0\}$ whose fiber $\pi_{\mathcal{A}_\bullet}^{-1}(F)$ is identified with the solutions of $F = 0$ in $(\mathbb{C}^\times)^n$. This viewpoint gives geometric structure to families of sparse polynomial systems whereby we may understand their constituents.

More difficult than counting solutions of polynomial systems is computing them. Over the last sixty years, mathematicians laid the groundwork for computational algebraic geometry, developing symbolic algorithms for studying and computing solutions of polynomials. More recently, techniques from numerical analysis joined algebraic geometry to form a novel computational paradigm known as numerical algebraic geometry. While symbolic algorithms use the algebraic properties of a polynomial system to study its solutions, numerical algebraic geometry studies varieties by computing numerical approximations of points on them, thus providing a predominantly geometric viewpoint toward computations in algebraic geometry.

Due to the geometric nature of numerical algebraic geometry, many definitions and proofs from geometry translate directly to numerical algorithms. For example, the definition of the monodromy group of a branched cover immediately suggests a numerical method to compute it (Algorithm 6.5.1). Another example is Huber and Sturmfels' proof of the Bernstein-Kushnirenko Theorem in [2] which produces the polyhedral homotopy algorithm (Algorithm 6.3.5) for computing all solutions of $F = 0$.

In Section 8 we give an algorithm which improves upon the polyhedral homotopy whenever the branched cover $\pi_{\mathcal{A}_\bullet}$ decomposes into a composition of branched covers. This decomposition happens if and only if the monodromy group of $\pi_{\mathcal{A}_\bullet}$ is imprimitive, a condition that Esterov [3] classified via computable conditions on \mathcal{A}_\bullet . Our algorithm (Algorithm 8.3.3) assesses whether or not $\pi_{\mathcal{A}_\bullet}$ decomposes and recursively computes fibers of the decomposition to compute a fiber of $\pi_{\mathcal{A}_\bullet}$, thus solving a sparse polynomial system with support \mathcal{A}_\bullet .

Conversely, algorithms in numerical algebraic geometry can extract information about Newton polytopes. In 2012, Hauenstein and Sottile suggested a numerical algorithm (Algorithm 7.1.2) which functions as a vertex oracle for the Newton polytope of the defining equation of a hypersurface. In Section 7, we explain how this algorithm is stronger than a vertex oracle and as a consequence, introduce the notion of a numerical oracle. Based on ideas from Hept and Theobald [4], we develop a tropical membership test (Algorithm 7.2.2) which relies on the algorithm of Hauenstein and Sottile as a subroutine. We analyze the convergence rates of each algorithm (Theorem 7.3.1) and explain our implementation of them in Section 7.4. Finally, we use our implementation to investigate the colossal Lüroth polytope (Section 7.6) and determine the implicit equation a hypersurface from algebraic vision (Section 7.5).

We provide all necessary background in Sections 2-6. Section 2 includes elementary results regarding polytopes, numerical oracles, mixed volumes, and subdivisions. In Section 3 we give a basic introduction to algebraic geometry necessary for the subsequent sections. In Section 4 we discuss branched covers, decomposable branched covers, and monodromy/Galois groups; we also give a proof that the monodromy/Galois group of a branched cover is imprimitive if and only if the branched cover is decomposable. In Section 5 we connect the previous sections by introducing Newton polytopes, sparse polynomial systems, and tropical algebraic geometry. Section 6 builds the theory of numerical algebraic geometry and contains an assembly of numerical algorithms, including Huber and Sturmfels' treatment of the polyhedral homotopy, as well as algorithms which use monodromy to solve polynomial systems.

2. POLYTOPES

We remark that a portion of the discussion of numerical oracles in this section also appears in the article [1] by the author*.

2.1 Describing polytopes

A subset $S \subset \mathbb{R}^n$ is **convex** if for any $p, q \in S$ the line segment between them $[p, q] = \{\lambda p + (1 - \lambda)q \mid 0 \leq \lambda \leq 1\}$ is also contained in S . The **convex hull** of S is

$$\text{conv}(S) = \bigcap \{S' \subset \mathbb{R}^n \mid S \subset S', S' \text{ convex}\}.$$

Lemma 2.1.1. *If $\mathcal{A} = \{\alpha_1, \dots, \alpha_k\} \subset \mathbb{R}^n$ is finite then*

$$\text{conv}(\mathcal{A}) = \left\{ \sum_{i=1}^k \lambda_i \alpha_i \mid \sum_{i=1}^k \lambda_i = 1, \lambda_i \in \mathbb{R}_{\geq 0} \right\}.$$

Proof. The forward containment is true since the right-hand-side is a convex set containing \mathcal{A} . Indeed, if $p = \sum_{i=1}^k \lambda_i \alpha_i$ and $q = \sum_{i=1}^k \nu_i \alpha_i$ are elements of the right-hand-side and $\gamma \in [0, 1]$, then

$$\gamma p + (1 - \gamma)q = \sum_{i=1}^k (\gamma \lambda_i + (1 - \gamma)\nu_i) \alpha_i$$

is as well.

The reverse containment for $k = 1$ or $k = 2$ is true by definition. Assume it is true for $k - 1$ and let $\alpha = \sum_{i=1}^k \lambda_i \alpha_i$ be an element of the right-hand-side. Without loss of generality, assume $\lambda_1 \neq 0$ so that

$$\alpha = \lambda_1 \alpha_1 + (1 - \lambda_1) \left(\frac{\lambda_2}{1 - \lambda_1} \alpha_2 + \dots + \frac{\lambda_k}{1 - \lambda_1} \alpha_k \right).$$

Since $p := \alpha_1$ and $q := \sum_{i=2}^k \frac{\lambda_i}{1 - \lambda_1} \alpha_i$ are points in $\text{conv}(\mathcal{A})$ by induction, the segment $[p, q]$ containing α must be in $\text{conv}(\mathcal{A})$ as well. \square

Definition 2.1.2. A **polytope** is any subset $P \subset \mathbb{R}^n$ that can be written as the convex hull of finitely many points. If these points can be taken to be in \mathbb{Z}^n , then P is called an **integral polytope**.

Example 2.1.3. For ease of reading, we will often encode points in \mathbb{R}^n as the columns of a matrix. Let $\mathcal{A} = \begin{pmatrix} 0 & 0 & 3/2 & 2 & 2 & 2 & 3 & 4 \\ 2 & 3 & 3/2 & 0 & 3 & 4 & 2 & 0 \end{pmatrix} \subset \mathbb{R}^2$. The polytope $Q = \text{conv}(\mathcal{A})$ shown in Figure 2.1 is an integral polytope since we may write $Q = \text{conv}\left(\begin{pmatrix} 0 & 0 & 2 & 2 & 4 \\ 2 & 3 & 0 & 4 & 0 \end{pmatrix}\right)$.

The **dimension** of a subset $S \subset \mathbb{R}^n$, denoted $\text{dim}(S)$, is the dimension of its **affine span**,

$$\mathbb{R}S = \left\{ \sum_{i=1}^k \lambda_i s_i \mid s_i \in S, \lambda_i \in \mathbb{R}, \sum_{i=1}^k \lambda_i = 1 \right\},$$

*Reprinted with permission from T. Brysiewicz, “Numerical Software to Compute Newton polytopes and Tropical Membership,” *Mathematics in Computer Science*, 2020. Copyright 2020 by Springer Nature.




Figure 2.1: An integral polytope $Q \subset \mathbb{R}^2$.

and the **codimension** of S is $\text{codim}(S) = n - \dim(S)$. **Polygons** are polytopes of dimension two. If S is compact, we define the **support function** of S as

$$h_S: \mathbb{R}^n \rightarrow \mathbb{R}$$

$$\omega \mapsto \max_{x \in S} \langle x, \omega \rangle.$$

Given $\omega \in \mathbb{R}^n$, the **subset of S exposed by ω** is

$$S_\omega = \{x \in S \mid \langle x, \omega \rangle = h_S(\omega)\}.$$

A **face \mathcal{F}** of a polytope $P \subset \mathbb{R}^n$ is any subset of P of the form $\mathcal{F} = \emptyset$ or $\mathcal{F} = P_\omega$ for some $\omega \in \mathbb{R}^n$. Faces of dimensions 0, 1, k , $\dim(P) - 1$ are called **vertices**, **edges**, **k -faces**, and **facets** respectively. The set of vertices is denoted $\text{vert}(P)$ and the set of facets is denoted $\text{facets}(P)$.

Example 2.1.4. Let Q be as in Example 2.1.3. The dimension of Q is 2 and its codimension is 0. Let $\omega_1 = (-1, 1)$, $\omega_2 = (2, 1)$, $\omega_3 = (-1, -2)$, and $\omega_4 = (0, 0)$. Then

$$h_Q(\omega_1) = 3, \quad h_Q(\omega_2) = 8, \quad h_Q(\omega_3) = -2, \quad h_Q(\omega_4) = 0,$$

and the faces exposed by $\omega_1, \omega_2, \omega_3$, and ω_4 are

$$Q_{\omega_1} = \{(0, 3)\}, \quad Q_{\omega_2} = \text{conv}(\{(2, 4), (4, 0)\}), \quad Q_{\omega_3} = \{(2, 0)\}, \quad Q_{\omega_4} = Q.$$

Figure 2.2 depicts these directions and faces. In total, Q has one empty face, five vertices, five facets (edges), and one 2-face. \diamond

Given a polytope $P \subset \mathbb{R}^n$, it is useful to collect directions $\omega \in \mathbb{R}^n$ which expose the same face into cones. A subset $C \subset \mathbb{R}^n$ is a **cone** if for any $p \in C$, we have that $\lambda p \in C$ for $\lambda \in \mathbb{R}_{\geq 0}$. A cone is a convex cone if it is closed under addition. Indeed if p and q are elements of a cone C which




Figure 2.2: Left: Directions $\omega_1, \omega_2, \omega_3$, and $\omega_4 = (0, 0)$. Right: The polytope $Q \subset \mathbb{R}^2$ and three of its proper faces exposed by ω_1, ω_2 , and ω_3 .

is closed under addition and $\lambda \in [0, 1]$ then $\lambda p + (1 - \lambda)q \in C$ since each summand is in C . The (outer) normal fan of a polytope P is the collection

$$\mathcal{N}(P) = \{C[\omega]\}_{\omega \in \mathbb{R}^n}$$

of convex cones

$$C[\omega] = \{\omega' \in \mathbb{R}^n \mid P_\omega \subseteq P_{\omega'}\}.$$

We denote the set of all $C[\omega]$ of codimension at least i by $\mathcal{N}^{(i)}(P)$.

Example 2.1.5. Figure 2.3 displays Q along with its normal fan $\mathcal{N}(Q)$ which has one zero-dimensional cone (the origin), five one-dimensional cones, and five two-dimensional cones. \diamond

Lemma 2.1.6. [5, Proposition 2.2] Every polytope may be written as

$$P = \text{conv}(\text{vert}(P)). \tag{2.1}$$

If $\mathcal{A} \subset \mathbb{R}^n$ is finite, then $\text{vert}(\text{conv}(\mathcal{A})) \subseteq \mathcal{A}$.

Lemma 2.1.7. [5, Proposition 2.3] Let F be a face of a polytope $P \subset \mathbb{R}^n$.

- (1) F is a polytope with $\text{vert}(F) = F \cap \text{vert}(P)$.
- (2) Every intersection of faces of P is a face of P .
- (3) The faces of F are exactly the faces of P that are contained in F .
- (4) $F = P \cap \mathbb{R}F$.

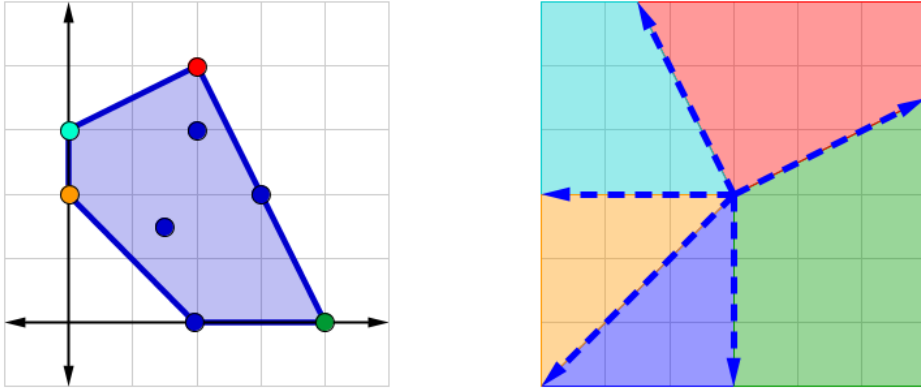


Figure 2.3: A polytope and its corresponding normal fan

Lemma 2.1.6 gives one way to canonically represent a polytope: as the convex hull of its vertices. This representation is called the **vertex representation** of a polytope. Halfspaces provide another way to represent polytopes. A **halfspace** of \mathbb{R}^n is any subset of the form

$$\mathbb{R}_{\omega,c}^n = \{x \in \mathbb{R}^n \mid \langle x, \omega \rangle \leq c\} \subset \mathbb{R}^n,$$

for some $\omega \in \mathbb{R}^n$ and $c \in \mathbb{R}$. Given a polytope $P \subset \mathbb{R}^n$ and any direction $\omega \in \mathbb{R}^n$, the halfspace $H_P(\omega) = \mathbb{R}_{\omega, h_P(\omega)}^n$ contains P . Note that $H_P(\omega) = H_P(\lambda \cdot \omega)$ for any $\lambda > 0$.

Lemma 2.1.8. [5, Theorem 2.15] Every polytope $P \subset \mathbb{R}^n$ may be written as

$$P = \mathbb{R}P \cap \left(\bigcap_{i=1}^m H_P(\omega_i) \right) \quad (2.2)$$

for any set $\{\omega_i\}_{i=1}^m \subset \mathbb{R}^n$ such that $\{P_{\omega_i}\}_{i=1}^m = \text{facets}(P)$. Conversely, any bounded intersection of halfspaces is a polytope.

If a polytope is n -dimensional, then it has a unique representation of the form (2.2) since each facet is $(n - 1)$ -dimensional and is exposed by its unique outer-normal ray. Note that these are the one-dimensional cones in the normal fan of a polytope. If a polytope has positive codimension, then it has a unique representation of the form (2.2) within its affine hull (the ω_i in (2.2) are taken to be parallel with the affine hull of P). We call such a unique representation the **halfspace representation** of a polytope.

Example 2.1.9. The polytope Q in Example 2.1.3 has the halfspace representation,

$$Q = H_Q(2, 1) \cap H_Q(0, -1) \cap H_Q(-1, -1) \cap H_Q(-1, 0) \cap H_Q(-1, 2).$$

Each of these halfspaces are displayed in Figure 2.4. ◇




Figure 2.4: Five halfspaces in \mathbb{R}^2 whose intersection is Q .

2.2 Oracles

While the vertex and halfspace representations are the most common ways of expressing a polytope, other representations come from functions called oracles. Colloquially, an oracle is an entity which provides prophetic insight whenever queried. Likewise, the **vertex oracle** for a polytope $P \subset \mathbb{R}^n$ is the function

$$\mathbb{V}_P: \mathbb{R}^n \rightarrow \mathbb{R}^n \cup \{\text{PFE}\}$$

$$\omega \mapsto \begin{cases} P_\omega & \dim(P_\omega) = 0 \\ \text{PFE} & \text{otherwise} \end{cases}$$

where PFE abbreviates the expression “Positive dimensional Face Exposed”. We remark that $\mathbb{V}_P(\omega) = \text{PFE}$ if and only if $\omega \in \mathcal{N}^{(1)}(P)$. The process of evaluating a vertex oracle is called **querying** the oracle.

Remark 2.2.1. When a vertex oracle query returns a vertex $\mathbb{V}_P(\omega) = v$, it implicitly returns the information that $h_P(\omega) = \langle v, \omega \rangle$ and therefore that $P \subset \mathbb{R}_{\omega, \langle v, \omega \rangle}^n = H_P(\omega)$.

Let $\mathbf{0}$ denote the all 0’s vector in \mathbb{R}^n , $\mathbf{1}$ denote the all 1’s vector in \mathbb{R}^n , and e_i denote the i -th coordinate vector in \mathbb{R}^n . For any $v \in \mathbb{R}^n$, let $|v|$ denote the sum of its coordinates. Given a polytope $P \subset \mathbb{R}^n$, let $\mathcal{L}(P) = P \cap \mathbb{Z}^n$ denote its set of **lattice points**.

Proposition 2.2.2. *If $P \subset \mathbb{R}^n$ is an integral polytope, then the vertex representation of P can be recovered from the vertex oracle for P .*

Proof. Let $P \subset \mathbb{R}^n$ be an integral polytope and \mathbb{V}_P its vertex oracle. To prove the proposition, we first bound P between two polytopes by querying the vertex oracle as follows.

Let $\omega^* = (\omega_1^*, \dots, \omega_n^*) \in \mathbb{R}_{>0}^n$ be a vector such that $\omega_1^*, \dots, \omega_n^*$ are rationally independent (i.e. $\langle x, \omega^* \rangle \neq 0$ for any $\mathbf{0} \neq x \in \mathbb{Z}^n$). Observe that $\mathbb{V}_P(\omega^*)$ must return a vertex: otherwise, there exist two vertices p_1, p_2 such that $\langle p_1, \omega^* \rangle = \langle p_2, \omega^* \rangle$ implying that $x = p_1 - p_2$ is an integer point whose dot product with ω^* is nonzero. A consequence of Remark 2.2.1 is that the halfspace $H_P(\omega^*)$ containing P is computed as well. Since ω is in the positive orthant, $H_P(\omega^*)$ bounds $P \cap \mathbb{R}_{\geq 0}^n$.

Similarly, for every vertex v of the hypercube $\text{cube}(n) = [-1, 1]^n$, we let $v \circ \omega^*$ denote the Hadamard (coordinate-wise) product so that the output $\mathbb{V}_P(v \circ \omega^*)$ is a vertex of P . Again, each oracle query bounds P in the corresponding orthant of \mathbb{R}^n so that the intersection

$$P^* = \bigcap_{v \in \text{cube}(n)} H_P(v \circ \omega^*),$$

is bounded, and thus by Lemma 2.1.8, is a polytope. Setting $P_* = \text{conv}(\{\mathbb{V}_P(v \circ \omega) \mid v \in \text{cube}(n)\})$ gives containments

$$P_* \subseteq P \subseteq P^*. \quad (2.3)$$

The proof proceeds algorithmically. Set $P^* = \text{conv}(\mathcal{L}(P^*))$ so that P^* is integral. Since P is integral, the containments (2.3) are still true. For every $p \in \text{vert}(P^*) \setminus P_*$, pick ω such that $\mathbb{V}_{P^*}(\omega) = p$. Since p is the unique point in P^* obtaining a maximum dot product with ω and $P \subseteq P^*$ then $p \in P$ if and only if $\mathbb{V}_P(\omega) = p$. We have three cases: either $p \in P$ and so $\mathbb{V}_P(\omega) = p$ (case (i)), or $\mathbb{V}_P(\omega)$ returns PFE (case (ii)) or $\mathbb{V}_P(\omega)$ returns another vertex $q \neq p$ (case (iii)).

Case (i): If $\mathbb{V}_P(\omega) = p$ then set $P_* = \text{conv}(P_* \cup p)$. Note that the containments (2.3) still hold and that the number of lattice points of P_* has increased.

Case (ii): If $\mathbb{V}_P(\omega) = \text{PFE}$, then $p \notin P$ and so we may set $P^* = \text{conv}(\mathcal{L}(P^*) \setminus p)$ while preserving (2.3). In this case, the number of lattice points of P^* has decreased.

Case (iii): If $\mathbb{V}_P(\omega) = q \neq p$, then we may set $P_* = \text{conv}(P_* \cup q)$ and $P^* = \text{conv}(\mathcal{L}(P^* \cap H_P(\omega)))$ while preserving (2.3). In this case, the number of lattice points of P_* may have increased depending on whether or not q was already in P_* , but it will always be the case that the number of lattice points of P^* has decreased.

Each oracle query involves one of the above cases and each case preserves the containments (2.3) while either increasing the number of lattice points in P_* or decreasing the number of lattice points in P^* . Thus, this process must terminate with $\text{vert}(P^*) \setminus P_* = \emptyset$, proving that these polytopes are equal to each other and so $P_* = P = P^*$. \square

Algorithm 2.2.3 (Vertex oracle \rightarrow vertex representation).

Input:

- The vertex oracle \mathbb{V}_P for an integral polytope $P \subset \mathbb{R}_{\geq 0}^n$

Output:

- The vertex representation for P

Steps:

```

0 Pick  $\omega^* = (\omega_1^*, \dots, \omega_n^*) \in \mathbb{R}_{>0}^n$  with rationally independent coordinates
1 set  $P_* = \emptyset$ , set  $P^* = \mathbb{R}^n$ 
2 for each vertex  $v \in \text{cube}(n)$  do
  2.1 set  $P_* = \text{conv}(P_* \cup \mathbb{V}_P(v \circ \omega^*))$ 
  2.2 set  $P^* = P^* \cap H_P(v \circ \omega^*)$ 
3 while  $\mathcal{L}(P_*) \neq \mathcal{L}(P^*)$  do
  3.1 set  $P^* = \text{conv}(\mathcal{L}(P^*))$ 
  3.2 Pick  $p \in \text{vert}(P^*) \setminus P_*$ 
  3.3 Find  $\omega \in \mathbb{R}^n$  such that  $\mathbb{V}_{P^*}(\omega) = p$ 
  3.4 if  $\mathbb{V}_P(\omega) = p$  then set  $P_* = \text{conv}(P_* \cup p)$ 
  3.5 if  $\mathbb{V}_P(\omega) = \text{PFE}$  then set  $P^* = \text{conv}(\mathcal{L}(P^*) \setminus p)$ 
  3.6 if  $\mathbb{V}_P(\omega) = q \neq p$  then
    3.6.1 set  $P^* = P^* \cap H_P(\omega)$ 
    3.6.2 if  $q \notin P_*$  then set  $P_* = \text{conv}(P_* \cup p)$ 
4 return  $\text{vert}(P_*)$ 

```

Example 2.2.4. Figure 2.5 displays the steps required to complete Algorithm 2.2.3 on Q from Example 2.1.3. We use $\omega^* = (1, \sqrt{2})$ in step (0) of the algorithm. Step (2) in Algorithm 2.2.3 is represented by the top-left graphic showing the four vertex oracle queries on the vectors ω^* , $-\omega^*$, $(-1, \sqrt{2})$, and $(1, -\sqrt{2})$. Each query reveals a vertex of Q and a halfspace containing Q . The intersection of all such halfspaces Q^* is depicted in grey in the first image along with Q_* in green and $\text{conv}(\mathcal{L}(Q^*))$ in red.

The next image (to the right) displays the oracle query $\mathbb{V}_Q(1, 0) = (4, 0)$, revealing a vertex which was already found. Thus, this oracle query does not increase the size of Q_* but it does establish that $(5, 1)$ (a previous vertex of Q^*) is not contained in Q and so the size of Q^* is reduced. The third image (bottom left) attempts to establish whether or not $(4, 2) \in Q$ by choosing $\omega = (2, 1)$ so that $\mathbb{V}_{Q^*}(\omega) = (4, 2)$ and querying $\mathbb{V}_Q(\omega) = \text{PFE}$. This does not find a new vertex of Q , nor does it find a new halfspace containing Q . It does, however, reveal that $(4, 2) \notin Q$ and so Q^* is again reduced to $\text{conv}(\mathcal{L}(Q^*) \setminus (4, 2))$. At this stage, $(0, 2)$ is the unique vertex of Q^* which is not in Q_* and $\mathbb{V}_Q(-2, 1) = (0, 2)$ reveals that it is a vertex of Q . The outer polytope Q^* is reduced again, the inner polytope Q_* grows, and Q^* becomes equal to Q_* , ending the algorithm. \diamond

Remark 2.2.5. Implementing Algorithm 2.2.3, as is, requires the representation of a rationally independent vector ω^* on a computer for step (2). Theoretically, a random $\omega \in \mathbb{R}^n$ will expose a vertex of P with probability one and so in practice, we replace steps (0) and (2) by randomly querying the oracle in each orthant until a vertex is returned. This process bounds P in a polytope P^* . We remark that probability one statements about the theory may not translate to probability one computations and we give a more detailed discussion in Remark 7.3.3 in Section 7. \diamond




Figure 2.5: A graphical interpretation of Algorithm 2.2.3 running on the polytope Q in Example 2.1.3.

We denote the standard full-dimensional simplex in \mathbb{R}^n by $\Delta_n = \text{conv}(\mathbf{0}, e_1, \dots, e_n)$ and the dilation of Δ_n by a factor of d by $d\Delta_n = \text{conv}(\mathbf{0}, d \cdot e_1, \dots, d \cdot e_n)$. The **degree** of a polytope $P \subset \mathbb{R}_{\geq 0}^n$, is $\text{deg}(P) = h_P(\mathbf{1})$. A polytope is **homogeneous** if $|p| = \text{deg}(P)$ for all $p \in P$ and the **homogenization** of P is $\tilde{P} = \{(p, \text{deg}(P) - |p|) \mid p \in P\} \subset \mathbb{R}^{n+1}$.

Definition 2.2.6. The **numerical oracle** for a polytope $P \in \mathbb{R}^n$ is the function

$$\mathcal{O}_P: \mathbb{R}^n \rightarrow \mathbb{R}^n \cup \{\text{EEP}\}$$

$$\omega \mapsto \begin{cases} P_\omega & \dim(P_\omega) = 0 \\ \min(P_\omega) & 0 < \dim(P_\omega) < \dim(P) \\ \text{EEP} & P_\omega = P \end{cases}$$

where $\min(P_\omega)$ is the coordinate-wise minimum of all points in P_ω .

The expression `EEP` abbreviates `Exposes Entire Polytope`. This oracle is dubbed “numerical” because it arises naturally from the numerical HS-algorithm (Algorithm 7.1.2 of Section 7).

Generally, one cannot distinguish whether the output of a numerical oracle for a polytope P is a vertex $v = P_\omega$ or the coordinate-wise minimum $w = \min(P_\omega)$ of a positive-dimensional face. For example, the numerical oracle query $\mathcal{O}_{\Delta_2}(1, 1)$ returns $\mathbf{0}$ not because $\mathbf{0}$ is a vertex, but because $\mathbf{0} = \min(\text{conv}(e_1, e_2))$. Thus, at first glance, a numerical oracle may seem weaker than a vertex oracle. However, when the polytope P is homogeneous of degree d these cases may be distinguished easily since the sum of the coordinates of a vector output of $\mathcal{O}_P(\omega)$ will be d if and only if the vector is a vertex and it will be less than d otherwise. Restricted to homogeneous

polytopes, a numerical oracle gives strictly more information than a vertex oracle, implying the following corollary to Proposition 2.2.2.

Corollary 2.2.7. *If P is a homogeneous integral polytope then the vertex representation of P may be recovered from its numerical oracle.*

Other oracles for polytopes exist and are well-studied. For example, Emiris et. al. [6] developed an algorithm similar to Algorithm 2.2.3 for oracles which are stronger than vertex oracles: instead of returning PFE, they return a vertex on the corresponding positive-dimensional face.

2.3 Mixed volume

We develop some of the theory of mixed volumes of polytopes and include multiple formulas and characterizations of mixed volume. We list them here for convenience.

- (1) Coefficient of a volume function (Definition 2.3.3).
- (2) Volume alternating sum formula (Lemma 2.3.6).
- (3) Axiomatic characterization (Lemma 2.3.8).
- (4) Lattice point alternating sum formula for integral polytopes (Lemma 2.3.9).
- (5) Sum of volumes of mixed cells formula (Lemma 2.4.3).

We give a sixth way of computing mixed volume in Section 5 via the Bernstein-Kushnirenko Theorem (Proposition 5.3.1).

We begin our discussion by introducing two natural operations on subsets of \mathbb{R}^n . Let $S_1, S_2 \subset \mathbb{R}^n$ and $\lambda \in \mathbb{R}_{\geq 0}$. The set

$$\lambda S_1 = \{\lambda s \mid s \in S_1\},$$

is the **scaling** of S_1 by λ . The set

$$S_1 + S_2 = \{s_1 + s_2 \mid s_1 \in S_1, s_2 \in S_2\},$$

is the **Minkowski sum** of S_1 and S_2 . The scaling of a polytope $P = \text{conv}(\mathcal{A})$ by $\lambda \in \mathbb{R}_{\geq 0}$ is clearly a polytope given as $\lambda P = \text{conv}(\lambda \mathcal{A})$. The following lemma proves an analogous result for Minkowski sums of polytopes.

Lemma 2.3.1. *Let $P, Q \subset \mathbb{R}^n$ be polytopes.*

- (1) *The support functions of P and Q are additive: $h_{P+Q} = h_P + h_Q$.*
- (2) *The Minkowski sum $P + Q$ is a polytope which may be written as $\text{conv}(\text{vert}(P) + \text{vert}(Q))$.*
- (3) *If $F \subset P + Q$ is a face, then there exist unique faces $F_P \subseteq P$ and $F_Q \subseteq Q$ such that $F = F_P + F_Q$.*
- (4) *If P and Q are integral, so is $P + Q$.*

Proof. Additivity of support functions is immediate since

$$h_{P+Q}(\omega) = \max_{s \in P+Q} \langle s, \omega \rangle = \max_{p \in P, q \in Q} \langle p + q, \omega \rangle = \max_{p \in P} \langle p, \omega \rangle + \max_{q \in Q} \langle q, \omega \rangle.$$

To show that $P + Q$ is a polytope, we first show $P + Q$ is convex. Let $a = p_1 + q_1$ and $b = p_2 + q_2$ for $p_1, p_2 \in P$ and $q_1, q_2 \in Q$. Then $v \in [a, b]$ implies

$$\begin{aligned} v &= \lambda a + (1 - \lambda)b \\ &= \lambda(p_1 + q_1) + (1 - \lambda)(p_2 + q_2) \\ &= (\lambda p_1 + (1 - \lambda)p_2) + (\lambda q_1 + (1 - \lambda)q_2) \in P + Q, \end{aligned}$$

proving that $P + Q$ is convex. To see that $P + Q \subset \text{conv}(\text{vert}(P) + \text{vert}(Q))$, suppose towards contradiction that there exists $v \in P + Q \setminus \text{conv}(\text{vert}(P) + \text{vert}(Q))$. Then there exists a halfspace of $\text{conv}(\text{vert}(P) + \text{vert}(Q))$ not containing v . In other words, there exists ω such that $\langle v, \omega \rangle = h_{P+Q}(\omega) > h_P(\omega) + h_Q(\omega)$, a contradiction by part (1). Thus,

$$\text{vert}(P) + \text{vert}(Q) \subset P + Q \subset \text{conv}(\text{vert}(P) + \text{vert}(Q)),$$

and taking the convex hull of this containment proves parts (2) and (4).

To prove part (3), observe that for any $\omega \in \mathbb{R}^n$ we have $(P + Q)_\omega = P_\omega + Q_\omega$ by part (1). Suppose

$$(P + Q)_\omega = P_{\omega'} + Q_{\omega''},$$

for some other $\omega', \omega'' \in \mathbb{R}^n$. The evaluation of $x \mapsto \langle x, \omega \rangle$ at any point on the right-hand-side must equal $h_P(\omega) + h_Q(\omega)$, implying that $P_{\omega'} = P_\omega$ and $Q_{\omega''} = Q_\omega$. \square

To fix notation, let $P_\bullet = \{P_1, \dots, P_n\}$ be a collection of n polytopes in \mathbb{R}^n . We denote the set $\{1, \dots, n\}$ by $[n]$. The following result is due to Minkowski when $d = 3$ [7].

Lemma 2.3.2 (H. Minkowski [7]). *The function*

$$V(P_\bullet): \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}$$

$$V(P_\bullet)(\lambda_1, \dots, \lambda_n) = \text{vol}(\lambda_1 P_1 + \dots + \lambda_n P_n)$$

is a homogeneous polynomial of degree n in $\mathbb{R}[\lambda_1, \dots, \lambda_n]$ where vol denotes the n -dimensional Euclidean volume.

Definition 2.3.3. The **mixed volume** of P_\bullet , denoted $\text{MV}(P_\bullet)$, is the coefficient of $\lambda_1 \lambda_2 \cdots \lambda_n$ in $V(P_\bullet)$.

Example 2.3.4. Consider $A = \text{conv}(\mathbf{0}, e_1, e_2, e_1 + e_2)$ and $B = \text{conv}(\mathbf{0}, e_1, e_2)$ as displayed in Figure 2.6. Then $V(A, B) = \lambda_1^2 + 2\lambda_1 \lambda_2 + \frac{1}{2}\lambda_2^2$ and so $\text{MV}(A, B) = 2$. \diamond

Lemma 2.3.5. *Let $P, P_1, \dots, P_n, Q \subset \mathbb{R}^n$ be polytopes and let $a \in \mathbb{R}_{\geq 0}$. Then,*

- (1) $\text{MV}(P, \dots, P) = n! \text{vol}(P)$.
- (2) MV is symmetric in its arguments.




Figure 2.6: A graphic expressing $\text{vol}(\lambda_1 A + \lambda_2 B)$ for two polygons $A, B \subset \mathbb{R}^2$.

(3) *MV is multilinear:*

$$\text{MV}(aP_1 + Q, P_2, \dots, P_n) = a \text{MV}(P_1, \dots, P_n) + \text{MV}(Q, P_2, \dots, P_n).$$

Proof. Note that $\text{vol}(\lambda_1 P + \dots + \lambda_n P) = \text{vol}((\lambda_1 + \dots + \lambda_n)P) = (\lambda_1 + \dots + \lambda_n)^n \text{vol}(P)$ and so the coefficient of $\lambda_1 \cdots \lambda_n$ is $n! \text{vol}(P)$. Part (2) is immediate from the definition of mixed volume. For a proof of part (3), see [8, Lemma 3.6]. \square

Lemma 2.3.6. [8, Theorem 3.7] *Given a collection of polytopes P_1, \dots, P_n ,*

$$\text{MV}(P_1, \dots, P_n) = \sum_{I \subset [n]} (-1)^{n-|I|} \text{vol} \left(\sum_{i \in I} P_i \right).$$

Proof. We restate the proof given in [8]. Due to precisely the properties of mixed volume in Lemma 2.3.5, we may treat the statement in the theorem as the polynomial equation

$$\begin{aligned} n!x_1 \cdots x_n &= (x_1 + \cdots + x_n)^n - \sum_{i=1}^n (x_1 + \cdots + x_{i-1} + x_{i+1} + \cdots + x_n)^n + \cdots \\ &\quad \cdots + (-1)^{n-2} \sum_{i < j} (x_i + x_j)^n + (-1)^{n-1} \sum_{i=1}^n x_i^n, \end{aligned} \quad (2.4)$$

where $x_{i_1} \cdots x_{i_N} \leftrightarrow \text{vol}(\lambda_{i_1} \cdot P_{i_1} + \cdots + \lambda_{i_N} \cdot P_{i_N})$. To verify (2.4), we may simply check how many times each monomial appears in the right-hand-side. The monomial x_i^n appears once in the first term, $n - 1$ times in the second, and so on to give a total of

$$1 - (n - 1) + \binom{n - 1}{2} - \cdots + (-1)^{n-2}(n - 1) + (-1)^{n-1} = (1 - 1)^{n-1} = 0.$$

Similarly, every term on the right-hand-side cancels except for the mixed term $x_1 \cdots x_n$ which appears $n!$ times. \square

Since the formula in Lemma 2.3.6 is short when $n = 2$, we state it as a corollary.

Corollary 2.3.7. *The mixed volume of two convex polygons $P_1, P_2 \subset \mathbb{R}^2$ is*

$$\text{MV}(P_1, P_2) = \text{vol}(P_1 + P_2) - \text{vol}(P_1) - \text{vol}(P_2).$$

Lemma 2.3.8. *The only function from n -tuples of polytopes to \mathbb{R} satisfying the properties in Lemma 2.3.5 is MV.*

Proof. The proof of the formula of Lemma 2.3.6 relied precisely on the properties in Lemma 2.3.5. Thus, any other function satisfying those properties will have the same formula. \square

When each polytope in a collection P_\bullet is integral, there is a discrete analog of Lemma 2.3.6 involving lattice point enumeration.

Lemma 2.3.9. [9, Corollary 3.10] *Given a collection of integral polytopes P_1, \dots, P_n ,*

$$\text{MV}(P_1, \dots, P_n) = (-1)^n + \sum_{\emptyset \neq I \subset [n]} (-1)^{n-|I|} \left| \mathcal{L} \left(\sum_{i \in I} P_i \right) \right|.$$

2.4 Subdivisions

Following [2] we give the notion of subdivisions of collections of finite subsets of \mathbb{R}^n . The combinatorial constructions in this section provide a fifth description of the mixed volume of a collection of polytopes and are fundamentally important for Algorithm 6.3.5 of Section 6.3.3.

Let $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ be a collection of finite subsets of \mathbb{R}^n whose union affinely spans \mathbb{R}^n . A **cell** of \mathcal{A}_\bullet is a tuple $\mathcal{C}_\bullet = (\mathcal{C}_1, \dots, \mathcal{C}_k)$ of nonempty subset $\mathcal{C}_i \subset \mathcal{A}_i$. We define

$$\begin{aligned} \text{type}(\mathcal{C}_\bullet) &= (\dim(\text{conv}(\mathcal{C}_1)), \dots, \dim(\text{conv}(\mathcal{C}_k))), \\ \text{conv}(\mathcal{C}_\bullet) &= \text{conv}(\mathcal{C}_1 + \dots + \mathcal{C}_k), \\ |\mathcal{C}_\bullet| &= |\mathcal{C}_1| + |\mathcal{C}_2| + \dots + |\mathcal{C}_k|, \\ \text{vol}(\mathcal{C}_\bullet) &= \text{vol}(\text{conv}(\mathcal{C}_\bullet)). \end{aligned}$$

Definition 2.4.1. A **subdivision** of \mathcal{A}_\bullet is a collection $S^\bullet = \{\mathcal{C}_\bullet^{(1)}, \dots, \mathcal{C}_\bullet^{(m)}\}$ of cells satisfying

- (1) $\dim(\text{conv}(\mathcal{C}_\bullet^{(i)})) = n$ for all $i = 1, \dots, m$.
- (2) $\text{conv}(\mathcal{C}_\bullet^{(i)}) \cap \text{conv}(\mathcal{C}_\bullet^{(j)})$ is a proper face of $\text{conv}(\mathcal{C}_\bullet^{(i)})$ and $\text{conv}(\mathcal{C}_\bullet^{(j)})$ for all $i \neq j \in [m]$.
- (3) $\bigcup_{i=1}^m \text{conv}(\mathcal{C}_\bullet^{(i)}) = \text{conv}(\mathcal{A}_\bullet)$.

If S^\bullet additionally satisfies

- (4) $|\text{type}(\mathcal{C}_\bullet^{(i)})| = n$ for all $i = 1, \dots, m$,

then we say it is a **mixed subdivision**. Even stronger, if S^\bullet additionally satisfies

$$(5) \sum_{i=1}^k \left(|\mathcal{C}_i^{(j)}| - 1 \right) = n \text{ for all } j = 1, \dots, m,$$

then we say it is a **fine mixed subdivision**.

A cell \mathcal{C}_\bullet of a subdivision S^\bullet is called a **mixed cell** when $\min(\text{type}(\mathcal{C}_\bullet)) > 0$ and a **fine mixed cell** if it additionally satisfies $\sum_{i=1}^k (|\mathcal{C}_i| - 1) = n$. When $k = n$, a cell \mathcal{C}_\bullet is mixed if $\text{type}(\mathcal{C}_\bullet) = \mathbf{1}$ and it is fine mixed if $|\mathcal{C}_i| = 2$ for $i = 1, \dots, k$.

Example 2.4.2. When $k = 1$, every subdivision of \mathcal{A}_\bullet is a mixed subdivision because parts (1) and (4) of Definition 2.4.1 become the same statement. The fine mixed subdivisions of \mathcal{A}_\bullet are those with the property that the convex hull of each cell is an n -simplex. Such subdivisions comprise a rich family of combinatorial objects called **triangulations** [10]. \diamond

The definitions above provide a new description of mixed volume.

Lemma 2.4.3. [2, Theorem 2.4] Suppose $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ is a collection of finite subsets of \mathbb{R}^n whose union affinely spans \mathbb{R}^n and let $P_i = \text{conv}(\mathcal{A}_i)$. If S^\bullet is a mixed subdivision of \mathcal{A}_\bullet and $r = (r_1, \dots, r_k) \subset \mathbb{N}^k$ such that $|r| = n$, then the mixed volume of

$$P = \left(\underbrace{P_1, \dots, P_1}_{r_1}, \underbrace{P_2, \dots, P_2}_{r_2}, \dots, \underbrace{P_k, \dots, P_k}_{r_k} \right),$$

is the sum of the volumes of the mixed cells in S^\bullet of type (r_1, r_2, \dots, r_k) :

$$\text{MV}(P) = \sum_{\substack{\mathcal{C}_\bullet \in S^\bullet \\ \text{type}(\mathcal{C}_\bullet) = (r_1, \dots, r_k)}} \text{vol}(\mathcal{C}_\bullet).$$

We describe a process which produces subdivisions from functions. Let $\mathcal{A} \subset \mathbb{R}^n$ be a finite set and let $\ell: \mathcal{A} \rightarrow \mathbb{R}$ be any function. Let $\Gamma_\ell: \mathcal{A} \rightarrow \mathbb{R}^{n+1}$ be the function $\Gamma_\ell(\alpha) = (\alpha, \ell(\alpha))$. We call ℓ a **lifting function** and we call the polytope

$$\text{conv}_\ell(\mathcal{A}) = \text{conv}(\Gamma_\ell(\mathcal{A})) \subset \mathbb{R}^{n+1},$$

the **lift of \mathcal{A} by ℓ** . Similarly, given a set of functions $\ell_\bullet = (\ell_1, \dots, \ell_k)$ with $\ell_i: \mathcal{A}_i \rightarrow \mathbb{R}$, let $\Gamma_{\ell_\bullet}: \mathcal{A}_\bullet \rightarrow \mathbb{R}^{n+1}$ be the function $\Gamma_{\ell_\bullet}(\alpha_1, \dots, \alpha_k) = \sum_{i=1}^k \Gamma_{\ell_i}(\alpha_i)$. Analogously, define

$$\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet) = \text{conv}(\Gamma_{\ell_\bullet}(\mathcal{A}_\bullet)) = \sum_{i=1}^k \text{conv}_{\ell_i}(\mathcal{A}_i).$$

For any polytope $P \subset \mathbb{R}^{n+1}$, the **lower hull** of P is the set

$$\text{hull}(P) = \{P_\omega \mid \omega \in \mathbb{R}^{n+1} \text{ and } \langle \omega, e_{n+1} \rangle < 0\}.$$

The $n + 1$ above is suggestive in that we will often take lower hulls of lifts of polytopes.

Lemma 2.4.4. Let $\mathcal{A} \subset \mathbb{R}^n$ be a finite collection of points and $\ell: \mathcal{A} \rightarrow \mathbb{R}$ a function. The projection of the lower hull of $\text{conv}_\ell(\mathcal{A})$ onto the first n coordinates is $\text{conv}(\mathcal{A})$.

Proof. Since $\text{conv}(\mathcal{A})$ is full-dimensional in its affine span, we may assume $\dim(\text{conv}(\mathcal{A})) = n$ and show that $\alpha \in \text{vert}(\text{conv}(\mathcal{A})) \implies \Gamma_\ell(\alpha) \in \underline{\text{hull}}(\text{conv}_\ell(\mathcal{A}))$.

Let $\alpha \in \text{vert}(\text{conv}(\mathcal{A}))$ and $\omega \in \mathbb{R}^n$ so that $\text{conv}(\mathcal{A})_\omega = \alpha$. Then $(\omega, 0)$ exposes $\Gamma_\ell(\alpha)$ and is in the interior of the $(n + 1)$ -dimensional cone $C[(\omega, 0)]$. Thus, there exists a direction with negative last coordinate which exposes $\Gamma_\ell(\alpha)$ implying that $\Gamma_\ell(\alpha) \in \underline{\text{hull}}(\text{conv}_\ell(\mathcal{A}))$. \square

Definition 2.4.5. Given a set ℓ_\bullet of lifting functions $\ell_i: \mathcal{A}_i \rightarrow \mathbb{R}$, let S^{ℓ_\bullet} be the set of maximal (with respect to inclusion) cells \mathcal{C}_\bullet of \mathcal{A}_\bullet satisfying

- (1) $\dim(\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)) = n$,
- (2) $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet) \in \underline{\text{hull}}(\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet))$.

We remark that the maximality condition in Definition 2.4.5 ensures that $\{\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)\}_{\mathcal{C}_\bullet \in S^{\ell_\bullet}}$ are distinct. Indeed if $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet) = \text{conv}_{\ell_\bullet}(\mathcal{C}'_\bullet)$ but $\mathcal{C}_\bullet \neq \mathcal{C}'_\bullet$ then the union $\mathcal{C}_\bullet \cup \mathcal{C}'_\bullet$ satisfies conditions (1) and (2) of Definition 2.4.5 and contains each cell, contradicting maximality.

Lemma 2.4.6. *The set S^{ℓ_\bullet} is a subdivision of \mathcal{A}_\bullet .*

Proof. If $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ is only n -dimensional, it must lie in a hyperplane implying that $S^{\ell_\bullet} = \mathcal{A}_\bullet$ is the trivial subdivision.

Let $\pi: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ be the projection onto the first n coordinates. Suppose $\mathcal{C}_\bullet \in S^{\ell_\bullet}$ and $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$ is exposed by $\omega \in \mathbb{R}^{n+1}$ where ω has negative last coordinate. Since $\dim(\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)) = n$, its projection under π has dimension at most n . Moreover, its projection has dimension less than n only if the affine span of $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$ contains a line which projects to a point under π . But no such line exists because ω has negative last coordinate and exposes $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$. Thus, S^{ℓ_\bullet} satisfies (1) of Definition 2.4.1.

Given distinct \mathcal{C}_\bullet and \mathcal{C}'_\bullet in S^{ℓ_\bullet} , both $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$ and $\text{conv}_{\ell_\bullet}(\mathcal{C}'_\bullet)$ are facets of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$, and so by part (2) of Lemma 2.1.7, their intersection is a face of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ as well. By part (3) of Lemma 2.1.7, that intersection is a face of both $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$ and $\text{conv}_{\ell_\bullet}(\mathcal{C}'_\bullet)$. It is proper since $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$ and $\text{conv}_{\ell_\bullet}(\mathcal{C}'_\bullet)$ are distinct. Thus, S^{ℓ_\bullet} satisfies (2) of Definition 2.4.1. Part (3) of Definition 2.4.1 follows from Lemma 2.4.4. \square

Any subdivision of the form S^{ℓ_\bullet} is called the **coherent subdivision** of \mathcal{A}_\bullet induced by ℓ_\bullet .

Example 2.4.7. Consider the set $\mathcal{A}_\bullet = \{\mathcal{A}_1\}$ where \mathcal{A}_1 consists of all lattice points in the 3-dilate of the unit square in \mathbb{R}^2 . Let $\ell_\bullet = \{\ell_1\}$ where $\ell_1: \mathcal{A}_1 \rightarrow \mathbb{R}$ is defined by

$$\ell_1(\alpha) = \begin{cases} \pi & \alpha \text{ is in the boundary of } \text{conv}(\mathcal{A}_1) \\ 1 & \text{otherwise} \end{cases}.$$

Then

$$\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet) = \text{conv}_{\ell_1}(\mathcal{A}_1) = \text{conv} \begin{pmatrix} 0 & 0 & 3 & 3 & 1 & 1 & 2 & 2 \\ 0 & 3 & 0 & 3 & 1 & 2 & 1 & 2 \\ \pi & \pi & \pi & \pi & 1 & 1 & 1 & 1 \end{pmatrix}.$$

The lower hull of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ consists of five facets exposed by the directions

$$(0, 0, -1), (0, 1 - \pi, -1), (1 - \pi, 0, -1), (0, \pi - 1, -1), (\pi - 1, 0, -1),$$

which project down to $\text{conv}(\mathcal{A}_1)$, producing a description of the subdivision $S^{\ell_\bullet} = \{\mathcal{C}_\bullet^{(1)}, \dots, \mathcal{C}_\bullet^{(5)}\}$.

The collection $\left\{ \text{conv} \left(\mathcal{C}_\bullet^{(i)} \right) \right\}_{i=1}^5$ consists of five quadrangles displayed in blue in Figure 2.7. \diamond




Figure 2.7: A lifting of a dilated square and the corresponding polyhedral subdivision.

Example 2.4.8. Let $\mathcal{A}_\bullet = \{\mathcal{A}_1, \mathcal{A}_2\}$ where

$$\begin{aligned}\mathcal{A}_1 &= \{(0, 0), (0, 1), (1, 0), (1, 1)\}, \\ \mathcal{A}_2 &= \{(0, 0), (1, 2), (2, 1)\}.\end{aligned}$$

Let $\ell_\bullet = (\ell_1, \ell_2)$ be the functions defined by

$$\begin{aligned}\ell_1(0, 0) &= 2, & \ell_1(0, 1) &= 3, & \ell_1(1, 0) &= 3, & \ell_1(1, 1) &= 3, \\ \ell_2(0, 0) &= 1, & \ell_2(1, 2) &= 1, & \ell_2(2, 1) &= 1.\end{aligned}$$

Figure 2.8 displays $\text{conv}(\mathcal{A}_1)$ and $\text{conv}(\mathcal{A}_2)$ along with the lower hulls of the convex hulls of their lifts in the first two images. The third image displays the lower hull of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ along with the two points of $\Gamma_{\ell_1}(\mathcal{A}_1) + \Gamma_{\ell_2}(\mathcal{A}_2)$ which do not belong to any facet in the lower hull. The third image also contains a depiction of the induced subdivision on \mathcal{A}_\bullet . The green parallelograms and the pink diamond are the mixed cells of the subdivision. The sum of their areas is equal to $4 = \text{MV}(\mathcal{A}_\bullet)$ verifying Lemma 2.4.3. Figure 2.9 shows the projections of these lower facets. \diamond

2.5 Monotonicity and positivity of mixed volume

The **defect** of a collection of polytopes $P = \{P_1, \dots, P_k\}$ is

$$d(P) = \dim \left(\sum_{i=1}^k P_i \right) - k.$$

We say P is **essential** if the defect of any nonempty subset of P is nonnegative.

Lemma 2.5.1. *A collection of polytopes $P_\bullet = \{P_1, \dots, P_n\}$ in \mathbb{R}^n has positive mixed volume if and only if P_\bullet is essential.*




Figure 2.8: A coherent fine mixed subdivision.




Figure 2.9: The projection of a coherent fine mixed subdivision.

Mixed volume is monotonic with respect to inclusion: if P_1 and Q_1, \dots, Q_n are polytopes in \mathbb{R}^n where $P_1 \subset Q_1$, then

$$\text{MV}(P_1, Q_2, \dots, Q_n) \leq \text{MV}(Q_1, \dots, Q_n). \quad (2.5)$$

On the other hand, $P_1 \subsetneq Q_1$ does not imply that the inequality (2.5) is strict.

Conditions for strict monotonicity were originally determined by Maurice Rojas [11] in 1994 but have since been rediscovered for the unmixed case [12] ten years later and again rediscovered and explained in the mixed case [13, 14] another ten years after that. The following version comes from [14].

A subset $S \subset P$ of a convex polytope *touches* a face F of P whenever $S \cap F$ is nonempty.

Lemma 2.5.2. [14, Proposition 3.2] *Let P_1 and $Q_\bullet = (Q_1, \dots, Q_n)$ where P_1, Q_1, \dots, Q_n are polytopes in \mathbb{R}^n such that $P_1 \subset Q_1$. Then $\text{MV}(P_1, Q_2, \dots, Q_n) = \text{MV}(Q_\bullet)$ if and only if P_1 touches every face $(Q_1)_\omega$ for ω in the set*

$$U = \{\omega \in \mathbb{R}^n \mid \{(Q_2)_\omega, \dots, (Q_n)_\omega\} \text{ is essential}\}.$$




Figure 2.10: Polytopes P_1, Q_1 , and Q_2 as in Example 2.5.3 along with a fine mixed subdivision displaying that $\text{MV}(P_1, Q_2) = \text{MV}(Q_1, Q_2)$.

Example 2.5.3. Let $Q_1 = Q_2 = \text{conv}(e_1, -e_1, e_2, -e_2)$ and let $P_1 = [0, 1]^2$. The collection U in Lemma 2.5.2 is the set $\{e_1 + e_2, e_1 - e_2, -e_1 + e_2, -e_1 - e_2\}$ of directions exposing the facets of Q_1 .

Indeed, for every $\omega \in U$, we have $P_1 \cap (Q_1)_\omega \neq \emptyset$, so P_1 touches each facet. Figure 2.10 displays $P_1 \subset Q_1$ and Q_2 along with a depiction of a mixed subdivision of each of the sums $P_1 + Q_2$ and $Q_1 + Q_2$. The mixed cells of each subdivision are the same and so the pairs of polytopes have the same mixed volume, $MV(P_1, Q_2) = MV(Q_1, Q_2) = 16$. \diamond

Given two collections of polytopes $P_\bullet = (P_1, \dots, P_n)$ and $Q_\bullet = (Q_1, \dots, Q_n)$ in \mathbb{R}^n with $P_i \subset Q_i$, one may either iterate Lemma 2.5.2 to determine strict monotonicity or use the following generalized version.

Lemma 2.5.4. [14, Theorem 3.3] *Let $P_\bullet = (P_1, \dots, P_n)$ and $Q_\bullet = (Q_1, \dots, Q_n)$ be collections of polytopes in \mathbb{R}^n such that $P_i \subset Q_i$ for $i = 1, \dots, n$. For $\omega \in \mathbb{R}^n$ let*

$$T_\omega = \{i \in [n] \mid P_i \text{ touches } (Q_i)_\omega\}.$$

Then

$$MV(P_\bullet) < MV(Q_\bullet)$$

if and only if there exists ω such that the collection $\{(Q_i)_\omega \mid i \in T_\omega\} \cup \{Q_i \mid i \in [n] \setminus T_\omega\}$ is essential.

