

Tropical and non-Archimedean curves

by

Ralph Elliott Morrison

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Bernd Sturmfels, Chair

Professor David Eisenbud

Professor Claire J. Tomlin

Professor Xinyi Yuan

Spring 2015

Tropical and non-Archimedean curves

Copyright 2015
by
Ralph Elliott Morrison

Abstract

Tropical and non-Archimedean curves

by

Ralph Elliott Morrison

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Bernd Sturmfels, Chair

Tropical geometry is a young field of mathematics that connects algebraic geometry and combinatorics. It considers “combinatorial shadows” of classical algebraic objects, which preserve information while being more susceptible to discrete methods. Tropical geometry has proven useful in such subjects as polynomial implicitization, scheduling problems, and phylogenetics. Of particular interest in this work is the application of tropical geometry to study curves (and other varieties) over non-Archimedean fields, which can be *tropicalized* to tropical curves (and other tropical varieties).

Chapter 1 presents background material on tropical geometry, and presents two perspectives on tropical curves: the *embedded* perspective, which treats them as balanced polyhedral complexes in Euclidean space, and the *abstract* perspective, which treats them as metric graphs. This chapter also presents the background on curves over non-Archimedean fields necessary for the rest of this work, including the *moduli space of curves* of a given genus and the *Berkovich analytic space* associated to a curve.

Chapters 2 and 3 study tropical curves embedded in the plane. Chapter 2 deals with tropical plane curves that intersect non-transversely, and opens with a result on which configurations of points in such an intersection can be lifted to intersection points of classical curves. It then moves on to present a joint work with Matthew Baker, Yoav Len, Nathan Pflueger, and Qingchun Ren that builds up a theory of bitangents of smooth tropical plane quartic curves in parallel to the classical theory.

Chapter 3 presents joint work with Sarah Brodsky, Michael Joswig, and Bernd Sturmfels, and is a study of which metric graphs arise as skeletons of smooth tropical plane curves. We begin by defining *the moduli space of tropical plane curves*, which is the tropical analog of Castryck and Voight’s space of *nondegenerate curves* in [CV09]. The first main theorem is that our space is full-dimensional inside of the tropicalization of the corresponding classical space, a result proved using *honeycomb curves*. The chapter proceeds to a computational study of the moduli space of tropical plane curves, and explicitly computes the spaces for genus up to 5. The chapter closes with both theoretical and computational results on tropical hyperelliptic curves that can be embedded in the plane.

Chapter 4 presents joint work with Qingchun Ren and is an algorithmic treatment of a special family of curves over a non-Archimedean field called *Mumford curves*. These are of particular interest in tropical geometry, as they are the curves whose tropicalizations can have genus-many cycles. We build up a family of algorithms, implemented in `sage` [S⁺13], for computing many objects associated to such a curve over the field of p -adic numbers, including its Jacobian, its Berkovich skeleton, and points in its canonical embedding.

Chapter 5 is joint work with Ngoc Tran, and is a departure from studying tropical curves. It considers what it means for matrix multiplication to commute tropically, both in the context of tropical linear algebra and by considering the tropicalization of the classical *commuting variety*, whose points are pairs of commuting matrices. We give necessary and sufficient conditions for small matrices to commute, and illustrate three different tropical spaces, each of which has some claim to being “the” space of tropical commuting matrices.

Acknowledgments

I would like to thank Bernd Sturmfels for his support and guidance throughout my time at Berkeley. I am deeply grateful for his mentorship and dedication to his students, and for his outstanding ability to set mathematical collaborations into motion. I would also like to express gratitude for many other scholars who helped shape my early mathematical career through guidance and helpful conversations, including Matthew Baker, Melody Chan, Maria Angelica Cueto, Michael Joswig, Madhusudan Manjunath, Hannah Markwig, Sam Payne, Joe Rabinoff, Annette Werner, and Josephine Yu. I thank my many coauthors (Matthew Baker, Sarah Brodsky, Michael Joswig, Yoav Len, Nathan Pflueger, Qingchun Ren, Bernd Sturmfels, and Ngoc Tran), who have kindly agreed to allow adaptations of our joint work to appear in this thesis.

Finally, I would like to thank my parents James and Ruth, my sisters Elizabeth and Jane, and the rest of my family for all their support. They have always been an important part of my mathematical life, from teaching me about Fibonacci numbers to putting up with me telling them about different sizes of infinity.

Contents

Contents	ii
List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Tropical geometry	1
1.2 Embedded tropical curves	3
1.3 Abstract tropical curves	5
1.4 Curves over non-Archimedean fields	9
2 Non-proper tropical intersections	11
2.1 Intersections of plane curves	11
2.2 Examples of non-proper intersections	17
2.3 Bitangents of tropical plane quartic curves	21
3 Smooth tropical plane curves	35
3.1 The moduli space of tropical plane curves	35
3.2 Honeycomb curves	38
3.3 Computational methods	41
3.4 Computations of $\mathbb{M}_g^{\text{planar}}$	46
3.5 Tropical hyperelliptic curves in the plane: theory	59
3.6 Tropical hyperelliptic curves in the plane: computation	65
4 Algorithms for Mumford curves	74
4.1 The theory of Mumford curves	74
4.2 Good fundamental domains	77
4.3 Computations from good starting data	82
4.4 Finding good starting data	95
5 The tropical commuting variety	100
5.1 Tropical linear algebra	100

5.2 Three tropical commuting spaces	106
Bibliography	113

List of Figures

1.1	A tropical quadric curve in \mathbb{R}^2	2
1.2	Inducing a subdivision on a Newton polygon, dual to a tropical curve	4
1.3	A non-regular triangulation of a polygon	5
1.4	Unimodular triangulation, tropical quartic, and skeleton	6
1.5	Three abstract tropical curves, of genus 1, 3, and 0	6
1.6	The graph of a rational function on an abstract tropical curve	7
1.7	An illustration of chip-firing	8
2.1	Tropical curves intersecting non-transversely	12
2.2	Possible images of intersection points under tropicalization	13
2.3	Graphs of tropical rational functions on a tropical intersection	14
2.4	Two quadrics intersecting non-transversely	17
2.5	The graph of a rational function on a quadric intersection	18
2.6	A moduli space of intersection configurations, with seven examples	19
2.7	The self-intersection of a tropical cubic	20
2.8	An intersection configuration on a cubic	21
2.9	The five trivalent graphs of genus 3	22
2.10	Smooth tropical plane quartic curves and the dual triangulations	23
2.11	The construction of a theta characteristic, and a corresponding bitangent line	26
2.12	Two divisors on genus 3 graphs	28
2.13	The form of triangulations, and a dual cycle	29
2.14	Genus 3 graphs, with hyperelliptic and nonhyperelliptic metrics	31
2.15	Edges in a triangulation and the dual cycle	32
2.16	The 7 theta characteristics for each genus 3 graph	34
3.1	A honeycomb triangulation, the dual tropical curve, and its skeleton	39
3.2	Full-dimensional honeycomb triangulations	41
3.3	Triangulations for genus 2 curves	45
3.4	The five trivalent graphs of genus 3, with edges labeled	47
3.5	A triangulation for full-dimensional genus 3 (000) graphs	49
3.6	Triangulations for lower-dimensional genus 3 (000) graphs	49
3.7	A triangulation for full-dimensional genus 3 (020) graphs	50

3.8	Triangulations for lower-dimensional genus 3 (020) graphs	50
3.9	Triangulations for full-dimensional genus 3 (111) graphs	51
3.10	Triangulations for lower-dimensional genus 3 (111) graphs	51
3.11	Triangulations for genus 3 (212) graphs	51
3.12	Nonhyperelliptic genus 4 polygons, and a triangulation	53
3.13	Triangulations of a square genus 4 graphs with flexible edge lengths	53
3.14	The 17 trivalent graphs of genus 4	54
3.15	The maximal genus 5 polygons	56
3.16	The 38 genus 5 graphs in $\mathbb{M}_5^{\text{planar}}$	57
3.17	A genus 5 graph realized by a single triangulation	57
3.18	The 33 genus 5 graphs not in $\mathbb{M}_5^{\text{planar}}$	58
3.19	A triangle forced by a sprawling graph	60
3.20	A graph that isn't crowded, but has a crowded subgraph	62
3.21	A 2-connected hyperelliptic graph that is not a chain	63
3.22	The standard embedding of a chain	64
3.23	A partial triangulation of a nonhyperelliptic polygon	65
3.24	The five maximal hyperelliptic polygons of genus 3	66
3.25	Hyperelliptic triangulations giving rise to the same metrics	67
3.26	The start of the construction of the triangulation Δ'	67
3.27	Several steps leading to Δ'	68
3.28	A triangulation Δ highlighting the point $(2, 1)$, with the corresponding σ 's in Δ and s 's in the dual tropical cycle.	68
3.29	A labelled skeleton of a chain tropical curve with $g = 6$	69
3.30	Triangulations for hyperelliptic metrics on the genus 3 (020)	72
3.31	Triangulations for hyperelliptic metrics on the genus 3 (111) and (212)	73
4.1	A tree giving rise to a dumbbell graph skeleton	89
4.2	A tree giving rise to a theta graph skeleton	89
4.3	A tree giving rise to a K_4 graph skeleton	90
4.4	The Newton polygon of a plane quartic, and the dual tropical curve	94
5.1	The action of a polytrope on \mathbb{TP}^2	102
5.2	The images of 3-dimensional polytropes	103
5.3	The 2×2 tropical commuting variety	108
5.4	The three tropical commuting spaces for 2×2 matrices	109
5.5	The three tropical commuting spaces for 3×3 matrices and beyond	109