3. ALGEBRAIC GEOMETRY

Algebraic geometry is the study of solution sets of polynomial equations. Such sets are called varieties and there is an intimate dictionary between the algebraic properties of collections of polynomials and the geometric properties of the varieties they define.

We explain a small subset of algebraic geometry relevant to this dissertation. For a more thorough treatment of algebraic geometry we invite the reader to consult [15, 16, 17, 18]. In particular, *Ideals, Varieties, and Algorithms* by Cox, Little, and O’Shea [15] takes a concrete and computational approach to solutions of polynomial equations that is suitable for undergraduates.

Throughout this section, we write $\mathbb{C}[x]$ for the polynomial ring $\mathbb{C}[x_1, \dots, x_n]$ in n variables with coefficients in \mathbb{C} . Given a collection $F \subset \mathbb{C}[x]$, we write $\langle F \rangle$ for the ideal in $\mathbb{C}[x]$ generated by all elements of F . When working in few variables, we use the more familiar variables of x, y, z , and w in that order.

3.1 Affine varieties

We denote n -dimensional complex affine space by

$$\mathbb{C}^n = \{(a_1, a_2, \dots, a_n) \mid a_i \in \mathbb{C}, \quad i = 1, \dots, n\}.$$

For any subset $F \subset \mathbb{C}[x]$, the **affine variety** defined by F is

$$\mathcal{V}(F) = \{(a_1, \dots, a_n) \in \mathbb{C}^n \mid f(a_1, \dots, a_n) = 0 \text{ for all } f \in F\} \subset \mathbb{C}^n.$$

We also refer to $\mathcal{V}(F)$ as the **vanishing locus** of F , the **zero set** of F , or the affine variety **cut out** by F . It is worth mentioning that many texts refer to $\mathcal{V}(F)$ as an “affine algebraic set” and reserve the term “affine variety” for a more specific object. If $f \in \mathbb{C}[x]$ and $f(a) = 0$ for some $a \in \mathbb{C}^n$, then we say that f **vanishes** at a . We sometimes will decorate the notation \mathbb{C}^n with subscripts to indicate the coordinates involved. For example, $\mathcal{V}(y - x^2) \subset \mathbb{C}_{x,y}^2$.

If $X \subset Y$ are both affine varieties, we say X is a **subvariety** of Y . The set \mathbb{C}^n is an affine variety cut out by $\{0\} \subset \mathbb{C}[x]$. We list some affine subvarieties of \mathbb{C}^2 in Figures 3.1-3.5.




Figure 3.1: The set $\mathcal{V}(0)$ defines $\mathbb{C}^2 \subset \mathbb{C}^2$.




Figure 3.2: The set $\mathcal{V}(1)$ defines $\emptyset \subset \mathbb{C}^2$.




Figure 3.3: The set $\mathcal{V}(x^2 + y^2 - 1)$ defines the unit circle in \mathbb{C}^2 .




Figure 3.4: The set $\mathcal{V}(x - y)$ defines a line in \mathbb{C}^2 .




Figure 3.5: The set $\mathcal{V}(x - a, y - b)$ defines a single point $(a, b) \subset \mathbb{C}^2$.

For any subset $S \subset \mathbb{C}^n$ (not necessarily a variety), we denote the set of all polynomials which vanish on S by

$$\mathcal{I}(S) = \{f \in \mathbb{C}[x] \mid f(s) = 0 \text{ for all } s \in S\}.$$

This set is an ideal in the polynomial ring $\mathbb{C}[x]$ since if $f, g \in \mathcal{I}(S)$ and $h \in \mathbb{C}[x]$, then $f + hg \in \mathcal{I}(S)$ because

$$f(s) + h(s)g(s) = 0 + h(s) \cdot 0 = 0 \quad \text{for all } s \in S.$$

Hence, we call $\mathcal{I}(S)$ the **ideal** of S . At this point, we may think of \mathcal{V} and \mathcal{I} as the functions,

$$\begin{aligned} \mathcal{V}: \{\text{subsets of } \mathbb{C}[x]\} &\rightarrow \{\text{subsets of } \mathbb{C}^n\} \\ \mathcal{I}: \{\text{subsets of } \mathbb{C}^n\} &\rightarrow \{\text{subsets of } \mathbb{C}[x]\}. \end{aligned}$$

Lemma 3.1.1. *The functions \mathcal{V} and \mathcal{I} are inclusion reversing:*

- (1) If $S_1 \subset S_2 \subset \mathbb{C}^n$ then $\mathcal{I}(S_2) \subset \mathcal{I}(S_1)$.
- (2) If $F_1 \subset F_2 \subset \mathbb{C}[x]$ then $\mathcal{V}(F_2) \subset \mathcal{V}(F_1)$.

Proof. Suppose $S_1 \subset S_2 \subset \mathbb{C}^n$. Then any polynomial vanishing on S_2 vanishes on the subset S_1 and so $\mathcal{I}(S_2) \subset \mathcal{I}(S_1)$. Suppose that $F_1 \subset F_2 \subset \mathbb{C}[x]$. Then if every element of F_2 vanishes at some $a \in \mathbb{C}^n$, then every element of the subset F_1 vanishes at a as well. \square

Lemma 3.1.2. *For any subset $F \subset \mathbb{C}[x]$, we have $\mathcal{V}(F) = \mathcal{V}(\langle F \rangle)$.*

Proof. If $g \in \langle F \rangle$, then

$$g = \sum_{f \in F} h \cdot f, \quad h \in \mathbb{C}[x], \tag{3.1}$$

and so evaluating the sum at a point $a \in \mathcal{V}(F)$ shows that $g(a) = \sum_{f \in F} h(a) \cdot 0 = 0$ and thus $\mathcal{V}(F) \subseteq \mathcal{V}(\langle F \rangle)$. Conversely, since $F \subset \langle F \rangle$, we have $\mathcal{V}(\langle F \rangle) \subseteq \mathcal{V}(F)$ proving equality. \square

Proposition 3.1.3 (Hilbert's Basis Theorem [19]). *Every ideal $I \in \mathbb{C}[x]$ may be written as $I = \langle f_1, \dots, f_k \rangle$ for some $k \in \mathbb{N}$ and $f_i \in \mathbb{C}[x]$.*

A more general version of Hilbert's Basis Theorem states that the polynomial ring $R[x]$ over any Noetherian ring R is Noetherian. Hilbert proved the case when R is either a field or the ring of integers [19]. Consequently, when studying affine varieties $X = \mathcal{V}(F) \subset \mathbb{C}^n$, we may assume that F is finite.

Given an affine variety $X = \mathcal{V}(f_1, \dots, f_k) \subset \mathbb{C}^n$, declare that the subvarieties of X of the form $X \cap \mathcal{V}(g_1, \dots, g_m)$ for some $g_1, \dots, g_m \in \mathbb{C}[x]$ are **closed**. Lemma 3.1.4 along with the facts that \emptyset and \mathbb{C}^n are affine varieties prove that this gives a topology on $X = \mathbb{C}^n$, which we call the **Zariski topology**.

Lemma 3.1.4. *Finite unions and arbitrary intersections of closed affine subvarieties of \mathbb{C}^n are closed affine subvarieties of \mathbb{C}^n .*

Proof. Let $F, G \subset \mathbb{C}[x]$ be finite generating sets for the ideals I and J respectively. Then

$$\mathcal{V}(I) \cap \mathcal{V}(J) = \mathcal{V}(I + J),$$

equivalently,

$$\mathcal{V}(F) \cap \mathcal{V}(G) = \mathcal{V}(F \cup G).$$

These intersections may be taken to be arbitrary by Hilbert's Basis Theorem. Finite unions are also varieties since,

$$\mathcal{V}(I) \cup \mathcal{V}(J) = \mathcal{V}(IJ),$$

or equivalently,

$$\mathcal{V}(F) \cup \mathcal{V}(G) = \mathcal{V}(\{f \cdot g \mid f \in F, g \in G\}),$$

completing the proof. □

Figures 3.6 and 3.7 display examples of unions and intersections of varieties.

The Zariski topology on a closed subvariety of \mathbb{C}^n is the subspace topology inherited from the Zariski topology on \mathbb{C}^n . Affine varieties come equipped with a second topology: the subspace topology inherited from the Euclidean topology on $\mathbb{C}^n \cong \mathbb{R}^{2n}$. The Zariski topology is weaker than the Euclidean topology in the sense that closed/open sets in the Zariski topology are closed/open in the Euclidean topology but the converse is very much not true.

For any subset $S \subset \mathbb{C}^n$, denote its closure in the Zariski topology by \overline{S} . The following lemma is dual to Lemma 3.1.2.

Lemma 3.1.5. *For any subset $S \subset \mathbb{C}^n$ we have $\mathcal{I}(S) = \mathcal{I}(\overline{S})$.*

Proof. We have $\mathcal{I}(S) \supset \mathcal{I}(\overline{S})$ immediately. Suppose $f \in \mathcal{I}(S)$ so that $f(s) = 0$ for all $s \in S$. If $f \notin \mathcal{I}(\overline{S})$ then there exists some point $s' \in \overline{S}$ such that $f(s') \neq 0$ implying that $\overline{S} \cap \mathcal{V}(f)$ is a variety which is strictly smaller than \overline{S} and contains S , a contradiction. □

Even when restricted to ideals and closed affine varieties, the functions \mathcal{V} and \mathcal{I} are not inverses of each other. It is true that $\mathcal{V}(\mathcal{I}(X)) = X$ for any closed affine variety $X \subset \mathbb{C}^n$, but it is not true that $\mathcal{I}(\mathcal{V}(I)) = I$ for any ideal $I \subset \mathbb{C}[x]$. For example $\mathcal{I}(\mathcal{V}(\langle x^2 \rangle)) = \langle x \rangle$. For \mathcal{V} and \mathcal{I} to be inverses of each other, we must restrict the domain of \mathcal{V} to the subset of ideals satisfying $f^m \in I \iff f \in I$, called **radical ideals**. For any ideal I , the set $\sqrt{I} = \{f \in \mathbb{C}[x] \mid f^m \in I \text{ for some } m \in \mathbb{N}\}$ is a radical ideal called the **radical** of I .




Figure 3.6: The set $\mathcal{V}(x^2 + y^2 - 1, x - y)$ defines two points.




Figure 3.7: The set $\mathcal{V}((x^2 + y^2 - 1) \cdot (x - y))$ defines the union of the unit circle and a line.

Proposition 3.1.6 (Hilbert's Nullstellensatz [20]). *Given an ideal $I \subset \mathbb{C}[x]$,*

$$\mathcal{I}(\mathcal{V}(I)) = \sqrt{I}.$$

The Nullstellensatz implies that with a further restriction to radical ideals, the functions

$$\mathcal{V}: \{\text{radical ideals in } \mathbb{C}[x]\} \rightarrow \{\text{closed affine subvarieties of } \mathbb{C}^n\}$$

$$\mathcal{I}: \{\text{closed affine subvarieties of } \mathbb{C}^n\} \rightarrow \{\text{radical ideals of } \mathbb{C}[x]\}$$

are inverses. As corollaries we have that $\mathcal{V}(I) = \mathcal{V}(\sqrt{I})$ and that $\mathcal{V}(I) \subset \mathbb{C}^n$ is empty if and only if $I = \mathbb{C}[x]$.

Every polynomial $f \in \mathbb{C}[x]$ defines a function

$$\begin{aligned} f: \mathbb{C}^n &\rightarrow \mathbb{C} \\ x &\mapsto f(x). \end{aligned}$$

A **regular function** on an affine variety $X \subset \mathbb{C}^n$ is the restriction of a polynomial function on \mathbb{C}^n to X . Two regular functions f and g on X are the same if and only if $f - g \in \mathcal{I}(X)$ and so the regular functions on X are identified with equivalence classes in the quotient ring $\mathbb{C}[X] = \mathbb{C}[x]/\mathcal{I}(X)$ called the **coordinate ring** of X . Just as regular functions on affine varieties are restrictions of polynomials, a **regular map** of affine varieties $X \subset \mathbb{C}^n, Y \subset \mathbb{C}^m$ is any function

$$\begin{aligned} \varphi: X &\rightarrow Y \\ x &\mapsto (\varphi_1(x), \dots, \varphi_m(x)) \end{aligned}$$

where each $\varphi_i: X \rightarrow \mathbb{C}$ is a regular function. We say φ is an **isomorphism** if it is bijective and its inverse is also a regular map.

A regular map $\varphi: X \rightarrow Y$ of affine varieties naturally induces a \mathbb{C} -algebra homomorphism on the coordinate rings of X and Y in the opposite direction:

$$\begin{aligned} \varphi^*: \mathbb{C}[Y] &\rightarrow \mathbb{C}[X] \\ f &\mapsto f \circ \varphi. \end{aligned}$$

Conversely, given any \mathbb{C} -algebra homomorphism $\phi: \mathbb{C}[Y] \rightarrow \mathbb{C}[X]$, with $\mathbb{C}[Y] = \mathbb{C}[y]/\mathcal{I}(Y)$ and $\mathbb{C}[X] = \mathbb{C}[x]/\mathcal{I}(X)$ let $[g_i] \in \mathbb{C}[X]$ be the image of $[y_i]$ under ϕ . The map,

$$\begin{aligned} \phi^\#: X &\rightarrow Y \\ x &\mapsto (g_1(x), \dots, g_m(x)), \end{aligned}$$

is a regular map of varieties. Note then that φ is an isomorphism of affine varieties if and only if φ^* is a \mathbb{C} -algebra isomorphism.

Example 3.1.7. Given an affine variety $\mathcal{V}(f_1, \dots, f_k) = X \subset \mathbb{C}^n$, subvarieties of X are not always closed. To see this, consider the open subset $U_f = X \setminus \mathcal{V}(f)$ for some $0 \neq f \in \mathbb{C}[X]$. While U_f cannot be expressed as $X \cap \mathcal{V}(g_1, \dots, g_r)$ for any collection $g_1, \dots, g_r \in \mathbb{C}[x]$ (U_f is not closed) it can still be given the structure of a variety in the following way.

Introduce a new variable z and consider $Y = \mathcal{V}(f_1, \dots, f_k) \cap \mathcal{V}(fz - 1) \subset \mathbb{C}^{n+1}$. Here, Y is a closed subvariety of the variety cut out by the same equations as X considered in a higher dimensional space. The coordinate ring $\mathbb{C}[Y]$ is isomorphic to $\mathbb{C}[X][\frac{1}{f}]$ via the map $z \mapsto \frac{1}{f}$. This gives U_f the structure of an affine variety and hence we call it a **principal affine open** subvariety of X .

3.2 Projective varieties

The fundamental theorem of algebra states that a univariate polynomial of degree d has d complex zeros, counted with multiplicity. This fact does not hold over the real numbers and so extending the notion of polynomial equations over \mathbb{R} to those over \mathbb{C} casts the real case into a larger picture which is better behaved. Similarly for varieties, we extend the notion of affine varieties to projective varieties. Doing so produces a more unified understanding of affine varieties.

We wish to keep the notation of $\mathbb{C}[x]$ for a polynomial ring in n variables, and so many of our statements will involve \mathbb{P}^{n-1} rather than \mathbb{P}^n . When we write this, we assume $n \geq 2$.

Definition 3.2.1. Define the equivalence \sim on the set $\mathbb{C}^n \setminus \{\mathbf{0}\}$ by setting $x = (x_1, \dots, x_n) \sim (y_1, \dots, y_n) = y$ if and only if $y = \lambda x$ for some $\lambda \in \mathbb{C} \setminus \{0\}$. **Projective $(n - 1)$ -space** is the quotient

$$\mathbb{P}^{n-1} = (\mathbb{C}^n \setminus \{\mathbf{0}\}) / \sim.$$

We write the equivalence class of (a_1, \dots, a_n) in \mathbb{P}^{n-1} as $[a_1 : \dots : a_n]$.

The zeros of a polynomial $f \in \mathbb{C}[x]$ are well-defined on \mathbb{P}^{n-1} whenever f satisfies the condition

$$f(x) = 0 \text{ if and only if } f(\lambda x) = 0, \text{ for any } \lambda \in \mathbb{C} \setminus \{0\}.$$

This property is equivalent to f being homogeneous. A polynomial

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x_1^{\alpha_1} \cdots x_n^{\alpha_n} \in \mathbb{C}[x], \quad c_\alpha \in \mathbb{C} \setminus \{0\}$$

is **homogeneous** of degree d if $|\alpha| = d$ for all $\alpha \in \mathcal{A}$. Denote the set of homogeneous polynomials of degree d by $\mathbb{C}[x]_d$. If a polynomial is not homogeneous, we say it is **inhomogeneous**.

One may erroneously guess that since the zeros of $f \in \mathbb{C}[x]$ are well-defined on \mathbb{P}^{n-1} if and only if f is homogeneous, then the zeros of $\{f_1, \dots, f_k\} \subset \mathbb{C}[x]$ are well-defined if and only if

f_1, \dots, f_k are homogeneous, but this is **not** necessary. For example, the zero set of $\{x^3 + x^2 + y^2 - z^2, x\}$ are the points $\{[0 : 1 : -1], [0 : -1 : 1]\} \in \mathbb{P}^2$. Due to the argument in Lemma 3.1.2, the zeros of $\{x^3 + x^2 + y^2 - z^2, x\}$ are the same as the zeros of $I = \langle x^3 + x^2 + y^2 - z^2, x \rangle = \langle x^2 + y^2 - z^2, x \rangle$. Ideals such as I which can be generated by homogeneous elements are called **homogeneous ideals**. The common zeros of a collection $F \subset \mathbb{C}[x]$ are well-defined on projective space exactly when $\langle F \rangle$ is a homogeneous ideal.

Definition 3.2.2. Let $F \subset \mathbb{C}[x]$ be a collection of polynomials such that $\langle F \rangle$ is a homogeneous ideal. The **projective variety** defined by F is

$$\mathcal{V}(F) = \{[a_1 : \dots : a_n] \in \mathbb{P}^{n-1} \mid f([a_1 : \dots : a_n]) = 0 \text{ for all } f \in F\} \subset \mathbb{P}^{n-1}.$$

Since \mathbb{P}^{n-1} is a quotient of $\mathbb{C}^n \setminus \{0\}$ with projection $\pi : \mathbb{C}^n \setminus \{0\} \rightarrow \mathbb{P}^{n-1}$, any subset $S \subset \mathbb{P}^{n-1}$ may be pulled back to the subset

$$\mathcal{C}X = \overline{\pi^{-1}(S)} \cup \mathbf{0} \subset \mathbb{C}^n,$$

called the **affine cone** over S . For any subset $S \subset \mathbb{P}^{n-1}$, we define the set $\mathcal{I}(S)$ to be the set of




Figure 3.8: Left: An affine cone over a quartic curve in \mathbb{P}^2 . Right: An affine cone over a circle in \mathbb{P}^2 .

polynomials which vanish on the cone over S . This is an ideal, and the following proves something stronger.

Lemma 3.2.3. *If $S \subset \mathbb{P}^n$, then the ideal $\mathcal{I}(S)$ is homogeneous.*

Proof. Suppose f is a non-homogeneous generator of $\mathcal{I}(S)$ given as

$$f = \sum_{i=0}^k f_i(x),$$

where each $f_i(x)$ is homogeneous. Since f vanishes on a subset of projective space, for any $s \in S$, we have $f(s) = 0$ if and only if $f(\lambda s) = 0$ for any $\lambda \in \mathbb{C} \setminus \{0\}$. On the other hand

$$f(\lambda s) = \sum_{i=0}^k \lambda^{\deg(f_i)} f_i(s),$$

is a polynomial in λ which must vanish whenever $s \in S$. Thus, thinking of λ as a variable, each coefficient $f_i(s)$ of λ^i must be zero. This implies that $f_i(x) \in \mathcal{I}(S)$ and in particular, $f_0 = 0$. Thus, f may be replaced as a generator of $\mathcal{I}(S)$ with the finite set $\{f_i\}_{i=1}^k$ since $f \in \langle f_1, \dots, f_k \rangle \subset \mathcal{I}(S)$. \square

The sets $\mathbb{P}^{n-1} = \mathcal{V}(0)$ and $\emptyset = \mathcal{V}(x_1, \dots, x_n) \subset \mathbb{P}^{n-1}$ are projective varieties. For any projective variety $X \subset \mathbb{P}^{n-1}$, declaring subvarieties of the form $X \cap \mathcal{V}(I) \subset X$ to be closed gives a topology by the same arguments as in the affine case. This topology is also called the **Zariski topology**. The ideal $\mathfrak{m}_0 = \langle x_1, \dots, x_n \rangle$ is called the **irrelevant ideal** since for any homogeneous ideal I , $\mathcal{V}(I \cdot \mathfrak{m}_0) \subset \mathbb{P}^{n-1}$ is the same as $\mathcal{V}(I)$. Since $\mathbf{0}$ is always contained in the cone over S the ideal $\mathcal{I}(S)$ is always contained in the irrelevant ideal.

The same arguments as in the affine case show that the following basic properties of the functions

$$\begin{aligned} \mathcal{V}: \{\text{homogeneous ideals in } \mathbb{C}[x]\} &\rightarrow \{\text{closed projective subvarieties of } \mathbb{P}^{n-1}\} \\ \mathcal{I}: \{\text{subsets of } \mathbb{P}^{n-1}\} &\rightarrow \{\text{homogeneous ideals in } \mathbb{C}[x] \text{ containing } \mathfrak{m}_0\} \end{aligned}$$

hold projectively.

- (1) \mathcal{V} and \mathcal{I} are inclusion reversing,
- (2) $\mathcal{V}(F) = \mathcal{V}(\langle F \rangle)$,
- (3) $\mathcal{I}(S) = \mathcal{I}(\overline{S})$,
- (4) $\sqrt{I \cdot \mathfrak{m}_0} = \mathcal{I}(\mathcal{V}(I))$ (projective Nullstellensatz).

Thus, the functions

$$\begin{aligned} \mathcal{V}: \{\text{homog. radical ideals in } \mathbb{C}[x] \text{ contained in } \mathfrak{m}_0\} &\rightarrow \{\text{closed projective subvarieties of } \mathbb{P}^{n-1}\} \\ \mathcal{I}: \{\text{closed projective subvarieties of } \mathbb{P}^{n-1}\} &\rightarrow \{\text{homog. radical ideals in } \mathbb{C}[x] \text{ contained in } \mathfrak{m}_0\} \end{aligned}$$

are inclusion-reversing inverses.

3.3 Charts on projective space

Consider the open set

$$U_i = \mathbb{P}^{n-1} \setminus \mathcal{V}(x_i).$$

Since any point in projective space has some nonzero coordinate, the U_i cover \mathbb{P}^{n-1} and every point in U_i has a unique representative of the form

$$\left(\frac{x_1}{x_i}, \dots, \frac{x_{i-1}}{x_i}, 1, \frac{x_{i+1}}{x_i}, \dots, \frac{x_n}{x_i} \right).$$

The maps

$$\begin{aligned} \varphi_i: U_i &\rightarrow \mathbb{C}^{n-1} \\ [x] &\mapsto \left(\frac{x_1}{x_i}, \dots, \frac{x_{i-1}}{x_i}, \frac{x_{i+1}}{x_i}, \dots, \frac{x_n}{x_i} \right) \end{aligned}$$

are charts for \mathbb{P}^{n-1} as a manifold. We call these the **standard affine open charts** for \mathbb{P}^{n-1} as they identify each U_i with an affine space. We sometimes will refer to the U_i themselves as charts.

3.4 Homogenizing and dehomogenizing

Let $X \subset \mathbb{P}^{n-1}$ be a projective variety. For $X = \mathcal{V}(f_1, \dots, f_k) \subset \mathbb{P}^{n-1}$ with $f_i \in \mathbb{C}[x]$, the affine cone of X is simply $\mathcal{C}X = \mathcal{V}(f_1, \dots, f_k) \subset \mathbb{C}^n$. For any variable x_i , the intersection of the affine cone of X with the hyperplane $H_i = \mathcal{V}(x_i - 1) \cong \mathbb{C}^{n-1}$ is a closed affine subvariety $\mathcal{C}X \cap \mathcal{V}(x_i - 1) \subset \mathbb{C}^{n-1}$ called the **dehomogenization** of X with respect to x_i . Identifying H_i with \mathbb{C}^{n-1} via the standard affine open chart φ_i , the dehomogenization of X with respect to x_i is the same as the image of $\varphi_i: X \cap U_i \rightarrow \mathbb{C}^{n-1}$.

Conversely, suppose X is an affine subvariety of \mathbb{C}^{n-1} . By introducing a new coordinate x_n we define the **projective closure** of X as

$$\overline{X} = \overline{\{[x : 1] \mid x \in X\}} \subset \mathbb{P}^{n-1}. \quad (3.2)$$

This is the same as taking the closure $\overline{\varphi_n^{-1}(X)}$ where φ_n is the standard affine open chart on \mathbb{P}^{n-1} . When considering the projective closure (3.2) of an affine variety, we call the hyperplane $H_n^\infty = \mathcal{V}(x_n)$ the **hyperplane at infinity**. Of course, the processes of projectively closing an affine variety and dehomogenizing a projective variety may be done with respect to any hyperplane $H \subset \mathbb{C}^n$ not passing through the origin, via the exact same geometric procedure. In these cases, the corresponding hyperplane at infinity is the hyperplane H^∞ through the origin with the same normal direction as H .

The dehomogenization of $\overline{X} \subset \mathbb{P}^{n-1}$ with respect to x_n is exactly X and writing equations for a dehomogenization is straightforward: if $F = \{f_1, \dots, f_k\} \subset \mathbb{C}[x]$ is a collection of homogeneous polynomials, then the dehomogenization of $\mathcal{V}(F)$ with respect to the variable x_i is the affine variety

$$\mathcal{V}(g_1, \dots, g_k) \subset \mathbb{C}^{n-1},$$

where $g_j := f_j(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n)$ is the **dehomogenization** of f_j with respect to x_i . The inverse task of producing the algebraic equations for \overline{X} from those for X is much more difficult. Given a polynomial

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x_1^{\alpha_1} \cdots x_{n-1}^{\alpha_{n-1}} \in \mathbb{C}[x_1, \dots, x_{n-1}], \quad c_\alpha \in \mathbb{C} \setminus \{0\}$$

of degree d , the **homogenization** of f with respect to a new variable x_n is the polynomial

$$\tilde{f} = \sum_{\alpha \in \mathcal{A}} c_{\alpha} x_1^{\alpha_1} \cdots x_{n-1}^{\alpha_{n-1}} \cdot x_n^{d-|\alpha|} \in \mathbb{C}[x].$$

Similarly, for a subset F of polynomials, let $\tilde{F} = \{\tilde{f}\}_{f \in F}$ be the **homogenization** of F . It is easy to see that dehomogenizing \tilde{f} with respect to x_n recovers f and so the dehomogenization of $\mathcal{V}(\tilde{F}) \subset \mathbb{P}^{n-1}$ with respect to x_n is $\mathcal{V}(F) \subset \mathbb{C}^{n-1}$. Unfortunately, $\mathcal{V}(\tilde{F}) \neq \overline{\mathcal{V}(F)}$ and so merely homogenizing equations for an affine variety does not produce equations for its projective closure. In order for this to work, we must homogenize the ideal $I = \langle F \rangle$ generated by F ; that is, $\mathcal{V}(\tilde{I}) = \overline{\mathcal{V}(F)}$. Thus, the homogenization of the ideal generated by a collection of polynomials is *not* the ideal generated by the homogenizations of those polynomials. We illustrate the failure of the naïve homogenization \tilde{F} to cut out the projective closure $\overline{\mathcal{V}(F)}$ in the following example.

Example 3.4.1. Let $F = \{xy - 1, z - x^2\} \subset \mathbb{C}[x, y, z]$ define the set $C \subset \mathbb{C}^3$ called the twisted cubic displayed in Figure 3.9. Let $I = \langle F \rangle$. The homogenization of I with respect to w is




Figure 3.9: A twisted cubic.

$\tilde{I} = \langle xy - w^2, zw - x^2, yz - xw \rangle$, but the ideal generated by the homogenization of F with respect to w is $\langle \tilde{F} \rangle = \langle xy - w^2, zw - x^2 \rangle \subset \tilde{I}$. Notice that the line $\{[0 : s : t : 0] \mid [s : t] \in \mathbb{P}^1\}$ is contained in $\mathcal{V}(xy - w^2, zw - x^2)$ but not in $\mathcal{V}(xy - w^2, zw - x^2, yz - xw)$. \diamond

3.5 Regular functions

We define the **homogeneous coordinate ring** of a projective variety $X \subset \mathbb{P}^{n-1}$ to be the graded quotient ring

$$\mathbb{C}[X] = \mathbb{C}[x]/\mathcal{I}(X).$$

Note that this is the coordinate ring of $\mathcal{C}X$. Contrary to the affine case, most polynomials are not functions on any projective variety $X \subset \mathbb{P}^{n-1}$, rather, the *only* polynomial functions on X are constants: if $f \in \mathbb{C}[x]$ is not constant, then for $\lambda \in \mathbb{C} \setminus \{0, 1\}$ we have $f(\lambda x) = \lambda^{\deg(f)} f(x) \neq f(x)$.

Despite there being almost no polynomial functions on a projective variety, there are still functions between projective varieties which are locally given by polynomials. Given a collection $\{f_1, \dots, f_k\} \subset \mathbb{C}[X]_d$ of polynomials of the same degree, the function

$$f: X \setminus \mathcal{V}(f_1, \dots, f_k) \rightarrow \mathbb{P}^{k-1}$$

$$x \mapsto [f_1(x) : \dots : f_k(x)]$$

is well-defined. If $\varphi: X \rightarrow \mathbb{P}^{k-1}$ is a function such that for every $x \in X$ there exist $f_1, \dots, f_k \in \mathbb{C}[x]$ of the same degree such that $x \notin \mathcal{V}(f_1, \dots, f_k)$ and

$$\varphi(y) = [f_1(y) : \dots : f_k(y)], \text{ for all } y \in X \setminus \mathcal{V}(f_1, \dots, f_k),$$

then we say the map φ is **regular**. Two projective varieties are isomorphic if there exist regular maps $\varphi: X \rightarrow Y$ and $\psi: Y \rightarrow X$ which are inverses of each other. Regular maps of affine/projective varieties are continuous maps under the Zariski topology.

3.6 Irreducibility and dimension

The term “variety” without the adjectives “affine” or “projective” refers to either an affine variety or a projective variety.

3.6.1 Irreducibility

A nonempty variety X is **irreducible** if it satisfies

$$X = X_1 \cup X_2 \implies X = X_1 \text{ or } X = X_2,$$

whenever X_1 and X_2 are closed subvarieties of X . Otherwise, we say it is **reducible**. A union $X = X_1 \cup X_2 \cup \dots \cup X_m$ of sets is **irredundant** if $X_i \not\subset X_j$ for any distinct $i, j \in [m]$. Note that if $X = X_1 \cup X_2$ is a witness for the reducibility of a variety X , then this union is irredundant.

Lemma 3.6.1. *Every nonempty variety X may be written as an irredundant union of finitely many irreducible closed subvarieties*

$$X = X_1 \cup X_2 \cup \dots \cup X_m.$$

Proof. Let X be a variety. If it is irreducible, the lemma is satisfied. Otherwise, it may be written as a union $X = Y_1 \cup Y^{(1)}$ of proper closed subvarieties. As a convention, we suppose that $Y^{(1)}$ is irreducible if Y_1 is, and otherwise we reorder them. Similarly, if $Y^{(1)}$ is reducible, we write $Y^{(1)} = Y_2 \cup Y^{(2)}$. Iteratively applying this process to $Y^{(j)}$ produces a proper infinite chain

$$X \supsetneq Y_1 \supsetneq Y_2 \supsetneq Y_3 \supsetneq \dots$$

of closed subvarieties $Y_i \subset X$. Applying \mathcal{I} to this chain gives an ascending chain of ideals which is proper since each Y_i is closed, contradicting Hilbert’s Basis Theorem. \square

Lemma 3.6.2. *Let X be a variety. If X admits two irredundant decompositions,*

$$X = X_1 \cup \dots \cup X_m, \text{ and } X = Y_1 \cup \dots \cup Y_{m'},$$

into irreducible closed subvarieties, then $m = m'$ and $\{X_1, \dots, X_m\} = \{Y_1, \dots, Y_{m'}\}$.

Proof. We will show that for all i , $Y_i = X_j$ for exactly one j . Consider

$$Y_i = X \cap Y_i = \bigcup_{j=1}^m X_j \cap Y_i.$$

Since Y_i is irreducible, one of the sets in the union must equal Y_i , or equivalently, $X_j \cap Y_i = Y_i$ for some $j \in [m]$, implying that $Y_i \subseteq X_j$. Applying this argument to X_j shows that $X_j \cap Y_k = X_j$ for some $k \in [m']$, implying that $X_j \subseteq Y_k$. Together, this implies $Y_i \subseteq Y_k$ and since the unions are irredundant, $Y_i = Y_k = X_j$. Iterating this argument on $X = \bigcup_{k \neq j} X_k$ and $Y = \bigcup_{k \neq i} Y_k$ proves the result. \square

We call the decomposition in Lemma 3.6.2 the **irreducible decomposition** of X .

Lemma 3.6.3. *Irreducible varieties are those whose ideals are prime.*

Proof. Suppose X is reducible, witnessed by $X_1 \cup X_2$, where X_1, X_2 are proper nontrivial closed subvarieties of X . Write $I_1 = \mathcal{I}(X_1)$ and $I_2 = \mathcal{I}(X_2)$ so that $I_1 I_2 = \mathcal{I}(X)$. Picking $f_1 \in I_1 \setminus I_2$ and $f_2 \in I_2 \setminus I_1$, we see that $f_1, f_2 \notin \mathcal{I}(X)$ but $f_1 f_2 \in \mathcal{I}(X)$ so $\mathcal{I}(X)$ is not prime. Conversely, suppose $\mathcal{I}(X)$ is not prime, witnessed by $f_1, f_2 \notin \mathcal{I}(X)$ yet $f_1 f_2 \in \mathcal{I}(X)$. Let $I_1 = \langle f_1 \rangle + \mathcal{I}(X)$ and $I_2 = \langle f_2 \rangle + \mathcal{I}(X)$. We claim that $X = X_1 \cup X_2$ where $X_1 = \mathcal{V}(I_1)$ and $X_2 = \mathcal{V}(I_2)$. Both $X_1, X_2 \subset X = \mathcal{V}(\mathcal{I}(X))$ since $\mathcal{I}(X) \subset I_1(X)$ and $\mathcal{I}(X) \subset I_2(X)$. Moreover, their union $X_1 \cup X_2$ is $\mathcal{V}(I_1 I_2) = \mathcal{V}(\mathcal{I}(X)) = X$. \square

3.6.2 Dimension

The **dimension** of an irreducible variety X is the longest length $\dim(X)$ of a proper chain of irreducible closed subvarieties

$$\emptyset = X_{-1} \subsetneq X_0 \subsetneq X_1 \subsetneq \cdots \subsetneq X_{\dim(X)} = X.$$

If X is not irreducible, then its **dimension** is the maximum dimension of its irreducible components. The **codimension** of a subvariety $X \subset Z$ is $\text{codim}_Z(X) = \dim(Z) - \dim(X)$. We will omit the subscript on codimension whenever $Z = \mathbb{C}^n$ or $Z = \mathbb{P}^n$ or we have specifically mentioned Z and so the subscript is clear from context. If X and Y are both subvarieties of Z and $\dim(X) = \text{codim}(Y)$, then we say that X and Y have **complementary dimension**.

A variety of codimension 1 is the zero set of a single polynomial and is called a **hypersurface**. Zero-dimensional varieties are finite collections of points. If X is a closed subvariety of an irreducible variety Z and $\dim(X) = \dim(Z)$ then $X = Z$.

Given a variety X one expects the intersection of X and a hypersurface to have dimension one less than X . The following lemma states that the dimension is lowered by at most one in the projective setting.

Lemma 3.6.4. [18, I.6.2 Corollary 5 of Theorem 4] *Let $f_1, \dots, f_k \in \mathbb{C}[x]$ be homogeneous polynomials and suppose $X \subset \mathbb{P}^{n-1}$ is a projective variety of dimension m . Then we have that $\dim(\mathcal{V}(f_1, \dots, f_k) \cap X) \geq m - k$.*

The affine analog of Lemma 3.6.4 gives a weaker conclusion.

Lemma 3.6.5. [18, I.6.2 Corollary 2 of Theorem 5] Let $f_1, \dots, f_k \in \mathbb{C}[x]$ and $X \subset \mathbb{C}^n$ an affine variety of dimension m . Every irreducible component of $\mathcal{V}(f_1, \dots, f_k) \cap X \subset \mathbb{C}^n$ has dimension at least $m - k$.

We distinguish the conclusions of Lemma 3.6.4 and Lemma 3.6.5 in the following example.

Example 3.6.6. Let $X = \mathbb{C}_{x,y}^2$ and $f_1 = xy - 1, f_2 = x$. Then $\mathcal{V}(f_1, f_2) = \emptyset$. While it is true that Lemma 3.6.5 guarantees that every irreducible component of $\mathcal{V}(f_1, f_2) \cap X$ has dimension at least 0, the variety $\mathcal{V}(f_1, f_2) \cap X$ has no irreducible components and so the lemma does not apply.




Figure 3.10: The affine varieties $\mathcal{V}(xy - 1)$ and $\mathcal{V}(x)$.

Naïvely homogenizing, take $\tilde{f}_1 = xy - z^2, \tilde{f}_2 = x$ and $X = \mathbb{P}^2$, so that

$$\mathcal{V}(\tilde{f}_1, \tilde{f}_2) \cap X = \{[0 : 1 : 0]\} \subset \mathbb{P}^2.$$

This is nonempty as guaranteed by Lemma 3.6.4.

Notice that with respect to this homogenization, the point $[0 : 1 : 0]$ is on the line at infinity. This aligns with our intuition as Figure 3.10 shows that the line $\mathcal{V}(x)$ and the hyperbola $\mathcal{V}(xy - 1)$ asymptotically approach each other along the y -axis. \diamond

Corollary 3.6.7. If $\emptyset \neq \mathcal{V}(f_1, \dots, f_k) \subset \mathbb{C}^n$ then $\mathcal{V}(f_1, \dots, f_k)$ has dimension at least $n - k$.

Lemma 3.6.8. [18, I.6.2 Theorem 6.] Let X and Y be subvarieties of \mathbb{C}^n (or \mathbb{P}^n) of dimensions m_1 and m_2 respectively. Then every irreducible component of $X \cap Y$ has dimension at least $m_1 + m_2 - n$. Moreover, if X and Y are projective and $m_1 + m_2 \geq n$ then $X \cap Y \neq \emptyset$.

3.7 Function fields and rational functions

When $X \subset \mathbb{C}^n$ is irreducible, $\mathcal{I}(X)$ is prime and so its coordinate ring $\mathbb{C}[X]$ is an integral domain. The field of fractions of $\mathbb{C}[X]$ is called the **function field** of X , denoted $\mathbb{C}(X)$, and consists of all rational functions $g/h: X \dashrightarrow \mathbb{C}$ such that $h \notin \mathcal{I}(X)$.

If X is an irreducible projective variety, the **function field** of X , denoted $\mathbb{C}(X)$, consists of rational functions $g/h: X \dashrightarrow \mathbb{C}$ such that g and h have the same degree and $h \notin \mathcal{I}(X)$. We use the dashed arrow notation to remind ourselves that rational functions are not defined everywhere but they are well-defined on the open subset $U = X \setminus \mathcal{V}(h)$. Indeed if $u \in U$, then

$$g(\lambda u)/h(\lambda u) = (\lambda^d g(u))/(\lambda^d h(u)) = g(u)/h(u)$$

for all $\lambda \in \mathbb{C} \setminus \{0\}$. Unlike the affine case, the function field of X is *not* the field of fractions of $\mathbb{C}[X]$, but rather, its 0-th graded piece. A **rational map** $\varphi: X \dashrightarrow \mathbb{P}^m$ from X to projective space is given as

$$\varphi = [\varphi_1 : \cdots : \varphi_{m+1}], \quad \varphi_i \in \mathbb{C}(X), \quad \text{for } i = 1, \dots, m+1, \quad (3.3)$$

where $\varphi_i = g_i/h_i$. The map φ is defined on the open set

$$V = X \setminus (\mathcal{V}(g_1, \dots, g_{m+1}) \cup \mathcal{V}(h_1 \cdots h_{m+1})).$$

We may always write a rational map (3.3) so that φ_i are polynomials. Since each φ_i is of the form $\varphi_i = g_i/h_i$, we simply clear denominators,

$$\varphi = [g_1/h_1 : \cdots : g_{m+1}/h_{m+1}] = [f_1 : \cdots : f_{m+1}] \quad (3.4)$$

where $f_i = (g_i/h_i) \cdot \prod_{j=1}^{m+1} h_j$. Even though the coordinates of every rational function may be written as polynomials, these are not regular functions because $\mathcal{V}(f_1, \dots, f_{m+1})$ may not be empty.

Two rational functions $g/h, g'/h' \in \mathbb{C}(X)$ are equal whenever $gh' - g'h \in \mathcal{I}(X)$. Of course, they may be defined on different open subsets $U = X \setminus \mathcal{V}(h)$ and $U' = X \setminus \mathcal{V}(h')$, but they agree on the dense open subset $U \cap U'$. Similarly, two rational maps

$$\varphi = [f_1 : \cdots : f_{m+1}] \quad \text{and} \quad \varphi' = [f'_1 : \cdots : f'_{m+1}],$$

written in the form (3.4), are the same if $f_i f'_j - f_j f'_i \in \mathcal{I}(X)$ for all $i, j \in [m+1]$. Equivalently, φ and φ' agree on a dense open subset of X . Thus, for any dense open subset $U \subset X$, rational maps on X are determined by their values on U . Hence, when U is an affine open subvariety of X , $\mathbb{C}(U) = \mathbb{C}(X)$.

3.8 Products, graphs, and the degree of a variety

Given a function $f: A \rightarrow B$ of sets, the graph of f is simply the set $\Gamma(f) = \{(a, b) \mid a \in A, b = f(a)\} \subset A \times B$. We may similarly define the graph of a regular or rational map of algebraic varieties, however, *a priori* these graphs do not come equipped with the structure of a variety. We obtain a variety structure on the graph of a map by developing a variety structure on products of varieties *vis-à-vis* Segre maps.

3.8.1 Segre maps

Given two projective spaces \mathbb{P}^{n-1} and \mathbb{P}^{m-1} , define the **Segre map**

$$\sigma_{n-1,m-1}: \mathbb{P}^{n-1} \times \mathbb{P}^{m-1} \rightarrow \mathbb{P}^{nm-1},$$

to be the function sending a pair of points $[x] \in \mathbb{P}^{n-1}$ and $[y] \in \mathbb{P}^{m-1}$ to the point whose coordinates are all possible pair-wise products of the coordinates of $[x]$ and $[y]$, namely,

$$\sigma_{n-1,m-1}([x_1 : \cdots : x_n], [y_1 : \cdots : y_m]) = [x_1y_1 : \cdots : x_iy_j : \cdots : x_ny_m].$$

Giving \mathbb{P}^{nm-1} coordinates $z_{i,j} = x_iy_j$, the image of the Segre map is

$$\Sigma_{n-1,m-1} = \mathcal{V}(z_{i,j}z_{k,l} - z_{i,l}z_{k,j}) \subset \mathbb{P}^{nm-1},$$

and is called the **Segre variety**.

3.8.2 Products

Defining the product of affine varieties is easy. If $X \subset \mathbb{C}^n$ and $Y \subset \mathbb{C}^m$, the Cartesian product $X \times Y = \{(x, y) \mid x \in X, y \in Y\}$ naturally lives in $\mathbb{C}^n \times \mathbb{C}^m \cong \mathbb{C}^{n+m}$ via the map $((x_1, \dots, x_n), (y_1, \dots, y_m)) \mapsto (x_1, \dots, x_n, y_1, \dots, y_m)$ and its structure as an affine variety comes from this realization of $X \times Y$ as a subvariety of \mathbb{C}^{n+m} .

Given two projective varieties $X \subset \mathbb{P}^{n-1}$ and $Y \subset \mathbb{P}^{m-1}$, from now on, whenever we write the **product** $X \times Y$ we will mean the image of the Cartesian product $X \times Y$ under the Segre map

$$X \times Y = \{\sigma_{n-1,m-1}(x, y) \mid x \in X, y \in Y\}.$$

The Segre map is injective and so we will write elements of $X \times Y$ as (x, y) where $x \in X$ and $y \in Y$. The projection maps $\pi_X: X \times Y \rightarrow X$ and $\pi_Y: X \times Y \rightarrow Y$ onto the first and second coordinates are regular maps. When $X \subset \mathbb{C}^{n-1}$ and $Y \subset \mathbb{P}^{m-1}$ we have $X \xrightarrow{\iota} \overline{X} \subset \mathbb{P}^{n-1}$ and so we take $X \times Y$ to be the variety $X \times Y = \sigma_{n-1,m-1}(\iota(X), Y) \subset \mathbb{P}^{nm-1}$.

3.8.3 Graphs

Given a regular function $\varphi: X \rightarrow Y$ define the **graph of φ** to be

$$\Gamma(\varphi) = \{(x, y) \mid x \in X, y = \varphi(x)\} \subset X \times Y.$$

This is a closed subvariety of $X \times Y$ and the projection maps are regular. When X or Y are projective, we will often first take affine open subsets so that φ is a map of affine varieties and the graph is an affine variety. When we do this, we may assume $X = \mathcal{V}(f_1, \dots, f_k) \subset \mathbb{C}^n$ and $Y \subset \mathbb{C}^m$ so the graph of φ is the subvariety of $\mathbb{C}^n \times \mathbb{C}^m \cong \mathbb{C}^{n+m}$ given explicitly as

$$\Gamma(\varphi) = \mathcal{V}(f_1, \dots, f_k, \varphi_1 - x_{n+1}, \dots, \varphi_m - x_{n+m}).$$

Lemma 3.8.1. *The closure of the image of an irreducible variety under a regular map is irreducible.*

Proof. Suppose $\varphi: X \rightarrow Y$ is a regular map with Y reducible, witnessed by $Y = Y_1 \cup Y_2$. Since φ is continuous with respect to the Zariski topology, $X = \varphi^{-1}(Y_1) \cup \varphi^{-1}(Y_2)$ is an irredundant union of proper nonempty closed subvarieties of X witnessing the reducibility of X . \square

Given a rational map $\varphi: X \dashrightarrow \mathbb{P}^m$ of projective varieties, let $U \subset X$ be its domain of definition. We define the **graph** of φ , denoted $\Gamma(\varphi)$, to be the closure of $\Gamma(\varphi|_U)$ in $X \times \mathbb{P}^m$ and we define the **image** of φ to be the image of $\Gamma(\varphi)$ under π_Y . The inverse image of a subvariety $Z \subset \mathbb{P}^m$ is $\varphi^{-1}(Z) = \pi_X(\pi_{\mathbb{P}^m}^{-1}(Z))$. Given $Y \subset \mathbb{P}^m$, a **rational map** $\varphi: X \dashrightarrow Y$ is any rational map $\varphi: X \dashrightarrow \mathbb{P}^m$ whose image is contained in Y .

Lemma 3.8.2. *The image of an irreducible variety under a rational map is irreducible.*

3.8.4 Dominant maps

Unfortunately, given two rational maps $\varphi: X \dashrightarrow Y$, and $\psi: Y \dashrightarrow Z$, the composition $\psi \circ \varphi: X \dashrightarrow Z$ is not always well-defined as shown in the following example.

Example 3.8.3. Let

$$\begin{array}{ccc} \varphi: \mathbb{P}^1 \rightarrow \mathbb{P}^3 & & \psi: \mathbb{P}^3 \dashrightarrow \mathbb{P}^2 \\ [u : v] \mapsto [u^3 : u^2v : uv^2 : v^3] & \text{and} & [x : y : z : w] \mapsto [xz - y^2 : yw - z^2 : xw - yz]. \end{array}$$

Then $\psi \circ \varphi([u : v]) = [0 : 0 : 0]$ is a point in \mathbb{P}^2 . ◇

The problem in Example 3.8.3 is that the image of φ is disjoint from the domain of definition of ψ . This motivates the definition of dominant maps, a subset of rational maps for which composition is always well-defined.

We say a rational map $\varphi: X \dashrightarrow Y$ of varieties is **dominant** if $\varphi(X)$ is dense in Y . If $\varphi: X \dashrightarrow Y$ is dominant with domain of definition U and $\psi: Y \dashrightarrow Z$ with domain of definition V , then the domain of definition of the composition $\psi \circ \varphi: X \dashrightarrow Z$ is $U \cap \varphi^{-1}(V)$.

In the same way that a regular map $\varphi: X \rightarrow Y$ of affine varieties induces a \mathbb{C} -algebra homomorphism $\varphi^*: \mathbb{C}[Y] \rightarrow \mathbb{C}[X]$, a dominant map $\varphi: X \dashrightarrow Y$ induces an *injective* \mathbb{C} -algebra homomorphism which (when X is irreducible) extends to the function field $\varphi^*: \mathbb{C}(Y) \rightarrow \mathbb{C}(X)$. Conversely, given an injective homomorphism $\phi: \mathbb{C}(Y) \rightarrow \mathbb{C}(X)$ of function fields, we obtain a dominant rational map $\phi^\#: X \dashrightarrow Y$.

Lemma 3.8.4. [18, I.6.3 Theorem 7] *Let $\varphi: X \rightarrow Y$ be a surjective regular map between irreducible varieties and that $\dim(X) = n$ and $\dim(Y) = m$. Then $m \leq n$ and*

- (1) $\dim(F) \geq n - m$ for any $y \in Y$ and for any component F of the fiber $\varphi^{-1}(y)$.
- (2) there exists a nonempty open subset $U \subset Y$ such that $\dim(\varphi^{-1}(y)) = n - m$ for $y \in U$.

Lemma 3.8.5. [18, I.6.3 Theorem 8] *Let $\varphi: X \rightarrow Y$ be a regular map between projective varieties with $\varphi(X) = Y$. Suppose that Y is irreducible, and that all the fibers $\varphi^{-1}(y)$ for $y \in Y$ are irreducible of the same dimension. Then X is irreducible.*

Proposition 3.8.6. [16, Proposition 7.16] *Given a dominant map $\pi: X \dashrightarrow Y$, there exists an open subset $U \subset Y$ such that the fiber $\pi^{-1}(u)$ is finite if and only if π^* expresses the field $\mathbb{C}(X)$ as a finite extension of the field $\mathbb{C}(Y)$. The number of points in a fiber over $u \in U$ is the degree of the field extension.*

Proof. We recount the proof from [16]. Without loss of generality, replace X and Y with affine open subsets so that π is a projection map $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_{n-1})$ of affine varieties. Thus, the function field $\mathbb{C}(X)$ is generated over $\mathbb{C}(Y)$ by x_n . If x_n is algebraic over $\mathbb{C}(Y)$ with minimal polynomial

$$g_{(x_1, \dots, x_{n-1})}(x_n) = a_d(x_1, \dots, x_{n-1})x_n^d + a_{d-1}(x_1, \dots, x_{n-1})x_n^{d-1} + \dots,$$

we may clear denominators so that the coefficients of g are regular functions. The discriminant D of g is a closed subset of the coefficient space since \mathbb{C} is algebraically closed and so outside of this locus the fibers of π consist of exactly d points.

Conversely, if x_n is transcendental, then any polynomial in $\mathcal{I}(X)$ written in $\mathbb{C}(x_1, \dots, x_{n-1})[x_n]$ must be identically zero as functions on Y . That is, the fiber $\pi^{-1}(y)$ for any $y \in Y$ contains infinitely many points. \square

We remark that the locus of points $x^* \in Y$ which do not have the generic fiber size in Proposition 3.8.6 come in three types:

- (1) The coefficient x^* belongs to the discriminant D because $g_{x^*}(x_n)$ has roots with multiplicity.
- (2) The coefficient x^* belongs to the discriminant D because $a_d(x^*) = 0$.
- (3) The rational coefficients $a_i(x_1, \dots, x_{n-1})$ are not defined at x^* .

A rational map $\pi: X \dashrightarrow Y$ satisfying Proposition 3.8.6 is called a **generically finite** map. The degree of the field extension is called the **degree** of the map.

Corollary 3.8.7. *Suppose $f: X \dashrightarrow Y$ is a dominant map of irreducible varieties of the same dimension. Then f satisfies Proposition 3.8.6.*

3.8.5 Degree of a variety

A variety cut out by linear equations is called a **linear** variety. The set of all linear subvarieties of \mathbb{P}^n of dimension k corresponds to the set of all $k + 1$ planes in \mathbb{C}^{n+1} through the origin. This space is called the **Grassmannian** of $(k + 1)$ -planes in \mathbb{C}^{n+1} and is denoted $\text{Gr}(k + 1, n + 1)$. The Grassmannian itself is a projective variety cut out by all relations amongst the minors of a $(k + 1) \times (n + 1)$ matrix. Similarly, a linear subvariety $L \subset \mathbb{C}^n$ of dimension k corresponds to the $(k + 1)$ -plane \bar{L} in \mathbb{P}^n . Thus, it makes sense to talk about subvarieties and open subsets of the space of linear spaces of a particular dimension.

Lemma 3.8.8. *Let X be an irreducible codimension m subvariety of \mathbb{C}^n or \mathbb{P}^n . There exists an open subset $V \subset \text{Gr}(k + 1, n + 1)$ with the property $L \in V \implies 0 < |L \cap X| < \infty$ if and only if $k = m$. When $k = m$, there exists a smaller open subset $V' \subset V$ for which the number of such intersection points is constant.*

Proof. The result is true for an affine variety if and only if it is true for its projective closure. Let X be projective and suppose such an open set $V \subset Y = \text{Gr}(k + 1, n + 1)$ exists. Consider the variety

$$Z = \{(x, L) \mid L \in Y, x \in L \cap X\} \subset X \times Y$$

with projections π_X and π_Y to X and Y respectively. By assumption, the image of π_Y contains V and the fibers of π_Y over a point $v \in V$ are finite. The fibers over π_X are all irreducible of dimension $\dim(Y) - (n - k)$ and so Z is irreducible of dimension $\dim(X) + \dim(Y) - (n - k)$ by Lemmas 3.8.4 and 3.8.5. If $\dim(X) < n - k$ then $\dim(Z) < \dim(Y) = \dim(V)$ and so $V \not\subset \pi_Y(Z)$, a contradiction. Thus, $k \geq n - \dim(X) = m$. On the other hand, if $k > m$ then by Lemma 3.6.4 the intersection $X \cap L$ is either empty or at least one-dimensional. We conclude $k = m$.

Conversely, if $k = m$,

$$\dim(Z) = \dim(X) + \dim(Y) - (n - m) = \dim(Y),$$

implying that π_Y is generically finite (such a V exists). By Lemma 3.8.6 there is an open subset $V' \subset V$ such that the number of points in a fiber of π_Y over V' is constant. \square

When $X, L \subset Z$, and $L \in V'$ as in the above lemma, then cardinality $|X \cap L|$ is some constant $d \in \mathbb{N}$. This number d is called the **degree** of X and is denoted $\deg(X)$. Given an irreducible polynomial $f \in \mathbb{C}[x]$ The degree of a hypersurface $\mathcal{V}(f)$ is the degree of f . The degree of a collection of d points is d . We give the **first Bertini theorem**.

Lemma 3.8.9. [18, II.6.1 Theorem 1] *Let X and Y be irreducible varieties defined over a field of characteristic 0 and $f: X \rightarrow Y$ a regular map such that $f(X)$ is dense in Y . Suppose that X remains irreducible over the algebraic closure $\overline{\mathbb{C}(Y)}$ of $\mathbb{C}(Y)$. Then there exists an open dense set $U \subset Y$ such that all the fibers $f^{-1}(y)$ over $y \in U$ are irreducible.*

Corollary 3.8.10. *Let X be a variety and H a general hyperplane. Then*

- (1) $\deg(X) = \deg(X \cap H)$.
- (2) $\dim(X) - 1 = \dim(X \cap H)$.
- (3) *If X is irreducible of dimension at least two, then $X \cap H$ is irreducible.*

Proof. Part (1) follows directly from the definition of the degree of a variety. For part (2), if X is irreducible and $\dim(X) = \dim(X \cap H)$ then H must contain X , but most hyperplanes do not contain a nonempty variety. For part (3), suppose that L is the normal line to the hyperplane H and consider the linear projection $f: X \rightarrow L$. Then for a general point $y \in L$, the fiber $f^{-1}(y)$ is irreducible and the hyperplane slice $X \cap H$ corresponds to one such fiber. \square

3.9 Singular points

Let $X = \mathcal{V}(F) \subset \mathbb{C}^n$ be an irreducible affine variety of dimension m such that $\langle F \rangle$ is a radical ideal. We say X is **smooth** at a point $p \in X$ if the rank of the **Jacobian matrix**

$$DF = \begin{bmatrix} \frac{\partial f_i}{\partial x_j} \end{bmatrix}$$

evaluated at p is $n - m$, otherwise we say p is **singular**. We say X is **smooth** if it is smooth at all of its points. The set $\text{Sing}(X)$ of singular points of X is a proper closed subvariety of X [17, Theorem 5.3] and so the set of smooth points of X is open and dense. If p is a point of a projective variety X , then p is smooth on X if p is smooth on $U_i \cap X$ for some affine chart containing p .

The following proposition is called the **second Bertini theorem**.

Proposition 3.9.1. [18, II.6.2 Theorem 2] Let $f: X \rightarrow Y$ be a regular dominant map with X smooth. There exists a dense open set $U \subset Y$ such that the fiber $f^{-1}(y)$ is nonsingular for every $y \in U$.

A corollary of the second Bertini theorem is fundamental to the theory of numerical algebraic geometry (Section 6).

Corollary 3.9.2. If X is a smooth variety and H is a general hyperplane then $X \cap H$ is smooth.

Proof. This follows by the same argument as in Corollary 3.8.10 replacing the first Bertini theorem with the second Bertini theorem. \square

Let $\mathcal{V}(F) \subset \mathbb{C}_x^n \times \mathbb{C}_t$ be an irreducible affine variety of dimension one such that the projection

$$\begin{aligned} \pi: \mathcal{V}(F) &\rightarrow \mathbb{C}_t \\ (t, x_1, \dots, x_n) &\mapsto t \end{aligned}$$

is dominant. The Jacobian DF encodes the points $t \in \mathbb{C}_t$ for which $\pi^{-1}(t)$ does not have the generic cardinality as in Proposition 3.8.6. Let $D_t F = \frac{\partial F}{\partial t}$ and $D_x F = \frac{\partial F}{\partial x}$ so that DF is the matrix whose first column is $D_t F$ and whose last n columns are $D_x F$. Given a point $p = (t^*, x^*) \in \mathcal{V}(F)$, p is smooth on $\mathcal{V}(F)$ when $\text{rank}(DF(t^*, x^*)) = n$. If $\text{rank}(D_x F(p)) = n - 1$, then p is singular ($\text{rank}(DF(p)) = n - 1$) on $\mathcal{V}(F)$ or the fiber $\pi^{-1}(t^*)$ has points with multiplicity. We depict this dichotomy in Figure 3.11.

Example 3.9.3. Consider the curve $\mathcal{V}(f)$ with

$$f = (x - 3)^2 - (t - 1)(t + 1)(t + 2)^2,$$

displayed in Figure 3.11. The rank of $D_x f$ is zero at the points $(-2, 3)$, $(-1, 3)$, and $(1, 3)$ on $\mathcal{V}(f)$. The rank of Df at these points is 0, 1, and 1 respectively. \diamond




Figure 3.11: Three points on a quartic curve in $\mathcal{V}(f) \subset \mathbb{C}_t \times \mathbb{C}_x$ such that the matrix $D_x f$ has rank zero when evaluated at these points.

4. BRANCHED COVERS AND GROUPS

Representing geometric objects as fibers of maps is a powerful method in geometry. For example, the simple problem of solving a quadratic equation $ax^2 + bx + c = 0$ for $a, b, c \in \mathbb{C}$ may be interpreted geometrically via a map

$$\begin{aligned} \pi: \{([a : b : c], x) \in \mathbb{P}_{a,b,c}^2 \times \mathbb{C} \mid ax^2 + bx + c = 0\} &\rightarrow \mathbb{P}_{a,b,c}^2 \\ ([a : b : c], x) &\mapsto [a : b : c], \end{aligned}$$

over the parameter space $\mathbb{P}_{a,b,c}^2$. We identify the solutions of a quadratic equation such as $3x^2 + 8x + 4 = 0$ with the fiber $\pi^{-1}([3 : 8 : 4]) = \{([3 : 8 : 4], -2), ([3 : 8 : 4], -\frac{2}{3})\}$. The subset $U \subset \mathbb{P}_{a,b,c}^2$ of parameters whose corresponding quadratic equation has two distinct solutions is the complement of the vanishing of the discriminant $B = \mathcal{V}(a(b^2 - 4ac))$ and comprises a dense open subset of $\mathbb{P}_{a,b,c}^2$. Since $a \neq 0$ for $[a : b : c] \in U$, rescaling to monic quadratic equations,

$$\begin{aligned} \pi|_{a=1}: \{(b, c, x) \in \mathbb{C}^3 \mid x^2 + bx + c = 0\} &\rightarrow \mathbb{C}_{b,c}^2 \\ (b, c, x) &\mapsto (b, c) \end{aligned}$$

gives a “branched cover” of affine varieties. In this framework, the solutions of $3x^2 + 8x + 4 = 0$ are identified with the fiber over the parameter $(\frac{8}{3}, \frac{4}{3}) \in \mathbb{C}_{b,c}^2$. Figure 4.1 depicts this parameter space along with the set $U|_{a=1} \subset \mathbb{C}_{b,c}^2$. Every point $(b, c) \in \mathbb{C}_{b,c}^2$ which is not on the dotted parabola in




Figure 4.1: The parameter space $\mathbb{C}_{b,c}^2$ along with the discriminant $\mathcal{V}(b^2 - 4c)$.

Figure 4.1 is in $U|_{a=1}$. Fibers over parameters in the red region, like $\pi|_{a=1}^{-1}(0, 1) = \pm\sqrt{-1}$, have two distinct (complex conjugate) nonreal points, and the fibers over points in the blue region have two distinct real points. Points on the parabola have fibers consisting of one real solution occurring with multiplicity two.

The variety $\mathcal{V}(b^2 - 4c)$ is a hypersurface in $\mathbb{C}_{b,c}^2$ and thus has (complex) codimension 1 and real codimension 2. Thus, U is a connected real manifold, even though it is disconnected when restricted to $\mathbb{R}_{b,c}^2$ as seen in Figure 4.1.

The discussion above distills the essence of the behavior of branched covers. We give an elementary treatment of branched covers and covering spaces in Section 4.1 and we provide background on permutation groups, monodromy groups, and Galois groups in Sections 4.2-4.3. We conclude in Section 4.4 with a discussion of decomposable branched covers.

4.1 Branched covers

An **(irreducible) branched cover** is a dominant map $\pi: X \dashrightarrow Z$ of irreducible varieties of the same dimension. We may assume that we restrict to an affine open subset of X so that $\pi: X \rightarrow Z$ is regular with $X \subset \mathbb{C}^n$ and $Z \subset \mathbb{C}^m$. Irreducible branched covers are generically finite in the sense of Proposition 3.8.6 and thus there exists a number d and a dense open set $U \subset Z$ such that for any $u \in U$, the fiber $\pi^{-1}(u)$ has cardinality d and π^* expresses the field $\mathbb{C}(X)$ as a degree d field extension of $\mathbb{C}(Z)$.

More generally, a **branched cover** is a map $\pi: X \rightarrow Z$ such that X is reducible and the restriction of π to some top dimensional component of X is an irreducible branched cover. The restriction of π to every other top dimensional component is either dominant or the image is a proper closed subvariety of Z . Let X_1, \dots, X_k be those components of X such that the restriction π_i of π to X_i is dominant. Suppose π_i has fibers of cardinality d_i over any point in the dense open subset $U_i \subset Z$. Then it is immediate that for any $u \in U = \bigcap_{i=1}^k U_i$, the cardinality of the fiber $\pi^{-1}(u)$ is $d = \sum_{i=1}^k d_i$.

Given a branched cover $X \xrightarrow{\pi} Z$ as above, d is the **degree** of π , U is the set of **regular values** of π , and the complement of U is the **branch locus** of π . We say π is **trivial** if $d = 1$. With respect to the real Euclidean topologies X and Z inherit from their ambient spaces, there exists an open cover $\{V_\beta\}$ of U such that for each β , the fiber $\pi^{-1}(V_\beta)$ is a disjoint union of d open sets in $\pi^{-1}(U)$, each of which is mapped homeomorphically onto V_β . Such a map $\pi|_U: \pi^{-1}(U) \rightarrow U$ is called a **d -sheeted covering space**.

Many properties of branched covers, like the well-definedness of degree and regular values, extend immediately from their irreducible restrictions. Therefore, in the interest of brevity, we use “branched cover” to refer to an irreducible branched cover, unless otherwise stated. We refrain from elaborating on branched covers which are not irreducible.

4.2 Permutation groups

We recall some terminology concerning permutation groups [21]. For $d \in \mathbb{N}$, the **symmetric group** S_d on d elements is the group of bijections from $[d]$ to $[d]$ under composition. Any subgroup $G \subset S_d$ of the symmetric group acts on the ordered set $\{1, 2, \dots, d\}$ by permuting its elements and is thus called a **permutation group**. A permutation group acts **transitively** if for every $i, j \in [d]$, there exists $g \in G$ such that $g(i) = j$. For now, we will assume that G acts transitively on $[d]$.

A **block** of G is a subset $B \subset [d]$ such that for every $g \in G$, either $gB = B$ or $gB \cap B = \emptyset$. The subsets \emptyset , $[d]$, and every singleton are blocks of every permutation group. If these trivial blocks are the only blocks, then G is **primitive** and otherwise it is **imprimitive**.

When G is imprimitive, we have a factorization $d = ab$ with $1 < a, b < d$ and there is a bijection $[a] \times [b] \leftrightarrow [d]$ such that G preserves the projection $[a] \times [b] \rightarrow [b]$. That is, the fibers $\{[a] \times \{i\} \mid i \in [b]\}$ are blocks of G , its action on this set of blocks gives a homomorphism $G \rightarrow S_b$

with transitive image, and the kernel acts transitively on each fiber $[a] \times \{i\}$. In particular, G is a subgroup of the wreath product $S_a \wr S_b = (S_a)^b \rtimes S_b$, where S_b acts on $(S_a)^b$ by permuting factors.

We observe a second characterization of imprimitive permutation groups G . Since G acts transitively, if $H \subset G$ is the stabilizer of a point $c \in [d]$, then H has index d in G and we may identify $[d]$ with the set G/H of cosets. If B is a nontrivial block of G containing c , then its stabilizer L is a proper subgroup of G that strictly contains H . Furthermore, using the map $G/H \rightarrow G/L$, we see that G is imprimitive if and only if the stabilizer of the point $eH \in G/H$ is not a maximal subgroup.

4.3 Monodromy groups and Galois groups

Let $\pi: X \rightarrow Z$ be a degree d branched cover so that the restriction $\pi^{-1}(U) \xrightarrow{\pi} U$ is a d -sheeted covering space. A **lift** of a continuous function $\gamma: Y \rightarrow U$ is a map $\tilde{\gamma}: Y \rightarrow X$ such that $\pi \circ \tilde{\gamma} = \gamma$. The **path lifting property** for a covering space says that for any path $\gamma: [0, 1] \rightarrow U$ and any lift \tilde{u}_0 of the point $u_0 = \gamma(0)$, there is a unique path $\tilde{\gamma}: [0, 1] \rightarrow X$ which lifts γ with the property that $\tilde{\gamma}(0) = \tilde{u}_0$ [22].

Since the cardinality of the fiber $\pi^{-1}(\gamma(0))$ is d , there are d paths $\{\tilde{\gamma}_i(t)\}_{i=1}^d$ lifting γ , giving a bijection m_γ from the fiber over $\gamma(0)$ to the fiber over $\gamma(1)$ defined by $m_\gamma(\tilde{\gamma}_i(0)) = \tilde{\gamma}_i(1)$. When $\gamma(0) = \gamma(1)$, we call γ a **(monodromy) loop based** at $\gamma(0)$. The set of all m_γ such that γ is a loop based at $u \in U$ forms a group $\mathcal{M}_{\pi,u}$ called the **monodromy group** of π based at u .

For any path γ in U , conjugation by m_γ gives an isomorphism $\mathcal{M}_{\pi,\gamma(0)} \cong \mathcal{M}_{\pi,\gamma(1)}$. Since X is irreducible, U is path-connected and so as a permutation group, the monodromy group is well-defined up to the relabelling of points in a fiber. We define the **monodromy group** of π , denoted \mathcal{M}_π , to be this group.

Lemma 4.3.1. *The monodromy group of a branched cover $X \xrightarrow{\pi} Z$ is transitive.*

Proof. Let $p, q \in \pi^{-1}(u)$ for some $u \in U$. The set $\pi^{-1}(U)$ is path-connected and so a path τ connecting p to q projects to a loop $\gamma = \pi \circ \tau$ with τ as a lift. Hence, $m_\gamma(p) = q$. \square

We define the **Galois group** G_π of π to be the Galois group of $K/\mathbb{C}(Z)$ where K is the Galois closure of $\mathbb{C}(X)/\mathbb{C}(Z)$. Harris [23] gave a modern proof of the following proposition, but this idea goes back at least to Hermite [24].

Proposition 4.3.2. [23, pg. 689] *The groups G_π and \mathcal{M}_π for a branched cover π are equal.*

4.4 Decomposable branched covers

A branched cover $\pi: X \rightarrow Z$ is **decomposable** if there is a dense open subset $V \subset Z$ over which π factors

$$\pi^{-1}(V) \xrightarrow{\varphi} Y \xrightarrow{\psi} V, \quad (4.1)$$

with φ and ψ both nontrivial branched covers. The fibers of φ over points of $\psi^{-1}(v)$ are blocks of the action of G_π on $\pi^{-1}(v)$, which implies that G_π is imprimitive. Pirola and Schlesinger [25] observed that decomposability of π is equivalent to imprimitivity of G_π . We give a proof, as we discuss the problem of computing a decomposition.

Proposition 4.4.1. *A branched cover is decomposable if and only if its Galois group is imprimitive.*

Proof. We need only to prove the reverse direction. As above, let $\mathbb{C}(Z)$, $\mathbb{C}(X)$, and K be the function fields of Z , X , and the Galois closure of $\mathbb{C}(X)/\mathbb{C}(Z)$, respectively, and let G_π be the Galois group of $K/\mathbb{C}(Z)$. Let H be the subgroup of G_π such that $\mathbb{C}(X) = K^H$, the fixed field of H . The set of Galois conjugates of $\mathbb{C}(X)$ forms the orbit G_π/H , and the number of conjugates is the degree of the branched cover $X \rightarrow Z$.

If G_π acts imprimitively, then the stabilizer L of a nontrivial block B containing $\mathbb{C}(X)$ is a proper subgroup properly containing H . Thus its fixed field $M = K^L$, which is the intersection of the conjugates of $\mathbb{C}(X)$ in the block B , is an intermediate field between $\mathbb{C}(Z)$ and $\mathbb{C}(X)$. For any variety Y' with function field M , there will be dense open subsets Y of Y' and V of Z such that (4.1) holds. \square

While imprimitivity is equivalent to decomposability, the proof does not address how to compute the variety Y of (4.1). One way is as follows. Replace Z and X by affine open subsets, if necessary, and let $y_1, \dots, y_m \in \mathbb{C}[X]$ be regular functions on X that generate M over $\mathbb{C}(Z)$. Let x_1, \dots, x_m be indeterminates and let $I \subset \mathbb{C}(Z)[x_1, \dots, x_m]$ be the kernel of the map $\mathbb{C}(Z)[x_1, \dots, x_m] \rightarrow \mathbb{C}(X)$ given by $x_i \mapsto y_i$. This is the zero-dimensional ideal of algebraic relations satisfied by y_1, \dots, y_m . Replacing Z by a dense affine open subset if necessary, we may choose generators g_1, \dots, g_r of I that lie in $\mathbb{C}[Z][x_1, \dots, x_m]$ —their coefficients are regular functions on Z . There is an open subset $V \subset Z$ such that the ideal I defines an irreducible variety $Y \subset V \times \mathbb{C}^m$ whose projection to V is a branched cover and whose function field is M . Restricting $X \rightarrow Z$ to V , we obtain the desired decomposition, with the map $X \rightarrow Y$ given by the functions y_1, \dots, y_m .

This does not address the practicality of computing Y , but it does indicate an approach. Given the subgroup L of G_π and a set of generators of $\mathbb{C}[X]$ over $\mathbb{C}[Z]$, if we apply the Reynolds averaging operator [26] for L to monomials in the generators, we obtain the desired generators y_1, \dots, y_m of M . One problem is that elements of G_π may not act on X , so their action on elements of $\mathbb{C}[X]$ may be hard to describe.

There is an exception to this. If $L \neq H$ normalizes H in G and $\pi: X \rightarrow Z$ is a covering space, then $\Gamma = L/H$ acts freely on X , preserving the fibers—it is a group of [deck transformations](#) of $X \rightarrow Z$ [27, Ch 13]. When Γ acts on the original branched cover, $Y = X/\Gamma$ is the desired space, and both Y and the map $X \rightarrow Y$ may be computed by applying the Reynolds operator for Γ to generators of $\mathbb{C}[X]$. The examples given in [28, Section 5] are of this form, and the authors use this approach to compute decompositions.

Example 4.4.2. Not all imprimitive groups have this property. Consider the wreath product $G = S_3 \wr S_3$, which acts imprimitively on the nine-element set $[3] \times [3]$. The stabilizer of the point $(3, 3)$ is the subgroup $H = ((S_3)^2 \times S_2) \rtimes S_2$, where $S_2 \subset S_3$ is the stabilizer of $\{3\}$. Then H is its own normalizer in G , as S_2 is its own normalizer in S_3 . \diamond

All imprimitive Galois groups in the Schubert calculus constructed in [29, Section 3] and in [30] have stabilizer H equal to its normalizer. For these, the decomposition of the branched cover follows from a deep structural understanding of the corresponding Schubert problem. There remain many Schubert problems whose Galois group is expected to be imprimitive, yet a decomposition (4.1) of the corresponding branched cover is unknown.

4.5 Real branched covers

The nonreal solutions of any univariate polynomial $f \in \mathbb{R}[x]$ come in complex conjugate pairs. Similarly, for a multivariate polynomial system $F = (f_1, \dots, f_k) \subset \mathbb{R}[x]$, any point $z \in \mathbb{C}^n$ satisfies $F(z) = 0$ if and only if its complex conjugate \bar{z} satisfies $F(\bar{z}) = 0$.

When a branched cover $\pi: X \rightarrow Z$ with $Z \subset \mathbb{C}^m$ has the property that for any $z \in Z \cap \mathbb{R}^m$ the ideal $\mathcal{I}(\varphi^{-1}(z))$ can be generated by real polynomials, we say π is a **real branched cover**. The set of real regular values of a real branched cover is possibly disconnected in \mathbb{R}^m , and we call these connected components **discriminant chambers**.

Lemma 4.5.1. *If $z, z' \in Z \cap \mathbb{R}^m$ are in the same discriminant chamber, then the number of real points in $\pi^{-1}(z)$ is equal to the number of real points in $\pi^{-1}(z')$.*

Proof. Let z, z' be in the same discriminant chamber D_z and let $\gamma: [0, 1] \rightarrow Z \cap D_z$ be a path from z to z' . For any point $\gamma(t^*)$ for $t^* \in [0, 1]$ the fiber $\pi^{-1}(t^*)$ consists of $\deg(\pi)$ distinct points. On the other hand, since nonreal points in the fiber come in complex conjugate pairs, the number of real points in a fiber over $\gamma([0, 1])$ changes only if either two real points come together and become complex or two complex points come together and become real. However, this cannot happen since points in each fiber over γ are distinct. \square

Example 4.5.2. Let

$$f = 4(\phi^2 x^2 - y^2)(\phi^2 y^2 - z^2)(\phi^2 z^2 - x^2) - (1 + 2\phi)(x^2 + y^2 + z^2 - 1^2)^2 \in \mathbb{C}[x, y, z],$$

where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. The surface $\mathcal{V}(f)$ is known as the **Barth sextic**. The projection

$$\pi: \mathbb{C}_{x,y,z}^3 \rightarrow \mathbb{C}_{x,y}^2$$

is a branched cover of degree 4.

The branch locus B of π is displayed in Figure 4.3 along with labels indicating the number of real points in any fiber of the corresponding discriminant chamber. Over \mathbb{Q} , B decomposes into two lines (purple and green) and two sextics (blue and red). Over \mathbb{R} , the blue sextic curve decomposes into the union of a conic and four lines. The red curve is irreducible over \mathbb{R} . The boxed region in Figure 4.3 contains a small discriminant chamber whose fibers have two real points. An enlarged depiction of this chamber is displayed in Figure 4.4. \diamond




Figure 4.2: The Barth sextic.




Figure 4.3: The discriminant of the projection of the Barth sextic onto $\mathbb{C}^2_{x,y}$ with the number of real points in the fiber of any point in each discriminant chamber indicated.




Figure 4.4: One of the small discriminant chambers not easily noticeable in Figure 4.3.

5. NEWTON POLYTOPES, SUPPORT, TROPICAL GEOMETRY, AND SPARSE POLYNOMIAL SYSTEMS

We introduce Newton polytopes, sparse polynomial systems, and tropical geometry. These connect ideas from Sections 2, 3, and 4. Material in Section 5.2 appears in the article [1] by the author*.

5.1 Newton polytopes

Let $\mathbb{C}^\times = \mathbb{C} \setminus \{0\}$ be the multiplicative group of nonzero complex numbers and $(\mathbb{C}^\times)^n$ be the n -dimensional complex torus. For each $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}^n$, the **(Laurent) monomial** with exponent α ,

$$x^\alpha = x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_n^{\alpha_n},$$

is a character (multiplicative map) $x^\alpha: (\mathbb{C}^\times)^n \rightarrow \mathbb{C}^\times$. Any finite linear combination

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha, \quad c_\alpha \in \mathbb{C},$$

of monomials is a **(Laurent) polynomial** which also defines a function $f: (\mathbb{C}^\times)^n \rightarrow \mathbb{C}$. When $c_\alpha \in \mathbb{C}^\times$ for all $\alpha \in \mathcal{A}$, we say that \mathcal{A} is the **support** of f and write $\text{supp}(f) = \mathcal{A}$. Otherwise, we say that f is **supported on** \mathcal{A} . We denote the vector space of all polynomials supported on \mathcal{A} by $\mathbb{C}^{\mathcal{A}}$. Consistent with the notation for polytopes, for any $\omega \in \mathbb{R}^n$ we set,

$$f_\omega = \sum_{\alpha \in \mathcal{A}_\omega} c_\alpha x^\alpha.$$

If \mathcal{A} is the support of f , then the support of $x^\beta f$ is $\beta + \mathcal{A}$, the translation of \mathcal{A} by β . As a monomial x^β for $\beta \in \mathbb{Z}^n$ is invertible on $(\mathbb{C}^\times)^n$, the polynomials f and $x^\beta f$ have the same sets of zeros in \mathbb{C}^\times . By translating the support of a polynomial by integer vectors we may assume that $\mathbf{0}$ is in the affine \mathbb{Z} -span of \mathcal{A} without changing any assertions about the zeros of f in $(\mathbb{C}^\times)^n$, thus we define $\mathbb{Z}\mathcal{A}$ to be the lattice generated by differences $\alpha - \beta$ for $\alpha, \beta \in \mathcal{A}$. For similar reasons, the results from Section 3 extend to this setting by shifting $\text{supp}(f)$ to the positive orthant so that f is polynomial.

The **Newton polytope** of f (or of $\mathcal{V}(f)$) is

$$\text{New}(f) = \text{New}(\mathcal{V}(f)) = \text{conv}(\text{supp}(f)).$$

We say f has **dense support** in $\text{New}(f)$ if $\text{supp}(f) = \mathcal{L}(\text{New}(f))$, the set of lattice points in $\text{New}(f)$. The Newton polytope of a polynomial and its support both encode a considerable amount of information about the polynomial and its zero set. Moreover, these combinatorial objects behave well under certain algebro-geometric transformations on polynomials and varieties.

*Reprinted with permission from T. Brysiewicz, “Numerical Software to Compute Newton polytopes and Tropical Membership,” *Mathematics in Computer Science*, 2020. Copyright 2020 by Springer Nature.

5.1.1 Basic observations about Newton polytopes

Let $f \in \mathbb{C}[x]$. Then the following observations are immediate from our definitions.

- (1) $\text{New}(\tilde{f}) = \widetilde{\text{New}(f)}$ where $\tilde{}$ denotes homogenization.
- (2) f is homogeneous if and only if $\text{New}(f)$ is homogeneous.
- (3) $\deg(f) = \deg(\text{New}(f))$.
- (4) $\text{New}(f)$ is an integral polytope.

For any $f, g \in \mathbb{C}[x]$, the Newton polytope of $f \cdot g$ is $\text{New}(f) + \text{New}(g)$. Indeed, Lemma 2.3.1 implies that the vertices of $\text{New}(f) + \text{New}(g)$ are uniquely represented as $\alpha' + \beta'$ for some $\alpha' \in \text{vert}(\text{New}(f))$ and $\beta' \in \text{vert}(\text{New}(g))$. Thus, the only term of the sum

$$f \cdot g = \sum_{\substack{\alpha \in \text{supp}(f) \\ \beta \in \text{supp}(g)}} c_\alpha c_\beta x^\alpha x^\beta$$

which has exponent $\alpha' + \beta'$ is $c_{\alpha'} c_{\beta'} x^{\alpha' + \beta'}$ which is in the support of $f \cdot g$ because $c_{\alpha'} \cdot c_{\beta'} \neq 0$.

Supports (and Newton polytopes) respect permutations of variables. For any permutation $\sigma \in S_n$, the support of

$$\sigma(f) = \sum_{\alpha \in \mathcal{A}} c_\alpha x_1^{\alpha_{\sigma(1)}} x_2^{\alpha_{\sigma(2)}} \cdots x_n^{\alpha_{\sigma(n)}}$$

is the set

$$\sigma(\mathcal{A}) = \{\sigma(\alpha) \mid \alpha \in \mathcal{A}\} = \{(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)}) \mid \alpha \in \mathcal{A}\}.$$

Consequently, $\text{New}(\sigma(f)) = \sigma(\text{New}(f)) = \text{conv}(\sigma(\mathcal{A}))$. Hyperplanes containing $\text{New}(f)$ correspond to scalings of the variables x_1, \dots, x_n which do not alter the variety $\mathcal{V}(f)$.

Lemma 5.1.1. *Let $f \in \mathbb{C}[x]$ be a polynomial with support \mathcal{A} and let $\omega \in \mathbb{R}^n$. Then \mathcal{A} is contained in the hyperplane $\langle \alpha, \omega \rangle - h_{\mathcal{A}}(\omega) = 0$ if and only if $\mathcal{V}(f) = \mathcal{V}(f(t^{\omega_1} x_1, \dots, t^{\omega_n} x_n))$ for all $t \in \mathbb{C}^\times$.*

Proof. The equality $\mathcal{V}(f) = \mathcal{V}(f(t^{\omega_1} x_1, \dots, t^{\omega_n} x_n))$ holds for all $t \in \mathbb{C}^\times$ if and only if for all $a \in \mathcal{V}(f)$,

$$\begin{aligned} 0 &= f(t^{\omega_1} a_1, \dots, t^{\omega_n} a_n) = \sum_{\alpha \in \mathcal{A}} c_\alpha t^{\langle \alpha, \omega \rangle} a^\alpha \\ &= \sum_{k=-h_{\mathcal{A}}(\omega)}^{h_{\mathcal{A}}(\omega)} \left(\sum_{\substack{\alpha \in \mathcal{A} \\ \langle \alpha, \omega \rangle = k}} t^k c_\alpha a^\alpha \right) \\ &= \sum_{k=-h_{\mathcal{A}}(\omega)}^{h_{\mathcal{A}}(\omega)} t^k g_k(a). \end{aligned} \tag{5.1}$$

The right-most-side is a polynomial in t and thus $g_k(a) = 0$ for all $k = -h_{\mathcal{A}}(\omega), \dots, h_{\mathcal{A}}(\omega)$ and all $a \in \mathcal{V}(f)$. However, this means that $\mathcal{V}(f) \subset \mathcal{V}(g_k)$ for all k . Since f is not identically zero, at least

one g_k is not. Suppose $g_j \neq 0$ for some j . Then containment of hypersurfaces implies $\deg(g_j) \geq \deg(f)$ and since $t^j g_j$ is a summand of (5.1), these degrees must be the same. Containment of hypersurfaces also implies that $g_j(x) = r(x) \cdot f(x)$ for some $r \in \mathbb{C}[x]$, but since the degrees of g_j and f are equal, r must be a constant implying $\mathcal{V}(f) = \mathcal{V}(g_j)$. Consequently, every other summand of (5.1) must be zero, proving that \mathcal{A} is contained in the hyperplane $\langle \alpha, \omega \rangle - h_{\mathcal{A}}(\omega) = 0$.

The converse is true since if $\langle \alpha, \omega \rangle = h_{\mathcal{A}}(\omega)$ for all $\alpha \in \mathcal{A}$, then $f(t^{\omega_1} x_1, \dots, t^{\omega_n} x_n) = t^{h_{\mathcal{A}}(\omega)} f(x)$, and thus cuts out the same variety as f for any $t \in \mathbb{C}^\times$. \square

Remark 5.1.2. Fix $r \in \mathbb{N}$ and $k = (k_1, \dots, k_r) \in \mathbb{N}^r$ and consider the grouping of variables $\left\{ \{x_{i,j}\}_{i=1}^{k_j} \right\}_{j=1}^r$. By definition of projective space, the zero set of a polynomial

$$f = \sum_{\alpha=(\alpha^{(1)}, \dots, \alpha^{(r)}) \in \mathcal{A}} c_\alpha x_{i,1}^{\alpha^{(1)}} \cdots x_{i,r}^{\alpha^{(r)}} \in \mathbb{C}[x_{i,j}]$$

with support \mathcal{A} is well-defined subvariety of $\mathbb{P}^{k_1} \times \dots \times \mathbb{P}^{k_r}$ if and only if it is invariant under scaling any of the variable groups: for each $j \in [r]$ and $t \in \mathbb{C}^\times$, the polynomial f is invariant under the action which multiplies each variable in the group $\{x_{i,j}\}_{i=1}^{k_j}$ by t . By Lemma 5.1.1, this is equivalent to the condition that for all $\alpha \in \mathcal{A}$ and $j \in [r]$, there exists d_j such that $|\alpha^{(j)}| = d_j$. The vector $d = (d_1, \dots, d_r)$ is called the **multidegree** of $\mathcal{V}(f)$.

Lemma 5.1.1 has strong implications when considering invariants. Fix some support $\mathcal{A} \subset \mathbb{Z}^n$ and suppose that $\mathcal{F} \subset \mathbb{C}[c_\alpha]_{\alpha \in \mathcal{A}}$ is a polynomial in the coefficient space $\mathbb{C}^{\mathcal{A}}$ of all polynomials

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha \in \mathbb{C}[x]$$

supported on \mathcal{A} . Observe that an action of a group $G \curvearrowright \mathbb{C}^n$ naturally induces an action $G \curvearrowright \mathbb{C}^{\mathcal{A}}$ on the coefficient space. If for all $f \in \mathcal{V}(\mathcal{F})$ and all $\sigma \in G$, we have that $\sigma \cdot f \in \mathcal{V}(\mathcal{F})$, then we say that \mathcal{F} is **invariant** under the action of G .

Proposition 5.1.3. *Suppose that $\mathcal{F} \in \mathbb{C}[c_\alpha]_{\alpha \in \mathcal{A}}$ is a homogeneous polynomial of degree D with variables in the coefficient space $\mathbb{C}^{\mathcal{A}}$ of all polynomials*

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha \in \mathbb{C}[x],$$

supported on $\mathcal{A} \subset \mathbb{Z}^n$. Suppose further that $|\alpha| = d$ for all $\alpha \in \mathcal{A}$. Let A be the $n \times |\mathcal{A}|$ matrix whose columns are points in \mathcal{A} and whose rows are $\{\omega_{x_1}, \dots, \omega_{x_n}\}$.

- (1) *If \mathcal{F} is invariant under the scaling $x_i \mapsto tx_i$ for some $i \in [n]$ and all $t \in \mathbb{C}^\times$, then $\mathcal{F}_{\omega_{x_i}} = F$.*
- (2) *Suppose \mathcal{F} is invariant under all scalings and permutations of the variables x_i and that $\mathbb{R}A$ is n dimensional. Then $p \in \text{New}(\mathcal{F})$ solves the linear equation*

$$\begin{pmatrix} A \\ \mathbf{1} \end{pmatrix} p = \left(\frac{dD}{n}, \dots, \frac{dD}{n}, D \right)^T.$$

In particular, $\text{New}(\mathcal{F}) \subset \mathbb{R}_p^{|\mathcal{A}|}$ is contained in an affine linear space of codimension n .

Proof. Given $f = \{c_\alpha\}_{\alpha \in \mathcal{A}}$, the action of $t \mapsto tx_1$ on \mathbb{C}^n induces the action

$$f_t = f(tx_1, x_2, \dots, x_n) = \{t^{\langle e_1, \alpha \rangle} c_\alpha\} = \{t^{(\omega_{x_1})_\alpha} c_\alpha\}$$

on the coefficients of f and thus the variables of \mathcal{F} . If \mathcal{F} is invariant under this action, then $f_t \in \mathcal{V}(\mathcal{F})$ if and only if $f \in \mathcal{V}(\mathcal{F})$ for all $t \in \mathbb{C}^\times$. Hence by Lemma 5.1.1 we have that $\mathcal{F}_{\omega_{x_1}} = \mathcal{F}$. The same argument applies for scaling any other variable.

If \mathcal{F} is invariant under scaling any of the variables x_1, \dots, x_n , then P is contained in the intersection $\bigcap_{i=1}^n H_i$ where

$$H_i = \{p \in \mathbb{R}_p^{|\mathcal{A}|} \mid \langle p, \omega_{x_i} \rangle = h_P(\omega_{x_i})\}$$

by part (1). Since \mathcal{F} is also invariant under the symmetric group S_n , the value of the support function $h = h_P(\omega_{x_i})$ does not depend on i . Since \mathcal{F} is homogeneous, P is also contained in the affine hyperplane

$$H_{\text{deg}} = \{p \in \mathbb{R}_p^{|\mathcal{A}|} \mid \langle p, \mathbf{1} \rangle = D\}.$$

Thus, the set

$$H = H_{\text{deg}} \cap \left(\bigcap_{i=1}^n H_i \right)$$

is the solution set of the matrix equation,

$$\begin{pmatrix} A \\ \mathbf{1} \end{pmatrix} p = (h, h, \dots, h, D)^T.$$

Note that since $|\alpha| = d$ for all $\alpha \in \mathcal{A}$, we have $(1, 1, \dots, 1, -d) \begin{pmatrix} A \\ \mathbf{1} \end{pmatrix} = \mathbf{0}$. Therefore, $hn - dD = 0$ and so $h = \frac{dD}{n}$. \square

5.1.2 Integer linear algebra and coordinate changes

Supports of polynomials do not maintain their structure under generic linear changes of coordinates: for a generic linear map $\phi: \mathbb{C}^n \rightarrow \mathbb{C}^n$, the composition $f(\phi(z))$ has dense support $\text{deg}(f)\Delta_n$. Supports do, however, respect partial evaluation in the following sense. Let $\pi_I: \mathbb{Z}^n \rightarrow \mathbb{Z}^{|I|}$ be the projection onto the coordinates indexed by $I \subset [n]$.

Lemma 5.1.4. *Let $f \in \mathbb{C}[x]$ be a polynomial with support \mathcal{A} and let $a_{k+1}, \dots, a_n \in \mathbb{C}^\times$ be general. Then the support of $f(x_1, \dots, x_k, a_{k+1}, \dots, a_n)$ is the projection $\pi_{[k]}(\mathcal{A})$.*

Supports of polynomials transform naturally under monomial changes of coordinates. Identifying the set $\text{Hom}((\mathbb{C}^\times)^n, \mathbb{C}^\times)$ of characters on $(\mathbb{C}^\times)^n$ with the free abelian group \mathbb{Z}^n , a homomorphism $\Phi: (\mathbb{C}^\times)^m \rightarrow (\mathbb{C}^\times)^k$ is determined by k characters of $(\mathbb{C}^\times)^m$, equivalently by a homomorphism (linear map) $\varphi: \mathbb{Z}^k \rightarrow \mathbb{Z}^m$ of free abelian groups. Note that φ is also the map pulling a character of $(\mathbb{C}^\times)^k$ back along Φ . In particular, an invertible map $\Phi: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$ (a monomial change of coordinates) pulls back to an invertible map $\varphi: \mathbb{Z}^n \rightarrow \mathbb{Z}^n$, identifying $\text{GL}(n, \mathbb{Z})$ with the group of possible monomial coordinate changes. We will write $\Phi = \varphi^*$ and $\varphi = \Phi^*$ for these, not to be confused with the notation for the homomorphism of coordinate rings induced by a regular map of varieties. If $\Phi(x) = (x^{\alpha_1}, \dots, x^{\alpha_n})$ where the integer span of $\{\alpha_1, \dots, \alpha_n\}$ is \mathbb{Z}^n , then the

map $\varphi = \Phi^*: \mathbb{Z}^n \xrightarrow{\sim} \mathbb{Z}^n$ sends the i -th standard basis vector e_i to α_i and is represented by the invertible matrix A whose i -th column is α_i .

Suppose that f is a polynomial on $(\mathbb{C}^\times)^n$ with support \mathcal{A} . Given a homomorphism $\Phi: (\mathbb{C}^\times)^m \rightarrow (\mathbb{C}^\times)^n$, the composition $f(\Phi(z))$ for $z \in (\mathbb{C}^\times)^m$ is a polynomial supported on $\varphi(\mathcal{A})$, where the coefficient of z^β is the sum of coefficients of x^α for $\alpha \in \varphi^{-1}(\beta) \cap \mathcal{A}$. For generic choices of coefficients of x^α , this sum is nonzero and so $f(\Phi(z))$ has support $\varphi(\mathcal{A})$.

5.1.3 Smith normal form

Let $\mathcal{A} = \{0, \alpha_1, \dots, \alpha_m\} \subset \mathbb{Z}^n$ be a collection of integer vectors. The sublattice $\mathbb{Z}\mathcal{A} \subset \mathbb{Z}^n$ that it generates is the image of a \mathbb{Z} -linear map $\mathbb{Z}^m \rightarrow \mathbb{Z}^n$ and is represented by a $n \times m$ integer matrix A whose columns are the vectors α_i . Suppose that $\mathbb{Z}\mathcal{A}$ has rank k . The **Smith normal form** of A is a factorization into integer matrices

$$A = PDQ, \quad (5.2)$$

where $P \in \text{GL}(n, \mathbb{Z})$ and $Q \in \text{GL}(m, \mathbb{Z})$ are invertible, and D is the rectangular matrix whose only nonzero entries are d_1, \dots, d_k along the diagonal of its principal $k \times k$ submatrix. These are the **invariant factors** of A and they satisfy $d_1 | d_2 | d_3 | \dots | d_k$. The sublattice $\mathbb{Z}\mathcal{A} \subset \mathbb{Z}^n$ has a basis given by the columns of the matrix PD . If we apply the coordinate change P^{-1} to \mathbb{Z}^n , then $\mathbb{Z}\mathcal{A}$ becomes the subset of the coordinate space $\mathbb{Z}^k \oplus \mathbf{0}^{n-k}$ given by $d_1\mathbb{Z} \oplus d_2\mathbb{Z} \oplus \dots \oplus d_k\mathbb{Z} \oplus \mathbf{0}^{n-k}$.

The Smith normal form is also useful in solving binomial equations over $(\mathbb{C}^\times)^n$. Fix a collection $F \subset \mathbb{C}[x]$ of binomials

$$a_1x^{\alpha_1} = b_1x^{\beta_1} \quad a_2x^{\alpha_2} = b_2x^{\beta_2} \quad \dots \quad a_nx^{\alpha_n} = b_nx^{\beta_n}$$

with $a_i, b_i \in \mathbb{C}^\times$ for $i = 1, \dots, n$. Recall that we can scale the equations and translate their support so that $a_i = 1$ and $\beta_i = 0$ for all $i = 1, \dots, n$. We now assume our system F is of the form

$$x^{\alpha_1} = b_1 \quad x^{\alpha_2} = b_2 \quad \dots \quad x^{\alpha_n} = b_n. \quad (5.3)$$

It is useful to use matrices as exponents. For example, we encode x^{α_1} as $(x_1, \dots, x_n)^{((\alpha_1)_1, \dots, (\alpha_1)_n)}$. Letting A be the matrix whose columns are $\alpha_1, \dots, \alpha_n$, we write $x^A = x^{(\alpha_1, \dots, \alpha_n)} = (x^{\alpha_1}, \dots, x^{\alpha_n})$ so that (5.3) is written as $x^A = b$.

Assume for simplicity that \mathcal{A} spans \mathbb{R}^n so that d_n of a Smith normal form $A = PDQ$ is nonzero. Then

$$(x^A)^{Q^{-1}} = b^{Q^{-1}}$$

and setting $z^{P^{-1}} = x$ gives

$$z^{P^{-1}AQ^{-1}} = z^D = b^{Q^{-1}}. \quad (5.4)$$

whose solutions are clearly the set $\mathcal{Z} = \{z \mid z_i \text{ is a } d_i\text{-th root of } b_i\}$ of $\prod_{i=1}^n d_i$ points. Taking $x = z^P$ expresses these solutions in terms of x .

5.1.4 Centroids and trace curves

Given an affine variety $X \subset \mathbb{C}^n$ and a generic linear space L of complementary dimension to X , the intersection $X \cap L$ is finite and consists of $\deg(X)$ points. The **centroid** of $X \cap L$, denoted $\mu(X \cap L)$ is the coordinate-wise average of those points. A family of linear spaces $\{L_t\}_{t \in \mathbb{C}}$ is a **pencil** if there exists a vector $v \in \mathbb{C}^n$ such that $L_t = t \cdot v + L_0$ for all $t \in \mathbb{C}$. The following lemma is the basis for the numerical algorithm known as the **trace test** (see Section 6.5.1).

Lemma 5.1.5. *Let $X \subset \mathbb{C}^n$ be an irreducible affine variety and let L_t be a general pencil of linear spaces of complementary dimension. The Zariski closure of the union*

$$\mu(X \cap L_t) = \bigcup_{t \in \mathbb{C}} \mu(X \cap L_t)$$

is an affine line.

Proof. Let $X \subset \mathbb{C}^n$ be an irreducible affine variety of dimension m . Observe that if L is a linear space of complementary dimension, then $\pi(\mu(X \cap L)) = \mu(\pi(X) \cap \pi(L))$ where $\pi: \mathbb{C}^n \rightarrow \mathbb{C}^{n-1}$ is any projection such that $\dim(\pi(L)) = \dim(L) - 1$. Projecting this way $n - m - 1$ times produces $\pi': \mathbb{C}^n \rightarrow \mathbb{C}^{m+1}$ so that $\dim(\pi'(L)) = 1$. Thus, $\pi'(X)$ is a hypersurface in \mathbb{C}^{m+1} and $\mu(X \cap L) \in \pi'^{-1}(\mu(\pi'(X) \cap \pi'(L)))$. Let v_1, \dots, v_{n-m} span L and define $\pi_i: \mathbb{C}^n \rightarrow \mathbb{C}^{m+1}$ to be the projection such that $\dim(\pi_i(v_j)) = 0$ whenever $i \neq j$. Then the intersection

$$\bigcap_{i=1}^{n-m} \pi_i^{-1}(\mu(\pi_i(X) \cap \pi_i(L)))$$

is the point $\mu(X \cap L)$. Thus, it is enough to prove the statement for when X is a hypersurface.

Let $X \subset \mathbb{C}^n$ be a hypersurface, let L_t be a general pencil of lines, and let $P = \bigcup_{t \in \mathbb{C}} L_t$. Consider $X' = X \cap P$. By Lemma 3.8.10, X' is a curve and so it is enough to prove the statement for plane curves.

Suppose $\mathcal{V}(f) = X \subset \mathbb{C}^2$ is a plane curve of degree d , and L_t a general pencil of lines. After an action by rotation, we may assume that L_t is the family $\mathcal{V}(x - t)$. This rotation is a generic linear change of coordinates because the family L_t is general and so the support of f must be $d\Delta_n$. Since scaling does not change the zero set, we assume that the coefficient of y^d is one. Then $X \cap L_t = X \cap \mathcal{V}(x - t)$ has points $\{(t, y_i(t))\}_{i=1}^d$ where $y_i(t)$ are the zeros of

$$f(t, y) = \prod_{i=1}^d (y - y_i(t)) = y^d - (y_1(t) + \dots + y_d(t))y^{d-1} + \dots$$

for some rational functions $y_i(t)$. On the other hand, the coefficient of y^{d-1} in $f \in \mathbb{C}[x][y]$ is $c_{(1,d-1)}x + c_{(0,d-1)}$ and so $-(y_1(t) + \dots + y_d(t)) = c_{(1,d-1)}x + c_{(0,d-1)}$. Since the y -coordinate of $\mu(X \cap L_t)$ is $\frac{1}{d}(y_1(t) + \dots + y_d(t))$, the points satisfying $-dy = c_{(1,d-1)}x + c_{(0,d-1)}$ are the points which are centroids of this family. In other words, the centroids are on the graph of the function

$$y = -\frac{1}{d}(c_{(1,d-1)}x + c_{(0,d-1)}). \quad (5.5)$$

□

The line of centroids guaranteed by Lemma 5.1.5 is called the **trace line** of X with respect to L_t .

Example 5.1.6. Let

$$f = 2 - 4x + x^3 + (-2 - 2x)y + (3 - x)y^2 + y^3$$

and let $L_t = \mathcal{V}(x - t)$ so that (5.5) computes the trace line of $\mathcal{V}(f)$ to be $\mathcal{V}(y - \frac{1}{3}x + 1)$. The cubic $\mathcal{V}(f)$ and its trace line are depicted in Figure 5.1. Notice that even though many lines L_t do not intersect $\mathcal{V}(f)$ in three real points, the centroids are still real. This is because the points $\mathcal{V}(f) \cap L_t$ must appear in complex conjugates and so their imaginary parts will cancel in the average. ◇




Figure 5.1: A plane cubic $\mathcal{V}(f)$ (blue), the trace line $\mu(\mathcal{V}(f, x - t))$ (red), and the specific centroids $\mu(\mathcal{V}(f, x + 3)), \mu(\mathcal{V}(f, x - 1)), \mu(\mathcal{V}(f, x - 2))$.

When the Newton polytope of a plane curve X of degree d is smaller than $d\Delta_2$ the family of lines $\mathcal{V}(x - t)$ is not generic with respect to X . Therefore, Equation 5.5 does not compute the curve of centroids. In particular, the closure of these centroids may not be a line. The following result gives a formula for the curve of centroids when the family $\mathcal{V}(x - t)$ is not generic with respect to X .

Lemma 5.1.7. *Suppose*

$$f = \sum_{(i,j) \in \mathcal{A}} c_{i,j} x^i y^j \in \mathbb{C}[x, y]$$

for $\mathcal{A} \subset \mathbb{Z}_{\geq 0}^2$ and $L_t = \mathcal{V}(x - t)$. Then

$$\overline{\bigcup_{t \in \mathbb{C}} \mu(\mathcal{V}(f) \cap L_t)} = \mathcal{V} \left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)-1} x^i + \deg_y(f) y \left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)} x^i \right) \right),$$

where $\deg_x(f) = \max_{(i,j) \in \mathcal{A}}(i)$ and $\deg_y(f) = \max_{(i,j) \in \mathcal{A}}(j)$.

Proof. As with the proof of Lemma 5.1.5, we take

$$f(t, y) = \left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)} t^i \right) y^d + \left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)-1} t^i \right) y^{\deg_y(f)-1} + \dots$$

and writing $f(t, y)$ as a monic polynomial tells us that

$$-(y_1(t) + \dots + y_{\deg_y(f)}(t)) = \frac{\left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)-1} t^i \right)}{\left(\sum_{i=0}^{\deg_x(f)} c_{i, \deg_y(f)} t^i \right)}.$$

Since $x = t$ and the y -coordinate of $\mu(X \cap L_t)$ is $\frac{1}{\deg_y(f)}(y_1(t) + \cdots + y_{\deg_y(f)}(t))$ we write this as

$$-\deg_y(f)y = \frac{\left(\sum_{i=0}^{\deg_x(f)} c_{i,\deg_y(f)-1}x^i\right)}{\left(\sum_{i=0}^{\deg_x(f)} c_{i,\deg_y(f)}x^i\right)},$$

and clearing denominators gives the result. \square

When the family L_t is not general as in Lemma 5.1.7, we define the **trace curve** of X with respect to L_t to be the closure of the set of centroids of $X \cap L_t$ for $t \in \mathbb{C}$.

Example 5.1.8. Consider the quartic curve

$$f = 1 - x + x^2 + (5 + x - 3x^2)y + (-3 + 3x - x^2)y^2$$

in \mathbb{C}^2 whose Newton polytope, support, and coefficients are depicted in Figure 5.2.




Figure 5.2: Left: The Newton polytope, support, and coefficients of f . Right: The curve $\mathcal{V}(f)$ (blue), the trace curve of $\mathcal{V}(f)$ with respect to $\mathcal{V}(x - t)$ (red), and five lines in the family L_t along with the centroids of their intersections with $\mathcal{V}(f)$.

The equation of the trace curve of $\mathcal{V}(f)$ with respect to the nongeneric family of lines $\mathcal{V}(x - t)$ is

$$g = (5 + x - 3x^2) + 2y(-3 + 3x - x^2).$$

Lemma 5.1.7 essentially states that the equation g can be read off from the coefficients of the top two rows of the polytope $\text{New}(f)$. \diamond

If $X = X_1 \cup \cdots \cup X_r \subset \mathbb{C}^n$ is a reducible affine variety and L is a generic linear space of complementary dimension, then $\mu(X \cap L) = \frac{1}{\deg(X)} \sum_{i=1}^r \mu(X_i \cap L) \cdot \deg(X_i)$ and so we have the following corollary.

Corollary 5.1.9. *Let $X \subset \mathbb{C}^n$ be an affine variety which is possibly reducible and let L_t be a general pencil of linear spaces of complementary dimension. The union of the centroids of the intersections $X \cap L_t$ is an affine line.*

5.2 Tropical geometry

Newton polytopes are intimately related to tropical geometry. We only begin to touch on the topic here and encourage the reader to reference [31] for a more extensive treatment.

The tropicalization of a variety depends on the choice of a valuation ν on the base field involved (in our case \mathbb{C}). Relevant to this document is the trivial valuation: $\nu(c) = 0$ for all $c \in \mathbb{C}^\times$. With this valuation, the **tropicalization** of a polynomial

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha, \quad \mathcal{A} = \text{supp}(f)$$

is the map

$$\begin{aligned} \text{trop}(f): \mathbb{R}^n &\rightarrow \mathbb{R} \\ \omega &\mapsto \max_{\alpha \in \mathcal{A}} \langle \alpha, \omega \rangle \end{aligned} \quad (5.6)$$

and the **tropicalization** of the hypersurface $\mathcal{V}(f)$ is

$$\text{trop}(\mathcal{V}(f)) = \{\omega \in \mathbb{R}^n \mid \text{the maximum in } \text{trop}(f)(\omega) \text{ is attained at least twice}\}. \quad (5.7)$$

By (5.6), $\text{trop}(f)$ is the same function as $h_{\text{New}(f)}$ and by (5.7), the tropicalization of $\mathcal{V}(f)$ is the codimension 1 part of the normal fan of the Newton polytope of f , namely $\mathcal{N}^{(1)}(\text{New}(f))$ (see Section 2.1).