List of Tables

3.1	Dimensions of the 1278 moduli cones in \mathbb{M}_{T_4}	48
3.2	Triangulation counts and dimensions for the graphs of genus 4	55
3.3	All cones from triangulations of the maximal nonhyperelliptic genus 4 polygons	56
3.4	Dimensions of the moduli cones for two genus 3 hyperelliptic polygons	71

Chapter 1

Introduction

In this chapter we review the basics of tropical geometry, focusing on tropical curves from both embedded and abstract viewpoints. We also review background on curves over non-Archimedean fields and the associated Berkovich spaces. For more background, see [BPR12] and [MS15].

1.1 Tropical geometry

Tropical geometry is concerned with polynomials over the *tropical semiring* $(\overline{\mathbb{R}}, \oplus, \odot)$. Here, $\overline{\mathbb{R}}$ is the usual real numbers together with an element ∞ , defined to be greater than any element of \mathbb{R} . The binary operations are defined by $a \oplus b := \min\{a, b\}$ and $a \odot b := a + b$, where $a \oplus \infty = \infty \oplus a = a$ and $a \odot \infty = \infty \odot a = \infty$ for any $a \in \overline{\mathbb{R}}$. This algebra is called a *semiring* since it has all the structure of a ring except there are no additive inverses. In fact, apart from additive inverses, it has all the structure of a field: 0 is the multiplicative identity, and any a in \mathbb{R} has $-a$ as a multiplicative inverse.

A polynomial in n variables over this semiring can be interpreted as a function from \mathbb{R}^n to \mathbb{R} given by the minimum of linear forms, one for each monomial of the polynomial. For instance, the polynomial

$$1 \odot x^{\odot 2} \oplus x \odot y \oplus 1 \odot y^{\odot 2} \oplus x \oplus y \oplus 1$$

can be thought of as a map from \mathbb{R}^2 to \mathbb{R} sending (x, y) to $\min\{1 + 2x, x + y, 1 + 2y, x, y, 1\}$.

Let f be a tropical polynomial in n variables. The *tropical vanishing locus* $V(f)$ is the set of all points in \mathbb{R}^n where the minimum in f is attained at least twice. This defines a *tropical hypersurface*. The combinatorics of the tropical hypersurface can be found from a subdivision of the Newton polytope of f . If f is the tropical quadric above, then $V(f)$ is the tropical plane curve pictured in Figure 1.1. The curve divides \mathbb{R}^2 into six unbounded regions, one for each of the six monomials in the polynomial achieving the minimum. The tropical curve consists of four vertices (at $(0, 0)$, $(-1, 0)$, $(0, -1)$, and $(1, 1)$), three bounded edges, and six rays.




Figure 1.1: A tropical quadric curve in \mathbb{R}^2

More generally, a *tropical prevariety* is the intersection of a finite number of tropical hypersurfaces. We will momentarily see that some tropical prevarieties arise from classical varieties, in which case we call them *tropical varieties*. A tropical variety has the structure of a weighted, balanced polyhedral complex. In Figure 1.1, this “balanced” property means that the sum at each vertex of the outgoing slopes, when written as primitive integer vectors, is equal to 0.

An alternate description of tropical geometry connects it with varieties over non-Archimedean fields. Let K be an algebraically closed field with non-trivial non-Archimedean valuation $\text{val} : K^* \rightarrow \mathbb{R}$. A common example throughout this work will be $K = \mathbb{C}\{\{t\}\}$, the field of Puiseux series over the complex numbers with indeterminate t . This is the algebraic closure of the field of Laurent series over \mathbb{C} , and can be defined as

$$\mathbb{C}\{\{t\}\} = \left\{ \sum_{i=k}^{\infty} a_i t^{i/n} : a_i \in \mathbb{C}, n, k \in \mathbb{Z}, n > 0 \right\},$$

with $\text{val}(\sum_{i=k}^{\infty} a_i t^{i/n}) = k/n$ if $a_k \neq 0$. In particular, $\text{val}(t) = 1/n$. Other important examples of non-Archimedean fields are the p -adic numbers \mathbb{Q}_p , where p is a prime number, and the complex p -adic numbers \mathbb{C}_p , a complete and algebraically closed field containing \mathbb{Q}_p . These are featured in Chapter 4.

The *tropicalization map* $\text{trop} : (K^*)^n \rightarrow \mathbb{R}^n$ sends points in the n -dimensional torus over K into Euclidean space under coordinate-wise valuation:

$$\text{trop} : (a_1, \dots, a_n) \mapsto (\text{val}(a_1), \dots, \text{val}(a_n)).$$

For a variety $X \subset (K^*)^n$, the *tropicalization of X* , denoted $\text{Trop}(X)$, is the Euclidean closure of the image of X under this map. Such a subset of Euclidean space is called a *tropical variety*. The tropicalization of a variety very much depends on its embedding.

The following theorem connects the two definitions of a tropical variety. For $f = \sum_v a_i x^v \in K[x_1, \dots, x_n]$, let the *tropicalization of f* be the tropical polynomial $\text{trop}(f) := \bigoplus_v \text{val}(a_i) \odot x^{\odot v}$.

Theorem 1.1.1 (The Fundamental Theorem of Tropical Geometry). *Let $I \subset k[x_1, \dots, x_n]$ be an ideal, and let $X = V(I) \cap (K^*)^n$. Then*

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

This result was proven in an unpublished manuscript by Mikhail Kapranov in the case of hypersurfaces. See [MS15, Theorem 3.2.5] for a proof of the general result.

If the intersection in Theorem 1.1.1 can be taken over a finite subset of I , we call that finite collection of polynomials a *tropical basis* for $\text{Trop}(X)$. Every tropical variety has a tropical basis, although it may contain more elements than a basis for the classical ideal I (see Section 5.2 for an example). It follows that a tropical variety is a tropical prevariety as defined above.

A one-dimensional tropical variety is called a *tropical curve*. Throughout this work, we will consider these objects at varying levels of abstraction. When we refer to an *embedded tropical curve* (Section 1.2), we are considering a one-dimensional balanced polyhedral complex sitting in some \mathbb{R}^n . When we refer to an *abstract tropical curve* (Section 1.3), we are considering the underlying metric graph.

A common theme throughout this work is a phenomenon that the *tropicalization* of a classical object might differ from the most natural *tropical analog* of that classical object. In Chapter 2, we see that curves with finitely many intersection points can correspond to tropical curves with infinitely many intersection points, even though the classical intersection tropicalizes to a finite set. In Chapter 3, we see that the moduli space of smooth plane quartic curves tropicalizes to a bigger set than the moduli space of tropical smooth plane quartic curves. And in Chapter 5, we study the difference between tropical commuting matrices and tropicalizations of commuting matrices. An important task in tropical geometry is to identify these discrepancies, and when possible to understand the underlying structure of them.

1.2 Embedded tropical curves

We will focus in this section on the special case of tropical curves that are embedded in the plane. Such a tropical curve C is a tropical hypersurface in \mathbb{R}^2 , defined by a single tropical polynomial

$$f(x, y) = \bigoplus_{(i,j) \in S} h_{ij} \odot x^i \odot y^j,$$

where $S \subset \mathbb{Z}^2$ is a finite index set and no h_{ij} is ∞ . Let P denote the *Newton polygon* of f , which is the convex hull of S , and write $A = P \cap \mathbb{Z}^2$ for the set of lattice points of the lattice polygon P . Note that $S \subset A$. Associated to the tropical polynomial f is a height function $h : A \rightarrow \mathbb{R}$, defined by $h(i, j) = h_{i,j}$ (where $h(i, j) = \infty$ if there is no $(i, j)^{th}$ term in f). As described in the previous section, the tropical curve C defined by this min-plus polynomial consists of all points $(x, y) \in \mathbb{R}^2$ for which the minimum among the quantities $i \cdot x + j \cdot y + h(i, j)$ is attained at least twice as (i, j) runs over A .