Let $P = \text{New}(f)$, and fix a monomial change of coordinates $\Phi: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$ with $\varphi = \Phi^*$ so that we have $Q = \varphi(P) = \text{New}(f \circ \Phi)$. The map φ induces a map in the opposite direction on functionals $\{\alpha \mapsto \langle \alpha, \omega \rangle \mid \omega \in \mathbb{R}^n\}$. Consequently, ω is an element of $\text{trop}(\mathcal{V}(f))$ if and only if $\varphi^{-1}(\omega) \in \text{trop}(\mathcal{V}(f \circ \Phi))$ and so

$$\varphi^{-1}(\text{trop}(\mathcal{V}(f))) = \text{trop}(\mathcal{V}(f \circ \Phi)),$$

or equivalently,

$$\text{trop}(\mathcal{V}(f)) = \varphi(\text{trop}(\mathcal{V}(f \circ \Phi))). \quad (5.8)$$

The **tropicalization** of $\mathcal{V}(I)$ for some ideal $I \subseteq \mathbb{C}[x_1, \dots, x_n]$ is the intersection

$$\text{trop}(\mathcal{V}(I)) = \bigcap_{f \in I} \text{trop}(\mathcal{V}(f)).$$

Hept and Theobald in [4], motivated by the results of Bieri and Groves in [32], investigated how to write $\text{trop}(\mathcal{V}(I))$ as an intersection of finitely many tropical hypersurfaces coming from projections. The following is a consequence of the proof of Theorem 1.1 in [4].

Theorem 5.2.1. *If $I \subseteq \mathbb{C}[x]$ is an m -dimensional prime ideal, and $\{\pi_i: \mathbb{R}^n \rightarrow \mathbb{R}^{m+1}\}_{i=0}^{n-m}$ are generic projections,*

$$\text{trop}(\mathcal{V}(I)) = \bigcap_{i=0}^{n-m} \pi_i^{-1}(\pi_i(\text{trop}(\mathcal{V}(I))))$$

where each $\pi_i^{-1}(\pi_i(\text{trop}(\mathcal{V}(I))))$ is a tropical hypersurface.

Coordinate projections are not always generic and it is possible that only the proper containment

$$\bigcap_{\substack{J \subseteq [n] \\ \text{codim}(\pi_J(\mathcal{V}(I)))=1}} \pi_J^{-1}(\pi_J(\mathcal{V}(I))) \subsetneq \text{trop}(\mathcal{V}(I))$$

holds where π_J is the projection onto the coordinates indexed by $J \subset [n]$.

Remark 5.2.2. The notion of genericity involved in Theorem 5.2.1 comes from that of a geometrically regular projection. Let Y be a union of m -dimensional linear subsets of \mathbb{R}^n . A projection $\pi: \mathbb{R}^n \rightarrow \mathbb{R}^{m+1}$ is **geometrically regular** with respect to $Y \subset \mathbb{R}^n$ if the image of k -dimensional linear subspaces of Y remain k -dimensional and π respects containments: $\pi(Y_1) \subset \pi(Y_2) \implies Y_1 \subset Y_2$. These properties form an open dense subset within the set of projections and taking π_1, \dots, π_{n-m} distinct such projections gives

$$Y = \bigcap_{i=1}^{n-m} \pi_i^{-1}(\pi_i(Y)).$$

A tropical variety is contained in a union of finitely many linear spaces, but requires one more projection π_0 in order to write it as the intersection of preimages; this projection determines which part of each linear space belongs to the tropical variety. \diamond

Example 5.2.3. The following is Example 4.2.11 in [33]. Let

$$\begin{aligned} I_1 &= \langle xz + 4yz - z^2 + 3x - 12y + 5z, xy - 4y^2 + yz + x + 2y - z \rangle, \\ I_2 &= \langle xy - 3xz + 3yz - 1, 3xz^2 - 12yz^2 + xz + 3yz + 5z - 1 \rangle. \end{aligned}$$

The varieties defined by these two ideals are curves in \mathbb{C}^3 whose tropicalizations are the rays from the origin to the positive (product of coordinates is positive) and negative vertices of the cube $[-1, 1]^3$ respectively. We display both curves in Figure 5.3. Notice that for any $\{i, j\} \subset \{1, 2, 3\}$,




Figure 5.3: (Reprinted from [1]) An example of two tropical curves which cannot be distinguished from their coordinate projections

we have that $\pi_{\{i,j\}}(\text{trop}(\mathcal{V}(I_1))) = \pi_{\{i,j\}}(\text{trop}(\mathcal{V}(I_2)))$ is the tropical plane curve whose rays are the positive span of the vertices of the square $[-1, 1]^2$. Therefore, these two tropical curves cannot be distinguished from their coordinate projections. Note that these projections are not geometrically regular with respect to the union of linear spaces containing each tropical curve. \diamond

Remark 5.2.4. By (5.8), we have that $\text{trop}(\mathcal{V}(f)) = \varphi(\text{trop}(\mathcal{V}(f \circ \Phi)))$ for any monomial change of coordinates Φ and so for any $f_1, \dots, f_m \in \mathbb{C}[x]$,

$$\varphi^{-1}(\text{trop}(\mathcal{V}(f_1, \dots, f_m))) = \text{trop}(\mathcal{V}(f_1 \circ \Phi, \dots, f_m \circ \Phi)),$$

where $\varphi = \Phi^*$. Projecting gives

$$\pi_{[k]} \varphi^{-1}(\text{trop}(\mathcal{V}(f_1, \dots, f_m))) = \pi_{[k]} \text{trop}(\mathcal{V}(f_1 \circ \Phi, \dots, f_m \circ \Phi)). \quad (5.9)$$

Thus, one way to produce a projection $A: \mathbb{R}^n \rightarrow \mathbb{R}^k$ on tropical varieties other than a coordinate projection is to write A as $\pi_{[k]} \circ \varphi^{-1}$ such that φ is an $n \times n$ matrix over \mathbb{Z} and apply (5.9). \diamond

5.3 Sparse polynomial systems

Given a collection $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ of nonempty finite subsets of \mathbb{Z}^n , write $\mathbb{C}^{\mathcal{A}_\bullet} = \mathbb{C}^{\mathcal{A}_1} \times \dots \times \mathbb{C}^{\mathcal{A}_n}$ for the vector space of n -tuples $F = (f_1, \dots, f_n)$ of polynomials, where f_i is supported on \mathcal{A}_i , for each i . An element $F \in \mathbb{C}^{\mathcal{A}_\bullet}$ corresponds to a system of polynomial equations

$$f_1(x_1, \dots, x_n) = f_2(x_1, \dots, x_n) = \dots = f_n(x_1, \dots, x_n) = 0,$$

called a **sparse polynomial system supported on \mathcal{A}_\bullet** . We write F to refer to these equations or to their vector of coefficients, depending on context. For $\omega \in \mathbb{R}^n$, we let $F_\omega = ((f_1)_\omega, \dots, (f_n)_\omega)$. Letting $P_\bullet = (P_1, \dots, P_n)$ where $P_i = \text{conv}(\mathcal{A}_i)$, we define the mixed volume $\text{MV}(\mathcal{A}_\bullet)$ of \mathcal{A}_\bullet to be $\text{MV}(P_\bullet)$.

5.3.1 Geometry of sparse polynomial systems

Given $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_n)$, consider the incidence variety

$$X_{\mathcal{A}_\bullet} = \{(F, x) \in \mathbb{C}^{\mathcal{A}_\bullet} \times (\mathbb{C}^\times)^n \mid F(x) = 0\}$$

equipped with projections $\pi_{\mathcal{A}_\bullet}: X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ and $p: X_{\mathcal{A}_\bullet} \rightarrow (\mathbb{C}^\times)^n$. For $F \in \mathbb{C}^{\mathcal{A}_\bullet}$, the fiber $\pi_{\mathcal{A}_\bullet}^{-1}(F)$ is identified with the set $\mathcal{V}(F)$ of solutions in $(\mathbb{C}^\times)^n$ to $F = 0$.

For any $x \in (\mathbb{C}^\times)^n$, the fiber $p^{-1}(x)$ is a codimension n vector subspace of $\mathbb{C}^{\mathcal{A}_\bullet}$. Indeed, for each $i = 1, \dots, n$, the condition that $f_i(x) = 0$ is a linear equation in the coefficients $\mathbb{C}^{\mathcal{A}_i}$ of f_i , and these n linear equations are independent. As a consequence $X_{\mathcal{A}_\bullet}$ is irreducible of dimension

$$\dim(\mathbb{C}^\times)^n + \dim \mathbb{C}^{\mathcal{A}_\bullet} - n = \dim \mathbb{C}^{\mathcal{A}_\bullet},$$

by Lemma 3.8.4 and Lemma 3.8.5.

Proposition 5.3.1 (Bernstein-Kushnirenko). *Let $F \in \mathbb{C}^{\mathcal{A}_\bullet}$ be a system of polynomials supported on \mathcal{A}_\bullet . The number of isolated solutions in $(\mathbb{C}^\times)^n$ to $F = 0$ is at most $\text{MV}(\mathcal{A}_\bullet)$. There is a dense open subset $U \subset \mathbb{C}^{\mathcal{A}_\bullet}$ consisting of systems with exactly $\text{MV}(\mathcal{A}_\bullet)$ solutions.*

Thus $\pi_{\mathcal{A}_\bullet} : X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ is a branched cover if and only if $MV(\mathcal{A}_\bullet) \neq 0$. When this is the case, we denote the Galois group of $\pi_{\mathcal{A}_\bullet}$ by $G_{\mathcal{A}_\bullet}$. We remark that Proposition 5.3.1 gives a different way to compute the mixed volume of a collection of polytopes P_\bullet than the formulas given in Section 2.3: solve a polynomial system whose Newton polytopes comprise the collection P_\bullet and count the solutions in the algebraic torus. A corollary of Proposition 5.3.1 is Bézout's theorem.

Corollary 5.3.2 (Bézout). *Let $\Delta_\bullet = (d_1\Delta_n, \dots, d_n\Delta_n)$ with $d_1, \dots, d_n \in \mathbb{N}$. Then π_{Δ_\bullet} is a branched cover of degree $\prod_{i=1}^n d_i$.*

6. NUMERICAL ALGEBRAIC GEOMETRY

Numerical algebraic geometry refers to a collection of theoretical and computational techniques for studying algebraic varieties using numerical methods. Contrary to symbolic algorithms which use the algebraic description of a variety as input, numerical methods represent varieties by computing approximations of points on them. This gives a computational paradigm which is almost entirely geometric, albeit, theoretically grounded in the algebra and geometry developed in Section 3.

At its core, numerical algebraic geometry uses tools from numerical analysis to compute approximate solutions of zero-dimensional polynomial systems. Computations on positive-dimensional varieties are performed numerically via their zero-dimensional intersections with general affine linear spaces of complementary dimension. The information of such an intersection comprises the fundamental data structure in numerical algebraic geometry: a witness set. When equipped with the method of homotopy continuation, a witness set may be used to efficiently extract information from a variety.

Understanding the basic concepts underlying numerical algebraic geometry does not require an extensive background in algebraic geometry, but the language from Section 3 illuminates many of the ideas involved. For example, we will see that homotopy methods are conveniently chosen branched covers, a clever interpretation of the fibers, and a special (but not too special!) fiber which can be computed.

We begin by briefly explaining the core numerical methods underlying the theory in Section 6.1 and then move on to an assortment of algorithms from numerical algebraic geometry, including the polyhedral homotopy (Algorithm 6.3.5) and the monodromy solve algorithm (Algorithm 6.5.3). We remark that Figure 6.8 appears in the article [1] by the author¹.

6.1 Core numerical methods

We discuss what it means to numerically solve a polynomial system and explain two core numerical algorithms: Euler’s method and Newton’s method. These algorithms may be used as the predictor and corrector subroutines of a predictor-corrector method.

6.1.1 Approximate solutions

Given a polynomial map $F: \mathbb{C}^n \rightarrow \mathbb{C}^n$, the system $F = 0$ is a collection of n polynomial equations in n variables and is thus called a **square system**. We suppose for now that $\mathcal{V}(F)$ is finite. For such a multivariate map, define

$$N_F(x) = x - (DF)^{-1} F(x),$$

where DF is the Jacobian matrix of F evaluated at x , $(DF)^{-1}$ is its inverse, and x and $F(x)$ are column vectors. We remark that $N_F(x)$ is only well-defined when DF is nonsingular at $x \in \mathbb{C}^n$. Applying N_F to a point $x_0 \in \mathbb{C}^n$ is called a **Newton step** on x_0 , or a **Newton iteration**. A **Newton sequence** is a sequence of points $\{x_0, x_1, \dots\}$ defined recursively from some initial point x_0 by $x_{i+1} = N_F(x_i)$. A sequence $\{x_0, x_1, \dots\}$ **converges quadratically** to a point $\xi \in \mathbb{C}^n$ if for all i

$$\|x_i - \xi\| \leq 2^{1-2^i} \|x_0 - \xi\|.$$

¹Reprinted with permission from T. Brysiewicz, “Numerical Software to Compute Newton polytopes and Tropical Membership,” *Mathematics in Computer Science*, 2020. Copyright 2020 by Springer Nature.

Newton's method is a root-finding algorithm which iteratively applies Newton steps to some point $x_0 \in \mathbb{C}^n$ with the hope that the Newton sequence $\{x_0, x_1, \dots\}$ converges to a solution of $F = 0$.

Algorithm 6.1.1 (Newton's Method).

Input:

- A point $x_0 \in \mathbb{C}^n$
- A square polynomial system F
- Some number of iterations, $m \in \mathbb{N}$

Output:

- The m -th Newton iteration, $N_F^m(x_0)$

Steps:

```

1 set  $i = 0$ 
2 while  $i < m$  do
    2.1 set  $x_{i+1} = x_i - (DF|_{x_i})^{-1}F(x_i)$ 
    2.2 set  $i = i + 1$ 
3 return  $x_m$ 

```

Lemma 6.1.2. [34, Theorem 3.5] *If x_0 is sufficiently near a smooth point $\xi \in \mathcal{V}(F)$ then a Newton sequence beginning with x_0 will converge quadratically to ξ .*

A point $x_0 \in \mathbb{C}^n$ is an **approximate zero** of $F = 0$ with **associated zero** $\xi \in \mathcal{V}(F)$ if the Newton sequence starting at x_0 converges quadratically to ξ . In this sense, x_0 is a numerical solution to $F = 0$.

Certifying that a point is a numerical solution is made possible through α -theory [35, Ch.8], developed by Smale [36] in the 1980's. We introduce the notation

$$\begin{aligned} \beta(F, x) &= \|x - N_F(x)\| = \|DF(x)^{-1}F(x)\|, \\ \gamma(F, x) &= \sup_{k \geq 2} \left\| \frac{DF(x)^{-1}D^k F(x)}{k!} \right\|^{\frac{1}{k-1}}, \\ \alpha(F, x) &= \beta(F, x) \cdot \gamma(F, x), \end{aligned}$$

where $D^k F(x)$ is the symmetric tensor comprised of the k -th order partial derivatives of f . Since $D^k F$ is a linear map from the k -fold symmetric power of \mathbb{C}^n to \mathbb{C}^n , so is $DF(x)^{-1}D^k F(x)$. The norm in the definition of $\gamma(F, x)$ is the operator norm induced by the standard norms on \mathbb{C}^n and the symmetric powers of \mathbb{C}^n . With this notation, we state a sufficient condition on quadratic convergence which forms the basis for α -theory.

Proposition 6.1.3. *A point $x_0 \in \mathbb{C}^n$ is an approximate solution of a square system $F = 0$ if $\alpha(F, x_0) < (13 - 3\sqrt{17})/4 \approx 0.15767078$.*

Given a point $x \in \mathbb{C}^n$ and a square polynomial system $F = 0$, software such as **alphaCertified** [37] and **NumericalCertification** [38] verify the inequality in Proposition 6.1.3 and can thus rigorously certify that x_0 is an approximate solution of $F = 0$.

6.1.2 Euler's method

Euler's method is a standard numerical method for solving a first order ordinary linear differential equation given an initial value. Fix an ordinary linear differential equation encoded via a matrix equation

$$\frac{\partial x}{\partial t} = F(t; x(t)), \quad x(t_0) = x_0,$$

where $F(t; x(t)): \mathbb{C}_t \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n$ is continuous near $(t_0; x_0)$ in $\mathbb{C}_t \times \mathbb{C}_x^n$. Fix a step size $h > 0$ and define

$$E_F(t; x) = x + hF(t; x).$$

Applying E_F to a point $(t_0; x_0)$ is called an **Euler step**. An **Euler sequence** is a sequence of points $\{(t_0; x_0), (t_1; x_1), \dots\}$ where $t_{i+1} = t_i - h$ and $x_{i+1} = E_F(t_i; x_i)$.

Analogous to Newton's method, given a step size h and a number of steps m , **Euler's method** attempts to compute an approximation x_m of $x(t_m)$.

Algorithm 6.1.4 (Euler's method).

Input:

- A first order linear differential equation $\frac{\partial x}{\partial t} = F(t; x)$
- An initial value $x(t_0) = x_0$
- A step size h
- A number of steps m

Output:

- An approximation x_m of $x(t_m)$

Steps:

```

1 set  $i = 0$ 
2 while  $i < m$  do
    2.1 set  $x_{i+1} = E_F(t_i; x_i)$ 
    2.2 set  $t_{i+1} = t_i - h$ 
3 return  $x_m$ 

```

Example 6.1.5. Figure 6.1 displays four branches of a curve $\mathcal{V}(F) \subset \mathbb{C}_{t,x}^2$ where

$$F(t; x) = 5(1-t)(x-0.1)(x-0.4)^2(x-0.6) + t(x-0.25)(x-0.5)(x-0.75)(x-0.05).$$

The branch containing the point $(1; 0.75)$ is the graph of some function $x(t): [0, 1] \rightarrow \mathbb{R}^2$ satisfying $F(t, x(t)) = 0$ for $t \in [0, 1]$ and thus satisfying the differential equation $DF(t; x(t)) = 0$. After applying the chain rule, this becomes,

$$\frac{\partial x}{\partial t} = -\frac{-4x^4 + 5.95x^3 - 3.1375x^2 + .671875x - .0433125}{-16tx^3 + 17.85tx^2 + 20x^3 - 6.275tx - 22.5x^2 + .671875t + 7.8x - .8}$$

We perform Algorithm 6.1.4 on this differential equation using the auxiliary input

$$x_0 = x(1) = 0.75, \quad h = 0.1, \quad \text{and} \quad m = 10,$$

so that $x_m = x(0)$. The computed points $\{(t_i; x_i)\}_{i=0}^m$ are shown in Figure 6.1 in green. ◇




Figure 6.1: Algorithm 6.1.4 with $h = 0.1$, $x(1) = 0.75$, and $m = 10$.

6.1.3 Predictor-corrector methods

Given a differential equation

$$\frac{\partial x}{\partial t} = F(t; x(t)) \quad (6.1)$$

and some starting point $x(t_0) = x_0$ satisfying (6.1), a **predictor-corrector method** attempts to analytically continue $x(t)$ as t goes from t_0 to some $t_m \in \mathbb{R}$ (taking $h = \frac{t_m - t_0}{m}$) by interspersing applications of a predictor method (like Euler's method) and a corrector method (like Newton's method). Combining both prediction and correction increases the accuracy of (t_m, x_m) dramatically over the use of Euler's method alone (see Example 6.1.7).

Predictor-corrector methods are versatile and depend on choices of

- (1) a differential equation,
- (2) a predictor method,
- (3) a corrector method,
- (4) the parameters involved in both the predictor and the corrector methods.

We give a predictor-corrector method below when the predictor and corrector steps are Euler's method and Newton's method respectively. Thus, this algorithm requires both a differential equation and a system of equations G satisfying $G(t; x(t)) = 0$ for $t \in [0, 1]$ as input. We remark that this is an extremely simple version of such an algorithm and in practice, predictor-corrector methods are often much more nuanced, using predictor methods with higher accuracy, applying Newton's method repeatedly, and adapting the step size throughout the process as needed.

Algorithm 6.1.6 (Predictor-Corrector).**Input:**

- A system of equations $G(t; x)$ such that $G(t; x(t)) = 0$ for all $t \in [0, 1]$
- A first order linear differential equation $\frac{\partial x}{\partial t} = F(t; x)$ satisfied by $x(t)$
- An initial value $x(t_0) = x_0$
- A step size h
- A target t -value, t'

Output:

- An approximate solution of $x(t')$

Steps:

```

0 set  $m = \left\lfloor \frac{(t_0 - t_m)}{h} \right\rfloor$  so that  $t_m - h < t' < t_m$ 
1 set  $i = 0$ 
2 while  $i < m$  do
    2.1 set  $x_{i+1} = E_F(t_i; x_i)$ 
    2.2 set  $t_{i+1} = t_i - h$ 
    2.3 set  $x_{i+1} = N_{G(t_{i+1}; x)}(x_{i+1})$ 
3 set  $x_{m+1} = E_F(t_m; x_m)$  using a stepsize of  $t_m - t'$ 
4 set  $x_{m+1} = N_{G(t'; x)}(x_{m+1})$ 
5 return  $x_{m+1}$ 

```

Example 6.1.7. Figure 6.2 illustrates the accuracy increase in Algorithm 6.1.6 compared to Euler's method alone. We list the numerical data in Table 6.1. \diamond

t -value	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0
True values	.75	.69914	.668991	.649216	.6354	.625304	.617667	.611729	.607003	.603166	.6
Eul. Only	.75	.68175	.641803	.615861	.598597	.587539	.580981	.577285	.575157	.573848	.572984
Eul. Newt.	.75	.70252	.669168	.649393	.635461	.625331	.61768	.611735	.607006	.603168	.600001
Eul. Err.	0	.01739	.027188	.033355	.036803	.037765	.036686	.034444	.031846	.029318	.027016
Eul. Newt. Err.	0	.00338	.000177	.000177	.000061	.000027	.000013	.000006	.000003	.000002	.000001

Table 6.1: Numerical data for Algorithm 6.1.6 on input from Example 6.1.5.

6.1.4 Numerical errors

Since Algorithm 6.1.6 is numerical, it is subject to numerical errors. Due to the limits of rational computations, numerical methods require the approximation of numbers up to some precision. Applying the linear maps relevant to Newton's method and Euler's method to these approximations could possibly increase their error. Algorithm 6.1.6 is especially prone to this whenever the matrix DF evaluated at the approximation x^* has a high **condition number**,

$$\kappa(DF(x^*)) = \|(DF(x^*))^{-1}\| \cdot \|DF(x^*)\|.$$




Figure 6.2: Algorithm 6.1.6 applied to the same differential equation, initial value, and step size as in Example 6.1.5

When this happens, we say that the path is **ill-conditioned** at x^* . One way to alleviate issues coming from error accumulation due to low precision is to use **adaptive precision**. Adaptive precision involves changing the precision used during the predictor-corrector process based on indicators of the conditioning of the path being followed.

Another problem which could occur during Algorithm 6.1.6 is **path-jumping**. Path-jumping occurs when the result (t^*, x^*) of an Euler step attempting to approximate $(t^*, x(t^*))$ is close enough to a solution $(t^*, \hat{x}(t^*)) \neq (t^*, x(t^*))$ so that a Newton sequence starting with (t^*, x^*) converges to $(t^*, \hat{x}(t^*))$. We display how this may occur in Figure 6.3. One way to avoid path-jumping is




Figure 6.3: A visual display of path-jumping.

to decrease the step size during the predictor-corrector process, particularly when DF is has high

condition number.

Developing robust or certifiable predictor-corrector methods is a goal of much current research [39, 40, 41]. For further information about these topics, we refer the reader to [42].

6.2 Homotopies

Homotopies make the idea of continuous deformations rigorous and are defined with respect to general topological spaces. For our purposes, we restrict ourselves to homotopies arising from polynomial systems. Let $H(s; x) \in \mathbb{C}[s][x]$ be a system of n polynomials in m parameters $s^{(1)}, \dots, s^{(m)}$ and n variables x . Suppose the projection

$$\pi: \mathcal{V}(H(s; x)) \rightarrow \mathbb{C}_s^m \quad (6.2)$$

is a degree d branched cover with regular values $U \subset \mathbb{C}_s^m$. The system $H(s; x)$ may also be thought of as a map

$$H(s; x): \mathbb{C}_s^m \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n. \quad (6.3)$$

By the path-lifting property of covering spaces, composing $H(s; x)$ with any continuous path $\tau: [0, 1]_t \rightarrow \mathbb{C}_s^m$ with $\tau(0, 1] \subset U$ produces a map

$$H(\tau(t); x): [0, 1]_t \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n \quad (6.4)$$

along with d lifts $\{x_i(t)\}_{i=1}^d$ over $\tau(0, 1]$ satisfying $H(\tau(t); x_i(t)) = 0$ for all $t \in (0, 1]$. Set $H_\tau(t; x) = H(\tau(t); x)$ so that for any $t^* \in (0, 1]$, the polynomial system $H_\tau(t^*; x) \in \mathbb{C}[x]$ has d solutions $\{x_i(t^*)\}_{i=1}^d$. We call H_τ a **homotopy** with **start system** $H_\tau(1; x) \in \mathbb{C}[x]$ and **target system** $H_\tau(0; x) \in \mathbb{C}[x]$. We call $\{x_i(t)\}_{i=1}^d$ the **paths** of the homotopy H_τ and the set $\{x_i(1)\}_{i=1}^d \subset \mathbb{C}^n$ the **start solutions** of H_τ . The isolated solutions of the target system $H_\tau(0; x)$ are called **target solutions**. A homotopy is called **regular** if additionally, $\tau(0) \in U$.

We omit the subscript on H_τ when convenient. When the limit $\lim_{t \rightarrow 0} x_i(t)$ of some path exists, we extend $x_i(t): (0, 1] \rightarrow \mathbb{C}^n$ continuously by setting $x_i(0) := \lim_{t \rightarrow 0} x_i(t)$. If $H(t; x)$ is a homotopy, then for any $\lambda \in \mathbb{C}^\times$, we say $\lambda H(t; x)$ and $H(t; x)$ are **equivalent** and write $H(t; x) \equiv \lambda H(t; x)$ since the zeros of $\lambda H(t; x)$ are the same as those of $H(t; x)$.

Lemma 6.2.1. *If $H(t; x)$ is a homotopy, then each target solution has the form $x_i(0)$ for some path $x_i(t): [0, 1] \rightarrow \mathbb{C}^n$ of the homotopy.*

Proof. Suppose $H(t; x) = H(\tau(t); x)$ for $H(s; x) \in \mathbb{C}[s][x]$ and $\tau: [0, 1] \rightarrow \mathbb{C}_s^m$. Let $U \subset \mathbb{C}_s^m$ be the set of regular values of $\mathcal{V}(H(s; x)) \xrightarrow{\pi} \mathbb{C}_s^m$ and let $p \in \mathcal{V}(H(0; x))$ be a target solution.

Since $\mathcal{V}(H(0; x))$ is nonempty, Corollary 3.6.7 implies that p belongs to an irreducible component C of $\mathcal{V}(H(s; x))$ of dimension at least m . But the dimension of C is at most m since p is isolated in its fiber over $t = 0$. Thus $\dim(C) = m$.

Since C has dimension m and the point p in the fiber of $\pi|_C: C \rightarrow \mathbb{C}_s^m$ over $t = 0$ is isolated in its fiber, $\pi(C)$ is open and dense in \mathbb{C}_s^m and thus the intersection of U and $\pi|_C(C)$ is open and dense. Considering the homotopy $H(t; x): [0, 1] \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n$ as a map, observe that since $\pi|_C(C)$ is open and dense in \mathbb{C}_s^m , the set $H^{-1}(\mathbf{0}) \cap (0, \epsilon) \times \mathbb{C}_x^n$ contains points in C for any $\epsilon > 0$. Such points must be of the form $x_i(\epsilon)$ for some path $x_i(t)$ of H and thus $\lim_{\epsilon \rightarrow 0} x_i(\epsilon)$ converges to $p \in C$. \square

Lemma 6.2.2. *Let $F(x), G(x) \in \mathbb{C}[x_1, \dots, x_n]$ be square systems. Let*

$$H(s; x) = (1 - s)F(x) + sG(x).$$

If $s = 1$ is a regular value of $\pi: \mathcal{V}(H(s; x)) \rightarrow \mathbb{C}_s$ then there exists a subset $S \subset \mathbb{C} \times \mathbb{C}$ of full measure such that for $\gamma = (\gamma_0, \gamma_1) \in S$,

$$H_{\tau_\gamma}(t; x) \equiv (1 - t)\gamma_0 F(x) + t\gamma_1 G(x) \tag{6.5}$$

is a homotopy, where

$$\begin{aligned} \tau_\gamma: [0, 1] &\rightarrow \mathbb{C}_s \\ t &\mapsto \frac{t\gamma_1}{t\gamma_1 + \gamma_0 - t\gamma_0}. \end{aligned}$$

If $s = 0$ is a regular value of π as well, then $H_{\tau_\gamma}(t; x)$ is regular.

Proof. Let $\tau_\gamma(t) = \frac{t\gamma_1}{t\gamma_1 + \gamma_0 - t\gamma_0}$. We claim that $H(\tau_\gamma(t); x)$ has the same solutions as the right-hand-side of (6.5) for any $t \in \mathbb{C} \setminus \mathcal{V}(t\gamma_1 + \gamma_0 - t\gamma_0)$. To see this, note that

$$\begin{aligned} H(\tau_\gamma(t); x) &= (1 - \tau_\gamma(t))F + \tau_\gamma(t)G \\ &= \left(1 - \frac{t\gamma_1}{t\gamma_1 + \gamma_0 - t\gamma_0}\right)F + \frac{t\gamma_1}{t\gamma_1 + \gamma_0 - t\gamma_0}G \end{aligned}$$

The denominator $t\gamma_1 + \gamma_0 - t\gamma_0$ is zero when $t = \frac{-\gamma_0}{\gamma_1 - \gamma_0}$. When $t \neq \frac{-\gamma_0}{\gamma_1 - \gamma_0}$, the denominator is nonzero. Thus, for γ chosen in a subset of $\mathbb{C} \times \mathbb{C}$ of full measure, we can clear denominators without changing the solutions:

$$\begin{aligned} &= (t\gamma_1 + \gamma_0 - t\gamma_0 - t\gamma_1)F + t\gamma_1 G \\ &= (1 - t)\gamma_0 F + t\gamma_1 G. \end{aligned}$$

The branch locus D of π has complex codimension 1 in \mathbb{C}_t , (i.e. D is a finite set of points d_1, \dots, d_k in \mathbb{C}_t). We claim that the set of ratios γ_0/γ_1 with the property that $\tau_\gamma(t) = d_i$ for some i and some $t \in [0, 1]$ has measure zero in $\mathbb{C} \cong \mathbb{R}^2$. Because scaling does not change solutions, we may assume that $\gamma_1 = 1$. Note that $\tau_\gamma(t) = \frac{t}{t + (1-t)\gamma_0} = 1 + \frac{1}{(1-t)\gamma_0}$ so $\tau_\gamma(t) = d_i$ if and only if $(d_i - 1)(1 - t) = \gamma_0^{-1}$ for some $t \in [0, 1]$. Thus, the only γ_0^{-1} for which $\tau_\gamma(t) = d_i$ for some $i = 1, \dots, k$ and $t \in (0, 1]$ are those whose inverses are contained on the finitely many half-open line segments $\{(d_i - 1)(1 - t)\}_{t \in (0, 1]}$. This set has measure zero. Thus, the subset $S' \subset \mathbb{C} \times \mathbb{C}$ inducing such ratios has measure zero in $\mathbb{C} \times \mathbb{C} \cong \mathbb{R}^4$ and its complement $S = \mathbb{C} \times \mathbb{C} \setminus S'$ has full measure in $\mathbb{C} \times \mathbb{C}$. Moreover, if $\tau_\gamma(0) \neq d_i$ for any $i = 1, \dots, k$, then $\tau_\gamma([0, 1]) \cap \{d_1, \dots, d_k\} = \emptyset$ for general $\gamma \in \mathbb{C} \times \mathbb{C}$ implying $H(t; x)$ is regular. \square

A homotopy of the form

$$H(t; x) = (1 - t)F(x) + tG(x) \tag{6.6}$$

is called a **straight-line homotopy**. Given two square polynomial systems $F(x)$ and $G(x)$, the construction (6.6) may not be a homotopy, however, if 1 is a regular value of $\pi: \mathcal{V}(H(t; x)) \rightarrow \mathbb{C}_t$, then (6.5) is a homotopy with probability one: under any probability measure on the space $\mathbb{C} \times \mathbb{C}$ of choices for γ in (6.5), the probability that γ is chosen so that $H(\tau_\gamma(t); x)$ is a homotopy is one. Replacing $H(t; x)$ with (6.5) is called the **γ -trick**.

Lemma 6.2.3. *Suppose that $H(s; x) \in \mathbb{C}[s][x]$ is a square system and that*

$$\pi: \mathcal{V}(H(s; x)) \rightarrow \mathbb{C}_s^m$$

is a branched cover with regular values $U \subset \mathbb{C}_s^m$. If $s_1 \in U$ and $s_0 \in \mathbb{C}_s^m$, then there exists a path $\tau: [0, 1] \rightarrow U$ such that $\tau(0) = s_0$ and $\tau(1) = s_1$ making H_τ a homotopy. If $s_0 \in U$ then H_τ is a regular homotopy.

Proof. Since $s_1 \in U$, the line connecting s_0, s_1 in \mathbb{C}_s^m intersects the branch locus of π in finitely many points. Parametrize this line by

$$\begin{aligned} \tau': \mathbb{C} &\rightarrow \mathbb{C}_s^m \\ q &\mapsto (1 - q)s_0 + qs_1. \end{aligned}$$

Composing τ' with the map γ from Lemma 6.2.2 for generic $(\gamma_0, \gamma_1) \in \mathbb{C} \times \mathbb{C}$ produces a path $\tau: (0, 1] \rightarrow U$ so that H_τ is a homotopy, and additionally, if $s_0 \in U$ then $\tau: [0, 1] \rightarrow U$ and H_τ is a regular homotopy. \square

In light of this result, given a branched cover $\pi: X \rightarrow \mathbb{C}_s^m$, a value $s_0 \in \mathbb{C}_s^m$, and a regular value $s_1 \in \mathbb{C}_s^m$, we will henceforth use the phrase “a homotopy from s_1 to s_0 ” assuming that we take a homotopy as in Lemma 6.2.3.

6.2.1 Homotopy continuation

Given a homotopy $H(t; x)$ and some path $x(t): (0, 1] \rightarrow \mathbb{C}^n$ of the homotopy for which $x(1)$ is known, the method of path tracking uses the predictor-corrector algorithm to analytically continue $x(t)$ as t goes from 1 toward 0. Producing a differential equation satisfied by $x(t)$ is simple. By definition, $H(t; x(t)) = 0$ and therefore,

$$DH(t; x(t)) = 0. \tag{6.7}$$

Applying the chain rule to (6.7) gives

$$D_t H + D_x H \cdot \frac{\partial x}{\partial t} = 0. \tag{6.8}$$

Reordering, this becomes the [Daivdenko differential equation](#) [43],

$$\frac{\partial x}{\partial t} = -(D_x H)^{-1} D_t H \tag{6.9}$$

which when used as the input to the predictor-corrector algorithm (Algorithm 6.1.6) produces a [path tracking algorithm](#) for regular homotopies.

Algorithm 6.2.4 (Path tracking for regular homotopies).

Input:

- A regular homotopy $H(t; x)$
- Approximate start solutions S_1 to $H(1; x) = 0$
- A step size h

Output:

- Approximate target solutions S_0

Steps:

- 1 **for** $s \in S_1$ **do**
 - 1.0 Let $x_s(t)$ be the path of $H(t; x)$ with $x_s(1) = s$
 - 1.1 **set** $x_s(0)$ equal to the output of Algorithm 6.1.6 using the input
 - Differential equation: $\frac{\partial x}{\partial t} = -(D_x H)^{-1} D_t H$
 - System of equations: $H(t; x) = 0$
 - Initial value: $x_s(1)$
 - Step size: h
 - Target t -value: 0
- 2 **return** $S_0 := \{x_s(0)\}_{s \in S}$

Euler and Newton steps of the path tracking algorithm at (t^*, x^*) are explicitly

$$\begin{aligned} (x^*, t^*) &\xrightarrow{E} (t^* - h, x^* + h(D_x H(t^*; x^*))^{-1} D_t H(t^*; x^*)) \\ (x^*, t^*) &\xrightarrow{N} (t^*; x^* - (D_x H(t^*; x^*))^{-1} H(t^*; x^*)). \end{aligned} \quad (6.10)$$

Equations (6.10) are only valid at the points $(t^*; x^*)$ where $D_x H(t^*; x^*)$ is invertible. This is the case at all points $(t; x(t))$ corresponding to a path $x(t)$ of the regular homotopy. When H is not a regular homotopy, these conditions fail at 0, but more importantly, they become computationally prohibitive near zero as described in Section 6.1.4.

6.2.2 Endgames

We tame difficulties of homotopies at $t = 0$ using endgame algorithms to produce the nonregular analog of Algorithm 6.2.4. Let $H(t; x)$ be a homotopy coming from lifting a path in \mathbb{C}_s^m to paths $\{x_i(t)\}_{i=1}^d$ with respect to the branched cover

$$\pi: X \rightarrow \mathbb{C}_s^m.$$

If $H(t; x)$ is not regular, a path $x_i(t)$ may exhibit wild behavior near $t = 0$ arising from one of two situations, each preventing the effective use of Algorithm 6.2.4 on $H(t; x)$.

- (1) As $t \rightarrow 0$, the path $x(t)$ diverges.
- (2) The matrix $D_x H$ is not invertible at $p = (0; x(0))$ because
 - (a) the rank of $DH|_p$ is $n - 1$,
 - (b) the rank of $DH|_p$ is n , but the rank of $D_x H|_p$ is $n - 1$,

(c) the rank of DH_p is less than $n - 1$.

Figure 6.4 displays a homotopy where each instance occurs (in order from top to bottom). The




Figure 6.4: A homotopy displaying possible behaviors at $t = 0$.

default practical solution for handling (1) is to simply truncate paths which seem to be diverging. This is assessed throughout the path tracking process by testing at each step whether $|x_i(t)| < N$ for some tolerance $N \gg 0$. If the test fails, the path is no longer tracked under the assumption that it is diverging. Another option is to homogenize the equations of X and take a random dehomogenization. This involves choosing some hyperplane at infinity, and so long as this (real codimension 2) hyperplane does not meet any of the homotopy paths (which have real dimension 1), no path will diverge in the corresponding affine chart.

The next section deals with second case.

6.2.3 Cauchy endgame

Each instance of case (2) may be handled the same way via the Cauchy endgame.

Let $x(t)$ be a path of a homotopy $H(t; x)$. We assume throughout this section that the function $H(t; x)$ extends from a function on the domain $[0, 1]_t \times \mathbb{C}_x^n$ to a function on $\mathbb{C}_t \times \mathbb{C}_x^n$ so that $x(t)$ extends to a map $x(t) : U \rightarrow (H(t; x))^{-1}(\mathbf{0})$ where U are the regular values of \mathbb{C}_t .

There exists $\epsilon > 0$ such that $0 \in \mathbb{C}_t$ is the only branch point of the homotopy in the disc $\Delta \subset \mathbb{C}_t$ of radius ϵ centered at 0 and the map $x(t)$ has a Puiseux expansion

$$x(t) = \left(f_1 \left(t^{\frac{1}{r}} \right), \dots, f_n \left(t^{\frac{1}{r}} \right) \right),$$

for some $r \in \mathbb{N}$ and complex analytic functions f_1, \dots, f_n on the disc $D = \epsilon^{1/r} \Delta$. The number r is called the **winding number** of $x(t)$. Figure 6.5 displays the graph of some $x(t): \Delta \rightarrow \mathbb{C}_x$, with winding number $r = 2$, projected onto the product of $\Delta \subset \mathbb{C}_t$ and the real axis of \mathbb{C}_x . Let




Figure 6.5: A depiction of the local behavior of a path $x(t)$ of a homotopy near a branch point with winding number 2.

$\theta: D \rightarrow \Delta$ be the map $\theta(z) = z^r$. Composing gives $f(z) = (f_1(z), \dots, f_n(z)) = x(\theta(z))$ which is holomorphic on D and has the property that $f(0) = x(0)$.

Lemma 6.2.5. *Suppose that g is a holomorphic function on a closed disc $D \subset \mathbb{C}_z$ centered at the origin. Then*

$$g(0) = \frac{1}{2\pi i} \int_{\partial D} \frac{g(z)}{z} dz.$$

Algorithm 6.2.6 (Cauchy Endgame).

Input:

- A path $x(t)$ of a homotopy $H(t; x)$
- An approximation of $x(\epsilon)$ such that $0 \in \Delta \subset \mathbb{C}_t$ is the only branch point in the disc Δ centered at 0 with radius ϵ

Output:

- A numerical approximation of $x(0)$
- The winding number of $x(t)$

Steps:

- 1 Use Algorithm 6.2.4 to track the point $x(\epsilon)$ around a parametrization of the boundary of Δ by $t(s) = \epsilon e^{\sqrt{-1}s}$ to produce the points $(t(s); x(t(s)))$ (and store them) until on the r -th loop, $(\epsilon, x(t(0))) = (x(u(2\pi)), t_\epsilon)$
- 2 Approximate $x' \approx x(0)$ via the path integral in the Cauchy integral formula using the stored values in step (1)
- 3 **return** (x', r)




Figure 6.6: A visual depiction of the full path tracking algorithm.

Equipped with the Cauchy endgame, we may now state the full path tracking algorithm.

Algorithm 6.2.7 (Path tracking).

Input:

- A homotopy $H(t; x)$
- Approximate start solutions S_1 to $H(1; x)$
- A tolerance $N \gg 0$ for determining divergence
- An endgame tolerance $\epsilon > 0$

Output:

- Approximate target solutions S_0

Steps:

- 1 **for** $s \in S_1$ **do**
 - 1.0 Let $x_s(t)$ be the path of $H(t; x)$ with $x_s(1) = s$
 - 1.1 Compute $x_s(\epsilon)$ using Algorithm 6.2.4
 - 1.2 **if** $|x_s(\epsilon)| \leq N$ and there are no signs of ill-conditioning of the path $x_s(t)$, then continue tracking to $t = 0$
 - 1.3 **if** $|x_s(\epsilon)| > N$ then set $x_s(0) := \infty$
 - 1.4 **else** Use Algorithm 6.2.6 to compute $x_s(0)$
- 2 **return** $S_0 := \{x_s(0)\}_{s \in S}$

We illustrate this algorithm in Figure 6.6.

6.3 Homotopy continuation methods

Given a zero-dimensional polynomial system $F \in \mathbb{C}[x]$, the process of [homotopy continuation](#) finds the isolated solutions of $F = 0$ by the following model.

- (1) If F is overdetermined (more equations than variables), construct a square polynomial system \hat{F} so that $\mathcal{V}(F) \subset \mathcal{V}(\hat{F})$.
- (2) Find a branched cover $\pi: H(t; x) \rightarrow \mathbb{C}_s^m$ so that the fiber $\pi^{-1}(y_0)$ is $\mathcal{V}(\hat{F})$.
- (3) Compute the $\deg(\pi)$ solutions in the fiber $\pi^{-1}(y_1)$ for some $y_1 \in U$.
- (4) Construct a homotopy $H_\tau(t; x)$ where τ is a path connecting $\tau(1) = y_1$ to $\tau(0) = y_0$.
- (5) Apply a path tracking algorithm to compute the target solutions $\mathcal{V}(H(0; x)) = \mathcal{V}(\hat{F}) = \pi^{-1}(y_0)$ from the start solutions $\mathcal{V}(H(1; x)) = \pi^{-1}(y_1)$.
- (6) Determine which points of $\mathcal{V}(\hat{F})$ are isolated points of $\mathcal{V}(F)$.

Step (1) is done by [squaring-up](#) the system F . If $F = (f_1, \dots, f_k) \subset \mathbb{C}[x]$, then for a generic matrix $A \in \mathbb{C}^{k \times n}$, the system

$$\hat{F} = \left\{ \sum_{i=1}^k a_{i,j} f_i \right\}_{i=j}^n$$

is a square polynomial system such that the isolated points of $\mathcal{V}(F)$ are isolated points of $\mathcal{V}(\hat{F})$. Step (6) is usually performed heuristically by checking if $F(s) \approx 0$ at each isolated point $s \in \mathcal{V}(\hat{F})$. If $F(s) \approx 0$ up to some numerical tolerance, then s is deemed to be an isolated solution of $F = 0$. Recently, methods have been developed for certifying solutions of overdetermined systems [44, 45]. In our following discussions we will assume that the polynomial systems involved are already square.

The most general homotopy method is that of a parameter homotopy [46, 47]. Common special cases of parameter homotopies include the Bézout homotopy [48], the polyhedral homotopy [2, 49], and the witness homotopy. We explain these in the following sections. For reference, we include the ingredients of steps (1) and (2) of each homotopy in Table 6.2 at the end of Section 6.4.1.

6.3.1 Parameter homotopies

Let $H(s; x) \in \mathbb{C}[s][x]$ be a square parametrized polynomial system. A [parameter homotopy](#) is any homotopy coming from the restriction of a branched cover

$$\pi: \mathcal{V}(H(s; x)) \rightarrow \mathbb{C}_s^m$$

to a path $\tau: [0, 1] \rightarrow \mathbb{C}_s^m$ such that $\tau(0, 1]$ is contained in the regular values U of π so that

$$H_\tau(t; x): [0, 1]_t \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n$$

is a homotopy. In other words, every homotopy is a parameter homotopy.

The [parameter homotopy method](#) constructs a fiber $\pi^{-1}(s^*)$ in an *ad hoc* fashion. This theoretically can always be done via the Bézout homotopy method, explained in the next section, but

often a more immediate or efficient construction is apparent. In either case, it is standard practice to move from $\pi^{-1}(s^*)$ to a fiber $\pi^{-1}(s_1)$ over a general $s_1 \in U$ via Algorithm 6.2.4. Once $\pi^{-1}(s_1)$ has been computed for a general $s_1 \in \mathbb{C}_s^m$, one may quickly solve for a fiber $\pi^{-1}(s_0)$ by taking τ to be a general path connecting s_1 to s_0 and applying Algorithm 6.2.7 to the homotopy $H_\tau(t; x)$. We note that when $s_0 \in U$, by definition, $H_\tau(t; x)$ is a regular homotopy.

6.3.2 The Bézout homotopy

The Bézout homotopy method solves a zero-dimensional polynomial system $\mathcal{V}(f_1, \dots, f_n)$ where $f_i \in \mathbb{C}[x]$ has degree d_i . In the language of sparse polynomial systems, this method solves for any fiber of the branched cover

$$\pi_{\Delta_\bullet}: X_{\Delta_\bullet} \rightarrow \mathbb{C}^{\Delta_\bullet}$$

where $\Delta_\bullet = (d_1\Delta_n, \dots, d_n\Delta_n)$. By Bézout's theorem, this branched cover has degree $d = \prod_{i=1}^n d_i$. The fiber over $G = \{x_i^{d_i} - 1\}_{i=1}^n$ consists exactly of the points (x_1, \dots, x_n) where x_i is any of the d_i -th roots of unity. Consequently, $|\pi_{\Delta_\bullet}^{-1}(G)| = d$ and so G is a regular value of π_{Δ_\bullet} .

Given a polynomial system $F \in \mathbb{C}^{\Delta_\bullet}$, if the path $\tau_\gamma: [0, 1] \rightarrow \mathbb{C}^{\Delta_\bullet}$ is given by $\tau(t) = \gamma_0(1 - t)F + \gamma_1 tG$ for some random $\gamma_0, \gamma_1 \in \mathbb{C}$, then by the γ -trick, the map $H_{\tau_\gamma}(t; x): [0, 1]_t \times \mathbb{C}_x^n \rightarrow \mathbb{C}^n$ is a homotopy. Since the parameters of π are linear, this homotopy is

$$H_{\tau_\gamma}(t; x) = \gamma_0(1 - t)F(x) + \gamma_1 tG(x).$$

If $F = 0$ has d solutions, then F is a regular value, making $H_\tau(t; x)$ a regular homotopy.

Algorithm 6.3.1 (Bézout homotopy).

Input:

- A square polynomial system $F = (f_1, \dots, f_n) \subset \mathbb{C}[x]$

Output:

- Approximations of the isolated solutions of $F = 0$

Steps:

- 0 **set** $d_i = \deg(f_i)$, $G = \{x_i^{d_i} - 1\}_{i=1}^n$, $S_1 = \{a \in \mathbb{C}^n \mid a_i^{d_i} = 1\}$, and $\gamma_0, \gamma_1 \in \mathbb{C}$ random complex numbers
- 1 **return** the output of Algorithm 6.2.7 on input homotopy $H(t; x) = \gamma_0(1 - t)F + \gamma_1 tG$ and start solutions S_1

It is best to perform path tracking between two polynomial systems where at least one of them is general. In practice, both the start system G of the Bézout homotopy, and the target system F could have special structure. For this reason it is common to apply Algorithm 6.3.1 to solve a random polynomial system $\widehat{G} \in \mathbb{C}^{\Delta_\bullet}$ and subsequently apply a straight-line homotopy from \widehat{G} to F . This process comprises the [Bézout homotopy method](#).

Algorithm 6.3.2 (Bézout homotopy method).

Input:

- A square polynomial system $F = (f_1, \dots, f_n) \subset \mathbb{C}[x]$

Output:

- Approximations of the isolated solutions to $F = 0$

Steps:

- 0 **set** $d_i = \deg(f_i)$, $\widehat{G} \in \mathbb{C}^{\Delta \bullet}$ random, $\gamma_0, \gamma_1 \in \mathbb{C}$ random, and $H(t; x) = \gamma_0(1 - t)F + \gamma_1 t \widehat{G}$
- 1 **set** \widehat{S}_1 to be the output of Algorithm 6.3.1 applied to \widehat{G}
- 2 **return** the output of Algorithm 6.2.7 on input homotopy $H(t; x)$ and start solutions \widehat{S}_1

6.3.3 The polyhedral homotopy

Generalizing the Bézout homotopy, the polyhedral homotopy understands a zero-dimensional polynomial system $F = \{f_1, \dots, f_n\}$ as a member of the family $\mathbb{C}^{\mathcal{A} \bullet}$ of sparse polynomial systems supported on $\mathcal{A} \bullet = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ where $\text{supp}(f_i) = \mathcal{A}_i$. The relevant branched cover in this scenario is $\pi_{\mathcal{A} \bullet}: X_{\mathcal{A} \bullet} \rightarrow \mathbb{C}^{\mathcal{A} \bullet}$. Unlike more basic homotopy methods, a start system is not immediately available, but must be constructed. Much of the notation in the subsequent discussion comes from Section 2.3.

Suppose $F \in \mathbb{C}^{\mathcal{A} \bullet}$ is general and let $\ell \bullet = (\ell_1, \dots, \ell_n)$ be a set of lifting functions $\ell_i: \mathcal{A}_i \rightarrow \mathbb{Z}_{\geq 0}$ such that the induced subdivision $S^{\ell \bullet}$ (Definition 2.4.5) is a fine mixed subdivision of $\mathcal{A} \bullet$. Define

$$f_{i, \ell_i}(t; x) = \sum_{\alpha \in \mathcal{A}_i} c_{i, \alpha} x^{\alpha} t^{\ell_i(\alpha)},$$

so that $\text{New}(f_{i, \ell_i}) = \text{conv}_{\ell_i}(\mathcal{A}_i)$ and similarly define the homotopy

$$F_{\ell \bullet}(t; x) = \{f_{i, \ell_i}(t; x)\}_{i=1}^n,$$

coming from a path in the branched cover $\pi_{\mathcal{A} \bullet}: X_{\mathcal{A} \bullet} \rightarrow \mathbb{C}^{\mathcal{A} \bullet}$ discussed in Section 5.3.1. When $t = 1$, we have $F_{\ell \bullet} = F$ and for a general value of t , this is a zero-dimensional polynomial system with support $\mathcal{A} \bullet$ and so $\pi: \mathcal{V}(F_{\ell \bullet}) \rightarrow \mathbb{C}_t$ is a branched cover with $\text{MV}(\mathcal{A} \bullet)$ branches. As $t \rightarrow 0$, there are often many solutions of $F_{\ell \bullet}(t; x) = 0$ which diverge, although some may not. We understand these paths $\{x_i(t): \mathbb{C}_t \rightarrow \mathbb{C}^n\}_{i=1}^{\text{MV}(\mathcal{A} \bullet)}$ near $t = 0$ by analyzing their Puiseux expansions and changing coordinates accordingly. We explain the process below.

The branches of $F_{\ell \bullet}(t; x)$ are functions $x = x(t)$ admitting a Puiseux expansion

$$x(t) = (z_1 t^{\nu_1}, \dots, z_n t^{\nu_n}) + \text{terms with higher powers of } t,$$

for some $z = (z_1, \dots, z_n) \in \mathbb{C}^n$ and $\nu = (\nu_1, \dots, \nu_n) \in \mathbb{Q}^n$. Taking the composition $F_{\ell \bullet}(t; x(t))$ yields

$$\{f_{i, \ell_i}(x(t))\}_{i=1}^n = \left\{ \sum_{\alpha \in \mathcal{A}_i} c_{i, \alpha} z^{\alpha} t^{(\nu, \alpha) + \ell_i(\alpha)} + \text{terms with higher powers of } t \right\}_{i=1}^n.$$

The solutions of the above system approach those of

$$F^\nu(t; z) = \left\{ \sum_{\alpha \in \mathcal{A}_i} c_{i,\alpha} z^\alpha t^{\langle \nu, \alpha \rangle + \ell_i(\alpha)} \right\}_{i=1}^n,$$

as $t \rightarrow 0$. Let $\omega = (-\nu, -1)$ and observe that the terms of

$$\sum_{\alpha \in \mathcal{A}_i} c_{i,\alpha} z^\alpha t^{\langle \nu, \alpha \rangle + \ell_i(\alpha)} = \sum_{\alpha \in \mathcal{A}_i} c_{i,\alpha} z^\alpha t^{\langle -\omega, \Gamma_i(\alpha) \rangle}$$

with lowest power of t are those α such that the inner product $\langle \omega, \Gamma_i(\alpha) \rangle$ is maximized. Equivalently, these are the vectors α such that $\Gamma_i(\alpha) \in (\Gamma_i(\mathcal{A}_i))_\omega$. Dividing by the lowest power of t occurring in each polynomial of $F^\nu(t; z)$ and evaluating at $t = 0$ yields the polynomial system $G^\nu = 0$ consisting of polynomials

$$f_i^\nu = \sum_{\Gamma_i(\alpha) \in (\Gamma_i(\mathcal{A}_i))_\omega} c_{i,\alpha} z^\alpha.$$

The solutions of G^ν in $(\mathbb{C}^\times)^n$ are the same as those of $F^\nu(0; z)$ in $(\mathbb{C}^\times)^n$. Moreover, since the face of a Minkowski sum is a Minkowski sum of faces, we have that

$$\text{conv} \left(\sum_{i=1}^n (\Gamma_i(\mathcal{A}_i))_\omega \right) = (\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet))_\omega$$

is a face of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$. Let $\mathcal{C}_i(\nu) = \text{supp}(f_i^\nu)$ and $\mathcal{C}_\bullet(\nu) = (\mathcal{C}_1(\nu), \dots, \mathcal{C}_n(\nu))$.

Lemma 6.3.3. *The system G^ν has a solution in the algebraic torus if and only if \mathcal{C}_\bullet is a fine mixed cell of the fine mixed subdivision S^{ℓ_\bullet} .*

Proof. Suppose G^ν has a solution in $(\mathbb{C}^\times)^n$. Since G^ν is a general sparse polynomial system, the Bernstein-Kushnirenko Theorem asserts that the number of solutions in the algebraic torus is the mixed volume of the supports of the f_i^ν . By Lemma 2.5.1, the polytopes $\{\text{conv}(\Gamma_i(\mathcal{A}_i))_\omega\}_{i=1}^n$ form an essential set and so the dimension of $(\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet))_\omega$ is n . Thus, it is a facet in the lower hull of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ and we conclude that $\mathcal{C}_\bullet(\nu)$ is a cell of S^{ℓ_\bullet} .

If for any $i \in [n]$ the polynomial f_i^ν is a monomial, then G^ν has no solutions in the algebraic torus. Thus, each $\text{conv}(\mathcal{C}_i(\nu))$ has dimension at least 1 and so $\text{conv}(\mathcal{C}_\bullet(\nu))$ is mixed. Since S^{ℓ_\bullet} is a fine mixed subdivision, $\mathcal{C}_\bullet(\nu)$ is fine mixed. \square

If ν exposes a fine mixed cell, then each f_i^ν is a binomial and the binomial system G^ν may be solved using the Smith normal form (see Section 5.1.3) to produce all $d_\nu = \text{vol}(\mathcal{C}_\bullet(\nu))$ solutions X^ν to G^ν . These d_ν solutions correspond to limits of paths of $F_{\ell_\bullet}(t; x)$ and may be tracked from $t = 0$ to $t = 1$ by first predicting their values at some $\epsilon > 0$ in the z -coordinates, applying a corrector method, and changing coordinates back to $F_{\ell_\bullet}(t; x)$ to complete the path tracking from $t = \epsilon$ to $t = 1$. This gives d_ν points X_ν of $\mathcal{V}(F) \cap (\mathbb{C}^\times)^n$. By Lemma 2.4.3, this comprises all of the $\text{MV}(\mathcal{A}_\bullet)$ branches of this homotopy.

We illustrate the polyhedral homotopy with an example.

Example 6.3.4. Let \mathcal{A}_\bullet be as in Example 2.4.8 and take

$$\begin{aligned} f_1 &= 3 + 4x - 2y + xy \\ f_2 &= 6 - 2xy^2 + x^2y \end{aligned}$$

and let

$$\begin{aligned} \ell_1(0,0) &= 2, & \ell_1(0,1) &= \ell_1(1,0) = \ell_1(1,1) = 3, \\ \ell_2(0,0) &= \ell_2(1,2) = \ell_2(2,1) = 1. \end{aligned}$$

so that

$$F(t; x) = \{3t^2 + 4xt^3 - 2yt^3 + xyt^6, 6t - 2xy^2t + x^2yt\}.$$

The three directions $\omega_i = (-\nu_i, -1)$ exposing facets in the lower hull of $\text{conv}_{\ell_\bullet}(\mathcal{A}_\bullet)$ which correspond to fine mixed cells are

$$\omega_1 = (2, 2, -1), \quad \omega_2 = (-2, 1, -1), \quad \omega_3 = (1, -2, -1),$$

and so $\nu_1 = (-2, -2), \nu_2 = (2, -1), \nu_3 = (-1, 2)$.




Figure 6.7: A cartoon describing the Polyhedral homotopy.

Construct $F^{\nu_1}(t; x(t)) = \{3t^2 + 4z_1t - 2z_2t + z_1z_2t^2, 6t - 2z_1z_2^2t^{-5} + z_1^2z_2t^{-5}\}$ and divide out by the lowest powers of t to produce $G^{\nu_1}(z) = \{3 + z_1z_2, -2z_1z_2^2 + z_1^2z_2\}$ which has $2 = \text{vol}(\mathcal{C}_\bullet(\nu_1))$ solutions: $X^{\nu_1} = \left\{ (\sqrt{-3/2}, \sqrt{-6}), (-\sqrt{-3/2}, -\sqrt{-6}) \right\}$.

Similarly, for ν_2 construct $F(t; x(t))^{\nu_2} = \{3t^2 + 4z_1t^5 - 2z_2t^2 + z_1z_2t^7, 6t - 2z_1z_2^2t + z_1^2z_2t^4\}$ and $G^{\nu_2}(z) = \{3 - 2z_2, 6 - 2z_1z_2^2\}$. The system $G^{\nu_2}(z)$ has $1 = \text{vol}(\mathcal{C}_\bullet(\nu_2))$ solution, namely $X^{\nu_2} = \{(4/3, 3/2)\}$.

Finally, for ν_3 we have $F^{\nu_3}(t; x(t)) = \{3t^2 + 4z_1t^2 - 2z_2t^5 + z_1z_2t^7, 6t - 2z_1z_2^2t^4 + z_1^2z_2t\}$ and $G^{\nu_3}(z) = \{3 + 4z_1, 6 + z_1^2z_2\}$ which has $1 = \text{vol}(\mathcal{C}_\bullet(\nu_3))$ solution: $X^{\nu_3} = \{(-3/4, -32/3)\}$.

Each solution set X^{ν_i} may be used to approximate d_{ν_i} solutions of $F_{\ell_\bullet}(t; x)$ at $t = \epsilon > 0$ via a predictor and corrector step followed by a coordinate change. Subsequently, we may track these solutions of $X_{\nu_i} \subset \mathcal{V}(F) \cap (\mathbb{C}^\times)^n$ via the homotopy $F_{\ell_\bullet}(t; x)$ as t goes from ϵ to 1. \diamond

Algorithm 6.3.5 (Polyhedral homotopy).

Input: A general sparse polynomial system $F \in \mathbb{C}^{\mathcal{A}_\bullet}$

Output: All solutions of $F = 0$ in the algebraic torus $(\mathbb{C}^\times)^n$

Steps:

- 0 **set** solutions = \emptyset
- 1 Choose lifting functions ℓ_\bullet such that S^{ℓ_\bullet} is a fine mixed subdivision of \mathcal{A}_\bullet
- 2 Compute the mixed cells $\mathcal{C}_\bullet^{(1)}, \dots, \mathcal{C}_\bullet^{(m)}$
- 3 **for** each mixed cell \mathcal{C}_\bullet **do**
 - 3.1 Compute the vector $(-\nu, -1)$ exposing $\text{conv}_{\ell_\bullet}(\mathcal{C}_\bullet)$
 - 3.2 Compute $X^\nu = \mathcal{V}(G^\nu) \cap (\mathbb{C}^\times)^n$ using Smith normal form
 - 3.3 Move the solutions X^ν from $t = 0$ to $t = \epsilon > 0$ via a prediction and correction step and change coordinates to x
 - 3.4 Track the solutions of $F_{\ell_\bullet}(\epsilon; x)$ in step (3.3) from $t = \epsilon$ to $t = 1$ (backwards) under the homotopy $F_{\ell_\bullet}(t; x)$ and append the resulting solutions X_ν to the list solutions
- 4 **return** solutions

Since we usually do not *a priori* know whether or not a sparse polynomial is general in the sense of Proposition 5.3.1, given $F \in \mathbb{C}^{\mathcal{A}_\bullet}$, we solve for the isolated solutions of F in the algebraic torus via the following method.

Algorithm 6.3.6 (Polyhedral homotopy method).

Input: A sparse polynomial system $F \in \mathbb{C}^{\mathcal{A}_\bullet}$

Output: All isolated solutions of $F = 0$ in the algebraic torus $(\mathbb{C}^\times)^n$

Steps:

- 0 Pick a general sparse polynomial system $G \in \mathbb{C}^{\mathcal{A}_\bullet}$
- 1 Apply Algorithm 6.3.5 to G to produce all isolated points $\mathcal{V}(G) \cap (\mathbb{C}^\times)^n$
- 2 Track the points $\mathcal{V}(G) \cap (\mathbb{C}^\times)^n$ to the points $\mathcal{V}(F) \cap (\mathbb{C}^\times)^n$ via the straight-line homotopy $H(t; x) = \gamma_0(1 - t)F + \gamma_1tG$
- 3 **return** $\mathcal{V}(F) \cap (\mathbb{C}^\times)^n$

6.4 Witness sets

Positive-dimensional varieties are represented in numerical algebraic geometry by slicing them with sufficiently many general hyperplanes which cut out degree-many points. Numerical approximations of these points are computed using homotopy methods and stored in the fundamental data structure of numerical algebraic geometry, a witness set.

Definition 6.4.1. Let X be an irreducible variety. A **witness set** for X is a triple (F, L, S) where

- F : a finite set of polynomials such that X is an irreducible component of $\mathcal{V}(F)$.
- L : a general affine linear space of complementary dimension to X .
- S : a set containing approximations of each of the points in $X \cap L$.

If X is reducible with top-dimensional components X_1, \dots, X_k then a witness set for X is (F, L, S) where $S = S_1 \cup \dots \cup S_k$ such that (F, L, S_i) is a witness set for X_i .

We refer to L as a **witness slice** and S as **witness points**. One immediate way to compute a witness set is by using Algorithm 6.3.1.

Algorithm 6.4.2 (Constructing a witness set).

Input:

- A polynomial system $F = (f_1, \dots, f_{n-m}) \subset \mathbb{C}[x]$ such that X is the union of irreducible components of $\mathcal{V}(F)$ of dimension $m = \dim(\mathcal{V}(F))$

Output:

- A witness set (F, L, S) for X

Steps:

- 1 Choose m random linear polynomials $L = \{\ell_1, \dots, \ell_m\} \subset \mathbb{C}[x]$
- 2 Apply Algorithm 6.3.1 to $F \cup \ell$ to produce S
- 3 **return** (F, L, S)

If $X = X_1 \cup \dots \cup X_k \subset \mathbb{C}^n$ is the irreducible decomposition of a variety whose components are not all of the same dimension, let $\text{Dim}(X)$ denote the set $\{\dim(X_i)\}_{i=1}^k$. A **witness superset** of X is a collection $\{(F, L_i, S_i)\}_{i \in \text{Dim}(X)}$ where (F, L_i, S_i) is a witness set for the union of all irreducible components of X of dimension i . Methods for computing witness supersets include “working dimension by dimension” and the “cascade algorithm”. These are discussed in [50, Ch. 9.3-9.4].

6.4.1 The witness cover

Given an irreducible variety $X \subset \mathbb{C}^n$ of dimension m , let

$$W(X) = \{(x, L) \mid x \in X \cap \mathcal{V}(L), \text{ and } L \in (\mathbb{C}^{\Delta_n})^m\} \subset X \times (\mathbb{C}^{\Delta_n})^m$$

be the incidence variety of points on X with linear varieties cut out by m linear polynomials. Then $W(X)$ is irreducible of dimension $m(n+1)$ and the map,

$$\pi_{W(X)}: W(X) \rightarrow (\mathbb{C}^{\Delta_n})^m,$$

is a degree $\deg(X)$ branched cover (by Lemma 3.8.8) called the **witness cover** of X . With this language, it is straightforward to describe how to “move” witness sets.

Algorithm 6.4.3 (Regular Witness homotopy).

Input:

- A witness set (F, L, S) for X
- A general linear space L' of dimension $n - m$

Output:

- A witness set (F, L', S') for X

Steps:

- 1 **set** $H(t; x) = \gamma_0(1 - t)[F|L] + \gamma_1 t[F|L']$
- 2 Track the witness points S via $H(t; x)$ to the solutions S' of $F = L' = 0$
- 3 **return** (F, L', S')

We may deform witness sets of a variety to special linear intersections via a similar algorithm.

Algorithm 6.4.4 (Witness homotopy).

Input:

- A polynomial system $F = (f_1, \dots, f_{n-m}) \subset \mathbb{C}[x]$ such that X is an irreducible component of $\mathcal{V}(F)$ of dimension $m = \dim(\mathcal{V}(F))$
- A witness set (F, L, S) for X
- A linear space L' of dimension $n - m$

Output:

- The points $S' = X \cap L'$

Steps:

- 1 **set** $H(t; x) = \gamma_0(1 - t)[F|L] + \gamma_1 t[F|L']$
- 2 Track the witness points S via $H(t; x)$ to the solutions S' of $F = L' = 0$
- 3 **return** S'

Since L and L' are regular values of the branched cover $\pi_{W(X)}$, Lemma 6.2.2 guarantees that Algorithm 6.4.4 computes a witness set (F, L', S') . We remark that Algorithm 6.4.4 functions just as well for reducible varieties $X = \bigcup_{i=1}^k X_i$ where the map $\pi_{W(X)}$ becomes a branched cover which is not irreducible.

Now that we have explained each of the four homotopy methods mentioned in the introduction, we provide a reference table (Table 6.2) for their ingredients. Parameter homotopies and the Bézout homotopy are implemented in most numerical algebraic geometry software including **Bertini**, **PHCPack**, **HOM4PS**, **homotopycontinuation.jl** and **NAG4M2** [49, 51, 52, 53, 54]. The polyhedral homotopy is implemented in **PHCPack** and **HOM4PS** [49, 53].

6.4.2 Witness sets for images of maps

Much of the strength of numerical algebraic geometry stems from the fact that witness sets (the fundamental data structure in numerical algebraic geometry) can often be computed even when their symbolic analogs, Gröbner bases, cannot. In some sense, this is because Gröbner bases transparently express so much information about a variety while witness sets do not. Rather,

Homotopy	Relevant systems	Branched Cover	Start system
Parameter Section 6.3.1	$F_s \in \mathbb{C}[s][x]$	$\mathcal{V}(F_s) \xrightarrow{\pi} \mathbb{C}_s^m$	$\mathcal{V}(F_{s_1})$
Bézout Section 6.3.2	$F \in \mathbb{C}^{\Delta_\bullet} = \mathbb{C}^{d_1 \Delta_n, \dots, d_n \Delta_n}$	$X_{\Delta_\bullet} \xrightarrow{\pi_{\Delta_\bullet}} \mathbb{C}^{\Delta_\bullet}$	$\{x_i^{d_i} - 1\}_{i=1}^n$
Polyhedral Section 6.3.3	$F \in \mathbb{C}^{\mathcal{A}_\bullet}$	$X_{\mathcal{A}_\bullet} \xrightarrow{\pi_{\mathcal{A}_\bullet}} \mathbb{C}^{\mathcal{A}_\bullet}$	Constructed from a fine mixed subdivision
Witness Section 6.4	$X \cap L$ where $\dim(X) = m = \text{codim}(L)$	$W(X) \xrightarrow{\pi_{W(X)}} (\mathbb{C}^{\Delta_n})^m$	Any witness set for X

Table 6.2: Ingredients for homotopy methods.

witness sets offer users the option to discover information as-needed, similar to the oracles from Section 2.

One particular instance where witness sets can be easily computed is when the variety of interest is a projection. Because witness sets are geometric in nature, they behave well with respect to projections.

Definition 6.4.5. A [pseudo-witness set](#) for an affine variety Z is a quadruple $(F, \varphi, \varphi^{-1}(L_1), S)$ where

- F : a finite set of polynomials such that X is the union of top dimensional components of $\mathcal{V}(F) \subset \mathbb{C}^N$.
- φ : a coordinate projection $\varphi: \mathbb{C}^N \rightarrow \mathbb{C}^n$ such that $Z = \overline{\varphi(X)}$ and $\dim(Z) = \dim(X)$.
- L_1 : a general affine linear space in \mathbb{C}^n of complementary dimension to Z .
- S : a set containing approximations of each of the points in $X \cap \varphi^{-1}(L_1)$.

Often, one desires a pseudo-witness set for the image $Z = \overline{\varphi(X)}$ of a map $X \xrightarrow{\varphi} \mathbb{C}^n$ where the dimension of X is larger than its image. When this is the case, one may take $\dim(X) - \dim(Z)$ generic linear equations $L_2 \subset \mathbb{C}[x_1, \dots, x_N]$ so that $X \cap L_2$ so that the image of $X \cap L_2 \xrightarrow{\varphi} \mathbb{C}^n$ is Z and $\dim(X \cap L_2) = \dim(Z)$. By factoring φ through its graph, it is enough to be able to compute witness sets for projections. We do this in the following way.




Figure 6.8: (Reprinted from [1]) Constructing a pseudo-witness set for a projection of a twisted cubic

Algorithm 6.4.6 (Constructing a pseudo-witness set).

Input:

- A witness set (F, L, S) for a variety X of dimension m
- A coordinate projection $\varphi: \mathbb{C}^N \rightarrow \mathbb{C}^n$ such that $Z = \overline{\varphi(X)}$ and $\dim(Z) = m$

Output:

- A pseudo-witness set $(F, \varphi, \varphi^{-1}(L^*), S)$ for Z

Steps:

- 0 Assume that $\varphi(x_1, \dots, x_N) = (x_1, \dots, x_n)$
- 1 Fix $\dim(Z)$ random linear forms $L^* \subset \mathbb{C}[x_1, \dots, x_n]$
- 2 Use Algorithm 6.4.4 to compute a witness set (F, L^*, S^*) for X by moving $(F, L, S) \rightarrow (F, L^*, S^*)$
- 3 **return** $(F, \varphi, \varphi^{-1}(L^*), S)$

Example 6.4.7. We give two examples which exhibit subtleties in pseudo-witness sets. The first is the twisted cubic $C \subset \mathbb{C}^3$ with the projection $\varphi: \mathbb{C}^3 \rightarrow \mathbb{C}^2$ such that the image is a parabola. During the homotopy which constructs a pseudo-witness, one of the three points of intersection with the twisted cubic diverges towards infinity.

The second example involves the necessity of a dimension reduction. The variety X in this case is the cylinder $\mathcal{V}(x^2 + y^2 - 1) \subset \mathbb{C}_{x,y,z}^3$ along with the projection $\pi: \mathbb{C}_{x,y,z}^3 \rightarrow \mathbb{C}_{x,y}$ whose image is the circle defined by the same equation in $\mathbb{C}[x, y]$. In this case, to construct a pseudo-witness set for the circle, we must first slice the cylinder by a hyperplane to produce the red curve C in Figure 6.9. The dimension of C is the same as the dimension of its image and so one may simply deform a witness set for C to be vertical with respect to π to produce a pseudo-witness set for the circle. \diamond

Any algorithm that may be performed using solely the witness cover of a variety may also be performed using pseudo-witness sets since these may be moved just as easily with respect to the a




Figure 6.9: Constructing a pseudo-witness set for a projection of a cylinder via slicing.

pseudo-witness cover. Suppose $\varphi: X \rightarrow Z$ is a projection. To construct a pseudo-witness cover for Z , replace X and Z with affine open sets, intersect X with a linear space L_2 of codimension $\dim(X) - \dim(Z)$, and relabel variables so that $\varphi(x_1, \dots, x_N) = (x_1, \dots, x_n)$ is a degree d branched cover of affine varieties $X \subset \mathbb{C}^N, Z \subset \mathbb{C}^n$ of dimension m . Take

$$PW(Z, \varphi) = \{(x, L_1) \mid x \in X \cap L_1, \text{ and } L_1 \in (\mathbb{C}^{\Delta_n})^m\}$$

to be the incidence variety of points on X with linear varieties cut out by m polynomials in the first n coordinates. Define the **pseudo-witness cover** of Z with respect to φ (and L_2) to be the map

$$\pi_{PW(Z, \varphi)}: PW(X) \rightarrow (\mathbb{C}^{\Delta_n})^m.$$

Any pseudo-witness set of the form $(F, \varphi, \varphi^{-1}(L_1), S)$ is a fiber $S = \pi_{PW(X)}^{-1}(L_1)$ of $\pi_{PW(X)}$ by construction. In particular, we see that Z has degree $|S|/d$ witnessed by the $|S|/d$ points $\varphi(S) \subset Z \cap L_1$. Note, that $|S|$ is not necessarily equal to the degree of X , as shown in the first part of Example 6.4.7. For more information about pseudo-witness sets, see [55].

6.5 Monodromy

Recall the background on monodromy groups of branched covers in Section 4.

Let $F(s; x) \subset \mathbb{C}[s][x]$ be a parametrized polynomial system in m parameters s and n variables x so that

$$\pi: \mathcal{V}(F(s; x)) \rightarrow \mathbb{C}_s^m$$

is a degree d branched cover with regular values $U \subset \mathbb{C}_s^m$. We do not assume π is an irreducible

branched cover. For $s_1, s_2 \in U$ and $c_1 \in \mathbb{C}$ let

$$\begin{aligned} \tau_{s_1, s_2, c_1} : [0, 1]_t &\rightarrow \mathbb{C}_s^m \\ t &\mapsto (1-t)s_1 + c_1 t s_2 \end{aligned}$$

be a path in \mathbb{C}_s^m . For this section, we will assume that c_1 is in the Euclidean dense subset of \mathbb{C} which satisfies $\tau_{s_1, s_2, c_1}([0, 1]) \subset U$. Applying the path tracking algorithm for regular homotopies (Algorithm 6.2.4) to the homotopy $H_{s_1, s_2, c_1}(t; x)$ produced by π and τ_{s_1, s_2, c_1} gives the bijection $m_{\tau_{s_1, s_2, c_1}}$ discussed in Section 4.3. Picking another generic complex number c_2 , the composition

$$m_{\tau_{s_2, s_1, c_2}} \circ m_{\tau_{s_1, s_2, c_1}} : \pi^{-1}(s_1) \rightarrow \pi^{-1}(s_1)$$

is a monodromy permutation of the d points in the fiber over s_1 . This permutation is the monodromy element m_γ where γ is the loop in \mathbb{C}_s^m formed by following the concatenation of the paths τ_{s_1, s_2, c_1} and τ_{s_2, s_1, c_2} .

This leads immediately to a heuristic algorithm for computing elements of the monodromy group of a branched cover.

Algorithm 6.5.1 (Extract monodromy group element).

Input:

- A parametrized polynomial system $F(s; x) \subset \mathbb{C}[s][x]$ such that $\pi : \mathcal{V}(F(s; x)) \rightarrow \mathbb{C}_s^m$ is an irreducible branched cover of degree d
- A fiber $S_1 = \pi^{-1}(s_1)$

Output:

- An element g of the monodromy group \mathcal{M}_π

Steps:

- 1 Label $S_1 = \{p_1, \dots, p_d\}$ so that p_i is identified with $i \in [d]$
- 2 Pick $s_2 \in U$ and generic $c_1, c_2 \in \mathbb{C}$
- 3 Track all points in S_1 along H_{s_1, s_2, c_1} to produce S_2
- 4 Track all points in S_2 along H_{s_2, s_1, c_2} to produce $(m_{\tau_{s_2, s_1, c_2}} \circ m_{\tau_{s_1, s_2, c_1}})(S)$
- 5 Determine $g = m_{\tau_{s_2, s_1, c_2}} \circ m_{\tau_{s_1, s_2, c_1}}(p_i)$ for each $i \in [n]$ to determine g
- 6 **return** g

Example 6.5.2. Rather than producing loops with only two parameter values in our examples, we use three parameters, s_1, s_2 , and s_3 , to clarify the ideas and images. Figure 6.10 shows a schematic of the computation of a monodromy group element using numerical algebraic geometry. Labeling the points of $\pi^{-1}(s_1)$ from bottom to top as $1, 2, \dots, 5$, the element $m_\gamma \in \mathcal{M}_\pi$ is $m_\gamma = (1, 2)(4, 5)$, written in cycle notation. Its cycles are depicted in distinct colors in Figure 6.10. \diamond

To determine the monodromy group \mathcal{M}_π of a branched cover, one may repeatedly extract group elements using Algorithm 6.5.1 until the group generated by these elements fails to grow after many runs of the algorithm. This is, of course, heuristic. A more rigorous way to compute the monodromy group is to restrict the parameter space \mathbb{C}_s to a generic line $\mathbb{C}_t \subset \mathbb{C}_s$. The branch locus of π restricted to \mathbb{C}_t consists of finitely many points b_1, \dots, b_k . A theorem of Zariski [56] implies that the monodromy group \mathcal{M}_π is generated by loops around each b_i . For more information about computing monodromy groups of branched covers using numerical algebraic geometry, we refer the reader to [57].




Figure 6.10: A schematic of a single monodromy loop tracked numerically.

6.5.1 Solving via monodromy

Recall that the monodromy group of an irreducible branched cover is transitive (Lemma 4.3.1). Thus, the observation that monodromy permutations can be explicitly computed using numerical algebraic geometry suggests one way to compute $\pi^{-1}(s_1)$ given some point $q \in \pi^{-1}(s_1)$: pick a random monodromy loop γ , use a homotopy to track q via a lift of γ thus computing $m_\gamma(q)$, and repeat. This is the naïve version of the [monodromy solve](#) algorithm.

Algorithm 6.5.3 (Naïve monodromy solver).

Input:

- A parametrized polynomial system $F(s; x) \subset \mathbb{C}[s][x]$ such that $\pi: \mathcal{V}(F(s; x)) \rightarrow \mathbb{C}_s^m$ is an irreducible branched cover
- A single point q in some fiber $\pi^{-1}(s_1)$

Output:

- All points in $\pi^{-1}(s_1)$

Steps:

- 1 **set** $S_1 = \{q\}$
- 2 **while** $\pi^{-1}(s_1)$ has not been fully computed **do**
 - 2.1 Pick $s_2, \in U$ and generic $c_1, c_2 \in \mathbb{C}$
 - 2.2 Track all points in S_1 along H_{s_1, s_2, c_1} to produce S_2
 - 2.3 Track all points in S_2 along H_{s_2, s_1, c_2} to produce

$$S'_1 = (m_{\tau_{s_2, s_1, c_2}} \circ m_{\tau_{s_1, s_2, c_1}})(S) \subset \pi^{-1}(s_1)$$

- 2.4 **set** $S_1 = S_1 \cup S'_1$
- 3 **return** S_1

Remark 6.5.4. Conditions for determining when the fiber has been “fully computed” in step (2) are called **stopping criteria** and are not obvious. When the degree of the cover is known, then a stopping criterion for Algorithm 6.5.3 is “stop when $\deg(\pi)$ points of $\pi^{-1}(s_1)$ have been computed”. We give an alternative stopping criterion when π is the witness cover of a variety. \diamond

By Lemma 5.1.5, the centroids of witness points on a pencil of witness slices of an irreducible affine variety lie on an affine line. A stronger result is true.

Lemma 6.5.5. *Let X be an irreducible affine variety and L_t a general pencil of linear spaces of complementary dimension. Given a subset of witness points $S_0 \subset L_0 \cap X$, the centroids of the paths starting at S_0 along the homotopy over L_t in the witness cover moves affine linearly if and only if $S_0 = L_0 \cap X$.*

Let X be an irreducible variety of dimension m . When performing Algorithm 6.5.3 on the witness cover

$$\pi_{W(X)}: W(X) \rightarrow (\mathbb{C}^{\Delta_n})^m$$

a stopping criterion is that the condition in Lemma 6.5.5 holds. This may be checked during the monodromy algorithms by moving a witness set to three slices in a pencil and numerically taking the centroids of the witness points. As witness points are only numerical approximations, this test relies on determining whether the midpoint m_{pq} of two centroids $p, q \in \mathbb{C}^n$ satisfies $\frac{p+q}{2} - m_{pq} = 0$: it requires assessing whether or not a numerical value is zero. Although extremely reliable in practice, this means the trace test does not certify the computation of all witness points. Developing an algorithm to certify the trace task is an important open task in numerical algebraic geometry.

6.5.2 Monodromy solving for real branched covers

Algorithm 6.5.3 is not optimal. For example, suppose the first two loops of Algorithm 6.5.3 are γ and γ' depicted in Figure 6.11. These loops induce a transitive subgroup of S_5 , but Algorithm 6.5.3 would only use each once, discovering a total of three points after the second loop.

A model for monodromy algorithms as well as a strategy-analysis for choosing monodromy loops are given in [58]. We propose improving Algorithm 6.5.3 by taking advantage of automorphisms of fibers guaranteed by the structure of π . For example, many polynomial systems of interest are defined over the real numbers, whose solutions come in complex conjugate pairs (see Section 4.5) and so computing a nonreal solution $x \in \pi^{-1}(s_1)$ immediately computes its conjugate $\bar{x} \in \pi^{-1}(s_1)$. This occurs, in particular, whenever π is a real branched cover and s_1 is real. Thus, we propose an additional step to Algorithm 6.5.3.




Figure 6.11: An example of two monodromy loops generating a transitive subgroup of a monodromy group.

Algorithm 6.5.6 (Monodromy solver for real branched covers).

Input:

- A parametrized polynomial system $F(s; x) \subset \mathbb{C}[s][x]$ such that $\pi: \mathcal{V}(F(s; x)) \rightarrow \mathbb{C}_s^m$ is an irreducible real branched cover
- A single point q in some fiber $\pi^{-1}(s_1)$ with s_1 real

Output:

- All points in $\pi^{-1}(s_1)$

Steps:

- 1 **set** $S_1 = \{q\}$
- 2 **while** $\pi^{-1}(s_1)$ has not been fully computed **do**
 - 2.1 Pick $s_2, \in U$ and generic $c_1, c_2 \in \mathbb{C}$
 - 2.2 Track all points in S_1 along H_{s_1, s_2, c_1} to produce S_2
 - 2.3 Track all points in S_2 along H_{s_2, s_1, c_2} to produce

$$S'_1 = (m_{\tau_{s_2, s_1, c_2}} \circ m_{\tau_{s_1, s_2, c_1}})(S) \subset \pi^{-1}(s_1)$$

- 2.4 **set** $S_1 = S_1 \cup S'_1 \cup \overline{S'_1}$
- 3 **return** S_1

Example 6.5.7. Figure 6.12 depicts a schematic for Algorithm 6.5.6 showing a monodromy loop along with complex conjugation generating a transitive subgroup of S_5 . Algorithm 6.5.6 computes all solutions in three steps using only one monodromy loop along with complex conjugation. \diamond

We remark that step (2.4) in the above algorithm may be replaced by any operation g preserving the fiber $\pi^{-1}(s_1)$. In particular, if g is a deck transformation of the cover π , then one may append the orbit gS_1 to S_1 in step (2.4) at a nominal cost.




Figure 6.12: A schematic describing step (2) of Algorithm 6.5.6.

6.5.3 Expected success of monodromy solving

The authors of [58] address the question of how many monodromy loops are necessary to induce a transitive subgroup of \mathcal{M}_π under the assumptions

- (1) \mathcal{M}_π is the full symmetric group.
- (2) Random choices of $s_2, c_1,$ and c_2 samples elements of \mathcal{M}_π uniformly at random.

For branched covers π of degree d , they prove a generalization of Dixon's theorem [59], a result which implies that the probability of two random elements of S_d generating a transitive subgroup of S_d approaches 1 as $d \rightarrow \infty$.

One may hope for an analogous result with respect to Algorithm 6.5.6. That is, given a branched cover π of degree d and a regular value s_1 whose fiber is known to be fixed under the action ι of complex conjugation, what is the probability that a random monodromy element along with complex conjugation generate a transitive subgroup of $\pi^{-1}(s_1)$?

As in [58], we must make decide how to model the action of complex conjugation on a fiber. A fiber of π whose points are fixed under complex conjugation may consist entirely of real points (in which case Algorithm 6.5.6 is no different than Algorithm 6.5.3) or entirely of nonreal points, or some number in between. Thus, we analyze the case when conjugation on a fiber is modeled by random involutions in S_d and the case that it is modeled by a fixed-point free involution. The latter case is relevant in applications as there are many instances where we can guarantee that a fiber contains no real points (such as computing witness sets for varieties which are compact over \mathbb{R}).

We fix some notation. Let $R_i \subset S_i$ be a subset of the symmetric group with the property that whenever the subgroup generated by $\sigma \in S_i$ and $\tau \in R_i$ has k_i orbits of size i , then $\tau \in (R_1)^{k_1} \times (R_2)^{k_2} \times \cdots \times (R_d)^{k_d}$ where each factor of R_i acts on an orbit of size i . Note that not every sequence of subsets of S_i has this property. For example, any sequence of subsets starting as

$$R_1 = \{(1)\}, \quad R_2 = \{(1)\}, \quad R_3 = \{(1, 2)\}, \dots$$

does not have this property since the permutation $(1, 2)$ has orbits of sizes $(k_1, k_2) = (1, 1)$, but $(1, 2) \notin R_1 \times R_2$.

Proposition 6.5.8. *Let R_i be a sequence of subsets of S_i such that whenever $(\sigma, \tau) \in S_d \times R_d$ has k_i orbits of size i , then $(\sigma, \tau) \in (S_1 \times R_1)^{k_1} \times \cdots \times (S_d \times R_d)^{k_d}$. Let t_d be the probability that $(\sigma, \tau) \in S_d \times R_d$ generates a transitive subgroup of S_d . Then the t_i satisfy the recursion*

$$d|R_d| = \sum_{i=1}^d it_i|R_i| \cdot |R_{d-i}|.$$

Proof. We use the same strategy as [58, 59] to determine a recursion for the probabilities t_d .

Let $K_d = \{\bar{k} \in \mathbb{N}^d \mid \sum ik_i = d\}$ be the set of number partitions of d . The number of set partitions which have parts corresponding to some \bar{k} is $\left(\frac{d!}{\prod_{i=1}^d (i!)^{k_i} k_i!}\right)$. This is since there are $d!$ ways to place the numbers $\{1, \dots, n\}$ into a sequence of cycles of sizes k_1, \dots, k_d , but we have over counted since each cycle can be permuted $i!$ ways, and cycles of the same size may be permuted as well. Let $(\sigma, \tau) \in S_d \times R_d$, then if $\langle \sigma, \tau \rangle$ is a subgroup whose orbits have sizes \bar{k} , then σ, τ must respect these partitions so $(\sigma, \tau) \in (S_1 \times R_1)^{k_1} \times \cdots \times (S_d \times R_d)^{k_d}$. So we may assume σ and τ have been uniformly chosen from $S_1^{k_1} \times \cdots \times S_d^{k_d}$ and $R_1^{k_1} \times \cdots \times R_d^{k_d}$ respectively. Therefore, using the probabilities t_i we may count the elements in $S_d \times R_d$ via

$$\begin{aligned} |S_d \times R_d| &= \sum_{\bar{k} \in K_d} \left(\frac{d!}{\prod_{i=1}^d (i!)^{k_i} k_i!} \right) \prod_{i=1}^d [t_i (i! \cdot |R_i|)]^{k_i} \\ &= d! \sum_{\bar{k} \in K_d} \prod_{i=1}^d \left(\frac{t_i i! |R_i|}{i!} \right)^{k_i} \frac{1}{k_i!} \\ &= d! \sum_{\bar{k} \in K_d} \prod_{i=1}^d \frac{(t_i \cdot |R_i|)^{k_i}}{k_i!} \end{aligned}$$

and since $|S_d \times R_d| = d! \cdot |R_d|$ we have

$$|R_d| = \sum_{\bar{k} \in K_d} \prod_{i=1}^d \frac{(t_i \cdot |R_i|)^{k_i}}{k_i!}. \quad (6.11)$$

Using the theory of generating functions, we extract a recursion on the numbers t_i . Let $\hat{F}(x)$ be the generating function of $F(d) = |R_d|$ and recall the formal identity

$$\exp \left(\sum_{i=1}^{\infty} y_i x^i \right) = \sum_{d=0}^{\infty} x^d \left(\sum_{\bar{k} \in K_d} \prod_{i=1}^d \frac{y_i^{k_i}}{k_i!} \right).$$

Applying this formula to $\hat{F}(x)$ using Equation (6.11) gives us

$$\exp \left(\sum_{i=1}^{\infty} t_i |R_i| x^i \right) = \hat{F}(x).$$

Now consider $\hat{F}'(x)$:

$$\begin{aligned}\sum_{d=1}^{\infty} d|R_d|x^{d-1} &= \frac{\partial}{\partial x}\hat{F}(x) \\ &= \frac{\partial}{\partial x}\exp\left(\sum_{i=1}^{\infty} t_i|R_i|x^i\right).\end{aligned}$$

When we apply the chain rule to differentiate this expression we get

$$\begin{aligned}&= \hat{F}(x) \cdot \left(\sum_{i=1}^{\infty} it_i|R_i|x^{i-1}\right) \\ &= \left[\sum_{d=1}^{\infty} |R_d|x^d\right] \cdot \left[\sum_{i=1}^{\infty} it_i|R_i|x^{i-1}\right] \\ &= \sum_{d=1}^{\infty} \left(\sum_{i=1}^{\infty} it_i|R_i| \cdot |R_d|x^{d+i-1}\right)\end{aligned}$$

making the substitution $d' = d + i$ yields the equality

$$\sum_{d=1}^{\infty} d|R_d|x^{d-1} = \sum_{d'=1}^{\infty} x^{d'-1} \left(\sum_{i=1}^{d'} it_i|R_i| \cdot |R_{d'-i}|\right).$$

Equating the coefficients of x^d gives a recursion

$$d|R_d| = \sum_{i=1}^d it_i|R_i| \cdot |R_{d-i}| \tag{6.12}$$

$$t_d = 1 - \sum_{i=1}^{d-1} \frac{i}{d} t_i \frac{|R_i| \cdot |R_{d-i}|}{|R_d|}, \tag{6.13}$$

completing the proof. □

Even though we followed exactly the same argument used in the results of [58, 59] in our proof of Proposition 6.5.8, we are not aware of this elementary result in the literature. We remark that when R_i is a subgroup of S_i , the analysis of the probabilities that random elements of $R_i \times S_i$ generating either A_i or S_i has been done [60].

As mentioned, there are three natural choices of R_i to analyze:

- (1) $R_i = S_i$
- (2) $R_i = T_i$
- (3) $R_{2i} = \{\text{all fixed point free involutions}\} =: \mathbb{T}_{2i}$.

d	1	2	3	4	5	10	20	30
S_d	1	0.75	0.722	0.739	0.768	0.881	0.946	0.965
T_d	1	0.75	0.583	0.575	0.546	0.607	0.731	0.792
\mathbb{T}_d	-	1	-	0.833	-	0.863	0.937	0.962

Table 6.3: Some probabilities of generating a transitive action by uniformly choosing from $S_d \times R_d$

The first case is the subject of Dixon's theorem. The second case corresponds to choosing a random involution to model complex conjugation on a fiber of the monodromy algorithm. The third case corresponds to modeling complex conjugation on a fiber where every solution is nonreal (in particular d is even). We list the first few terms of the probabilities t_i in each case in Table 6.3.

We remark that the only property of a sequence $\{R_i\}_{i \in \mathbb{N}}$ determining the probabilities t_i are the cardinalities $|R_i|$.

Corollary 6.5.9. *For $d = 2n$, the probability that a fixed point free involution and a random element of S_d generate a transitive subgroup of S_d approaches 1 as $d \rightarrow \infty$.*

Proof. Let $p_d = 1 - t_d$ be the probability that a random element of S_d and a fixed point free involution do *not* generate a transitive subgroup of S_d . Note that if d is odd, then there are no fixed point free involutions. We let a_j be the number of fixed point free involutions on a set of cardinality $2j$ so that

$$a_j = (2j - 1)!! = \frac{(2j)!}{j!2^j}.$$

These may be recursively defined by $a_j = (2j - 1)a_{j-1}$ and $a_1 = 1$.

Set $d = 2n$ so that (6.12) becomes

$$2n \cdot a_n = \sum_{j=1}^n (2j) \cdot t_{2j} \cdot a_j \cdot a_{n-j}$$

which we may rearrange so that

$$t_{2n} = 1 - \sum_{j=1}^{n-1} \frac{2j}{2n} t_{2j} \frac{a_j \cdot a_{n-j}}{a_n} = 1 - \sum_{j=1}^{n-1} \frac{j}{n} t_{2j} \frac{a_j \cdot a_{n-j}}{a_n}.$$

Consequently, to show that $p_d \rightarrow 0$ as $d \rightarrow \infty$ we show that

$$\lim_{n \rightarrow \infty} p_{2n} = \lim_{n \rightarrow \infty} \sum_{j=1}^{n-1} \frac{j}{n} t_{2j} \frac{a_j \cdot a_{n-j}}{a_n} = 0.$$

Let $m = \lfloor \frac{n-1}{2} \rfloor$ and observe that since the t_{2j} are probabilities, they are bounded by 1 so

$$p_d \leq \sum_{j=1}^{n-1} \frac{j}{n} \frac{a_j \cdot a_{n-j}}{a_n}.$$

By symmetry, if n is even, we have

$$p_{2n} \leq \frac{1}{2} \frac{a_n^2}{a_n} + \sum_{j=1}^m \frac{a_j \cdot a_{n-j}}{a_n} \leq \left(\frac{1}{2}\right)^m + \sum_{j=1}^m \frac{a_j \cdot a_{n-j}}{a_n}$$

and so p_{2n} will approach 0 as $n \rightarrow \infty$ if and only if

$$\sum_{j=1}^m \frac{a_j \cdot a_{n-j}}{a_n}$$

does. If n is odd, this bound also holds. Since the a_j satisfy the recursion $a_j = (2j - 1)a_{j-1}$, we know that

$$a_j a_{n-j} = a_{j+1} a_{n-j-1} \frac{2n - 2j}{2j + 1} > a_{j+1} a_{n-j-1}$$

for all $1 \leq j \leq m$. Thus,

$$\begin{aligned} \lim_{n \rightarrow \infty} \sum_{j=1}^m \frac{a_j \cdot a_{n-j}}{a_n} &\leq \left(\lim_{n \rightarrow \infty} \frac{a_1 \cdot a_{n-1}}{a_n} + \sum_{j=2}^m \frac{a_2 \cdot a_{n-2}}{a_n} \right) \\ &= \lim_{n \rightarrow \infty} \left(\frac{1}{2n-1} + \sum_{j=2}^m \frac{3}{(2n-1)(2n-3)} \right) \\ &= \lim_{n \rightarrow \infty} \frac{1}{2n-1} + \frac{3(m-1)}{(2n-1)(2n-3)} = 0 \end{aligned}$$

showing that $\lim_{n \rightarrow \infty} p_{2n} \rightarrow 0$ so $\lim_{n \rightarrow \infty} t_{2n} = 1$. □

7. NEWTON POLYTOPES AND TROPICAL MEMBERSHIP VIA NUMERICAL ALGEBRAIC GEOMETRY

A major theme of numerical algebraic geometry is the extraction of information about a variety X using witness sets. For varieties arising as the image of a map, the algebraic information of generators of the ideal $\mathcal{I}(X)$ may not be readily available. Finding these generators is the problem of implicitization. While this may be done using symbolic methods involving Gröbner bases, this technique is often computationally prohibitive for moderate to large problems. Even when X is a hypersurface, its defining polynomial may be so large that it is not human-readable. Thus, one naturally desires a coarser description of the polynomial, such as its Newton polytope.

In 2012, Hauenstein and Sottile [61] sketched a numerical algorithm which functions as a vertex oracle for the Newton polytope of a hypersurface, relying only on the computation of witness sets. In Section 7.1, we explain how this algorithm, which we call the HS-algorithm, actually functions as a numerical oracle and is therefore stronger than originally anticipated.

Following ideas from Hept and Theobald, we extend the HS-algorithm to an algorithm for computing the tropicalization of an ideal in Section 7.2. In Section 7.3 we analyze the convergence of the HS-algorithm. We discuss our implementation in the **Macaulay2** [62] package **NumericalNP.m2** [63] in Section 7.4 and give large examples showcasing our software in Sections 7.5 and 7.6. Much of this material is contained in the article [1] by the author*.

7.1 The HS-Algorithm

Let $\mathcal{H} \subseteq \mathbb{C}^n$ be a degree d hypersurface defined by

$$f = \sum_{\alpha \in \mathcal{A}} c_{\alpha} x^{\alpha} \in \mathbb{C}[x] \quad c_{\alpha} \neq 0, \mathcal{A} \subseteq \mathbb{N}^n, |\mathcal{A}| < \infty$$

so that $\text{supp}(f) = \mathcal{A}$. Suppose that $a, b \in (\mathbb{C}^{\times})^n$ are general so that the line parametrized by $s \mapsto (a_1 s - b_1, \dots, a_n s - b_n)$ intersects \mathcal{H} at d points, making the map

$$\begin{array}{c} \mathcal{W}(\mathcal{H}) = \{(p_1, \dots, p_n, s) \mid f(p_1(a_1 s - b_1), \dots, p_n(a_n s - b_n)) = 0\} \\ \pi \downarrow \\ \mathbb{C}_p^n \end{array}$$

a degree d branched cover. For any direction $\omega \in \mathbb{R}^n$, the path $t \mapsto (t^{\omega_1}, \dots, t^{\omega_n})$ in \mathbb{C}_p^n corresponds to a family of lines \mathcal{L}_t parametrized by

$$\begin{array}{c} \mathbf{L}_t: \mathbb{C}_s \rightarrow \mathbb{C}^n \\ s \mapsto (t^{\omega_1}(a_1 s - b_1), \dots, t^{\omega_n}(a_n s - b_n)). \end{array}$$

This family of lines lifts to d paths $\{s_i(t)\}_{i=1}^d$ in $\mathcal{W}(\mathcal{H})$ via π , each corresponding to an intersection point of \mathcal{L}_t and \mathcal{H} . In other words, these paths comprise the data of witness points and so π essentially functions as a witness cover.

*Reprinted with permission from T. Brysiewicz, “Numerical Software to Compute Newton polytopes and Tropical Membership,” *Mathematics in Computer Science*, 2020. Copyright 2020 by Springer Nature.

Each path $s_i(t)$ may be tracked numerically with respect to the homotopy

$$\begin{aligned} f(\mathbf{L}_t(s)) &: (0, 1] \times \mathbb{C}_s \rightarrow \mathbb{C} \\ (t^{-1}, s) &\mapsto f(\mathbf{L}_t(s)) \end{aligned}$$

in the variables t^{-1} and s . We remark that the use of t^{-1} is a definitional technicality and that we will mostly work in t using the map

$$\begin{aligned} \pi_\omega &: f(\mathbf{L}_t(s))^{-1}(0) \rightarrow (1, \infty) \\ (t^{-1}, s) &\mapsto t \end{aligned} \tag{7.1}$$

so that the fiber of π_ω over t is identified with $\{s_i(t)\}_{i=1}^d$. The following lemma is the basis of the HS-algorithm.

Lemma 7.1.1. *As $t \rightarrow \infty$ the s -coordinates of the fibers $\pi_\omega^{-1}(t)$ converge to the solutions of $f_\omega(\mathbf{L}_1(s))$.*

Proof. Using the notation $(as - b)^\alpha = (a_1s - b_1)^{\alpha_1} \dots (a_ns - b_n)^{\alpha_n}$, observe that

$$\begin{aligned} f(\mathbf{L}_t) &= \sum_{\alpha \in \mathcal{A}} c_\alpha [t^{\omega_1}(a_1s - b_1)]^{\alpha_1} \dots [t^{\omega_n}(a_ns - b_n)]^{\alpha_n} \\ &= \sum_{\alpha \in \mathcal{A}} t^{\langle \omega, \alpha \rangle} c_\alpha (as - b)^\alpha \\ &= \sum_{\alpha \in \mathcal{A}_\omega} t^{h_{\mathcal{A}}(\omega)} c_\alpha (as - b)^\alpha + \sum_{\alpha \in \mathcal{A}_\omega^c} t^{\langle \omega, \alpha \rangle} c_\alpha (as - b)^\alpha \end{aligned} \tag{7.2}$$

Since t is not zero, we may scale (7.2) by $t^{-h_{\mathcal{A}}(\omega)}$ without changing its zeros. Thus the solutions of (7.2) are the same as those of

$$\sum_{\alpha \in \mathcal{A}_\omega} c_\alpha (as - b)^\alpha + \sum_{\alpha \in \mathcal{A}_\omega^c} t^{\langle \omega, \alpha \rangle - h_{\mathcal{A}}(\omega)} c_\alpha (as - b)^\alpha \tag{7.3}$$

where \mathcal{A}_ω^c is the complement of \mathcal{A}_ω in \mathcal{A} . Note that $\pi_\omega^{-1}(t) = \{(t, s_i(t))\}_{i=1}^d$ where $\{s_i(t)\}_{i=1}^d$ are the solutions of (7.3). Finding the values of each $s_i(t)$ as $t \rightarrow \infty$ is the same as finding the values of $s_i(t^{-1})$ as $t \rightarrow 0$ and since these paths are continuous, we substitute t^{-1} for t in (7.3) and take the limit as $t \rightarrow 0$:

$$\begin{aligned} &\sum_{\alpha \in \mathcal{A}_\omega} c_\alpha (as - b)^\alpha + \sum_{\alpha \in \mathcal{A}_\omega^c} (t^{-1})^{\langle \omega, \alpha \rangle - h_{\mathcal{A}}(\omega)} c_\alpha (as - b)^\alpha \\ &\sum_{\alpha \in \mathcal{A}_\omega} c_\alpha (as - b)^\alpha + \sum_{\alpha \in \mathcal{A}_\omega^c} (t)^{-\langle \omega, \alpha \rangle + h_{\mathcal{A}}(\omega)} c_\alpha (as - b)^\alpha. \end{aligned} \tag{7.4}$$

Note that $\langle -\omega, \alpha \rangle + h_{\mathcal{A}}(\omega) > 0$ for all $\alpha \in \mathcal{A}_\omega^c$ by definition, and so evaluating (7.4) at $t = 0$ gives $\mathcal{V}(\sum_{\alpha \in \mathcal{A}_\omega} c_\alpha (as - b)^\alpha) = \mathcal{V}(f_\omega(\mathbf{L}_1(s)))$. \square

Suppose a hypersurface $\mathcal{H} \subset \mathbb{C}^m$ is the image of a map $\varphi: X \rightarrow \mathcal{H}$. Recall that a (pseudo)-witness set for \mathcal{H} can be computed by computing a witness set for the graph of φ and applying Algorithm 6.4.6. Consequently, since fibers of π_ω are essentially witness points, we may compute them without access to the defining equation for \mathcal{H} . The HS-algorithm follows from the above observations.

We remind the reader that \tilde{f} denotes the homogenization of a polynomial f and \mathcal{O}_P denotes a numerical oracle for a polytope P .

Algorithm 7.1.2 (HS-Algorithm).

Input:

- A witness set, or pseudo-witness set, W for a hypersurface $\mathcal{H} \subseteq \mathbb{C}^n$
- A direction $\omega \in \mathbb{R}^n$

Output:

- $\mathcal{O}_{\text{New}(\tilde{f})}(\omega)$ where $\mathcal{H} = \mathcal{V}(f)$

Steps:

- 1 Pick random $a, b \in (\mathbb{C}^\times)^n$ and construct \mathbf{L}_t
- 2 Track the witness points in W to the intersection $\mathcal{H} \cap \mathcal{L}_1$
- 3 Track all points $\{s_i(1)\}_{i=1}^d$ along (7.1) from $t = 1 \rightarrow 2$.
- 4 If none of the points move, **return** EEP
- 5 Initialize $\beta = (0_1, 0_2, \dots, 0_n, 0_\infty) \in \mathbb{N}^{n+1}$
- 6 **for** i from 1 to d **do**
 - 6.1 Track the point $s_i(1)$ along (7.1) as $t \rightarrow \infty$
 - 6.2 If $s_i(t)$ converges or diverges, stop tracking it
 - 6.2.1 If $s_i(t)$ converged to ρ_i , increment β_i by one
 - 6.2.2 If $s_i(t)$ diverged, increment β_∞ by one
- 7 **return** β

Proof of correctness: We claim that Algorithm 7.1.2 is a numerical oracle (see Definition 2.2.6) for $\text{New}(f)$. We consider three situations, dependent on $\omega \in \mathbb{R}^n$, which result in different behaviors of the set $\{s_i(t)\}_{i=1}^d$ as $t \rightarrow \infty$.

- (1) (ω **exposes a single point**): This means that $f_\omega(\mathbf{L}_1(s)) = c_\beta(as - b)^\beta$ is a monomial with exponent $\beta = (\beta_1, \dots, \beta_n) \in \mathcal{A}$. This clearly has roots of $\rho_i = b_i/a_i$ appearing with multiplicity β_i . Note that if $|\beta| < d$ then there are $\beta_\infty = d - |\beta|$ paths which diverge as $t \rightarrow \infty$. One way to see this is to observe that β_∞ is the exponent of the homogenizing variable in the term \tilde{f}_ω .
- (2) (ω **exposes \mathcal{A}**): If ω exposes the entire polytope defined by \mathcal{A} , then the roots $\{s_i(t)\}_{i=1}^d$ remain constant as t varies since $f(\mathbf{L}_t) = t^{h_{\mathcal{A}}(\omega)} f(\mathbf{L}_1)$.
- (3) (ω **exposes a proper subset of \mathcal{A} consisting of more than one point**): If ω exposes a proper non-singleton subset of \mathcal{A} , then there is more than one term in f_ω . We remark this happens exactly when $\omega \in \text{trop}(\mathcal{V}(f))$. The terms of f_ω will have a common factor of $\prod_{i=1}^n (a_i s - b_i)^{m_i}$ where the vector m is the coordinate-wise minimum of the points in \mathcal{A}_ω .

Therefore, m_i roots will converge to ρ_i and $m_\infty = \min_{\beta \in \mathcal{A}_\omega} (d - |\beta|)$ points will diverge. All other roots will converge somewhere else in \mathbb{C} .

In each case, the output is that of a numerical oracle. □

7.2 Tropical membership

A direct consequence of the HS-algorithm is a tropical membership algorithm for hypersurfaces. Recall that for a polynomial $f \in \mathbb{C}[x]$ of degree d , a direction $\omega \in \mathbb{R}^n$ is an element of $\text{trop}(f)$ if and only if ω exposes a positive-dimensional face of $\text{New}(\tilde{f})$. Equivalently, recall that $\omega \in \text{trop}(f)$ if and only if a numerical oracle outputs a vector $v = \mathcal{O}_{\text{New}(\tilde{f})}(\omega) \in \mathbb{Z}^{n+1}$ satisfying $|v| < d$. Thus, since the HS-algorithm functions as a numerical oracle, it may be used as a tropical membership algorithm for hypersurfaces.

Algorithm 7.2.1 (Tropical Membership for Hypersurfaces).

Input:

- A witness set, or pseudo-witness set, W for a hypersurface $\mathcal{H} \subseteq \mathbb{C}^n$
- A direction $\omega \in \mathbb{R}^n$

Output:

- `true` if $\omega \in \text{trop}(\mathcal{H})$ and `false` otherwise.

Steps:

```

0 set  $d = \deg(\mathcal{H})$  and set  $\beta$  to be the output of the HS-algorithm on input  $W$  and  $\omega$ 
1 if  $|\beta| < d$  then return true, else return false

```

Given an arbitrary variety $\mathcal{V}(I) \subset \mathbb{C}^n$, recall that the tropicalization of $\mathcal{V}(I)$ may be realized as the intersection of preimages of projections of $\text{trop}(\mathcal{V}(I))$ (Lemma 5.2.1). When the coordinate projections are sufficiently generic, Algorithm 7.2.1 extends immediately to a tropical membership algorithm for the tropicalization of $\mathcal{V}(I)$. When they are not, this algorithm can only yield false positives. To handle this, we may obtain new projections of $\text{trop}(\mathcal{V}(I))$ by taking the coordinate projections of $\text{trop}(\mathcal{V}(I))$ after a linear change of coordinates on $\text{trop}(\mathcal{V}(I))$. We recall that a linear change of coordinates on $\text{trop}(\mathcal{V}(I))$ amounts to a monomial change of coordinates on I (see Remark 5.2.4) which will likely increase the degree of $\mathcal{V}(I)$. While this makes the computation of a witness set more difficult, it is often still manageable.

Algorithm 7.2.2 (Tropical Membership).**Input:**

- An m -dimensional variety $X = \mathcal{V}(I) \subseteq \mathbb{C}^n$
- A direction $\omega \in \mathbb{R}^n$

Output:

- `true` if $\omega \in \text{trop}(\mathcal{H})$ and `false` otherwise

Steps:

- 1 Replace I with its image under a generic monomial map Φ so that the coordinate projections of $\mathcal{V}(I)$ are generic
- 2 Replace ω with $\varphi^{-1}\omega$ where $\varphi = \Phi^*$
- 3 Compute a witness set W for X .
- 4 **for** each coordinate projection $\{\pi_J\}_{J \subseteq [n]}$ with $|J| = n - m - 1$ **do**
 - 4.1 Compute a pseudo-witness set W_J for $\pi_J(X)$
 - 4.2 **if** Algorithm 7.2.1 returns `false` on input $(W_J, \pi_J(\omega))$, **then** STOP and **return** `false`
- 5 **return** `true`

7.3 Convergence rates of the HS-Algorithm

Theorem 8 of [61] gives an analysis of the convergence of the HS-algorithm whenever ω exposes a vertex. We generalize this result to include the case where $\omega \in \text{trop}(\mathcal{V}(f))$. First, we introduce some notation. As before, let

$$f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha$$

be a polynomial with support \mathcal{A} and let $\omega \in \mathbb{R}^n$. The polynomial f_ω may be written as $f_\omega = x^m \cdot g(x)$ for some polynomial $g(x) \in \mathbb{C}[x]$ whose terms have no common monomial factor. After choosing generic points $a, b \in \mathbb{C}^n$, we write $f_\omega(\mathbf{L}_1)$ in its factored form as

$$f_\omega(\mathbf{L}_1) = (as - b)^m g(\mathbf{L}_1) = (as - b)^m \cdot K \cdot (s - \tau)^k.$$

Note that the $\tau = (\tau_1, \dots, \tau_{n'})$ are the n' complex roots of $g(\mathbf{L}_t)$ and so k is some n' -tuple satisfying $|k| < d - |m|$. In the case that $\text{New}(f)_\omega = \beta$ is a vertex, we have $m = \beta$, $g(x) = 1$, and $K = c_\beta$. We define the following constants based on the coefficients c_α , the support \mathcal{A} , and the constant K .

$$C = \max \left\{ \frac{|c_\alpha|}{|K|} \mid \alpha \in \mathcal{A} \right\}, \quad d_\omega = h_{\mathcal{A}}(\omega) - h_{\mathcal{A}^c}(\omega)$$

$$a_{\min} = \min\{1, |a_i| \mid i = 1, \dots, n\}, \quad a_{\max} = \max\{1, |a_i| \mid i = 1, \dots, n\}$$

Finally, based on the positions of the $\rho_i = \frac{b_i}{a_i}$ and the τ_j appearing as roots of $g(\mathbf{L}_1)$ we define the following constants for any $z \in \{\rho_i\}_{i=1}^n \cup \{\tau_j\}_{j=1}^{n'}$.

$$\gamma_z = \min \left\{ a_{\min}, \frac{|z - \hat{z}|}{2} \mid \hat{z} \in \{\rho_i\}_{i=1}^n \cup \{\tau_j\}_{j=1}^{n'} \setminus z \right\}$$

$$\Gamma_z = \max \left\{ \frac{2}{a_{\max}}, |z - \rho_i| \mid i = 1, \dots, n \right\}$$

The constant d_ω describes how close ω is to exposing a positive-dimensional face of $\text{New}(f)$. The constant γ_z is defined so that any point inside the circle of radius γ_z centered at z is closer to z than any other point in $\{\rho_i\}_{i=1}^{n'} \cup \{\tau_j\}_{j=1}^{n'}$. We include helpful graphics describing this notation in Figure 7.1 and Figure 7.2.




Figure 7.1: An example of locations of $\{\rho_i\}_{i=1}^4$ and $\{\tau_j\}_{j=1}^3$ in \mathbb{C}_s . The smaller circle has radius γ_{τ_1} and the larger circle has radius Γ_{τ_1} .

Theorem 7.3.1. *Suppose $\omega \in \mathbb{R}^n$. Let $s(t)$ be a path of the HS-algorithm converging to z as $t \rightarrow \infty$ and let β be the number of such paths converging to z . Let $t_1 \geq 0$ be a number such that if $t > t_1$ then $|s(t) - z| \leq \gamma_z$. Then for all $t > t_1$*

$$|s(t) - z|^\beta \leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot \left(\frac{a_{\max}}{a_{\min}} \left(1 + \frac{\Gamma_z}{\gamma_z} \right) \right)^d.$$

Proof. Recall from (7.3) we have

$$t^{-h_{\mathcal{A}}(\omega)} \cdot f(\mathbf{L}_t) = \sum_{\alpha \in \mathcal{A}_\omega} c_\alpha (as - b)^\alpha + \sum_{\alpha \in \mathcal{A}_\omega^c} t^{(\omega, \alpha) - h_{\mathcal{A}}(\omega)} c_\alpha (as - b)^\alpha. \quad (7.5)$$

Suppose $s(t): (1, \infty) \rightarrow \mathbb{C}_s$ is a continuous path in (7.1) so that $f(\mathbf{L}_t(s(t))) = 0$ for all $t > 1$. Then (7.5) gives

$$|f_\omega(as(t) - b)| = \left| \sum_{\alpha \in \mathcal{A}_\omega^c} t^{(\omega, \alpha) - h_{\mathcal{A}}(\omega)} c_\alpha (as(t) - b)^\alpha \right| \quad (7.6)$$

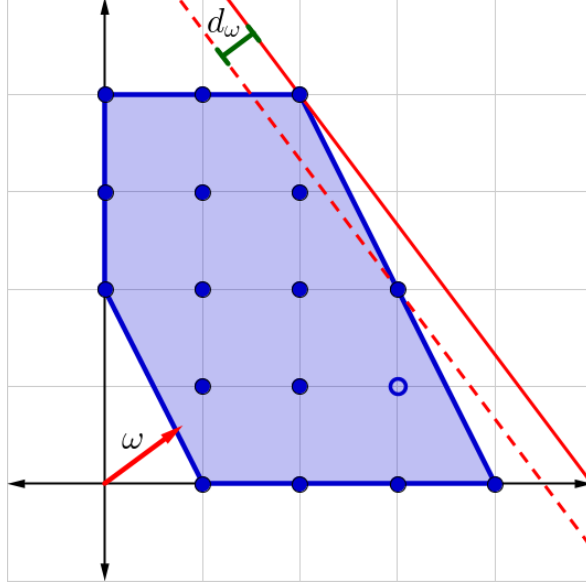


Figure 7.2: A unit vector ω and a geometric description of d_ω .

and so after dividing through by K and extracting the largest power of t from the sum,

$$|(as - b)^m \cdot (s - \tau)^k| \leq t^{-d_\omega} \cdot \sum_{\alpha \in \mathcal{A}_\omega^c} \left| \frac{c_\alpha}{K} \right| \cdot |(as(t) - b)^\alpha| \quad (7.7)$$

$$\leq t^{-d_\omega} \cdot C \cdot \sum_{\alpha \in \mathcal{A}_\omega^c} |(as(t) - b)^\alpha|. \quad (7.8)$$

Recalling that $|s(t) - z| \leq \gamma_z$ by hypothesis, we bound the right-hand summands,

$$\begin{aligned} |a_j s(t) - b_j| &= |a_j| \cdot |s(t) - \rho_j| \leq a_{\max} \cdot |s(t) - z + z - \rho_j| \\ &\leq a_{\max} \cdot (|s(t) - z| + |z - \rho_j|) \\ &\leq a_{\max} \cdot (\gamma_z + \Gamma_z) \end{aligned}$$

and so since $2 \leq a_{\max} \Gamma_z$,

$$|(as(t) - b)^\alpha| \leq (a_{\max}(\gamma_z + \Gamma_z))^{|\alpha|} \leq (a_{\max}(\gamma_z + \Gamma_z))^d. \quad (7.9)$$

Substituting (7.9) into (7.7) gives

$$|(as - b)^m \cdot (s - \tau)^k| \leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot (a_{\max} \cdot (\gamma_z + \Gamma_z))^d.$$

We now bound the factors on the left-hand-side of (7.7),

$$\begin{aligned} |s(t)a_j - b_j| &= |a_j| \cdot |s(t) - \rho_j| = |a_j| \cdot |s(t) - z + z - \rho_j| \geq a_{\min} \left| |z - \rho_j| - |s(t) - z| \right| \\ &\geq a_{\min} \cdot (2\gamma_z - \gamma_z) = a_{\min} \gamma_z. \end{aligned}$$

Similarly,

$$\begin{aligned} |s(t) - \tau_j| &= |s(t) - z + z - \tau_j| \geq \left| |z - \tau_j| - |s(t) - z| \right| \\ &\geq 2\gamma_z - \gamma_z = \gamma_z \geq a_{\min}\gamma_z \end{aligned}$$

and since $a_{\min}\gamma_z \leq 1$ and $|m| + |k| \leq d$ we have that

$$|(as(t) - b)^m (s(t) - \tau)^k| \geq (a_{\min}\gamma_z)^d.$$

So for either $z = \tau_j$ or $z = \rho_i$ we have

$$\begin{aligned} \left| \frac{(as(t) - b)^m (s(t) - \tau)^k}{(s(t) - \tau_j)^{k_j}} \right| &\geq (a_{\min}\gamma_{\tau_j})^{d-k_j}, \\ \left| \frac{(as(t) - b)^m (s(t) - \tau)^k}{(a_i s(t) - b_i)^{m_i}} \right| &\geq (a_{\min}\gamma_{\rho_i})^{d-m_i}, \end{aligned}$$

respectively.

We now suppose that $z = \rho_i$ and essentially recover Theorem 8 of [61]. Note that,

$$|s(t) - \rho_i|^{m_i} = \frac{1}{|a_i|^{m_i}} \cdot \left| \frac{(as(t) - b)^m (s(t) - \tau)^k}{\left(\prod_{j \neq i} (a_j s(t) - b_j)^{m_j}\right) \cdot (s(t) - \tau)^k} \right|$$

and so putting our bounds together gives

$$|s(t) - \rho_i|^{m_i} \leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot (a_{\max} \cdot (\gamma_{\rho_i} + \Gamma_{\rho_i}))^d \cdot \frac{1}{a_{\min}^{m_i}} \cdot \frac{1}{(a_{\min}\gamma_{\rho_i})^{d-m_i}}.$$

Recall that $1 \geq a_{\min} \geq \gamma_{\rho_i}$ so

$$|s(t) - \rho_i|^{m_i} \leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot \left(\frac{a_{\max}}{a_{\min}} \left(1 + \frac{\Gamma_{\rho_i}}{\gamma_{\rho_i}} \right) \right)^d.$$

On the other hand, if $z = \tau_j$, we have

$$|s(t) - \tau_j|^{k_j} \leq \frac{(as(t) - b)^m (s(t) - \tau)^k}{\left(\prod_{j \neq i} (s(t) - \tau_j)^{k_j}\right) \cdot (as(t) - b)^m}$$

and so putting our bounds together gives

$$|s(t) - \tau_j|^{k_j} \leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot (a_{\max} \cdot (\gamma_{\tau_j} + \Gamma_{\tau_j}))^d \cdot \frac{1}{(a_{\min}\gamma_{\tau_j})^{d-k_j}}.$$

Since $1 \geq a_{\min} \geq \gamma_{\tau_j}$, we obtain

$$\begin{aligned} |s(t) - \tau_j|^{k_j} &\leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot (a_{\max} \cdot (\gamma_{\tau_j} + \Gamma_{\tau_j}))^d \cdot \frac{1}{(a_{\min}\gamma_{\tau_j})^d} \\ &\leq t^{-d_\omega} \cdot C \cdot |\mathcal{A}_\omega^c| \cdot \left(\frac{a_{\max}}{a_{\min}} \left(1 + \frac{\Gamma_{\tau_j}}{\gamma_{\tau_j}} \right) \right)^d, \end{aligned}$$

completing the proof. □

We give an example displaying the convergence rates in Theorem 7.3.1.

Example 7.3.2. Consider the plane curve (hypersurface) given by $\mathcal{V}(f) \subset \mathbb{C}^2$ where

$$f = x + 20x^2 - 4x^3 + x^4 - 4xy + 10x^2y + y^2 + 8xy^2 + 4x^2y^2 + x^3y^2 - 4y^3 - 6xy^3 + 4x^2y^3 + 4y^4 - 4xy^4 + x^2y^4.$$

Its Newton polytope and tropicalization are displayed in Figure 7.3. Note that the only lattice point




Figure 7.3: The Newton polytope and tropicalization of a hypersurface.

of the Newton polytope not appearing in the support of f is the point $(3, 1)$. Figure 7.4 displays the convergence rates in Theorem 7.3.1 as follows. For a uniform sample of unit vectors $\omega \in S^2 \subset \mathbb{R}^2$, we draw a ray in the direction of ω with length equal to the minimum of 1 and $\frac{1}{d_\omega}$, the exponent appearing in Theorem 7.3.1. We remark that setting the length of the rays to be the minimum of $d^{-\omega}$ and 1 models the feature that when this algorithm is used in practice, the user must specify a tolerance describing how far to track t to see convergence. We also point out that the ridges indicated in Figure 7.4 occur because d_ω depends not only on the vertices of $\text{New}(f)$ but also on the monomials in $\text{supp}(f)$. \diamond

Example 7.3.2 shows that in practice, the numerical oracle for the Newton polytope of a hypersurface coming from the HS-algorithm comes with a cost associated to inputs near the tropicalization of the hypersurface: as the input of the HS-algorithm approaches the tropical hypersurface, the convergence rate becomes arbitrarily slow. Due to this feature, pairing Algorithm 7.1.2 with Algorithm 2.2.3 may be too computationally expensive for computing large Newton polytopes.

Remark 7.3.3. Figure 7.4 exposes an important drawback of the algorithms involved in this dissertation: many of our algorithms require a blind random choice of parameters avoiding some forbidden set of measure zero (in this case, the tropical variety).

The first issue with this is that choosing parameters *near* the forbidden set can cause computations to take arbitrarily long. Consequently, the true space of parameters which we want to avoid in our computations has positive measure. In the case of Figure 7.4, this is the set of directions which

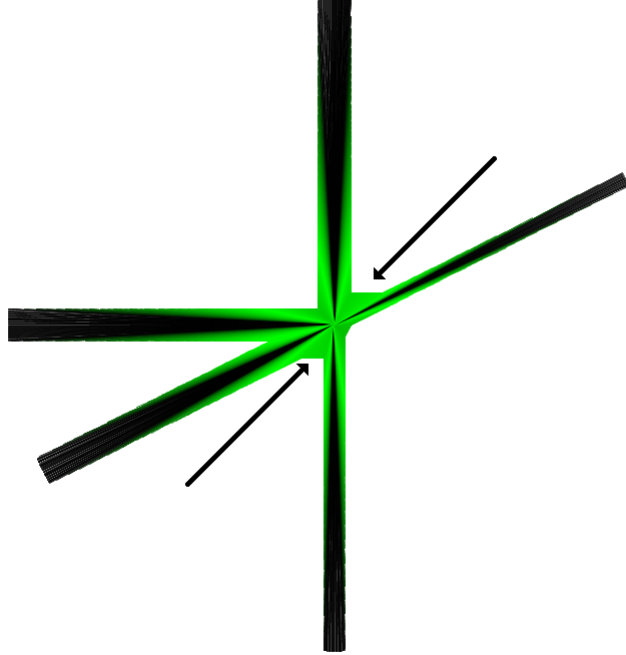


Figure 7.4: For directions $\omega \in S^1$ we draw the ray in direction ω with length $\min(1, d^{-\omega})$ describing the convergence rate proven in Theorem 7.3.1.

correspond to a black ray. Another example is the choice of γ in Lemma 6.2.2 and Lemma 6.2.3. If γ is chosen near the set of measure zero, then the condition number involved in path tracking can become large. This will either cause an instance of path-jumping, or if one is using adaptive precision, can cause the computation to take arbitrarily long.

The second issue is that our computations inherently work over a subset of rational numbers with bounded height. This technically causes problems with observations such as Remark 2.2.5 where the set of directions which do not expose a vertex form a finite subset of the finite set of rational numbers with bounded height; that is, a set of positive measure.

Nonetheless, even with positive measure, the forbidden sets involved in our computations remain heuristically small and in practice the algorithms remain effective. \diamond

7.4 Implementation of Algorithms 7.1.2 and 7.2.2

We describe our implementation of Algorithm 7.1.2 and Algorithm 7.2.2 along with the relevant supporting functions in our **Macaulay2** package **NumericalNP.m2** [63]. This package contains four main user functions, the first three of which implement the HS-algorithm and the last implements the tropical membership algorithm. All numerical computations are piped to **Bertini** [51] through the package **Bertini.m2** [64].

Function 7.4.1 computes a pseudo-witness set for the image of a variety $X \subseteq \mathbb{C}^N$ under a projection $\pi: \mathbb{C}^N \rightarrow \mathbb{C}^n$.

Function 7.4.1. `witnessForProjection`**Input:**

- `I`: Ideal defining $X \subseteq \mathbb{C}^N$
- `ProjCoord`: List of coordinates which are forgotten by π
- `OracleLocation` (option): Path in which to create witness files

Output:

- A subdirectory `/OracleLocation/WitnessSet` containing
 - `witnessPointsForProj`: Preimages of witness points of $\overline{\pi(X)}$
 - `projectionFile`: List of coordinates in `ProjCoord`
 - `equations`: List of equations defining $X' \subseteq X$ such that $\pi|_{X'}$ is generically finite and that $\overline{\pi(X')} = \overline{\pi(X)}$

Given a hypersurface \mathcal{H} , Function 7.4.2, `witnessToOracle`, creates all necessary **Bertini** files to track the witness set $\mathcal{H} \cap \mathcal{L}_t$ as $t \rightarrow \infty$ for any $\omega \in \mathbb{R}^n$. These files treat ω as a parameter so that the user who wants to query many directions needs only to produce these files once.

Function 7.4.2. `witnessToOracle`**Input:**

- `OracleLocation`: Path containing the directory `/WitnessSet`

Optional Input:

- `PointChoice`: Prescribes a and b explicitly (see Algorithm 7.1.2)
- `TargetChoice`: Prescribes targets b_i/a_i
- `NPConfigs`: List of **Bertini** path tracking configurations

Output:

- A subdirectory `/OracleLocation/Oracle` containing all necessary files to run the homotopy described in Algorithm 7.1.2.

Function 7.4.2, by default, chooses $a, b \in \mathbb{C}^n$ such that $\rho_i = a_i/b_i$ are the n -th roots of unity. One may choose to either specify a and b (`PointChoice`), or $\rho_i = a_i/b_i$ (`TargetChoice`) or request that these choices are random. When random, the function ensures that the points ρ_i are far from each other so that convergence to ρ_i is easily distinguished from convergence to ρ_j . **Bertini** is called to track the points in `/OracleLocation/WitnessSet` to the points $\overline{\pi(X)} \cap \mathcal{L}_1$. These become start solutions of the homotopy described in Algorithm 7.1.2 with parameters ω and t . There are many numerical choices for **Bertini**'s native path-tracking algorithms which can be specified via `NPConfigs`.

The fundamental function of **NumericalNP.m2** is `oracleQuery`. It runs the homotopy described in the HS-algorithm on a hypersurface $\mathcal{H} = \mathcal{V}(f)$, monitors convergence, and outputs the result of the numerical oracle.

Function 7.4.3. `oracleQuery`**Input:**

- `OracleLocation` (Option): Location containing the directory `/Oracle`
- ω : A vector in \mathbb{R}^n

Optional Input:

- `Certainty` • `Epsilon` • `MinTracks` • `MaxTracks` • `StepResolution`
- `MakeSageFile`

Output:

- $\mathcal{O}_{\text{New}(\tilde{f})}(\omega)$ or Reached `MaxTracks`
- A subdirectory `/OracleLocation/OracleCalls/Call#` containing
 - `SageFile`: Sage code animating the paths $s(t)$
 - `OracleCallSummary`: a human-readable file summarizing the results

To monitor convergence of points $s(t)$ the software tracks $t \rightarrow \infty$ in discrete steps. The option `StepResolution` specifies these step sizes. In each step and for each path $s_i(t)$, a numerical derivative is computed to heuristically determine convergence or divergence of the solution. If the solution is large and the numerical derivative exceeds $10^{\text{Certainty}}$ in two consecutive steps, then the path is declared to diverge, and if the numerical derivative is below $10^{-\text{Certainty}}$ in two consecutive steps, then the point is declared to converge. If a converged point is at most `Epsilon` from some ρ_i , then the software deems that it has converged to ρ_i . When a point is declared to converge or diverge, it is not tracked further. The option `MaxTracks` allows the user to specify




Figure 7.5: (Reprinted from [1]) Left: Values of t (magnitude of rays) such that `queryOracle` finishes for different ω (direction of rays) on a hypersurface with Newton polytope (center) and normal fan (right).

how long to wait for convergence of the paths $s(t)$.

Example 7.4.4. Figure 7.5 shows the Newton polytope of the same plane sextic as in Example 7.3.2. It also shows the convergence rate of the algorithm on different directions $\omega \in S^1$. The length of each green ray is proportional to the number of steps required for `oracleQuery` to finish and the black rays indicate that this convergence was not observed within the limit specified

by `MaxTracks`. We note that the striking resemblance of Figure 7.4 and 7.5 indicates that the value of t at which our implementation recognizes convergence is approximately proportional to the convergence rate we prove in Theorem 7.3.1. We include the image of the tropicalization of this curve to illustrate how the convergence rate involved in the HS-algorithm slows as ω approaches directions in the tropical variety. Nonetheless, we remind the reader that this slow convergence rate does not occur when ω is in the tropical variety. \diamond

One may also specify `MinTracks` which indicates the step at which convergence begins to be monitored. The option to create a Sage [65] animation (see Figure 7.6) of the solution paths helps the user recognize pathological behavior in the numerical computations and fine-tune parameters such as `Certainty`, `StepResolution`, or `Epsilon` accordingly.

Example 7.4.5. Consider the curve in $X \subseteq \mathbb{C}^3$ defined by

$$I = \langle xyt - (x - y - t)^2 + 3x + t, x + y^2 + t^2 \rangle \subseteq \mathbb{C}[x, y, t]$$

and let π be the projection forgetting the t coordinate. The following code written in **Macaulay2** computes a witness set for $\mathcal{C} = \overline{\pi(X)}$, prepares oracle files for the HS-algorithm and then runs the HS-algorithm in the direction $(3, 2)$. The software returns the list $\{2, 4, 0\}$ indicating that $\text{New}(\overline{\pi(X)})_{(3,2)} = (2, 4)$.

```
i1: loadPackage("NumericalNP");
i2: R=CC[x,y,t];
i3: I=ideal(x*y*t-(x-y-t)^2+3*x+t,x+y^2+t^2);
i4: witnessForProjection(I,{2},OracleLocation=>"Example");
i5: witnessToOracle("Example");
i6: time oracleQuery({3,2},OracleLocation=>"Example",MakeSageFile=>true)
    -- used 0.178448 seconds
o6: {2,4,0}
```

The equation of $\pi(X)$ is the polynomial in Example 7.3.2 and so its Newton polytope is displayed in Figure 7.5. Snapshots of the Sage animation created by `queryOracle` are shown in Figure 7.6. There, the circles are centered at $\rho_1 = 1$ and $\rho_2 = -1$ and have radius `epsilon`. The first image shows the intersections (in the s -coordinates) of the sextic $\overline{\pi(X)}$ with \mathcal{L}_1 in the complex plane \mathbb{C}_s . The second image is a snapshot showing two points converging to $s = 1$ and the third image shows four other points converging to $s = -1$. \diamond

Given an ideal I , the fourth function `tropicalMembership` computes a pseudo-witness set for each coordinate projection $\pi(\mathcal{V}(I))$ whose image is a hypersurface. The algorithm subsequently checks that `oracleQuery` indicates that $\pi(\omega) \in \text{trop}(\pi(\mathcal{V}(I)))$. If this is true for each coordinate projection, the algorithm returns `true` and otherwise returns `false`. The numerical options fed to `tropicalMembership` are passed along to `oracleQuery`.




Figure 7.6: Three snapshots of Sage animation from example with viewing window $[-4, 4]^2$

Function 7.4.6. `tropicalMembership`

Input:

- I : Ideal defining $X \subseteq \mathbb{C}^n$
- ω : A vector in \mathbb{R}^n

Optional Input:

- `Certainty` • `Epsilon` • `MinTracks` • `MaxTracks` • `StepResolution`
- `MakeSageFile`

Output:

- A list of oracle queries of $\pi(X)$ in directions $\pi(\omega)$ where π runs through all coordinate projections such that $\pi(X)$ is a hypersurface.
- `true` if all oracle queries exposed positive-dimensional faces and `false` otherwise

Example 7.4.7. We return to Example 5.2.3 of two tropical space curves which are different, yet have the same tropicalized coordinate projections. We depict these tropical curves again in Figure 7.7 and illustrate their behavior with our software.

```
i1 : loadPackage("NumericalNP");
i2 : R=QQ[x,y,z];
i3 : I_1=ideal {x*z+4*y*z-z^2+3*x-12*y+5*z, x*y-4*y^2+y*z+x+2*y-z};
i4 : I_2=ideal{x*y-3*x*z+3*y*z-1, 3*x*z^2-12*y*z^2+x*z+4*y*z+5*z-1};
i5 : I_1==I_2
o5 = false
i6 : directions:={{1,1,1},{1,1,-1},{1,-1,1},
{1,-1,-1},{-1,1,1},{-1,1,-1},{-1,-1,1},{-1,-1,-1}};
i7 : apply(directions,d->tropicalMembership(I_2,d))
o7 = {true, true, true, true, true, true, true, true}
i8 : apply(directions,d->tropicalMembership(I_1,d))
o8 = {true, true, true, true, true, true, true, true}
```

◇




Figure 7.7: (Reprinted from [1]) Two tropical space curves with the same tropical coordinate projections

Every projection of every vertex of $\text{cube}(3)$ is in the tropicalization of the corresponding projection of $\mathcal{V}(I_1)$ and $\mathcal{V}(I_2)$. Nonetheless, the tropicalizations of $\mathcal{V}(I_1)$ and $\mathcal{V}(I_2)$ are disjoint subsets of the vertices of the cube, exemplifying that an output of `tropicalMembership` is not a certification of membership in the tropical variety. Unfortunately, we cannot *a priori* decide whether or not our coordinate projections are generic.

Example 7.4.7 (continued). Consider the monomial change of coordinates Φ given by

$$\Phi(x) = xyz, \quad \Phi(y) = y, \quad \text{and} \quad \Phi(z) = z,$$

and let $\Phi^* = \varphi : \mathbb{Z}^3 \rightarrow \mathbb{Z}^3$ be the linear map corresponding to Φ . Let F and G be the generators used in the above code of I_1 and I_2 respectively. By Equation (5.8) of Section 5.2 we have that

$$\text{trop}(\mathcal{V}(F)) = \varphi(\text{trop}(\mathcal{V}(F \circ \Phi))), \quad \varphi = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & -1 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}^{-1}.$$

The linear transformation φ produces generic coordinate projections in the sense of Theorem 5.2.1 and the function `tropicalMembership` is able to distinguish $\text{trop}(\mathcal{V}(I_1))$ from $\text{trop}(\mathcal{V}(I_2))$.

```
i9 : I'_1=ideal apply((I_1)_*, f->sub(f, {x=>x*y*z, y=>y, z=>z}));
i10 : I'_2=ideal apply((I_2)_*, f->sub(f, {x=>x*y*z, y=>y, z=>z}));
i11 : directions'=apply(directions, d->{d#0-d#1-d#2, d#1, d#2})
o11 = {{-1, 1, 1}, {1, 1, -1}, {1, -1, 1}, {3, -1, -1},
      {-3, 1, 1}, {-1, 1, -1}, {-1, -1, 1}, {1, -1, -1}}
i12 : apply(directions', d->tropicalMembership(I'_1, d))
o12 = {false, true, true, false, true, false, true, true}
i13 : apply(directions', d->tropicalMembership(I'_2, d))
o13 = {true, false, false, true, false, true, true, false}
```

7.5 A hypersurface from algebraic vision

The following example is a hypersurface in the space of $3 \times 2 \times 2$ tensors coming from a multi-view variety of a pinhole camera and a two slit camera. This example can be found in Proposition

7.5 of [66], where the authors computed the polynomial symbolically via elimination with respect to another variety in the space of $3 \times 3 \times 3$ tensors. Although this computation is not new, it serves to demonstrate the strength of our implementation.

Consider the matrix

$$[A \ B \ C] = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & b_{1,1} & b_{1,2} & c_{1,1} & c_{1,2} \\ a_{2,1} & a_{2,2} & a_{2,3} & b_{2,1} & b_{2,2} & c_{2,1} & c_{2,2} \\ a_{3,1} & a_{3,2} & a_{3,3} & b_{3,1} & b_{3,2} & c_{3,1} & c_{3,2} \\ a_{4,1} & a_{4,2} & a_{4,3} & b_{4,1} & b_{4,2} & c_{4,1} & c_{4,2} \end{bmatrix}.$$

The matrix A represents a pinhole camera and (B, C) , a two slit camera. The corresponding multi-view variety X is a hypersurface in \mathbb{P}^{11} . Let $f_{i,j,k}$ be the minor corresponding to the submatrix which ignores columns a_i, b_j , and c_k . Then X is parametrized by these twelve minors

$$F: \mathbb{C}^{28} \rightarrow \mathbb{C}^{12}$$

$$\begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & b_{1,1} & b_{1,2} & c_{1,1} & c_{1,2} \\ a_{2,1} & a_{2,2} & a_{2,3} & b_{2,1} & b_{2,2} & c_{2,1} & c_{2,2} \\ a_{3,1} & a_{3,2} & a_{3,3} & b_{3,1} & b_{3,2} & c_{3,1} & c_{3,2} \\ a_{4,1} & a_{4,2} & a_{4,3} & b_{4,1} & b_{4,2} & c_{4,1} & c_{4,2} \end{bmatrix} \xrightarrow{F} [f_{i,j,k}]_{i \in \{1,2,3\}, j, k \in \{1,2\}}.$$

This map has 17-dimensional fibers. Rather than taking generic linear slices in \mathbb{C}^{28} , we find constant replacements for 17 of the variables under the condition that the Jacobian of F does not drop rank. This substitution gives a new map $\mathcal{F}: \mathbb{C}^{11} \rightarrow \mathbb{C}^{12}$ whose image is X .

We order the $f_{i,j,k}$ variables lexicographically,

$$(f_{111}, f_{112}, f_{121}, f_{122}, f_{211}, f_{212}, f_{221}, f_{222}, f_{311}, f_{312}, f_{321}, f_{322}).$$

The polynomial f which cuts out X is homogeneous of degree 6 in 12 variables, giving an *a priori* upper bound of 12, 376 possible monomials appearing in $\mathcal{A} = \text{supp}(f)$. There is a group action of $G \cong S_3 \times S_2 \times S_2$ on the coordinates in \mathbb{C}^{12} taking $f_{i,j,k} \rightarrow f_{\sigma(i), \tau(j), \nu(k)}$ which extends to an action on the vertices of the polytope. This action is transitive on $\{f_{i,j,k}\}_{i \in \{1,2,3\}, j, k \in \{1,2\}}$ and so to get a bound on the size of any coordinate $\alpha \in \mathcal{A}$, it is enough to bound one. An oracle query in the $(1, 0, \dots, 0)$ direction returns the vector $(2, 0, \dots, 0)$ along with four points which converge somewhere other than a target. As such, $\text{New}(f) \subset \bigcap_{i=1}^{12} \mathbb{R}_{e_i, 2}^{12} = 2 \cdot [0, 1]^{12}$. This reduces the possible number of lattice points in $\text{New}(f)$ to 8, 074. Querying the oracle in the independent directions

$$(1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0) \quad (0, 0, 0, 0, 1, 1, 1, 1, 0, 0, 0, 0)$$

$$(1, 0, 1, 0, 1, 0, 1, 0, 1, 0, 1, 0) \quad (1, 1, 0, 0, 1, 1, 0, 0, 1, 1, 0, 0) \quad (1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1)$$

returns Exposes entire polytope. Thus, $\text{New}(f)$ is a subset of a 7-dimensional subspace of \mathbb{R}^{12} . The following four directions expose four vertices of $\text{New}(f)$ which, after applying symmetries of G , become 60 vertices V of a 7-dimensional polytope $P_* \subset \text{New}(f)$ containing $60 + 6$

lattice points.

$$\begin{aligned}
\mathcal{O}_{\text{New}(f)}(6, -3.5, -1, 0.4, 0.16, .6, 0.2, 1.33, .66, .9, 4, -4) \\
&= (2, 0, 0, 0, 0, 0, 2, 0, 1, 1, 0) \\
\mathcal{O}_{\text{New}(f)}(.31, -.31, -.31, .31, -.31, .09, -.31, .31, .31, -.31, .09, -.31) \\
&= (1, 0, 0, 1, 0, 1, 0, 1, 1, 0, 1, 0) \\
\mathcal{O}_{\text{New}(f)}(-.31, -.31, .31, .09, -.31, .31, .31, -.31, .09, .31, -.31, -.31) \\
&= (0, 0, 1, 1, 0, 1, 1, 0, 1, 1, 0, 0) \\
\mathcal{O}_{\text{New}(f)}(.19, -.39, .13, .19, .04, .08, -.33, .04, .25, -.20, -.13, .71) \\
&= (1, 0, 1, 0, 1, 1, 0, 0, 0, 0, 0, 2).
\end{aligned}$$

We conclude that $\text{New}(f)$ is 7-dimensional. Two more oracle queries,

$$\begin{aligned}
\mathcal{O}_{\text{New}(f)}(-11, -3, -3, 5, -11, -3, -3, 5, 1, 9, 9, -31) &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0) \\
\mathcal{O}_{\text{New}(f)}(-5, 3, 3, -5, -5, 3, 3, -5, -5, 3, 3, -5) &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)
\end{aligned}$$

imply that a positive-dimensional face of $\text{New}(f)$ is exposed in each of these directions. The facets of P_* are also exposed by these directions but no other pair of points within $P^* = \bigcap_{i=1}^{12} \mathbb{R}_{e_i, 2}^{12}$ are exposed. Thus, $\text{New}(f) = P_*$.

Knowing the support of f , interpolation successfully recovers the polynomial computed in [66]:

$$\begin{aligned}
f &= f_{111}^2 f_{212} f_{221} f_{322}^2 - f_{111}^2 f_{212} f_{222} f_{321} f_{322} - f_{111}^2 f_{221} f_{222} f_{312} f_{322} + \\
&f_{111}^2 f_{222}^2 f_{312} f_{321} - f_{111} f_{112} f_{211} f_{221} f_{322}^2 + f_{111} f_{112} f_{211} f_{222} f_{321} f_{322} - \\
&f_{111} f_{112} f_{212} f_{221} f_{321} f_{322} + f_{111} f_{112} f_{212} f_{222} f_{321}^2 + f_{111} f_{112} f_{221}^2 f_{312} f_{322} + \\
&f_{111} f_{112} f_{221} f_{222} f_{311} f_{322} - f_{111} f_{112} f_{221} f_{222} f_{312} f_{321} - f_{111} f_{112} f_{222}^2 f_{311} f_{321} - \\
&f_{111} f_{121} f_{211} f_{212} f_{322}^2 + f_{111} f_{121} f_{211} f_{222} f_{312} f_{322} + f_{111} f_{121} f_{212}^2 f_{321} f_{322} - \\
&f_{111} f_{121} f_{212} f_{221} f_{312} f_{322} + f_{111} f_{121} f_{212} f_{222} f_{311} f_{322} - f_{111} f_{121} f_{212} f_{222} f_{312} f_{321} + \\
&f_{111} f_{121} f_{221} f_{222} f_{312}^2 - f_{111} f_{121} f_{222}^2 f_{311} f_{312} + f_{111} f_{122} f_{211} f_{212} f_{321} f_{322} + \\
&f_{111} f_{122} f_{211} f_{221} f_{312} f_{322} - 2f_{111} f_{122} f_{211} f_{222} f_{312} f_{321} - f_{111} f_{122} f_{212}^2 f_{321}^2 - \\
&2f_{111} f_{122} f_{212} f_{221} f_{311} f_{322} + 2f_{111} f_{122} f_{212} f_{221} f_{312} f_{321} + f_{111} f_{122} f_{212} f_{222} f_{311} f_{321} - \\
&f_{111} f_{122} f_{221}^2 f_{312}^2 + f_{111} f_{122} f_{221} f_{222} f_{311} f_{312} + f_{112}^2 f_{211} f_{221} f_{321} f_{322} - \\
&f_{112}^2 f_{211} f_{222} f_{321}^2 - f_{112}^2 f_{221}^2 f_{311} f_{322} + f_{112}^2 f_{221} f_{222} f_{311} f_{321} \\
&+ f_{112} f_{121} f_{211}^2 f_{322}^2 - f_{112} f_{121} f_{211} f_{212} f_{321} f_{322} - f_{112} f_{121} f_{211} f_{221} f_{312} f_{322} - \\
&2f_{112} f_{121} f_{211} f_{222} f_{311} f_{322} + 2f_{112} f_{121} f_{211} f_{222} f_{312} f_{321} + 2f_{112} f_{121} f_{212} f_{221} f_{311} f_{322} - \\
&f_{112} f_{121} f_{212} f_{222} f_{311} f_{321} - f_{112} f_{121} f_{221} f_{222} f_{311} f_{312} + f_{112} f_{121} f_{222}^2 f_{311}^2 \\
&- f_{112} f_{122} f_{211}^2 f_{321} f_{322} + f_{112} f_{122} f_{211} f_{212} f_{321}^2 + f_{112} f_{122} f_{211} f_{221} f_{311} f_{322} -
\end{aligned}$$

$$\begin{aligned}
& f_{112}f_{122}f_{211}f_{221}f_{312}f_{321} + f_{112}f_{122}f_{211}f_{222}f_{311}f_{321} - f_{112}f_{122}f_{212}f_{221}f_{311}f_{321} + \\
& f_{112}f_{122}f_{221}^2f_{311}f_{312} - f_{112}f_{122}f_{221}f_{222}f_{311}^2 + f_{121}^2f_{211}f_{212}f_{312}f_{322} - \\
& f_{121}^2f_{211}f_{222}f_{312}^2 - f_{121}^2f_{212}^2f_{311}f_{322} + f_{121}^2f_{212}f_{222}f_{311}f_{312} - \\
& f_{121}f_{122}f_{211}^2f_{312}f_{322} + f_{121}f_{122}f_{211}f_{212}f_{311}f_{322} - f_{121}f_{122}f_{211}f_{212}f_{312}f_{321} + \\
& f_{121}f_{122}f_{211}f_{221}f_{312}^2 + f_{121}f_{122}f_{211}f_{222}f_{311}f_{312} + f_{121}f_{122}f_{212}^2f_{311}f_{321} - \\
& f_{121}f_{122}f_{212}f_{221}f_{311}f_{312} - f_{121}f_{122}f_{212}f_{222}f_{311}^2 + f_{122}^2f_{211}f_{312}f_{321} - \\
& f_{122}^2f_{211}f_{212}f_{311}f_{321} - f_{122}^2f_{211}f_{221}f_{311}f_{312} + f_{122}^2f_{212}f_{221}f_{311}^2
\end{aligned}$$

7.6 The Lüroth invariant

7.6.1 The Lüroth invariant, hypersurface, and polytope.

Let \mathbb{C}_q^{15} be the vector space spanned by all homogeneous quartic plane curves with coefficients $\{q_{ijk}\}_{i+j+k=4}$ so that a quartic $Q \in \mathbb{C}_q^{15}$ is written as

$$Q = \sum_{i+j+k=4} q_{ijk} x^i y^j z^k.$$

Such a quartic $\mathcal{V}(Q) \subset \mathbb{P}^2$ is called **Lüroth** if it passes through the ten intersection points of five lines in \mathbb{P}^2 . We display one such quartic in Figure 7.8.




Figure 7.8: A Lüroth quartic.

The set of all Lüroth quartics \mathbb{L} is a hypersurface of degree 54 in \mathbb{P}_q^{14} called the **Lüroth hy-**

persurface. The group $\mathrm{PGL}(3, \mathbb{C})$ of all projective linear transformations of \mathbb{P}^2 acts on a plane quartic $\mathcal{V}(Q)$ by some element $A \in \mathrm{PGL}(3, \mathbb{C})$ in the natural way: $\mathcal{V}(Q) \mapsto A \cdot \mathcal{V}(Q) = \mathcal{V}(Q(A^*(x, y, z)))$. This action preserves intersection points and so if $\mathcal{V}(Q)$ is a Lüroth quartic, so is $A \cdot \mathcal{V}(Q)$. The defining equation Λ of the Lüroth hypersurface is called the **Lüroth invariant**.

The Lüroth hypersurface is parametrized by the coefficients of five homogeneous linear polynomials $\ell_i = a_i x + b_i y + c_i z \in \mathbb{C}[x, y, z]$. This parametrization is

$$\varphi: \mathbb{P}((\mathbb{C}^3)^5) \dashrightarrow \mathbb{P}^{14} \quad (7.10)$$

$$(\ell_1, \dots, \ell_5) \mapsto \sum_{j=1}^5 \prod_{i \neq j} \ell_i = \sum_{i+j+k=4} q_{ijk} x^i y^j z^k. \quad (7.11)$$

Finding Λ using symbolic elimination algorithms is computationally infeasible. Moreover, it is expected that Λ in its expanded form is not human-readable. Thus, we attempt to determine the **Lüroth polytope**, $\mathfrak{P} = \mathrm{New}(\Lambda) \subset \mathbb{R}^{15}$ using Algorithm 7.1.2. Before discussing computations, we explain some reductions to the problem.

Corollary 7.6.1. *Every point p in the Lüroth polytope \mathfrak{P} solves the linear equation*

$$\begin{pmatrix} 4 & 3 & 3 & 2 & 2 & 2 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 2 & 1 & 0 & 3 & 2 & 1 & 0 & 4 & 3 & 2 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 2 & 0 & 1 & 2 & 3 & 0 & 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} p = \begin{pmatrix} 72 \\ 72 \\ 72 \\ 54 \end{pmatrix}.$$

Proof. The Lüroth invariant is a homogeneous polynomial of degree 54 in the coefficient space $\mathbb{C}_q^{4\Delta_3} \cong \mathbb{P}^{14}$ which is invariant under permutations and scalings of variables x, y , and z . Observing that $\frac{54 \cdot 4}{3} = 72$ and applying Lemma 5.1.3 gives the result. \square

We order the coordinates of the space \mathbb{R}_p^{15} containing \mathfrak{P} as follows

$$\{p_{400}, p_{310}, p_{301}, p_{220}, p_{211}, p_{202}, p_{130}, p_{121}, p_{112}, p_{103}, p_{040}, p_{031}, p_{022}, p_{013}, p_{004}\}. \quad (7.12)$$

We may identify coordinates by their numerical bijection with $0, 1, \dots, 14$ as listed above. For example, p_{301} may be written as p_2 . Under this bijection, the permutation group S_3 acting on coordinates of the subscripts of p_{ijk} induces the following involutions,

$$\begin{aligned} \sigma_{xy} &= (0, 10)(1, 6)(2, 11)(4, 7)(5, 12)(9, 13) \\ \sigma_{yz} &= (1, 2)(3, 5)(6, 9)(7, 8)(10, 14)(11, 13) \\ \sigma_{xz} &= \sigma_{xy} \circ \sigma_{yz} \circ \sigma_{xy} = (0, 14)(1, 13)(2, 9)(3, 12)(4, 8)(6, 11), \end{aligned}$$

written in cycle notation. We write G for this subgroup $S_3 \hookrightarrow S_{15}$. Corollary 7.6.1 gives the *a priori* bounds of

$$p_{400} \in [0, 18], \quad p_{310} \in [0, 24], \quad p_{220} \in [0, 36], \quad p_{211} \in [0, 36] \quad (7.13)$$

$$\begin{array}{ccc}
\mathbb{C}_a^{14} \times \mathbb{C}_s & \xrightarrow{\pi_s} & \mathbb{C}_s \\
\pi_a \downarrow & & \downarrow \mathbf{L}_t \\
\mathbb{C}_a^{14} & \xrightarrow{\varphi} & \mathbb{C}_q^{15}
\end{array}$$

Figure 7.9: A fiber product construction to compute witness points.

on the sizes of each coordinate p_{ijk} of a point in \mathfrak{P} . In Section 7.6.3, we use the HS-algorithm to show that these bounds are not sharp.

7.6.2 Computational setup

To perform computations, we dehomogenize the domain $\mathbb{P}((\mathbb{C}^3)^5)$ of the parametrization (7.10) with respect to a random linear polynomial and work with the restricted map $\varphi: \mathbb{C}^{14} \rightarrow \mathbb{C}_q^{15}$.

We parametrize the lines \mathcal{L}_t in the HS-algorithm by $t \mapsto \{t^{\omega_i}(a_i s - b_i)\}_{i=0}^{14}$ and chose $a, b \in (\mathbb{C}^\times)^{15}$ so that the target points $\rho_i = b_i/a_i$ are the 15-th roots of unity $\{\zeta_{15}^i\}_{i=0}^{14}$ where $\zeta_{15} = e^{2\pi\sqrt{-1}/15}$. We do this under the heuristic assumption that choosing targets far from one another decreases the chances of the implementation **NumericalNP.m2** misattributing which target is the true limit of a path.

To compute the intersection points of $\mathcal{L}_t \cap \Lambda$ in the parameters s , we employ the fiber product construction in Figure 7.9 and solve the equations

$$F_t = \{q_{ijk}(a) - t^{\omega_i}(a_i s - b_i)\}_{i=0}^{14} \quad (7.14)$$

in $\mathbb{C}_a^{14} \times \mathbb{C}_s$, where $ijk \leftrightarrow \iota$ is the identification of $\{(i, j, k) \mid i, j, k \geq 0, i + j + k = 4\}$ with $\{0, \dots, 14\}$ given in (7.12). During the homotopy process, we project a solution $(a, s) \in \mathbb{C}_a^{14} \times \mathbb{C}_s$ to $s \in \mathbb{C}_s$ to monitor convergence.

Whenever querying the numerical oracle in a direction $\omega \in S^{14} \subset \mathbb{R}^{15}$, by Theorem 7.3.1 it is best to attempt to maximize d_ω . Given that we do not *a priori* know the polytope \mathfrak{P} , this is generally difficult. Nonetheless, we always project ω onto the kernel of the matrix in Corollary 7.6.1 and rescale this projection to be a unit vector. Given that $|\omega| = 1$, this process increases d_ω thus increasing the convergence rate of the HS-algorithm.

Given the size of these computations, we expect numerical errors to occur. However, assessing whether something went wrong during the path tracking process can be done in several ways.

Remark 7.6.2. Suppose that a numerical implementation of the HS-algorithm returns a vertex $\mathcal{O}_{\mathfrak{P}}(\omega) = v$ on the direction $\omega \in \mathbb{R}^{15}$. A numerical error has occurred if any of the following are true.

- (1) $v_\infty \neq 0$, (Since Λ is homogeneous, no points in the HS-algorithm will diverge)
- (2) $\mathcal{O}_{\mathfrak{P}}(\sigma(\omega)) \neq \sigma(v)$ for any $\sigma \in G$, (Λ is invariant under G)
- (3) v does not solve the matrix equation in Corollary 7.6.1.

- (4) $v \notin H_{\mathfrak{P}}(\nu)$ where $H_{\mathfrak{P}}(\nu)$ is the halfspace containing \mathfrak{P} implied by an oracle call $\mathcal{O}_{\mathfrak{P}}(\nu)$ we have already performed on $\nu \in \mathbb{R}^{15}$, (Since $\mathfrak{P} \subset H_{\mathfrak{P}}(\nu)$ for any $\nu \in \mathbb{R}^{15}$).

In the last case, an error has occurred on either $\mathcal{O}_{\mathfrak{P}}(\omega)$ or $\mathcal{O}_{\mathfrak{P}}(\nu)$. ◇

7.6.3 Vertices of the Lüroth polytope

Using **NumericalNP.m2**, we reproduce the result of [61] that

$$\mathfrak{P}_{(3,-5,3,2,3,-2,-1,4,-3,-2,3,1,-5,3,-5)} = (6, 0, 6, 0, 0, 0, 0, 30, 0, 0, 0, 0, 0, 12, 0),$$

indicating that $q_{400}^6 q_{301}^6 q_{121}^{30} q_{013}^{12}$ is a monomial in the support of Λ . Acting on the exponents by G reveals that

$$\begin{aligned} & q_{040}^6 q_{301}^6 q_{211}^{30} q_{103}^{12}, \quad q_{400}^6 q_{310}^6 q_{112}^{30} q_{031}^{12}, \quad q_{040}^6 q_{130}^6 q_{112}^{30} q_{301}^{12}, \\ & q_{004}^6 q_{013}^6 q_{211}^{30} q_{130}^{12}, \quad q_{004}^6 q_{130}^6 q_{121}^{30} q_{310}^{12}, \end{aligned}$$

are also monomials of Λ .

Figures 7.10-7.12 display snapshots from the Sage [65] animation (produced by Function 7.4.3 for $t = 1, 4, 8, 20, 30$, and 75 respectively) of the paths $\{s_i(t)\}_{i=1}^{54}$ in the HS-algorithm. The second image shows a clustering of 12 points toward ζ_{15}^{13} . The third image shows a clustering of six points toward ζ_{15}^0 and 1 point converging to ζ_{15}^2 . Next, a large cluster of 30 points move toward ζ_{15}^7 and five points move toward ζ_{15}^2 . They converge in the last image. As $t \rightarrow \infty$, there are two clusters




Figure 7.10: (Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 1$ and $t = 4$ resp.) of Algorithm 7.1.2.

of points which converge to ζ_{15}^2 : one of size one and another of size five. This suggests that these paths have winding numbers one and five respectively (see the Cauchy endgame in Section 6.2.2). We do not have any conjectures about what this means for the polytope \mathfrak{P} .




Figure 7.11: (Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 8$ and $t = 20$ resp.) of Algorithm 7.1.2.




Figure 7.12: (Reprinted from [1]) Two snapshots of a Sage animation of the paths $\{s_i(t)\}_{i=1}^{54}$ (for $t = 30$ and $t = 75$ resp.) of Algorithm 7.1.2.

Querying the oracle in the coordinate directions with respect to $p_{400}, p_{310}, p_{220}$, and p_{211} returns $18e_{400}, 24e_{310}, 28e_{220}$, and $32e_{211}$ respectively where e_{ijk} is the standard basis vector in the coordinate p_{ijk} . This gives new bounds of

$$p_{400} \in [0, 18], \quad p_{310} \in [0, 24], \quad p_{220} \in [0, 28], \quad p_{211} \in [0, 32] \quad (7.15)$$

improving the bounds (7.13) for p_{220} and p_{211} .

The initial upper bound on the number of terms in f based on homogeneity and degree is $\binom{54+15-1}{54} = 123, 234, 279, 768, 160$. Taking into account the linear space containing \mathfrak{P} we see that the p_{400}, p_{040} , and p_{004} coordinates of a point in \mathfrak{P} are determined by the rest. Thus, the number of points in \mathbb{Z}^{15} subject to the bounds of (7.15) and the matrix equation of Corollary 7.6.1 is

$$[x^{54}] \left(\frac{1-x^{24}}{1-x} \right)^6 \left(\frac{1-x^{28}}{1-x} \right)^3 \left(\frac{1-x^{32}}{1-x} \right)^3 = 879, 008, 719, 165$$

which improves the initial bound by a factor of ≈ 140 . Nonetheless, this number remains so large that interpolation is infeasible.

In total, we have found 1713 vertices, belonging to 1, 1, 28, and 271 orbits of sizes 1, 2, 3, and 6 using our implementation **NumericalNP.m2**. The orbits of size one and two which we found are

$$\{q_{400}^{18} q_{040}^{18} q_{004}^{18}\} \quad \text{and} \quad \{q_{301}^{18} q_{130}^{18} q_{013}^{18}, q_{103}^{18} q_{031}^{18} q_{310}^{18}\},$$

respectively. We list the other orbits in Table 7.1 along with an orbit representative, the number of times a representative of each orbit was found in our search, and the number of elements of each orbit. The two orbits colored in blue were found by Hauenstein and Sottile [61] in their computation of the Newton polytope of the hypersurface of even Lüroth quartics which has five vertices, two having a G -orbit of size one and three belonging to a G -orbit of size three. We did not find the vertex $(4, 0, 0, 14, 0, 14, 0, 0, 0, 0, 4, 0, 14, 0, 4)$. Up-to-date computations regarding the Lüroth polytope can be found at the author's webpage [63].

Vertex v	#	$ G \cdot v $
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 18, 0, 0, 0, 12)	360	6
(4, 0, 0, 0, 0, 28, 0, 0, 0, 0, 18, 0, 0, 0, 4)	126	3
(0, 0, 0, 0, 0, 25, 22, 0, 0, 0, 0, 0, 3, 0, 4)	117	6
(0, 0, 18, 0, 0, 0, 0, 0, 0, 18, 18, 0, 0, 0, 0)	113	3
(0, 0, 0, 0, 32, 0, 0, 0, 0, 8, 10, 0, 0, 0, 4)	106	6
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 7, 0, 22, 0, 1)	88	6
(0, 0, 0, 0, 0, 25, 22, 0, 0, 0, 0, 0, 0, 6, 1)	88	6
(0, 0, 18, 0, 0, 0, 18, 0, 0, 0, 0, 0, 0, 18, 0)	78	2
(0, 0, 0, 0, 32, 0, 0, 0, 0, 8, 6, 0, 8, 0, 0)	72	6
(0, 0, 8, 0, 0, 24, 0, 0, 0, 0, 18, 0, 0, 0, 4)	64	6
(0, 0, 0, 0, 30, 0, 0, 0, 0, 12, 6, 6, 0, 0, 0)	64	6
(0, 0, 0, 0, 0, 24, 22, 0, 0, 2, 0, 0, 0, 6, 0)	64	6
(18, 0, 0, 0, 0, 0, 0, 0, 0, 0, 18, 0, 0, 0, 18)	63	1
(4, 0, 0, 0, 28, 0, 0, 0, 0, 0, 11, 0, 0, 0, 11)	60	3
(0, 0, 0, 0, 30, 0, 0, 0, 0, 12, 10, 0, 0, 2, 0)	54	6
(0, 0, 0, 0, 0, 25, 21, 1, 0, 0, 0, 0, 0, 7, 0)	50	6
(0, 0, 12, 0, 0, 0, 22, 0, 0, 14, 0, 0, 0, 6, 0)	49	6
(0, 0, 12, 0, 0, 0, 16, 0, 0, 20, 6, 0, 0, 0, 0)	48	6
(0, 0, 0, 0, 30, 0, 6, 0, 0, 6, 0, 0, 12, 0, 0)	46	3
(0, 0, 12, 0, 0, 0, 24, 0, 0, 12, 0, 0, 0, 0, 6)	45	6
(0, 0, 0, 0, 0, 24, 16, 0, 0, 8, 6, 0, 0, 0, 0)	45	6
(0, 0, 0, 0, 0, 28, 0, 16, 0, 0, 10, 0, 0, 0, 0)	44	3
(0, 0, 0, 0, 0, 27, 0, 18, 0, 0, 9, 0, 0, 0, 0)	43	3
(0, 0, 8, 24, 0, 0, 0, 0, 0, 0, 6, 0, 0, 0, 16)	41	6
(0, 0, 0, 0, 0, 28, 16, 0, 0, 0, 6, 0, 0, 0, 4)	41	6
(0, 0, 8, 0, 24, 0, 0, 0, 0, 0, 12, 0, 0, 0, 10)	40	6
(0, 0, 0, 0, 24, 6, 12, 0, 0, 0, 0, 0, 0, 12, 0)	40	6
(0, 0, 0, 0, 29, 0, 0, 0, 0, 14, 10, 1, 0, 0, 0)	39	6
(0, 0, 12, 0, 18, 0, 0, 0, 0, 0, 10, 0, 0, 14, 0)	37	6
(0, 0, 12, 0, 0, 0, 20, 0, 0, 16, 0, 0, 6, 0, 0)	36	6
(0, 0, 1, 24, 0, 0, 0, 0, 0, 21, 0, 8, 0, 0, 0)	31	6
(0, 0, 13, 0, 0, 0, 24, 0, 0, 9, 0, 0, 0, 0, 8)	29	6
(0, 0, 0, 0, 28, 1, 0, 0, 0, 14, 11, 0, 0, 0, 0)	29	6
(0, 0, 12, 0, 0, 0, 18, 0, 0, 18, 0, 6, 0, 0, 0)	28	6
(0, 0, 0, 0, 32, 0, 0, 0, 0, 8, 6, 4, 0, 4, 0)	27	6
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 17, 0, 0, 4, 9)	26	6
(0, 0, 2, 0, 0, 21, 24, 0, 0, 0, 0, 0, 0, 0, 7)	25	6
(0, 0, 1, 0, 27, 0, 15, 0, 0, 0, 0, 0, 0, 0, 11)	25	6
(0, 0, 14, 0, 0, 0, 21, 0, 9, 0, 0, 0, 0, 0, 10)	24	6
(0, 0, 11, 0, 9, 0, 21, 0, 0, 0, 0, 0, 0, 0, 13)	21	6

Table 7.1: Vertices of the Lüroth polytope found.

Vertex v	#	$ G \cdot v $
(0, 0, 21, 0, 0, 0, 0, 0, 9, 16, 0, 0, 8, 0)	18	6
(0, 0, 18, 6, 0, 0, 0, 0, 6, 0, 20, 0, 0, 4)	18	6
(0, 0, 0, 0, 30, 0, 6, 0, 0, 6, 0, 8, 0, 0, 4)	18	6
(0, 0, 20, 0, 0, 0, 12, 0, 0, 0, 9, 0, 0, 0, 13)	17	6
(0, 0, 12, 0, 0, 0, 18, 0, 18, 0, 0, 0, 0, 0, 6)	17	6
(0, 0, 4, 28, 0, 0, 0, 0, 0, 4, 4, 0, 0, 0, 14)	17	6
(3, 0, 0, 0, 0, 16, 0, 28, 0, 0, 4, 0, 0, 0, 3)	16	3
(0, 0, 17, 0, 0, 0, 21, 0, 0, 0, 0, 0, 0, 9, 7)	16	6
(0, 0, 8, 0, 24, 0, 0, 0, 0, 0, 6, 4, 0, 12, 0)	16	6
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 15, 4, 0, 0, 11)	15	6
(0, 0, 8, 18, 0, 0, 12, 0, 0, 0, 0, 0, 0, 0, 16)	14	6
(0, 0, 4, 16, 0, 0, 0, 0, 28, 0, 3, 0, 0, 0, 3)	14	6
(0, 0, 0, 4, 0, 28, 0, 8, 0, 0, 12, 0, 0, 0, 2)	14	6
(0, 0, 4, 28, 0, 0, 0, 0, 4, 0, 3, 0, 0, 0, 15)	13	6
(0, 0, 0, 0, 20, 12, 8, 0, 0, 0, 0, 0, 14, 0, 0)	12	6
(0, 0, 19, 0, 0, 0, 6, 0, 0, 9, 0, 18, 0, 0, 2)	11	6
(0, 0, 0, 0, 0, 24, 19, 0, 0, 5, 0, 3, 3, 0, 0)	11	6
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 6, 2, 21, 0, 1)	10	6
(0, 0, 0, 3, 0, 22, 22, 0, 0, 0, 0, 0, 0, 0, 7)	10	6
(0, 0, 0, 0, 32, 0, 2, 0, 0, 6, 4, 4, 0, 6, 0)	10	6
(0, 12, 12, 0, 0, 0, 0, 0, 0, 0, 1, 0, 28, 0, 1)	9	3
(0, 0, 8, 0, 0, 20, 0, 0, 0, 8, 18, 0, 0, 0, 0)	9	3
(0, 0, 0, 0, 28, 0, 4, 0, 0, 12, 2, 8, 0, 0, 0)	9	6
(0, 0, 0, 0, 0, 28, 16, 0, 0, 0, 0, 6, 3, 0, 1)	9	6
(0, 0, 0, 0, 0, 24, 22, 0, 0, 2, 0, 0, 3, 0, 3)	9	6
(0, 0, 23, 0, 0, 0, 3, 0, 0, 0, 9, 9, 0, 0, 10)	8	6
(0, 0, 8, 12, 12, 0, 0, 0, 0, 0, 0, 12, 0, 0, 10)	8	6
(0, 0, 0, 0, 28, 0, 0, 0, 8, 8, 6, 4, 0, 0, 0)	8	6
(0, 0, 0, 0, 27, 0, 9, 0, 0, 9, 0, 0, 9, 0, 0)	8	3
(0, 0, 0, 0, 16, 8, 4, 0, 20, 0, 6, 0, 0, 0, 0)	8	6
(0, 0, 24, 0, 0, 0, 0, 0, 0, 0, 6, 13, 0, 9, 2)	7	6
(0, 0, 17, 0, 0, 0, 0, 21, 0, 0, 6, 0, 0, 6, 4)	7	6
(0, 0, 6, 0, 24, 0, 6, 0, 0, 0, 4, 0, 0, 14, 0)	7	6
(0, 0, 2, 20, 0, 0, 3, 0, 23, 0, 0, 0, 0, 0, 6)	7	6
(0, 0, 1, 23, 0, 0, 0, 0, 0, 23, 6, 0, 1, 0, 0)	7	6
(0, 0, 12, 0, 0, 0, 2, 24, 0, 10, 4, 0, 0, 2, 0)	6	6
(0, 0, 8, 0, 0, 8, 8, 24, 0, 0, 0, 0, 0, 0, 6)	6	6
(0, 0, 0, 12, 0, 20, 0, 0, 8, 0, 10, 0, 0, 0, 4)	6	6

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 0, 2, 0, 24, 20, 0, 0, 0, 0, 0, 8, 0)	6	6
(0, 0, 0, 0, 26, 2, 12, 0, 0, 4, 0, 0, 0, 10, 0)	6	6
(0, 0, 12, 0, 0, 0, 6, 18, 0, 12, 0, 6, 0, 0, 0)	5	3
(0, 0, 8, 0, 0, 21, 6, 0, 0, 0, 0, 18, 0, 0, 1)	5	6
(0, 0, 6, 0, 26, 0, 2, 0, 0, 0, 4, 4, 0, 12, 0)	5	6
(0, 0, 2, 0, 24, 0, 6, 0, 0, 12, 0, 10, 0, 0, 0)	5	6
(0, 0, 1, 0, 0, 24, 21, 0, 0, 0, 0, 0, 1, 7, 0)	5	6
(0, 0, 0, 0, 26, 6, 8, 0, 0, 0, 0, 4, 0, 10, 0)	5	6
(0, 0, 0, 0, 26, 6, 8, 0, 0, 0, 0, 0, 8, 6, 0)	5	6
(0, 0, 0, 0, 24, 8, 0, 0, 8, 0, 6, 0, 8, 0, 0)	5	6
(2, 0, 0, 0, 0, 19, 0, 26, 0, 0, 5, 0, 0, 0, 2)	4	3
(0, 0, 20, 6, 0, 0, 0, 0, 0, 0, 6, 12, 0, 0, 10)	4	6
(0, 0, 8, 6, 18, 0, 0, 0, 0, 0, 0, 10, 0, 12, 0)	4	6
(0, 0, 6, 0, 24, 0, 6, 0, 0, 0, 0, 6, 0, 12, 0)	4	6
(0, 0, 4, 0, 0, 18, 24, 0, 0, 0, 0, 0, 0, 0, 8)	4	6
(0, 0, 3, 28, 0, 0, 0, 0, 7, 0, 0, 3, 0, 0, 13)	4	3
(0, 0, 0, 12, 0, 20, 0, 0, 0, 8, 12, 0, 0, 0, 2)	4	6
(0, 0, 0, 10, 0, 22, 0, 8, 0, 0, 0, 12, 0, 0, 2)	4	6
(0, 0, 0, 4, 0, 28, 8, 0, 0, 0, 10, 0, 0, 0, 4)	4	6
(0, 0, 0, 0, 30, 0, 6, 0, 0, 6, 0, 6, 0, 6, 0)	4	3
(0, 0, 0, 0, 28, 0, 6, 0, 0, 10, 0, 8, 0, 2, 0)	4	6
(0, 0, 0, 0, 24, 0, 6, 0, 12, 6, 4, 0, 0, 2, 0)	4	6
(0, 0, 0, 0, 20, 12, 0, 0, 8, 0, 10, 0, 0, 4, 0)	4	6
(0, 0, 0, 0, 0, 28, 14, 2, 0, 0, 0, 6, 4, 0, 0)	4	6
(0, 0, 0, 0, 0, 24, 16, 0, 6, 2, 0, 6, 0, 0, 0)	4	6
(0, 0, 0, 0, 0, 24, 14, 2, 8, 0, 0, 6, 0, 0, 0)	4	6
(3, 0, 18, 0, 0, 0, 0, 0, 0, 6, 3, 20, 0, 0, 4)	3	6
(3, 0, 0, 0, 29, 0, 0, 0, 0, 2, 7, 5, 0, 0, 8)	3	6
(0, 2, 10, 0, 0, 0, 0, 26, 0, 10, 4, 0, 0, 2, 0)	3	3
(0, 0, 12, 0, 0, 0, 12, 0, 24, 0, 0, 4, 0, 0, 2)	3	6
(0, 0, 8, 6, 0, 18, 0, 0, 0, 0, 0, 20, 0, 0, 2)	3	6
(0, 0, 6, 0, 20, 0, 12, 0, 2, 0, 0, 0, 0, 14, 0)	3	6
(0, 0, 6, 0, 0, 12, 6, 24, 0, 0, 0, 0, 0, 6, 0)	3	6
(0, 0, 2, 0, 30, 0, 6, 0, 0, 0, 0, 6, 0, 6, 4)	3	6
(0, 0, 2, 0, 0, 27, 12, 0, 0, 0, 0, 12, 0, 0, 1)	3	6
(0, 0, 2, 0, 0, 22, 0, 22, 0, 0, 7, 0, 0, 0, 1)	3	6
(0, 0, 2, 0, 0, 20, 18, 0, 0, 8, 0, 6, 0, 0, 0)	3	6
(0, 0, 1, 22, 0, 1, 0, 0, 0, 23, 7, 0, 0, 0, 0)	3	6

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 1, 21, 2, 0, 0, 0, 0, 23, 7, 0, 0, 0, 0)	3	6
(0, 0, 0, 4, 0, 21, 0, 22, 0, 0, 5, 0, 0, 0, 2)	3	6
(0, 0, 0, 3, 0, 28, 0, 10, 0, 0, 11, 0, 0, 2, 0)	3	6
(0, 0, 0, 0, 32, 0, 2, 0, 0, 6, 4, 0, 8, 2, 0)	3	6
(0, 0, 0, 0, 29, 0, 0, 1, 0, 13, 10, 0, 0, 1, 0)	3	6
(0, 0, 0, 0, 28, 0, 6, 0, 0, 10, 0, 6, 4, 0, 0)	3	6
(0, 0, 0, 0, 24, 0, 2, 10, 0, 12, 4, 2, 0, 0, 0)	3	6
(0, 0, 0, 0, 18, 6, 4, 0, 20, 0, 4, 2, 0, 0, 0)	3	6
(0, 0, 0, 0, 8, 16, 8, 0, 16, 0, 6, 0, 0, 0, 0)	3	6
(0, 0, 0, 0, 6, 22, 14, 2, 0, 0, 0, 0, 10, 0, 0)	3	6
(0, 0, 0, 0, 4, 22, 0, 20, 0, 0, 7, 0, 0, 0, 1)	3	6
(0, 0, 0, 0, 3, 21, 19, 0, 0, 5, 0, 0, 6, 0, 0)	3	6
(0, 0, 0, 0, 0, 28, 10, 6, 0, 0, 0, 10, 0, 0, 0)	3	3
(0, 0, 0, 0, 0, 28, 8, 8, 0, 0, 6, 0, 4, 0, 0)	3	6
(0, 0, 0, 0, 0, 25, 19, 0, 3, 0, 3, 0, 0, 0, 4)	3	6
(0, 0, 0, 0, 0, 24, 18, 6, 0, 0, 0, 0, 0, 6, 0)	3	6
(0, 3, 18, 0, 0, 0, 3, 0, 0, 6, 0, 20, 0, 0, 4)	2	6
(0, 0, 23, 0, 0, 0, 3, 0, 0, 0, 3, 9, 12, 0, 4)	2	6
(0, 0, 14, 0, 0, 0, 0, 21, 0, 9, 6, 2, 0, 0, 2)	2	6
(0, 0, 12, 7, 0, 0, 2, 20, 0, 0, 3, 0, 0, 0, 10)	2	6
(0, 0, 12, 7, 0, 0, 0, 22, 0, 0, 3, 0, 0, 2, 8)	2	6
(0, 0, 12, 6, 0, 12, 0, 0, 0, 0, 0, 20, 0, 0, 4)	2	6
(0, 0, 12, 3, 0, 0, 0, 26, 0, 4, 3, 0, 0, 2, 4)	2	6
(0, 0, 12, 0, 0, 0, 16, 0, 18, 2, 0, 0, 0, 6, 0)	2	6
(0, 0, 12, 0, 0, 0, 6, 24, 0, 6, 0, 2, 0, 0, 4)	2	6
(0, 0, 9, 4, 2, 0, 0, 22, 0, 11, 2, 2, 2, 0, 0)	2	6
(0, 0, 8, 0, 20, 4, 0, 0, 0, 0, 10, 0, 0, 12, 0)	2	6
(0, 0, 8, 0, 20, 0, 0, 8, 0, 0, 6, 0, 0, 12, 0)	2	6
(0, 0, 6, 24, 0, 0, 0, 0, 0, 6, 0, 8, 0, 0, 10)	2	6
(0, 0, 5, 27, 0, 0, 0, 0, 0, 3, 3, 2, 0, 0, 14)	2	6
(0, 0, 4, 27, 0, 0, 0, 0, 6, 0, 0, 4, 0, 0, 13)	2	3
(0, 0, 4, 24, 4, 0, 0, 4, 0, 0, 3, 0, 0, 0, 15)	2	6
(0, 0, 4, 24, 0, 4, 4, 0, 0, 0, 3, 0, 0, 0, 15)	2	6
(0, 0, 4, 24, 0, 0, 8, 0, 0, 4, 0, 0, 0, 0, 14)	2	6
(0, 0, 1, 0, 27, 0, 0, 0, 1, 14, 11, 0, 0, 0, 0)	2	6
(0, 0, 0, 12, 0, 18, 0, 0, 0, 12, 12, 0, 0, 0, 0)	2	6
(0, 0, 0, 10, 0, 22, 4, 4, 0, 0, 0, 4, 10, 0, 0)	2	3

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 0, 10, 0, 20, 12, 0, 0, 0, 0, 8, 0, 4)	2	6
(0, 0, 0, 6, 18, 8, 0, 0, 8, 0, 0, 6, 8, 0, 0)	2	6
(0, 0, 0, 6, 0, 26, 8, 0, 0, 0, 0, 8, 6, 0, 0)	2	3
(0, 0, 0, 4, 0, 23, 0, 18, 0, 0, 5, 0, 4, 0, 0)	2	3
(0, 0, 0, 2, 8, 20, 12, 0, 0, 0, 0, 0, 12, 0, 0)	2	6
(0, 0, 0, 0, 32, 0, 4, 0, 0, 4, 2, 0, 10, 0, 2)	2	3
(0, 0, 0, 0, 28, 4, 0, 0, 0, 8, 10, 0, 0, 4, 0)	2	6
(0, 0, 0, 0, 28, 1, 14, 0, 0, 0, 0, 0, 1, 0, 10)	2	6
(0, 0, 0, 0, 26, 6, 6, 0, 2, 0, 0, 6, 0, 8, 0)	2	6
(0, 0, 0, 0, 26, 2, 6, 0, 10, 0, 0, 6, 0, 0, 4)	2	6
(0, 0, 0, 0, 26, 0, 2, 0, 12, 6, 4, 4, 0, 0, 0)	2	6
(0, 0, 0, 0, 26, 0, 0, 10, 0, 10, 6, 0, 0, 2, 0)	2	6
(0, 0, 0, 0, 24, 6, 0, 0, 0, 12, 12, 0, 0, 0, 0)	2	6
(0, 0, 0, 0, 24, 0, 4, 0, 14, 6, 4, 2, 0, 0, 0)	2	6
(0, 0, 0, 0, 24, 0, 2, 12, 0, 10, 4, 0, 0, 2, 0)	2	6
(0, 0, 0, 0, 22, 10, 8, 0, 0, 0, 4, 0, 0, 10, 0)	2	6
(0, 0, 0, 0, 20, 4, 8, 0, 16, 0, 0, 4, 0, 0, 2)	2	6
(0, 0, 0, 0, 16, 16, 0, 8, 0, 0, 6, 0, 8, 0, 0)	2	6
(0, 0, 0, 0, 16, 11, 0, 2, 16, 0, 9, 0, 0, 0, 0)	2	3
(0, 0, 0, 0, 12, 16, 0, 16, 0, 0, 6, 0, 0, 4, 0)	2	6
(0, 0, 0, 0, 2, 22, 20, 0, 4, 0, 0, 0, 0, 6, 0)	2	6
(0, 0, 0, 0, 0, 28, 15, 1, 0, 0, 0, 6, 3, 1, 0)	2	6
(0, 0, 0, 0, 0, 24, 8, 8, 8, 0, 6, 0, 0, 0, 0)	2	6
(3, 0, 4, 0, 0, 10, 0, 28, 0, 0, 4, 0, 0, 0, 5)	1	6
(3, 0, 2, 0, 27, 0, 0, 0, 0, 0, 10, 0, 0, 5, 7)	1	6
(3, 0, 2, 0, 27, 0, 0, 0, 0, 0, 3, 9, 0, 6, 4)	1	6
(3, 0, 0, 0, 25, 0, 0, 0, 8, 2, 3, 9, 0, 0, 4)	1	6
(2, 0, 4, 26, 0, 0, 0, 0, 0, 0, 4, 0, 2, 0, 16)	1	6
(2, 0, 0, 0, 30, 0, 0, 0, 0, 4, 4, 7, 0, 5, 2)	1	6
(1, 0, 0, 4, 4, 13, 0, 26, 0, 0, 2, 0, 0, 0, 4)	1	6
(0, 9, 11, 0, 0, 0, 12, 0, 0, 0, 0, 9, 0, 0, 13)	1	6
(0, 4, 4, 0, 24, 0, 0, 0, 0, 0, 3, 0, 16, 0, 3)	1	3
(0, 3, 20, 0, 0, 0, 3, 0, 0, 0, 6, 12, 0, 0, 10)	1	6
(0, 2, 9, 2, 0, 0, 0, 28, 0, 7, 2, 0, 0, 2, 2)	1	6
(0, 0, 20, 0, 6, 0, 0, 0, 0, 0, 12, 6, 0, 0, 10)	1	6
(0, 0, 20, 0, 6, 0, 0, 0, 0, 0, 6, 0, 21, 0, 1)	1	6
(0, 0, 20, 0, 3, 0, 6, 0, 0, 0, 0, 15, 0, 6, 4)	1	6
(0, 0, 20, 0, 0, 0, 0, 0, 12, 0, 0, 11, 0, 2, 0, 9)	1	6
(0, 0, 19, 0, 2, 0, 0, 11, 0, 0, 12, 0, 0, 0, 10)	1	6

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 18, 0, 0, 2, 0, 14, 0, 0, 11, 0, 0, 0, 9)	1	6
(0, 0, 17, 0, 0, 0, 12, 0, 9, 0, 0, 9, 0, 0, 7)	1	6
(0, 0, 16, 9, 0, 0, 0, 0, 6, 0, 0, 16, 0, 0, 7)	1	3
(0, 0, 15, 0, 0, 0, 0, 23, 0, 4, 6, 0, 0, 2, 4)	1	6
(0, 0, 15, 0, 0, 0, 0, 21, 0, 6, 6, 2, 0, 0, 4)	1	6
(0, 0, 14, 0, 9, 0, 12, 0, 0, 0, 0, 9, 0, 0, 10)	1	6
(0, 0, 14, 0, 2, 0, 0, 16, 10, 0, 7, 0, 0, 0, 5)	1	6
(0, 0, 12, 18, 0, 0, 0, 0, 0, 0, 4, 0, 0, 20, 0)	1	6
(0, 0, 12, 9, 0, 0, 0, 18, 0, 0, 3, 2, 0, 0, 10)	1	6
(0, 0, 12, 3, 0, 0, 0, 24, 2, 4, 3, 0, 2, 0, 4)	1	6
(0, 0, 12, 0, 9, 0, 6, 0, 12, 0, 0, 11, 0, 0, 4)	1	6
(0, 0, 12, 0, 2, 0, 0, 22, 0, 10, 6, 0, 0, 2, 0)	1	6
(0, 0, 12, 0, 0, 2, 0, 23, 0, 9, 6, 0, 0, 2, 0)	1	6
(0, 0, 12, 0, 0, 0, 16, 0, 20, 0, 0, 0, 0, 4, 2)	1	6
(0, 0, 12, 0, 0, 0, 4, 18, 0, 14, 6, 0, 0, 0, 0)	1	6
(0, 0, 12, 0, 0, 0, 2, 20, 2, 12, 6, 0, 0, 0, 0)	1	6
(0, 0, 11, 3, 0, 2, 0, 25, 0, 4, 4, 0, 0, 0, 5)	1	6
(0, 0, 11, 2, 0, 0, 2, 26, 0, 7, 2, 0, 0, 2, 2)	1	6
(0, 0, 10, 9, 2, 0, 0, 20, 0, 0, 3, 0, 0, 0, 10)	1	6
(0, 0, 10, 5, 0, 0, 0, 17, 0, 15, 7, 0, 0, 0, 0)	1	6
(0, 0, 10, 0, 14, 0, 14, 0, 0, 0, 0, 0, 0, 16, 0)	1	6
(0, 0, 9, 7, 0, 0, 0, 21, 0, 10, 4, 0, 0, 0, 3)	1	6
(0, 0, 9, 7, 0, 0, 0, 18, 0, 13, 4, 0, 3, 0, 0)	1	6
(0, 0, 9, 4, 2, 0, 0, 26, 0, 7, 2, 0, 0, 2, 2)	1	6
(0, 0, 8, 24, 0, 0, 0, 0, 0, 0, 1, 0, 0, 20, 1)	1	6
(0, 0, 8, 3, 0, 8, 0, 26, 0, 0, 3, 0, 0, 2, 4)	1	6
(0, 0, 8, 0, 24, 0, 0, 0, 0, 0, 10, 0, 0, 8, 4)	1	6
(0, 0, 8, 0, 24, 0, 0, 0, 0, 0, 8, 0, 8, 0, 6)	1	6
(0, 0, 6, 6, 20, 0, 0, 2, 0, 0, 0, 8, 0, 12, 0)	1	6
(0, 0, 6, 4, 22, 0, 2, 0, 0, 0, 0, 8, 0, 12, 0)	1	6
(0, 0, 6, 0, 22, 0, 2, 8, 0, 0, 4, 0, 0, 12, 0)	1	6
(0, 0, 4, 25, 0, 0, 6, 0, 4, 0, 0, 0, 0, 0, 15)	1	6
(0, 0, 4, 24, 0, 4, 0, 4, 0, 0, 4, 0, 0, 0, 14)	1	6
(0, 0, 4, 21, 4, 0, 6, 4, 0, 0, 0, 0, 0, 0, 15)	1	6
(0, 0, 4, 21, 0, 4, 10, 0, 0, 0, 0, 0, 0, 0, 15)	1	6
(0, 0, 4, 16, 0, 0, 0, 0, 25, 3, 3, 0, 0, 3, 0)	1	6
(0, 0, 3, 20, 0, 0, 0, 0, 0, 23, 8, 0, 0, 0, 0)	1	6
(0, 0, 3, 18, 0, 0, 8, 0, 0, 19, 0, 0, 6, 0, 0)	1	6

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 3, 7, 0, 22, 0, 5, 0, 0, 0, 16, 0, 0, 1)	1	6
(0, 0, 2, 0, 26, 4, 6, 0, 0, 0, 0, 6, 0, 10, 0)	1	6
(0, 0, 2, 0, 0, 20, 24, 0, 0, 2, 0, 0, 0, 0, 6)	1	6
(0, 0, 1, 28, 0, 0, 0, 0, 1, 12, 1, 0, 0, 11, 0)	1	6
(0, 0, 1, 27, 0, 0, 2, 0, 1, 12, 0, 0, 0, 11, 0)	1	6
(0, 0, 1, 23, 0, 0, 0, 0, 0, 23, 5, 2, 0, 0, 0)	1	6
(0, 0, 1, 22, 0, 0, 2, 0, 0, 23, 5, 0, 1, 0, 0)	1	6
(0, 0, 1, 0, 0, 27, 15, 0, 0, 0, 0, 6, 4, 1, 0)	1	6
(0, 0, 0, 14, 0, 18, 8, 0, 0, 0, 0, 0, 8, 4, 2)	1	6
(0, 0, 0, 14, 0, 18, 4, 4, 0, 0, 0, 0, 10, 4, 0)	1	6
(0, 0, 0, 14, 0, 18, 4, 0, 0, 4, 0, 4, 10, 0, 0)	1	6
(0, 0, 0, 14, 0, 17, 0, 10, 0, 0, 0, 8, 0, 0, 5)	1	6
(0, 0, 0, 12, 0, 18, 12, 0, 0, 0, 0, 0, 0, 12, 0)	1	6
(0, 0, 0, 10, 4, 18, 4, 4, 0, 0, 0, 0, 14, 0, 0)	1	6
(0, 0, 0, 10, 0, 20, 8, 4, 0, 0, 0, 0, 10, 0, 2)	1	6
(0, 0, 0, 10, 0, 19, 0, 14, 0, 0, 0, 8, 0, 0, 3)	1	6
(0, 0, 0, 9, 0, 21, 0, 12, 0, 0, 0, 10, 0, 0, 2)	1	6
(0, 0, 0, 8, 0, 22, 4, 8, 0, 0, 2, 0, 10, 0, 0)	1	6
(0, 0, 0, 8, 0, 22, 0, 12, 0, 0, 4, 0, 8, 0, 0)	1	3
(0, 0, 0, 6, 20, 6, 0, 2, 6, 0, 0, 8, 0, 6, 0)	1	6
(0, 0, 0, 6, 8, 18, 8, 0, 0, 0, 0, 0, 14, 0, 0)	1	6
(0, 0, 0, 6, 0, 26, 8, 0, 0, 0, 0, 12, 0, 0, 2)	1	6
(0, 0, 0, 6, 0, 19, 0, 22, 0, 0, 4, 0, 0, 0, 3)	1	6
(0, 0, 0, 6, 0, 19, 0, 22, 0, 0, 3, 0, 0, 4, 0)	1	6
(0, 0, 0, 4, 0, 28, 8, 0, 0, 0, 2, 8, 4, 0, 0)	1	3
(0, 0, 0, 1, 0, 25, 20, 0, 0, 0, 0, 1, 0, 7, 0)	1	6
(0, 0, 0, 0, 29, 0, 1, 0, 0, 13, 10, 0, 0, 0, 1)	1	6
(0, 0, 0, 0, 28, 4, 6, 0, 0, 2, 0, 6, 0, 8, 0)	1	6
(0, 0, 0, 0, 28, 0, 2, 8, 0, 6, 4, 0, 0, 6, 0)	1	6
(0, 0, 0, 0, 28, 0, 0, 8, 0, 8, 6, 0, 0, 4, 0)	1	6
(0, 0, 0, 0, 28, 0, 0, 6, 0, 10, 6, 0, 4, 0, 0)	1	6
(0, 0, 0, 0, 27, 0, 7, 0, 0, 11, 6, 0, 0, 0, 3)	1	6
(0, 0, 0, 0, 27, 0, 5, 0, 0, 13, 3, 6, 0, 0, 0)	1	6
(0, 0, 0, 0, 26, 6, 2, 0, 6, 0, 4, 4, 0, 6, 0)	1	6
(0, 0, 0, 0, 26, 2, 0, 0, 6, 10, 10, 0, 0, 0, 0)	1	6
(0, 0, 0, 0, 26, 0, 6, 4, 4, 6, 0, 0, 8, 0, 0)	1	3
(0, 0, 0, 0, 26, 0, 0, 8, 2, 10, 6, 0, 2, 0, 0)	1	6

Table 7.1 Continued.

Vertex v	#	$ G \cdot v $
(0, 0, 0, 0, 26, 0, 0, 8, 0, 12, 6, 2, 0, 0, 0)	1	6
(0, 0, 0, 0, 24, 8, 0, 0, 0, 8, 6, 8, 0, 0, 0)	1	6
(0, 0, 0, 0, 24, 6, 6, 0, 6, 0, 0, 0, 12, 0, 0)	1	6
(0, 0, 0, 0, 24, 2, 0, 8, 0, 12, 8, 0, 0, 0, 0)	1	6
(0, 0, 0, 0, 24, 0, 7, 0, 6, 11, 3, 3, 0, 0, 0)	1	6
(0, 0, 0, 0, 24, 0, 4, 8, 0, 12, 5, 0, 0, 0, 1)	1	6
(0, 0, 0, 0, 24, 0, 4, 8, 0, 12, 2, 4, 0, 0, 0)	1	6
(0, 0, 0, 0, 22, 4, 6, 12, 0, 2, 0, 0, 0, 8, 0)	1	6
(0, 0, 0, 0, 22, 2, 10, 0, 14, 0, 0, 2, 0, 0, 4)	1	6
(0, 0, 0, 0, 20, 4, 12, 0, 12, 0, 0, 0, 0, 4, 2)	1	6
(0, 0, 0, 0, 20, 4, 6, 14, 4, 0, 0, 0, 0, 2, 4)	1	6
(0, 0, 0, 0, 18, 8, 6, 14, 0, 0, 0, 0, 0, 8, 0)	1	6
(0, 0, 0, 0, 18, 6, 14, 0, 0, 10, 0, 0, 6, 0, 0)	1	6
(0, 0, 0, 0, 16, 16, 8, 0, 0, 0, 2, 0, 12, 0, 0)	1	6
(0, 0, 0, 0, 16, 16, 0, 8, 0, 0, 10, 0, 0, 0, 4)	1	6
(0, 0, 0, 0, 16, 8, 10, 0, 14, 0, 0, 0, 6, 0, 0)	1	6
(0, 0, 0, 0, 14, 10, 7, 17, 0, 0, 0, 1, 0, 0, 5)	1	6
(0, 0, 0, 0, 12, 20, 6, 0, 0, 2, 0, 14, 0, 0, 0)	1	6
(0, 0, 0, 0, 12, 20, 0, 8, 0, 0, 10, 0, 0, 4, 0)	1	6
(0, 0, 0, 0, 12, 12, 6, 18, 0, 0, 0, 0, 0, 6, 0)	1	6
(0, 0, 0, 0, 8, 20, 14, 0, 2, 0, 0, 0, 10, 0, 0)	1	6
(0, 0, 0, 0, 6, 21, 12, 6, 0, 0, 0, 0, 9, 0, 0)	1	6
(0, 0, 0, 0, 2, 22, 14, 0, 10, 0, 0, 6, 0, 0, 0)	1	6
(0, 0, 0, 0, 1, 24, 21, 0, 1, 0, 0, 0, 0, 7, 0)	1	6
(0, 0, 0, 0, 0, 28, 15, 1, 0, 0, 3, 3, 0, 4, 0)	1	6
(0, 0, 0, 0, 0, 27, 9, 9, 0, 0, 0, 9, 0, 0, 0)	1	3
(0, 0, 0, 0, 0, 25, 20, 2, 0, 0, 0, 0, 4, 0, 3)	1	6
(0, 0, 0, 0, 0, 25, 19, 3, 0, 0, 0, 3, 0, 0, 4)	1	6
(0, 0, 0, 0, 0, 25, 16, 0, 6, 0, 0, 6, 0, 0, 1)	1	6
(0, 0, 0, 0, 0, 24, 20, 2, 2, 0, 0, 0, 0, 6, 0)	1	6
(0, 0, 0, 0, 0, 24, 18, 6, 0, 0, 0, 0, 3, 0, 3)	1	6
(0, 0, 0, 0, 0, 24, 16, 3, 0, 5, 0, 6, 0, 0, 0)	1	6
(0, 0, 0, 0, 0, 24, 15, 9, 0, 0, 0, 3, 0, 0, 3)	1	6
(0, 0, 0, 0, 0, 24, 12, 9, 0, 3, 0, 6, 0, 0, 0)	1	6

Table 7.1 Continued.

8. SOLVING SPARSE DECOMPOSABLE SYSTEMS

We describe how to use a numerical homotopy to solve polynomial systems corresponding to fibers of decomposable branched covers. Recall that a branched cover $\pi: X \rightarrow Z$ is decomposable if there is a dense open subset $V \subset Z$ over which π factors as

$$\pi^{-1}(V) \longrightarrow Y \longrightarrow V \tag{8.1}$$

with φ and ψ both nontrivial branched covers. As discussed in Section 4.4, a result of Pirola and Schlesinger [25] states that the Galois group G_π acts imprimitively if and only if π is decomposable.

Améndola and Rodriguez [28] explained how to use an explicit decomposition to compute fibers $\pi^{-1}(z)$ using monodromy. They also showed how several examples from the literature involve a decomposable branched cover; for these, the variety Y and intermediate maps were determined using invariant theory as there was a finite group acting as automorphisms of $\pi: X \rightarrow Z$. In general, it is nontrivial to determine a decomposition (8.1) of a branched cover $\pi: X \rightarrow Z$ with imprimitive Galois group, especially when the cover has trivial automorphism group.

Esterov [3] determined which systems of sparse polynomials have an imprimitive Galois group. One goal was to classify those which are solvable by radicals. He identified two simple structures which imply that the system is decomposable. In these cases, the decomposition is transparent. He also showed that the Galois group is full symmetric when neither structure occurs. We use Esterov’s classification to give a recursive numerical homotopy continuation algorithm for solving decomposable sparse systems.

The first such structure is when a polynomial system is composed with a monomial map. For example, if $f(x) = g(x^3)$ then to solve $f(x) = 0$, first solve $g(y) = 0$ and then for each solution y , extract its third roots. The second structure is when the system is triangular, such as

$$f(x, y) = g(y) = 0.$$

To solve this, first solve $g(y) = 0$ and then for each solution y , solve $f(x, y) = 0$.

In general, Esterov’s classification leads to a sequence of branched covers, each corresponding to a sparse system with symmetric monodromy or to a monomial map. Our algorithm identifies this structure and uses it to recursively solve a decomposable system. We give some examples which demonstrate that, despite its overhead, this algorithm is a significant improvement over a direct use of the polyhedral homotopy (Algorithm 6.3.6).

By the Bernstein-Kushnirenko Theorem (Proposition 5.3.1), a general system of sparse polynomials has the same number of solutions as a system whose supports have the same convex hull. When the system supported on the vertices is decomposable, we propose using it as a start system in a parameter homotopy to solve the original system. This is similar in spirit to the Bézout homotopy (Algorithm 6.3.2).

We remind the reader of the general background we developed in Section 4 on Galois groups of branched covers as well as our explanation of the relation between decompositions of branched covers and imprimitivity of the corresponding Galois groups.

In Section 8.1, we explain Esterov’s classification and describe how to compute the correspond-

ing decompositions in Section 8.2. We present our algorithms for solving sparse decomposable systems in Section 8.3, and give an application to furnish start systems for parameter homotopies. Section 8.5 gives timings and information on the performance of our algorithm. Much of the material in this section appears in the paper of the same name [67] with Rodriguez, Sottile, and Yahl.

8.1 Decompositions of sparse polynomial systems

Let $\mathcal{A}_\bullet = (\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n)$ be a collection of supports $\mathcal{A}_i \subset \mathbb{Z}^n$. We describe two properties that a collection \mathcal{A}_\bullet may have, lacunary and (strictly) triangular, and then state Esterov's theorem about the Galois group $G_{\mathcal{A}_\bullet}$. We then present explicit decompositions of the projection $\pi: X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ when \mathcal{A}_\bullet is lacunary and when \mathcal{A}_\bullet is triangular. These form the basis for our algorithms.

Assume that $MV(\mathcal{A}_\bullet) > 1$. We say that \mathcal{A}_\bullet is **lacunary** if $\mathbb{Z}\mathcal{A}_\bullet \neq \mathbb{Z}^n$ (it has rank n as $MV(\mathcal{A}_\bullet) \neq 0$). We say that \mathcal{A}_\bullet is **triangular** if there is a nonempty proper subset $\emptyset \neq I \subsetneq [n]$ such that $\text{rank}(\mathbb{Z}\mathcal{A}_I) = |I|$, or equivalently, the defect of the collection of polytopes $\{\text{conv}(\mathcal{A}_i)\}_{i \in I}$ is zero. As we explain in Section 8.2, we may change coordinates and assume that $\mathbb{Z}\mathcal{A}_I \subset \mathbb{Z}^{|I|}$ so that $MV(\mathcal{A}_I)$ is defined using $\text{conv}(\mathcal{A}_i) \subset \mathbb{R}^{|I|}$ for $i \in I$. A system \mathcal{A}_\bullet of triangular supports is **strictly triangular** if for some $\emptyset \neq I \subsetneq [n]$ with $\text{rank}(\mathbb{Z}\mathcal{A}_I) = |I|$, we have $1 < MV(\mathcal{A}_I) < MV(\mathcal{A}_\bullet)$. It is elementary that if \mathcal{A}_\bullet is either lacunary or strictly triangular, then the branched cover $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ is decomposable and therefore $G_{\mathcal{A}_\bullet}$ is an imprimitive permutation group. We do this explicitly in Sections 8.1.1 and 8.1.2.

Proposition 8.1.1 (Esterov [3]). *Let \mathcal{A}_\bullet be a collection of supports with $MV(\mathcal{A}_\bullet) \neq 0$. The Galois group $G_{\mathcal{A}_\bullet}$ is equal to the symmetric group $S_{MV(\mathcal{A}_\bullet)}$ if and only if \mathcal{A}_\bullet is neither lacunary nor strictly triangular.*

8.1.1 Lacunary support

Let us begin with an example when $n = 2$. Let

$$\mathcal{A}_1 = \begin{pmatrix} 0 & 0 & 3 & 6 & 12 \\ 0 & 4 & 3 & 6 & 0 \end{pmatrix} \quad \text{and} \quad \mathcal{A}_2 = \begin{pmatrix} 0 & 3 & 6 & 9 & 9 \\ 0 & 7 & 2 & 1 & 5 \end{pmatrix}$$

be supports in \mathbb{Z}^2 . Then $\mathbb{Z}\mathcal{A}_\bullet$ has index 12

in \mathbb{Z}^2 as the map $\varphi(a, b)^T = (3a, 4b - a)^T$ is an isomorphism $\varphi: \mathbb{Z}^2 \xrightarrow{\sim} \mathbb{Z}\mathcal{A}_\bullet$, and $\det\begin{pmatrix} 3 & 0 \\ -1 & 4 \end{pmatrix} = 12$. If we set $\mathcal{B}_i = \varphi^{-1}(\mathcal{A}_i)$, then

$$\mathcal{B}_1 = \begin{pmatrix} 0 & 0 & 1 & 2 & 4 \\ 0 & 1 & 1 & 2 & 1 \end{pmatrix} \quad \text{and} \quad \mathcal{B}_2 = \begin{pmatrix} 0 & 1 & 2 & 3 & 3 \\ 0 & 2 & 1 & 1 & 2 \end{pmatrix}.$$

We display $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1$, and \mathcal{B}_2 in Figure 8.1. Then the map $\Phi = \varphi^*: (\mathbb{C}^\times)^2 \twoheadrightarrow (\mathbb{C}^\times)^2$ is given by $\Phi(x, y) = (x^3y^{-1}, y^4) = (z, w)$. If

$$\begin{aligned} f_1 &= 1 + 2y^4 + 4x^3y^3 + 8x^6y^6 + 16x^{12} \\ f_2 &= 3 + 5x^3y^7 + 7x^6y^2 + 11x^9y + 13x^9y^5, \end{aligned}$$




Figure 8.1: The lacunary image $(\mathcal{A}_1, \mathcal{A}_2)$ of the support $(\mathcal{B}_1, \mathcal{B}_2)$ under the map φ .

which is a polynomial system with support \mathcal{A}_\bullet , then $f_i = g_i \circ \Phi$, where

$$\begin{aligned} g_1 &= 1 + 2w + 4zw + 8z^2w^2 + 16z^4w \\ g_2 &= 3 + 5zw^2 + 7z^2w + 11z^3w + 13z^3w^2, \end{aligned}$$

is a polynomial system with support \mathcal{B}_\bullet . Therefore, the branched cover $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ factors as $X_{\mathcal{A}_\bullet} \rightarrow X_{\mathcal{B}_\bullet} \rightarrow \mathbb{C}^{\mathcal{B}_\bullet} = \mathbb{C}^{\mathcal{A}_\bullet}$ with the map $X_{\mathcal{A}_\bullet} \rightarrow X_{\mathcal{B}_\bullet}$ induced by Φ . Consequently, this implies that $G_{\mathcal{A}_\bullet} \subset (\mathbb{Z}/12\mathbb{Z})^{10} \rtimes S_{10}$, as $\mathbb{Z}^2/\mathbb{Z}\mathcal{A}_\bullet \simeq \mathbb{Z}/12\mathbb{Z}$, \mathcal{B}_\bullet is neither lacunary nor triangular, and $\text{MV}(\mathcal{B}_\bullet) = 10$.

We generalize this example. Suppose that $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ is lacunary. Then $\mathbb{Z}\mathcal{A}_\bullet$ has rank n but $\mathbb{Z}\mathcal{A}_\bullet \neq \mathbb{Z}^n$. Let $\varphi: \mathbb{Z}^n \xrightarrow{\sim} \mathbb{Z}\mathcal{A}_\bullet$ be an isomorphism. Then the corresponding map $\Phi = \varphi^*: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$ is a surjection with kernel $\text{Hom}(\mathbb{Z}^n/\mathbb{Z}\mathcal{A}_\bullet, \mathbb{C}^\times)$. For each $i = 1, \dots, n$, set $\mathcal{B}_i = \varphi^{-1}(\mathcal{A}_i)$. Then $\mathcal{B}_\bullet = (\mathcal{B}_1, \dots, \mathcal{B}_n)$ is a collection of supports with $\mathbb{Z}\mathcal{B}_\bullet = \mathbb{Z}^n$. Since φ is a bijection, we identify $\mathbb{C}^{\mathcal{B}_i}$ with $\mathbb{C}^{\mathcal{A}_i}$ and $\mathbb{C}^{\mathcal{B}_\bullet}$ with $\mathbb{C}^{\mathcal{A}_\bullet}$. Given a system $F \in \mathbb{C}^{\mathcal{A}_\bullet}$, let $\iota(F) \in \mathbb{C}^{\mathcal{B}_\bullet}$ be the corresponding system with support \mathcal{B}_\bullet .

Lemma 8.1.2. *Suppose that \mathcal{A}_\bullet is lacunary, $\varphi: \mathbb{Z}^n \xrightarrow{\sim} \mathbb{Z}\mathcal{A}_\bullet$ is an isomorphism with corresponding surjection $\Phi: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$. Let $\mathcal{B}_\bullet = \varphi^{-1}(\mathcal{A}_\bullet)$ and suppose that $\text{MV}(\mathcal{B}_\bullet) > 1$. Then the branched cover $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ is decomposable and $X_{\mathcal{A}_\bullet} \rightarrow X_{\mathcal{B}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet} = \mathbb{C}^{\mathcal{B}_\bullet}$ is a nontrivial decomposition of branched covers induced by the map Φ .*

Proof. If g is a polynomial with support $\mathcal{B} \subset \mathbb{Z}^n$, then the composition $g \circ \Phi$ is a polynomial with support $\varphi(\mathcal{B})$, with the coefficient of x^β in g equal to the coefficient of $x^{\varphi(\beta)}$ in $g \circ \Phi$. Since $\varphi(\mathcal{B}_i) = \mathcal{A}_i$, this gives the natural identifications $\iota: \mathbb{C}^{\mathcal{A}_i} \xrightarrow{\sim} \mathbb{C}^{\mathcal{B}_i}$ and $\iota: \mathbb{C}^{\mathcal{A}_\bullet} \xrightarrow{\sim} \mathbb{C}^{\mathcal{B}_\bullet}$ mentioned before the lemma. Under this identification, we have $\iota(f)(\Phi(x)) = f(x)$.

Since $\text{MV}(\mathcal{B}_\bullet) > 1$, the branched cover $X_{\mathcal{B}_\bullet} \rightarrow \mathbb{C}^{\mathcal{B}_\bullet}$ is nontrivial by definition. The identification $\iota: \mathbb{C}^{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{B}_\bullet}$ extends to a commutative diagram

$$\begin{array}{ccc} X_{\mathcal{A}_\bullet} & \xrightarrow{\iota \times \Phi} & X_{\mathcal{B}_\bullet} \\ \pi \downarrow & & \downarrow \pi \\ \mathbb{C}^{\mathcal{A}_\bullet} & \xrightarrow{\iota} & \mathbb{C}^{\mathcal{B}_\bullet} \end{array} \quad (8.2)$$

where $\iota \times \Phi$ is the restriction of the map $\iota \times \Phi: \mathbb{C}^{\mathcal{A}_\bullet} \times (\mathbb{C}^\times)^n \rightarrow \mathbb{C}^{\mathcal{B}_\bullet} \times (\mathbb{C}^\times)^n$ to $X_{\mathcal{A}_\bullet}$. The map $\iota \times \Phi: X_{\mathcal{A}_\bullet} \rightarrow X_{\mathcal{B}_\bullet}$ is a map of branched covers with $\ker \Phi$ acting freely on the fibers. If

we restrict the diagram (8.2) to the open subset V of $\mathbb{C}^{\mathcal{B}_\bullet}$ over which $X_{\mathcal{B}_\bullet} \rightarrow \mathbb{C}^{\mathcal{B}_\bullet}$ is a covering space, we obtain a composition of covering spaces with $\ker \Phi$ acting as deck transformations on $\pi^{-1}(V) \subset X_{\mathcal{A}_\bullet}$. Thus $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ is decomposable. \square

8.1.2 Triangular support

This requires more discussion before we can state the analog of Lemma 8.1.2. Let us begin with an example when $n = 3$. Suppose that

$$\mathcal{A}_1 = \mathcal{A}_2 = \mathcal{A} = \begin{pmatrix} 0 & 1 & 1 & 1 & 2 & 2 & 2 & 3 \\ 0 & 0 & 1 & 2 & 0 & 1 & 2 & 1 \\ 0 & 1 & 2 & 3 & 2 & 3 & 4 & 4 \end{pmatrix} \quad \text{and} \quad \mathcal{A}_3 = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 2 & 4 & 5 & 3 & 4 \end{pmatrix}.$$

The span $\mathbb{Z}\mathcal{A}$ of the first two supports is isomorphic to \mathbb{Z}^2 , with $\varphi(a, b)^T \mapsto (a, b, a + b)^T$ an isomorphism $\varphi: \mathbb{Z}^2 \xrightarrow{\sim} \mathbb{Z}\mathcal{A}_\bullet$. Set $\mathcal{B} = \varphi^{-1}(\mathcal{A})$. We display \mathcal{A} , \mathcal{A}_3 , and \mathcal{B} in the horizontal plane together on the left in Figure 8.2, and \mathcal{B} on the right. Consider the polynomial system

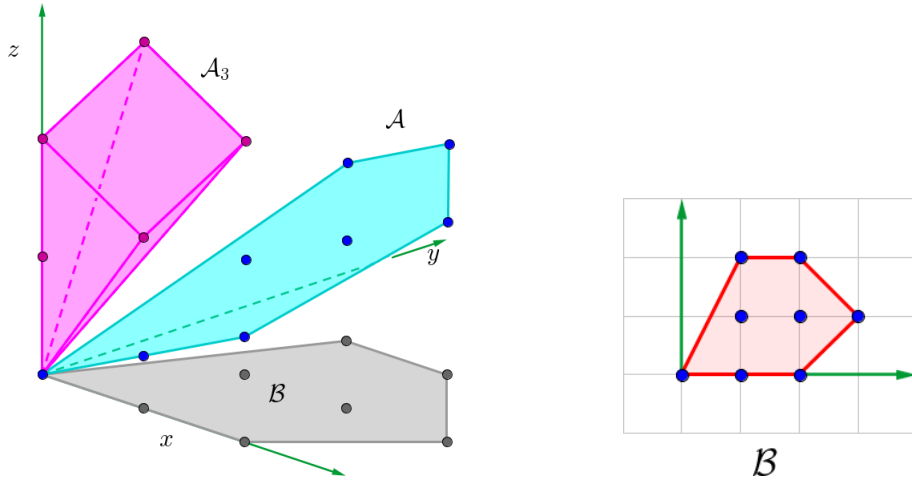


Figure 8.2: An example of triangular support.

$F = (f_1, f_2, f_3) \in \mathbb{C}[x, y, z]$ with support \mathcal{A}_\bullet ,

$$\begin{aligned} f_1 &= 1 + 2xz + 3xyz^2 + 4xy^2z^3 + 5x^2z^2 + 6x^2yz^3 + 7x^2y^2z^4 + 8x^3yz^4 \\ f_2 &= 2 + 3xz + 5xyz^2 + 7xy^2z^3 + 11x^2z^2 + 13x^2yz^3 + 17x^2y^2z^4 + 19x^3yz^4 \\ f_3 &= 1 + 3z^2 + 9z^4 + 27yz^5 + 81xz^3 + 243xyz^4. \end{aligned}$$

Let $\Phi: (\mathbb{C}^\times)^3 \rightarrow (\mathbb{C}^\times)^2$ be given by $\Phi(x, y, z) = (xz, yz) = (u, v)$. If

$$\begin{aligned} g_1 &= 1 + 2u + 3uv + 4uv^2 + 5u^2 + 6u^2v + 7u^2v^2 + 8u^3v \\ g_2 &= 2 + 3u + 5uv + 7uv^2 + 11u^2 + 13u^2v + 17u^2v^2 + 19u^3v, \end{aligned}$$

then $f_i = g_i \circ \Phi$ for $i = 1, 2$. To compute $\mathcal{V}(F)$, we first may compute $\mathcal{V}(g_1, g_2)$ which consists of eight points. For each solution $(u_0, v_0) \in \mathcal{V}(g_1, g_2)$, we may identify the fiber $\Phi^{-1}(u_0, v_0)$ with \mathbb{C}^\times by $z \mapsto (u_0 z^{-1}, v_0 z^{-1}, z)$. Then the restriction of f_3 to this fiber is

$$1 + (3 + 81u_0 + 243u_0v_0)z^2 + (9 + 27v_0)z^4,$$

which is a lacunary univariate polynomial with support $\{0, 2, 4\}$, and has four solutions (counted with multiplicity) when $v_0 \neq -1/3$.

This example generalizes to all triangular systems. Suppose that $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ is triangular. Let $\emptyset \neq I \subsetneq [n]$ be a proper subset witnessing the triangularity, so that $\text{rank}(\mathbb{Z}\mathcal{A}_I) = |I|$. Set $J = [n] \setminus I$. Let

$$\mathbb{Z}^I = \mathbb{Q}\mathcal{A}_I \cap \mathbb{Z}^n = \{v \in \mathbb{Z}^n \mid \exists m \in \mathbb{N} \text{ with } mv \in \mathbb{Z}\mathcal{A}_I\},$$

be the saturation of $\mathbb{Z}\mathcal{A}_I$, which is a free abelian group of rank $|I|$. As it is saturated, $\mathbb{Z}_J = \mathbb{Z}^n / \mathbb{Z}^I$ is free abelian of rank $n - |I| = |J|$.

Applying $\text{Hom}(\bullet, \mathbb{C}^\times)$ to the short exact sequence $\mathbb{Z}^I \hookrightarrow \mathbb{Z}^n \twoheadrightarrow \mathbb{Z}_J$ gives the short exact sequence of tori (whose characters are $\mathbb{Z}_J, \mathbb{Z}^n$, and \mathbb{Z}^I) with indicated maps,

$$(\mathbb{C}^\times)^{|J|} \simeq \mathbb{T}_J := \text{Hom}(\mathbb{Z}_J, \mathbb{C}^\times) \longleftarrow (\mathbb{C}^\times)^n \xrightarrow{\Phi} \mathbb{T}^I := \text{Hom}(\mathbb{Z}^I, \mathbb{C}^\times) \simeq (\mathbb{C}^\times)^{|I|}. \quad (8.3)$$

A polynomial f with support in \mathbb{Z}^I determines polynomial functions on $(\mathbb{C}^\times)^n$ and on \mathbb{T}^I with the first the pullback of the second. Let f be a polynomial on $(\mathbb{C}^\times)^n$ with support $\mathcal{A} \subset \mathbb{Z}^n$. Then its restriction to a fiber $\Phi^{-1}(y_0)$ of Φ is a regular function \bar{f} on the fiber, which is a coset of \mathbb{T}_J . Choosing an identification of $\mathbb{T}_J \simeq \Phi^{-1}(y_0)$, we obtain a polynomial \bar{f} on \mathbb{T}_J whose support is the image $\bar{\mathcal{A}}$ of \mathcal{A} in $\mathbb{Z}_J = \mathbb{Z}^n / \mathbb{Z}^I$. This polynomial \bar{f} depends upon the identification of the fiber with \mathbb{T}_J . Let $\bar{\mathcal{A}}_J$ be the image in \mathbb{Z}_J of the collection \mathcal{A}_J of supports. Then we have the product formula (see [68, Lem. 6] or [3, Thm. 1.10])

$$\text{MV}(\mathcal{A}_\bullet) = \text{MV}(\mathcal{A}_I) \cdot \text{MV}(\bar{\mathcal{A}}_J). \quad (8.4)$$

Since $\mathcal{A}_\bullet = \mathcal{A}_I \sqcup \mathcal{A}_J$, we have the identification $\mathbb{C}^{\mathcal{A}_\bullet} = \mathbb{C}^{\mathcal{A}_I} \oplus \mathbb{C}^{\mathcal{A}_J}$. Suppose that $F \in \mathbb{C}^{\mathcal{A}_\bullet}$ is a polynomial system with support \mathcal{A}_\bullet . Write $F_I \in \mathbb{C}^{\mathcal{A}_I}$ for its restriction to the indices in I , and the same for F_J . We have the diagram

$$\begin{array}{ccc} X_{\mathcal{A}_\bullet} & \xrightarrow{p_I \times \Phi} & X_{\mathcal{A}_I} \\ \pi \downarrow & & \downarrow \pi \\ \mathbb{C}^{\mathcal{A}_\bullet} & \xrightarrow{p_I} & \mathbb{C}^{\mathcal{A}_I} \end{array} \quad (8.5)$$

where $p_I \times \Phi$ is the restriction of the map $p_I \times \Phi: \mathbb{C}^{\mathcal{A}_\bullet} \times (\mathbb{C}^\times)^n \rightarrow \mathbb{C}^{\mathcal{A}_I} \times \mathbb{T}^I$ to $X_{\mathcal{A}_\bullet}$.

Let $V_{\mathcal{A}_\bullet} \subset \mathbb{C}^{\mathcal{A}_\bullet}$ be the dense open subset over which $X_{\mathcal{A}_\bullet}$ is a covering space. This is the set of polynomial systems F with support \mathcal{A}_\bullet which have exactly $\text{MV}(\mathcal{A}_\bullet)$ solutions in $(\mathbb{C}^\times)^n$. Similarly, let $V_{\mathcal{A}_I} \subset \mathbb{C}^{\mathcal{A}_I}$ be the subset where $X_{\mathcal{A}_I} \rightarrow \mathbb{C}^{\mathcal{A}_I}$ is a covering space. We will show that under the projection $\mathbb{C}^{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_I}$, the image of $V_{\mathcal{A}_\bullet}$ is a subset of $V_{\mathcal{A}_I}$. Define $Y_{\mathcal{A}_\bullet} \rightarrow V_{\mathcal{A}_\bullet}$ to be the restriction of $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ to the dense open set $V_{\mathcal{A}_\bullet}$. Also define $Y_{\mathcal{A}_I} \rightarrow V_{\mathcal{A}_I}$ to be the pullback of

$X_{\mathcal{A}_I} \rightarrow \mathbb{C}^{\mathcal{A}_I}$ along the map $V_{\mathcal{A}_\bullet} \rightarrow V_{\mathcal{A}_I}$. Write $\Phi: Y_{\mathcal{A}_\bullet} \rightarrow Y_{\mathcal{A}_I}$ for the map induced by Φ .

Lemma 8.1.3. *Suppose that \mathcal{A}_\bullet is a triangular set of supports in \mathbb{Z}^n witnessed by $I \subsetneq [n]$. Then $Y_{\mathcal{A}_\bullet} \rightarrow Y_{\mathcal{A}_I} \rightarrow V_{\mathcal{A}_\bullet}$ a composition of covering spaces. If $1 < \text{MV}(\mathcal{A}_I) < \text{MV}(\mathcal{A}_\bullet)$, then this decomposition is nontrivial, so that $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ is decomposable.*

Furthermore, each fiber of the map $Y_{\mathcal{A}_\bullet} \rightarrow Y_{\mathcal{A}_I}$ may be identified with the set of solutions of a polynomial system with support $\overline{\mathcal{A}_J}$.

Proof. Let $F \in V_{\mathcal{A}_\bullet}$. Then its number of solutions is $\#\mathcal{V}(F) = \text{MV}(\mathcal{A}_\bullet)$. If $x \in \mathcal{V}(F)$, then $\Phi(x) \in \mathbb{T}^I$ is a solution of $f_i = 0$ for $i \in I$. Thus $\Phi(\mathcal{V}(F)) \subset \mathcal{V}(F_I)$, the latter being the solutions of F_I on \mathbb{T}^I . For any $y \in \mathcal{V}(F_I)$, if we choose an identification $\mathbb{T}_J \simeq \Phi^{-1}(y)$ of the fiber, then the restriction of F to $\Phi^{-1}(y)$ is the system $\overline{F}_J = \{\overline{f}_j \mid j \in J\}$. By the Bernstein-Kushnirenko Theorem, this has at most $\text{MV}(\overline{\mathcal{A}_J})$ solutions. By the product formula (8.4) and our assumption on $\#\mathcal{V}(F)$, we conclude that the system F_I has $\text{MV}(\mathcal{A}_I)$ solutions, and for each $y \in \mathcal{V}(F_I)$, the system \overline{F}_J has $\text{MV}(\overline{\mathcal{A}_J})$ solutions.

In particular, this implies that the image of $V_{\mathcal{A}_\bullet}$ in $\mathbb{C}^{\mathcal{A}_I}$ is a subset of $V_{\mathcal{A}_I}$. As $V_{\mathcal{A}_\bullet}$ is open and dense in $\mathbb{C}^{\mathcal{A}_\bullet}$, its image contains an open dense subset. This proves the assertion that $Y_{\mathcal{A}_\bullet} \rightarrow Y_{\mathcal{A}_I} \rightarrow V_{\mathcal{A}_\bullet}$ is a decomposition of covering spaces. We have already shown that each fiber of the map $Y_{\mathcal{A}_\bullet} \rightarrow Y_{\mathcal{A}_I}$ is a polynomial system with support $\overline{\mathcal{A}_J}$ with exactly $\text{MV}(\overline{\mathcal{A}_J})$ solutions. Thus when $1 < \text{MV}(\mathcal{A}_I) < \text{MV}(\mathcal{A}_\bullet)$, we have $\text{MV}(\overline{\mathcal{A}_J}) > 1$, which shows that this decomposition is nontrivial. \square

8.2 Computing the decompositions

We show how to compute the decompositions of $X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$ from Section 8.1 when \mathcal{A}_\bullet is either lacunary or strictly triangular.

Let us consider the Smith normal form (see Section 5.1.3),

$$\mathcal{A} = PDQ, \tag{8.6}$$

when \mathcal{A} is the matrix whose columns are the vectors in \mathcal{A}_\bullet and $\text{MV}(\mathcal{A}_\bullet) > 0$. Then $d_n > 0$ as $\mathbb{Z}\mathcal{A}_\bullet$ has rank n , and \mathcal{A}_\bullet is lacunary when $d_n > 1$. In this case, an identification $\varphi: \mathbb{Z}^n \xrightarrow{\sim} \mathbb{Z}\mathcal{A}_\bullet$ is given by PD_n , where D_n is the principal $n \times n$ submatrix of D . Recall that the corresponding surjection $\varphi^* = \Phi: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$ has kernel $\text{Hom}(\mathbb{Z}^n/\mathbb{Z}\mathcal{A}_\bullet, \mathbb{C}^\times)$. Let $\psi = P^{-1}$. Then $\psi \circ \varphi = D_n$, so that if we set $\Psi = \psi^*$, then $\Phi \circ \Psi: (\mathbb{C}^\times)^n \rightarrow (\mathbb{C}^\times)^n$ is diagonal,

$$\Phi \circ \Psi(x_1, \dots, x_n) = (x_1^{d_1}, \dots, x_n^{d_n}). \tag{8.7}$$

Let $y = (y_1, \dots, y_n) \in (\mathbb{C}^\times)^n$. If we set $\rho_i = |y_i|$ and $\zeta_i = \arg(z_i)$ so that $y_i = \rho_i e^{\sqrt{-1}\zeta_i}$, then $(\Phi \circ \Psi)^{-1}(y)$ is the set

$$\left\{ \left(\rho_1^{1/d_1} e^{\sqrt{-1}\theta_1}, \dots, \rho_n^{1/d_n} e^{\sqrt{-1}\theta_n} \right) \mid \theta_i = \frac{\zeta_i + 2\pi j}{d_i} \text{ for } j = 0, \dots, d_i - 1 \right\} \tag{8.8}$$

as explained in Section 5.1.3.

Suppose that \mathcal{A}_\bullet is triangular, and let us use the notation of Section 8.1.2. We suppose that $I = [k] = \{1, \dots, k\}$ and $J = \{k+1, \dots, n\}$. Given a polynomial f on $(\mathbb{C}^\times)^n$, its restriction \overline{f} to a fiber of $\Phi: (\mathbb{C}^\times)^n \rightarrow \mathbb{T}^I$ is a regular function on the fiber, which is isomorphic to \mathbb{T}_J . To

represent \bar{f} as a polynomial on \mathbb{T}_J depends on the choice of a point in that fiber. Indeed, suppose that $f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha$. Let $y \in \mathbb{T}^I$ and $y_0 \in \Phi^{-1}(y)$ be a point in the fiber above y , so that $\mathbb{T}_J \ni z \mapsto y_0 z \in \Phi^{-1}(y)$ parameterizes $\Phi^{-1}(y)$. If we write $\bar{\alpha}$ for the image of $\alpha \in \mathbb{Z}^n$ in $\mathbb{Z}_J = \mathbb{Z}^n / \mathbb{Z}^I$, then

$$\bar{f}(z) = \sum_{\alpha \in \mathcal{A}} c_\alpha (y_0 z)^\alpha = \sum_{\beta \in \bar{\mathcal{A}}} z^\beta \left(\sum_{\alpha \in \mathcal{A} \text{ with } \bar{\alpha}=\beta} c_\alpha y_0^\alpha \right). \quad (8.9)$$

A uniform choice of a point in each fiber is given by fixing a splitting $\mathbb{T}^I \hookrightarrow (\mathbb{C}^\times)^n$ of the map $\Phi: (\mathbb{C}^\times)^n \rightarrow \mathbb{T}^I$. This gives an identification $(\mathbb{C}^\times)^n = \mathbb{T}^I \times \mathbb{T}_J$. Then points $y \in \mathbb{T}^I$ are canonical representatives of cosets of \mathbb{T}_J . As $k = |I|$, we may further fix isomorphisms $\mathbb{T}^I \simeq (\mathbb{C}^\times)^k$ giving $\mathbb{Z}^I \simeq \mathbb{Z}^k$ and $\mathbb{T}_J \simeq (\mathbb{C}^\times)^{n-k}$ giving $\mathbb{Z}_J \simeq \mathbb{Z}^{n-k}$.

Suppose now that $\mathcal{A} = \mathcal{A}_I$, and we compute a decomposition (8.6). Since $\mathbb{Z}\mathcal{A}_I$ has rank k , the diagonal matrix D has k nonzero invariant factors. The saturation L of $\mathbb{Z}\mathcal{A}_I$ is the image of PI_k , where I_k is the $n \times n$ matrix whose only nonzero entries are in its principal $k \times k$ submatrix, which forms an identity matrix. Then $\varphi = PI_k$ and $\Phi = \varphi^*$. Applying the coordinate change $\psi = P^{-1}$ to \mathbb{Z}^n identifies this saturation as the coordinate plane $\mathbb{Z}^k \oplus \mathbf{0}^{n-k}$ and the lattice $\mathbb{Z}\mathcal{A}_I$ as $d_1\mathbb{Z} \oplus d_2\mathbb{Z} \oplus \cdots \oplus d_k\mathbb{Z} \oplus \mathbf{0}^{n-k}$. As in Section 8.1.2, this identifies \mathbb{Z}/L with the complementary coordinate plane, $\mathbf{0}^k \oplus \mathbb{Z}^{n-k}$. Setting $\Psi = \psi^*$, the composition $\Phi \circ \Psi$ is the projection to the first k coordinates,

$$\Phi \circ \Psi : (\mathbb{C}^\times)^n \longrightarrow (\mathbb{C}^\times)^k \quad (8.10)$$

and we identify $\mathbb{T}_J = 1^k \times (\mathbb{C}^\times)^{n-k}$ and $\mathbb{T}^I = (\mathbb{C}^\times)^k \times 1^{n-k}$.

8.3 Solving decomposable sparse systems

We describe algorithms that use Esterov's conditions to solve sparse decomposable systems and suggest an application for computing a start system for solving a general (not necessarily decomposable) sparse polynomial system. In each, we let `SOLVE` be an arbitrary algorithm for solving a polynomial system. We assume that the system F to be solved is general given its support \mathcal{A}_\bullet in that it has $MV(\mathcal{A}_\bullet)$ solutions in $(\mathbb{C}^\times)^n$. If not, then one may instead solve a general polynomial system with support \mathcal{A}_\bullet and then use a parameter homotopy together with endgames to compute $\mathcal{V}(F)$. Recall the identification in (8.2) for Algorithm 8.3.1 and the notation $F_I \in \mathbb{C}^{\mathcal{A}_I}$ used in (8.5) for Algorithm 8.3.2.

Algorithm 8.3.1 (SolveLacunary).**Input:**

- A general polynomial system F whose support \mathcal{A}_\bullet is lacunary.

Output:

- All solutions $\mathcal{V}(F) \subset (\mathbb{C}^\times)^n$

Steps:

- 1 Compute the Smith normal form (8.6) of \mathcal{A}_\bullet , giving $\varphi = PD_n$, $\Phi = \varphi^*$, $\psi = P^{-1}$, and $\Psi = \psi^*$, so that $\Phi \circ \Psi$ is diagonal (8.7)
- 2 Use SOLVE to compute $\mathcal{V}(\iota(F)) \subset (\mathbb{C}^\times)^n$
- 3 Using the formula (8.8) to compute $(\Phi \circ \Psi)^{-1}(y)$ for $y \in \mathcal{V}(\iota(F))$, **return**

$$\left\{ \Psi(w) \mid w \in \bigcup_{z \in \mathcal{V}(\iota(F))} (\Phi \circ \Psi)^{-1}(z) \right\}$$

Proof of Correctness. By Lemma 8.1.2, $\mathcal{V}(F) = \Phi^{-1}(\mathcal{V}(\iota(F)))$. We apply Ψ to the points of $(\Phi \circ \Psi)^{-1}(z)$ for $z \in \mathcal{V}(\iota(F))$ to obtain points of $\mathcal{V}(F)$ in their original coordinates. \square

Algorithm 8.3.2 (SolveTriangular).**Input:**

- A general polynomial system F whose support \mathcal{A}_\bullet is triangular, witnessed by $0 < k < n$ such that $\text{rank}(\mathbb{Z}\mathcal{A}_{[k]}) = k$

Output:

- All solutions of $\mathcal{V}(F) \subset (\mathbb{C}^\times)^n$

Steps:

- 1 Compute the Smith normal form (8.6) of $\mathcal{A}_{[k]}$, giving $\varphi = PI_k$, $\Phi = \varphi^*$, $\psi = P^{-1}$, and $\Psi = \psi^*$, so that $\Phi \circ \Psi$ is the projection (8.10)
- 2 Use SOLVE to compute $\mathcal{V}(F_{[k]}) \subset (\mathbb{C}^\times)^k$
- 3 Choose $y_0 \in \mathcal{V}(F_{[k]})$ Use SOLVE to compute the points of the fiber $(\Phi \circ \Psi)^{-1}(y_0)$ in $Y_{\mathcal{A}_\bullet}$, which are $\mathcal{V}(\overline{F}_J) \subset \{y_0\} \times (\mathbb{C}^\times)^{n-k}$, where \overline{F}_J has support $\overline{\mathcal{A}}_J$ and $J = [n] \setminus [k]$
- 4 **for** each $y \in \mathcal{V}(F_{[k]})$ use a parameter homotopy with start system $\mathcal{V}(\overline{F}_J)$ to compute $(\Phi \circ \Psi)^{-1}(y)$ and **return**

$$\left\{ \Psi(w) \mid w \in \bigcup_{y \in \mathcal{V}(F_{[k]})} (\Phi \circ \Psi)^{-1}(y) \right\}$$

Proof of Correctness. By Lemma 8.1.3, every solution $x \in \mathcal{V}(F)$ lies over a solution $y = \Phi(x)$ to $F_{[k]}$ in $(\mathbb{C}^\times)^k$. As explained in Section 8.2, the map $\Phi \circ \Psi$ is a coordinate projection and $(\Phi \circ \Psi)^{-1}(y) = \mathcal{V}(\overline{F}_J)$. Here, $\overline{F}_J = (\overline{f}_{k+1}, \dots, \overline{f}_n)$ where \overline{f}_j has support $\overline{\mathcal{A}}_j$ and is computed using (8.9). We apply Ψ to convert these points to the original coordinates. \square

Our main algorithm takes a sparse system and checks Esterov’s criteria for decomposability. If the system is decomposable, the algorithm calls Algorithm 8.3.1 (if lacunary) or Algorithm 8.3.2 (if triangular), and in each of these algorithms calls to the solver SOLVE are assumed to be recursive calls back to Algorithm 8.3.3. If the polynomial system is indecomposable, then Algorithm 8.3.3 calls a black box solver BLACKBOX.

Algorithm 8.3.3 (SolveDecomposable).

Input:

- A generic polynomial system F with support \mathcal{A}_\bullet .

Output:

- All solutions of $\mathcal{V}(F) \subset (\mathbb{C}^\times)^n$

Steps:

- 1 Compute the Smith normal form PDQ (8.6) of \mathcal{A}_\bullet .
- 2 **if** $d_n > 1$, then **return** SolveLacunary(F)
- 3 **if** $d_n = 1$, then
 - 3.1 **for** all $\emptyset \neq I \subsetneq [n]$ compute the Smith normal form $PD_I Q$ (8.6) of \mathcal{A}_I
 - 3.2 **if** $\text{rank}(D_I) = |I|$ for some I , reorder so $I = [k]$ and **return** SolveTriangular(F, k)
 - 3.3 **else** neither of Esterov’s conditions hold and **return** BLACKBOX(F)

Proof of Correctness. First note that if the algorithm halts, then it returns the solutions $\mathcal{V}(F)$. Halting is clear in Case (3), but the other cases involve recursive calls back to Algorithm 8.3.3. In Case (1), SolveLacunary will call Algorithm 8.3.3 on a system $\iota(F)$ whose mixed volume is less than $MV(\mathcal{A}_\bullet)$. In Case (2), SolveTriangular will call Algorithm 8.3.3 on systems $F_{[k]}$ and \overline{F}_J , each involving fewer variables than F . Thus, in each recursive call back to Algorithm 8.3.3, either the mixed volume or the number of variables decreases, which proves that the algorithm halts. \square

8.4 Start systems

The start system in the Bézout homotopy (Algorithm 6.3.1) is a highly decomposable sparse polynomial system consisting of supports which are subsets of the original support of F , but have the same mixed volume. We propose a generalization, in which Algorithm 8.3.3 is used to compute a start system.

Example 8.4.1. Suppose that we have supports $\mathcal{A}_1 = \mathcal{A}_2 = \mathcal{A}$, shown in Figure 8.3 which are given by the columns of the matrix $\begin{pmatrix} 0 & 0 & 1 & 1 & 2 & 3 & 3 & 3 & 4 & 5 & 5 & 6 \\ 0 & 2 & 0 & 1 & 3 & 0 & 1 & 4 & 2 & 3 & 4 & 4 \end{pmatrix}$. Then $MV(\mathcal{A}_1, \mathcal{A}_2) = 2! \text{vol}(\text{conv}(\mathcal{A})) = 30$. Let $\mathcal{B}_1 = \mathcal{B}_2 = \begin{pmatrix} 0 & 0 & 3 & 3 & 6 \\ 0 & 2 & 0 & 4 & 4 \end{pmatrix}$ be the set of vertices of $\text{conv}(\mathcal{A})$. Given a general system $F \in \mathbb{C}^{\mathcal{A}_\bullet}$, let $G \in \mathbb{C}^{\mathcal{B}_\bullet} \subset \mathbb{C}^{\mathcal{A}_\bullet}$ be obtained from F by restriction to the monomials in \mathcal{B} . (That is, we set coefficients of monomials x^α in F to zero if $\alpha \notin \mathcal{B}$.) Then \mathcal{B}_\bullet is lacunary with the map $\Phi(x_1, x_2) = (x_1^3, x_2^2)$, and $\iota(G)$ has five solutions. We may use Algorithm 8.3.3 (more specifically, Algorithm 8.3.1) to compute $\mathcal{V}(G)$, and then compute $\mathcal{V}(F)$ using the straight-line homotopy $H(t; x)$ with start system $G = H(1; x)$ and tracking from the solutions $\mathcal{V}(G)$ at $t = 1$. \diamond

Example 8.4.1 motivates our final algorithm. For a collection $\mathcal{A}_\bullet = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ of supports, let $\text{vert}(\mathcal{A}_\bullet) = (\text{vert}(\mathcal{A}_1), \dots, \text{vert}(\mathcal{A}_n))$ where $\text{vert}(\mathcal{A}_i) = \text{vert}(\text{conv}(\mathcal{A}_i))$. Note that if $G \in$




Figure 8.3: A support \mathcal{A} such that $(\mathcal{A}, \mathcal{A})$ is neither triangular nor lacunary.

$\mathbb{C}^{\text{vert}(\mathcal{A}_\bullet)}$ is a regular value of the branched cover $\pi|_{X_{\text{vert}(\mathcal{A}_\bullet)}} : X_{\text{vert}(\mathcal{A}_\bullet)} \rightarrow \mathbb{C}^{\text{vert}(\mathcal{A}_\bullet)}$ then G is also a regular value of $\pi : X_{\mathcal{A}_\bullet} \rightarrow \mathbb{C}^{\mathcal{A}_\bullet}$. As such, G may be taken as a start system for a straight-line homotopy and used to compute $\mathcal{V}(F)$ for any $F \in \mathbb{C}^{\mathcal{A}_\bullet}$ with $\mathcal{V}(F)$ finite. The benefit of this approach is that $\pi|_{X_{\text{vert}(\mathcal{A}_\bullet)}}$ decomposes if π does. Therefore, as seen in Example 8.4.1, $\pi|_{X_{\text{vert}(\mathcal{A}_\bullet)}}$ is more likely (and no less likely) than π to be decomposable.

Algorithm 8.4.2 (Decomposable Start System).

Input:

- A set \mathcal{A}_\bullet of supports

Output:

- A start system G for a homotopy coming from $\pi_{\mathcal{A}_\bullet}$ and start solutions $\mathcal{V}(G)$

Steps:

- 1 Choose a general system $G \in \mathbb{C}^{\text{vert}(\mathcal{A}_\bullet)}$
- 2 Compute $\mathcal{V}(G)$ using Algorithm 8.3.3
- 3 **return** the pair $(G, \mathcal{V}(G))$

Proof of Correctness. As $G \in \mathbb{C}^{\text{vert}(\mathcal{A}_\bullet)}$ is general, it has $\text{MV}(\text{vert}(\mathcal{A}_\bullet))$ solutions. Since for each i , $\text{conv}(\mathcal{A}_i) = \text{conv}(\text{vert}(\mathcal{A}_i))$, we have $\text{MV}(\text{vert}(\mathcal{A}_\bullet)) = \text{MV}(\mathcal{A}_\bullet)$. Finally, $\mathbb{C}^{\text{vert}(\mathcal{A}_\bullet)}$ is the subspace of $\mathbb{C}^{\mathcal{A}_\bullet}$ where the coefficients of non-extreme monomials in each polynomial are zero. Thus $G \in \mathbb{C}^{\mathcal{A}_\bullet}$, which shows that $(G, \mathcal{V}(G))$ is a start system for \mathcal{A}_\bullet . \square

Remark 8.4.3. The Bézout homotopy motivated Algorithm 8.4.2. However, if we apply Algorithm 8.4.2 to the system of supports \mathcal{A}_\bullet , where \mathcal{A}_i consists of all monomials of degree at most d_i , then we will not get the start system for the Bézout homotopy. For example, when $n = 2$, $d_1 = 2$, and $d_2 = 3$, the supports are as shown in Figure 8.4. Here, \mathcal{B}_1 and \mathcal{B}_2 are the supports of the start system for the Bézout homotopy.

We leave open the challenge of finding a simple, general method to replace each set \mathcal{A}_i by a subset (or superset) \mathcal{B}_i of \mathcal{A}_i , so that $\text{MV}(\mathcal{A}_\bullet) = \text{MV}(\mathcal{B}_\bullet)$ and $\pi : X_{\mathcal{B}_\bullet} \rightarrow \mathbb{C}^{\mathcal{B}_\bullet}$ is decomposable.

A possible first step would be to take advantage of the results on monotonicity developed in Section 2.5. For example, if $\mathcal{A}_1 = \begin{pmatrix} 1 & 3 & 1 & 3 & 2 \\ 1 & 1 & 3 & 3 & 4 \end{pmatrix}$ and $\mathcal{A}_2 = \mathcal{B}_1 = \mathcal{B}_2 = \begin{pmatrix} 2 & 0 & 2 & 4 \\ 0 & 2 & 4 & 2 \end{pmatrix}$ then $\mathcal{A}_\bullet = (\mathcal{A}_1, \mathcal{A}_2)$ is neither lacunary nor triangular, but $\mathcal{B}_\bullet = (\mathcal{B}_1, \mathcal{B}_2)$ is lacunary. Moreover, $\text{MV}(\mathcal{A}_\bullet) = \text{MV}(\mathcal{B}_\bullet) = 8$ and so a general sparse polynomial system supported on \mathcal{A}_\bullet corresponds to a regular value of $\pi_{\mathcal{B}_\bullet}$.




Figure 8.4: Dense support $(\mathcal{A}_1, \mathcal{A}_2)$, the support $(\text{vert}(\mathcal{A}_1), \text{vert}(\mathcal{A}_2))$, and the support of the Bézout start system.

Thus, one may solve a general sparse decomposable system on \mathcal{B}_\bullet , and subsequently solve a system supported on \mathcal{A}_\bullet via a parameter homotopy. \diamond

8.5 A computational experiment

We explored the computational cost of using Algorithm 8.3.3 to solve sparse decomposable systems, comparing timings to **PHCPack** [49, 69] on a family of related systems.

Let $\mathcal{A}_1 = \begin{pmatrix} 0 & 1 & 2 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \end{pmatrix}$, $\mathcal{A}_2 = \begin{pmatrix} 1 & 0 & 1 & 2 & 1 \\ 0 & 1 & 1 & 1 & 2 \end{pmatrix}$, $\mathcal{B}_1 = \begin{pmatrix} 0 & 2 & 0 & 2 \\ 0 & 0 & 1 & 3 \end{pmatrix}$, and $\mathcal{B}_2 = \begin{pmatrix} 0 & 1 & 2 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 & 1 & 2 \end{pmatrix}$. We display these supports and their convex hulls in Figure 8.5. Let $\mathcal{C} = \{0, 1\}^5$ be the vertices of




Figure 8.5: The four supports involved in a computational experiment.

the five-dimensional cube. We construct sparse decomposable systems from $\mathcal{A}_\bullet = (\mathcal{A}_1, \mathcal{A}_2)$, $\mathcal{B}_\bullet = (\mathcal{B}_1, \mathcal{B}_2)$, and \mathcal{C} as follows.

Choose two injections $\iota, j: \mathbb{Z}^2 \rightarrow \mathbb{Z}^5$ such that $\iota(\mathbb{Z}^2) \cap j(\mathbb{Z}^2) = \{0\}$. For example, choose four linearly independent vectors $\iota_1, \iota_2, j_1, j_2 \in \mathbb{Z}^5$, and define $\iota(a, b) = a\iota_1 + b\iota_2$, and the same for j . Let us set

$$\mathcal{A}(\iota, j) = (\iota(\mathcal{A}_1), \iota(\mathcal{A}_2), j(\mathcal{B}_1), j(\mathcal{B}_2), \mathcal{C}).$$

Example 8.5.1. We now illustrate Algorithm 8.3.3 in detail on $\mathcal{A}(\iota, j)$ by considering the case when $\iota_1, \iota_2, j_1, j_2$ are the first four standard unit vectors e_1, \dots, e_4 . Suppose $F = (f_1, f_2, g_1, g_2, h)$ is a system of polynomials $\mathbb{C}[x_1, x_2, y_1, y_2, z]$ with support $\mathcal{A}(e_1, e_2, e_3, e_4)$. We use superscripts to

distinguish different calls of the same algorithm. When $\text{SolveDecomposable}^{(1)}(F)$ is called, it first checks if F is lacunary (it is not as $\mathbb{ZC} = \mathbb{Z}^5$), and then recognizes that F is triangular witnessed by (f_1, f_2) . As such, it calls $\text{SolveTriangular}^{(1)}(F, 2)$ which computes the $MV(\mathcal{A}_\bullet) = 5$ solutions p_1, \dots, p_5 to $\mathcal{V}(f_1, f_2)$ with **PHCPack**, our choice of BLACKBOX.

As its penultimate task, $\text{SolveTriangular}^{(1)}$ computes a fiber of the first solution p_1 by performing the substitution $(x_1, x_2) = p_1$ in g_1, g_2 and h , and recursively calls $\text{SolveDecomposable}^{(2)}$ on the system $(g_1(p_1, y, z), g_2(p_1, y, z), h(p_1, y, z)) \in \mathbb{C}[y_1, y_2, z]$. This system is recognized to be triangular witnessed by (g_1, g_2) and $\text{SolveTriangular}^{(2)}(g_1, g_2)$ computes the $MV(\mathcal{B}_\bullet) = 10$ solutions q_1, \dots, q_{10} using **PHCPack**. Next, $\text{SolveTriangular}^{(2)}$ computes a fiber above q_1 by performing the substitution $y = (y_1, y_2) = q_1$ in $h(p_1, y, z)$ producing the univariate polynomial $h(p_1, q_1, z)$ of degree 1 which has solution (p_1, q_1, z_1) . Finally, $\text{SolveTriangular}^{(2)}$ performs a parameter homotopy from q_1 to q_i to populate the fibers above each q_i . Thus $\text{SolveTriangular}^{(1)}$ populates the fiber above p_1 consisting of $10 \cdot 1 = 10$ solutions. As its final step, $\text{SolveTriangular}^{(1)}$ uses parameter homotopies from p_1 to p_i to populate all fibers producing all $5 \cdot 10 = 50$ solutions of $\mathcal{V}(F)$.

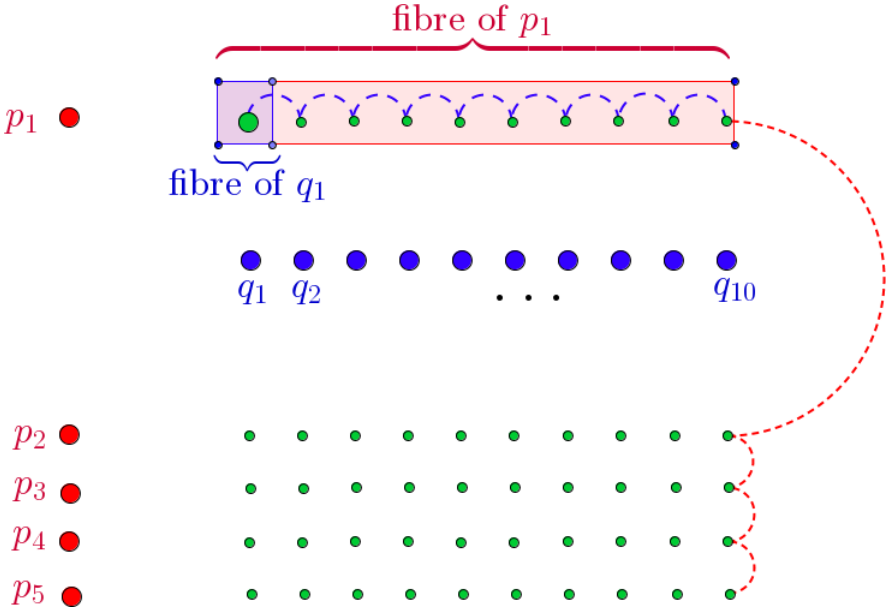


Figure 8.6: A schematic of the process in Example 8.5.1.

Figure 8.6 depicts a schematic of this process. Red objects correspond to the function $\text{solveDecomposable}^{(1)}$ and blue objects correspond to $\text{solveDecomposable}^{(2)}$. The largest points represent solutions which were computed directly. The dotted lines represent the use of monodromy to move fibers. \diamond

The overhead of this algorithm includes the computation of Smith normal forms and the search for subsets witnessing triangularity. Additionally, it often requires more path-tracking than a direct

use of **PHCPack**. Nonetheless, the overhead seems nominal, and compared to the paths tracked in **PHCPack**, the paths tracked in our algorithm either involve fewer variables or polynomials of smaller degree.

For example, in Example 8.5.1, our algorithm called **PHCPack** to solve two sparse polynomial systems with 5 and 10 solutions respectively. A parameter homotopy was called $10 - 1 = 9$ times on a system with 1 solution, then a different parameter homotopy was called $5 - 1 = 4$ times on a system with 10 solutions. In total, $5 + 10 + 9 + 40 = 64$ individual paths were tracked. In contrast, a direct use of **PHCPack** involves tracking exactly $MV(\mathcal{A}(e_1, e_2, e_3, e_4)) = 50$ paths, albeit in a higher dimensional space.

For more general ι and j , the recursive structure of our computation is similar to Example 8.5.1. Some notable differences include

- (1) $\iota(\mathcal{A}_\bullet)$ or $j(\mathcal{B}_\bullet)$ may be lacunary which induces further decompositions.
- (2) Monomial changes must be computed as $\iota(\mathcal{A}_\bullet)$ or $j(\mathcal{B}_\bullet)$ could involve all variables.
- (3) For most ι, j the univariate polynomial obtained from h has degree 5 and is solved by computing eigenvalues of its companion matrix.

For example, if we choose $e_1 - e_2, e_2 - e_3, e_3 - e_4, e_4 - e_5$ for $\iota_1, \iota_2, j_1, j_2$, then again, no system in the algorithm is lacunary, but the univariate polynomial obtained from h has support $\{0, 1, 2, 3, 4, 5\}$, so that $MV(\mathcal{A}(\iota, j)) = 250$.




Figure 8.7: Scatter plot of timings

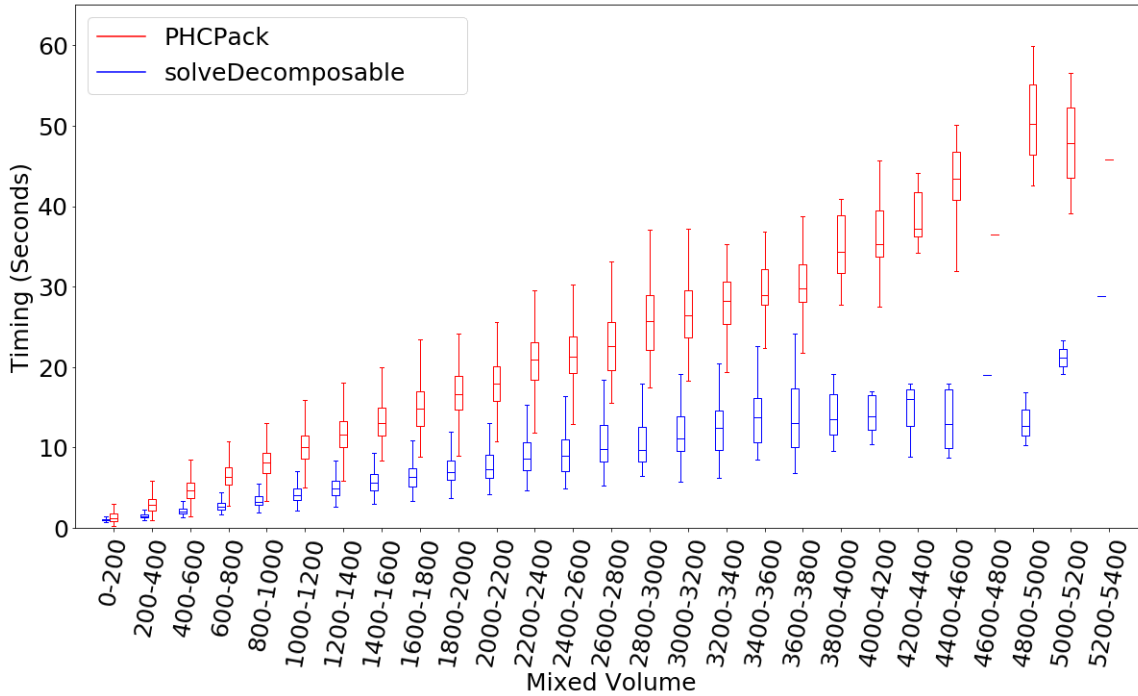


Figure 8.8: Box plot of timings

In our computational experiment, we produced 13563 instances of $\mathcal{A}(l, j)$ and solved each instance using our implementation of Algorithm 8.3.3 as well as with **PHCPack**. Due to ill-conditioning and heuristic choices of tolerances, some computations failed to produce all solutions. We only include the 10962 instances such that both **PHCPack** and Algorithm 8.3.3 computed all solutions.

We give a scatter plot of the elapsed timings in Figure 8.7 with respect to the mixed volume of the system. Figure 8.8 displays box plots of the timings of each algorithm grouped by sizes of mixed volumes. The boxes range from the first quartile q_1 to the third quartile q_3 of the group data with whiskers extending to the smallest and largest data points which are not outliers. Outliers are the data points which are smaller than $q_1 - 1.5I$ or larger than $q_3 + 1.5I$ where I is the length of the interquartile range (q_1, q_3).

A more detailed account of these computations, along with our implementation in **Macaulay2**, may be found at the website [70].

9. SUMMARY

Newton polytopes provide a rich combinatorial structure by which we may delineate polynomials and thus polynomial systems. The geometric nature of numerical algebraic geometry lends itself to algorithms which can extract and use this combinatorial data via the HS-algorithm and polyhedral homotopy algorithm, respectively. We augment both of these algorithms.

Using the HS-algorithm as a subroutine, we develop a tropical membership algorithm. We implement both the HS-algorithm and the tropical membership algorithm in a computer algebra system and analyze their convergence rates. We use the HS-algorithm to completely identify a large polynomial defining a hypersurface from algebraic vision. With the same software, we also determine many vertices of the Lüroth polytope.

We augment the polyhedral homotopy by developing and implementing an algorithm which recognizes when a sparse polynomial system is decomposable. It then uses this decomposition to numerically and recursively solve the sparse system. We compare timings of our software against the use of a polyhedral homotopy.

REFERENCES

- [1] T. Brysiewicz, “Numerical Software to Compute Newton Polytopes and Tropical Membership,” *Mathematics in Computer Science*, 2020.
- [2] B. Huber and B. Sturmfels, “A polyhedral method for solving sparse polynomial systems,” *Mathematics of Computation*, vol. 64, no. 212, pp. 1541–1555, 1995.
- [3] A. Esterov, “Galois theory for general systems of polynomial equations,” *Compositio Mathematica*, vol. 155, no. 2, pp. 229–245, 2019.
- [4] K. Hept and T. Theobald, “Tropical bases by regular projections,” *Proceedings of the American Mathematical Society*, vol. 137, no. 7, pp. 2233–2241, 2009.
- [5] G. M. Ziegler, *Lectures on polytopes*. New York: Springer, 1995.
- [6] I. A. Emiris, V. Fisikopoulos, C. Konaxis, and L. Penaranda, “An oracle-based, output-sensitive algorithm for projections of resultant polytopes,” *International Journal of Computational Geometry and Applications*, vol. 23, no. 04n05, 2013.
- [7] H. Minkowski, “Theorie der konvexen körper, insbesondere begründung ihres oberflächenbegriffs,” *Gesammelte Abhandlungen*, vol. II, pp. 131–229, 1911.
- [8] G. Ewald, *Combinatorial convexity and algebraic geometry*, vol. 168 of *Graduate Texts in Mathematics*. Springer, New York, 1996.
- [9] R. J. Steffens, *Mixed volumes, mixed Ehrhart theory and applications to tropical geometry and linkage configurations*. PhD thesis, Goethe Universität, Goethe Universität Frankfurt, 2009.
- [10] J. De Loera, J. Rambau, and F. Santos, *Triangulations: Structures for Algorithms and Applications*. Algorithms and Computation in Mathematics, Springer Berlin Heidelberg, 2010.
- [11] J. Rojas, “A convex geometric approach to counting the roots of a polynomial system,” *Theoretical Computer Science*, vol. 133, pp. 105–140, 10 1994.
- [12] A. Esterov, “Indices of 1-forms and newton polyhedra,” *Mathematicheskii Sbornik*, vol. 197, no. 7, pp. 1085–1108, 2006.
- [13] T. Chen, “Unmixing the mixed volume computation,” *Discrete & Computational Geometry*, vol. 62, pp. 55–86, 2019.
- [14] B. Frédéric and I. Soprunov, “Criteria for strict monotonicity of the mixed volume of convex polytopes,” *Advances in Geometry*, 02 2017.
- [15] D. Cox, J. Little, and D. O’Shea, *Ideals, Varieties and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer, 1991.
- [16] J. Harris, *Algebraic Geometry: A First Course*. Graduate Texts in Mathematics, Springer, 1992.
- [17] R. Hartshorne, *Algebraic Geometry*. Graduate Texts in Mathematics, Springer New York, 2013.
- [18] M. Reid and I. Shafarevich, *Basic Algebraic Geometry I*. Springer Berlin Heidelberg, 2013.

- [19] D. Hilbert, “über die theorie der algebraischen formen,” *Mathematische Annalen*, vol. 36, pp. 473–530, 1890.
- [20] D. Hilbert, “über die vollen invariantensysteme,” *Mathematische Annalen*, vol. 42, pp. 313–373, 1893.
- [21] H. Wielandt, *Finite permutation groups*. Translated from the German by R. Bercov, Academic Press, New York-London, 1964.
- [22] A. Hatcher, *Algebraic topology*. Cambridge: Cambridge University Press, 2002.
- [23] J. Harris, “Galois groups of enumerative problems,” *Duke Mathematical Journal*, vol. 46, no. 4, pp. 685–724, 1979.
- [24] C. Hermite, “Sur les fonctions algébriques,” *Comptes rendus de l’Académie des Sciences (Paris)*, vol. 32, pp. 458–461, 1851.
- [25] G. P. Pirola and E. Schlesinger, “Monodromy of projective curves,” *Journal of Algebraic Geometry*, vol. 14, no. 4, pp. 623–642, 2005.
- [26] H. Derksen and G. Kemper, *Computational invariant theory*. Invariant Theory and Algebraic Transformation Groups, I, Springer, Berlin, 2002. Encyclopaedia of Mathematical Sciences, 130.
- [27] J. R. Munkres, *Topology: a first course*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1975.
- [28] C. Améndola and J. Rodriguez, “Solving parameterized polynomial systems with decomposable projections,” 2016. arXiv:1612.08807.
- [29] A. Martín del Campo-Sanchez, F. Sottile, and R. Williams, “Classification of Schubert Galois groups in $Gr(4, 9)$.” arXiv.org/1902.06809, 2019.
- [30] F. Sottile, R. Williams, and L. Ying, “Galois groups of compositions of Schubert problems.” arXiv.org/1910.06843, 2019.
- [31] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*, vol. 161. Providence, RI: American Mathematical Society, 2015.
- [32] R. Bieri and J. Groves, “The geometry of the set of characters induced by valuations.,” *Journal für die reine und angewandte Mathematik*, vol. 347, pp. 168–195, 1984.
- [33] A. Chan, “Gröbner bases over fields with valuation and tropical curves by coordinate projections,” *Ph.D. thesis, University of Warwick*, 2013.
- [34] J. Nocedal and S. J. Wright, *Numerical Optimization*. New York, NY, USA: Springer, second ed., 2006.
- [35] L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and real computations*. Springer, New York, 1998.
- [36] S. Smale, “Newton’s method estimates from data at one point,” in *The merging of disciplines: new directions in pure, applied, and computational mathematics*, pp. 185–196, Springer, New York, 1986.
- [37] J. D. Hauenstein and F. Sottile, “Algorithm 921: alphacertified: certifying solutions to polynomial systems,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 38, no. 4, p. 28, 2012.

- [38] K. Lee, “Numerical certification.” Distributed with Macaulay 2.
- [39] C. Beltrán and A. Leykin, “Certified numerical homotopy tracking,” *Experimental Mathematics*, vol. 21, no. 1, pp. 69–83, 2012.
- [40] M. Burr, C. Yap, and J. Xu, “An approach for certifying homotopy continuation paths: Univariate case,” *In Proceedings of the 43rd International Symposium on Symbolic and Algebraic Computation*, pp. 399–406, 2018.
- [41] S. Telen, M. Van Barel, and J. Verschelde, “A robust numerical path tracking algorithm for polynomial homotopy continuation,” 2019. arXiv:1909.04984 .
- [42] P. Bürgisser and F. Cucker, *Condition: The geometry of numerical algorithms*, vol. 349. Springer Science & Business Media, 2013.
- [43] D. Davidenko, “Ob odnom novom metode chislennoy resheniya sistem nelineynykh uravnenii,” *Doklady Akademii Nauk SSR*, vol. 87, no. 4, pp. 601–602, 1953.
- [44] T. Akoglu, J. D. Hauenstein, and A. Szanto, “Certifying solutions to overdetermined and singular polynomial systems over \mathbb{Q} ,” *Journal of Symbolic Computation*, vol. 84, pp. 147–171, 2018.
- [45] T. Duff, N. Hein, and F. Sottile, “Certification for polynomial systems via square subsystems,” 2019. arXiv:1812.02851.
- [46] T. Y. Li, T. Sauer, and J. A. Yorke, “The cheater’s homotopy: an efficient procedure for solving systems of polynomial equations,” *SIAM Journal on Numerical Analysis*, vol. 26, no. 5, pp. 1241–1251, 1989.
- [47] A. P. Morgan and A. J. Sommese, “Coefficient-parameter polynomial continuation,” *Applied Mathematics and Computation*, vol. 29, no. 2, part II, pp. 123–160, 1989.
- [48] C. B. Garcia and W. I. Zangwill, “Finding all solutions to polynomial systems and other systems of equations,” *Mathematical Programming*, vol. 16, no. 1, pp. 159–176, 1979.
- [49] J. Verschelde, “Algorithm 795: PHCpack: A general-purpose solver for polynomial systems by homotopy continuation,” *ACM Transactions on Mathematical Software*, vol. 25, no. 2, pp. 251–276, 1999. Available at <http://www.math.uic.edu/~jan>.
- [50] D. J. Bates, J. D. Hauenstein, A. J. Sommese, and C. W. Wampler, *Numerically solving polynomial systems with Bertini*. SIAM, 2013.
- [51] D. J. Bates, J. D. Hauenstein, A. J. Sommese, and C. W. Wampler, “Bertini: Software for numerical algebraic geometry.” Available at bertini.nd.edu with permanent doi: [dx.doi.org/10.7274/R0H41PB5](https://doi.org/10.7274/R0H41PB5).
- [52] P. Breiding and S. Timme, “Homotopycontinuation.jl - a package for solving systems of polynomial equations in julia,” *Mathematical Software ICMS 2018, Lecture Notes in Computer Science*, 2018. Available at juliahomotopycontinuation.org.
- [53] T.-L. Lee, T. Li, and C. Tsai, “Hom4ps-2.0: A software package for solving polynomial systems by the polyhedral homotopy continuation method,” *Computing*, vol. 83, pp. 109–133, 2008.
- [54] A. Leykin, “Numerical algebraic geometry for macaulay2.” <http://people.math.gatech.edu/aleykin3/NAG4M2>.

- [55] J. D. Hauenstein and A. J. Sommese, “Witness sets of projections,” *Applied Mathematics and Computation*, vol. 217, no. 7, pp. 3349–3354, 2010.
- [56] O. Zariski, “A theorem on the poincaré group of an algebraic hypersurface,” *Annals of Mathematics*, vol. 38, no. 1, pp. 131–141, 1937.
- [57] J. D. Hauenstein, J. Rodriguez, and S. F., “Numerical computation of galois groups,” *Foundations of Computational Mathematics*, vol. 18, pp. 867–890, 2018.
- [58] T. Duff, C. Hill, A. Jensen, K. Lee, A. Leykin, and J. Sommars, “Solving polynomial systems via homotopy continuation and monodromy,” *IMA Journal of Numerical Analysis*, 2018.
- [59] J. Dixon, “The probability of generating the symmetric group,” *Math Z*, vol. 110, pp. 199–205, 1969.
- [60] T. Hayes and L. Babai, “The probability of generating the symmetric group when one of the generators is random,” *Publicationes Mathematicae Debrecen*, vol. 69, pp. 271–280, 10 2006.
- [61] J. D. Hauenstein and F. Sottile, “Newton polytopes and witness sets,” *Mathematics in Computer Science*, vol. 8, no. 2, pp. 235–251, 2012.
- [62] D. R. Grayson and M. E. Stillman, “Macaulay2, a software system for research in algebraic geometry.” Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [63] T. Brysiewicz, “Numerical computations of Newton polytopes.” Available at <http://www.math.tamu.edu/~brysiewicz/NumericalNP>, 2018.
- [64] D. J. Bates, E. Gross, A. Leykin, and J. Rodriguez, “Bertini for Macaulay2,” Oct. 2013.
- [65] W. Stein *et al.*, *Sage Mathematics Software (Version x.y.z)*. The Sage Development Team, 2017. <http://www.sagemath.org>.
- [66] J. Ponce, B. Sturmfels, and M. Trager, “Congruences and concurrent lines in multi-view geometry,” *Advances in Applied Mathematics*, vol. 88, pp. 62–91, 2017.
- [67] T. Brysiewicz, J. Rodriguez, F. Sottile, and T. Yahl, “Solving Decomposable Sparse Systems,” *arXiv:2001.04228*, 2019.
- [68] R. Steffens and T. Theobald, “Mixed volume techniques for embeddings of Laman graphs,” *Computational Geometry. Theory and Applications*, vol. 43, no. 2, pp. 84–93, 2010.
- [69] E. Gross, S. Petrović, and J. Verschelde, “Interfacing with PHCpack,” *The Journal of Software for Algebra and Geometry*, vol. 5, pp. 20–25, 2013.
- [70] T. Brysiewicz, J. Rodriguez, F. Sottile, and T. Yahl, “Software for decomposable sparse polynomial systems,” 2020. <https://www.math.tamu.edu/~thomasjyahl/research/DSS/DSSsite.html>.