The curve C is dual to the *regular subdivision* Δ of A induced by h , which we now describe. Lift each lattice point $a \in A$ to the height $h(a)$, then take the lower convex hull of the lifted points in \mathbb{R}^3 . We project back to \mathbb{R}^2 by omitting the height, thereby obtaining Δ . This is illustrated in Figure 1.2. The duality between the subdivision and the tropical curve is more easily seen by turning the tropical curve upside-down.




Figure 1.2: Inducing a subdivision on a Newton polygon, dual to a tropical curve

Each line segment making up C has rational slope and therefore a natural *lattice length* with respect to the lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$. If a line segment intersects \mathbb{Z}^2 exactly at its two endpoints, it has lattice length 1; all other lengths are determined by scaling and translating such segments. Equivalently, a line segment with slope p/q , where p and q are relatively prime integers, has length equal to the Euclidean length of the edge divided by $\sqrt{p^2 + q^2}$. For an alternate definition of the length of edges in a tropical curve, see Equation (3.3.1) in Section 3.3. Whenever we speak of *length* or *metrics* with respect to embedded tropical curves, we are referring to lattice length.

By no means does the subdivision Δ uniquely determine the tropical curve: different height functions can induce the same triangulation, and will give different lengths to the edges of the tropical curve even though the combinatorics will be the same. For instance, the tropical quadric curves in Figures 1.2 and 1.1 are dual to the same Δ , but are drawn with different lengths on the bounded edges. The set of all height functions h which induce the same subdivision Δ is a relatively open polyhedral cone in \mathbb{R}^A . Its closure is called the *secondary cone* and is denoted $\Sigma(\Delta)$. The collection of all secondary cones $\Sigma(\Delta)$ is a complete polyhedral fan in \mathbb{R}^A , which is called the *secondary fan* of A .




Figure 1.3: A non-regular triangulation of a polygon

A subdivision Δ is a *triangulation* if all maximal cells are triangles. The maximal cones in the secondary fan $\Sigma(\Delta)$ correspond to the regular triangulations Δ of A . We say the regular triangulation Δ of P is *unimodular* if each triangle in Δ has area $1/2$, or equivalently that the subdivision contains the maximal number of cells possible. We say the tropical curve C is *smooth* if the corresponding subdivision of the Newton polygon is a unimodular triangulation Δ . It is worth noting that there are some triangulations of polygons that are *not* regular, meaning that they are not induced by any height function h , and have no corresponding tropical curve. See [DLRS10] for more background, and Figure 1.3 for a non-regular triangulation.

The *genus* of any tropical curve is its genus as a topological space, i.e. its first Betti number. In the case of a smooth tropical plane curve C with Newton polygon P , the genus $g = g(C)$ is the number of interior lattice points of P . The curve C contains a subdivision of a metric graph of genus g with vertices of valency ≥ 3 as in [BPR12], and this subdivision is unique for $g \geq 2$. The underlying graph G is planar and has g distinguished cycles, one for each interior lattice point of P . We call G the *skeleton* of C . It is the smallest subspace of C to which C admits a deformation retract. Although the metric on G depends on C , the graph is determined by Δ . Abstract metric graphs will be discussed in more detail in Section 1.3.

For an illustration, see Figure 1.4, where the tropical curve is drawn upside-down to highlight the duality. The triangulation Δ on the left defines a family of smooth tropical plane curves of degree four. Such a curve has genus $g = 3$. Its skeleton G is shown on the right.

1.3 Abstract tropical curves

Following the conventions of [GK08], a graph for us means a finite and connected multigraph where loops are allowed. Given a graph Γ , the sets of edges and vertices are denoted $E(\Gamma)$ and $V(\Gamma)$, respectively, and $\text{val}(P)$ denotes the valency of a vertex $P \in V(\Gamma)$. When we refer to the genus $g(\Gamma)$ of the graph, we mean its genus as a topological space, i.e. the dimension




Figure 1.4: Unimodular triangulation, tropical quartic, and skeleton

of $H_1(\Gamma, \mathbb{R})$. We say Γ is n -connected if there do not exist $n - 1$ edges that if deleted would disconnect the graph. If removing a single edge disconnects a graph, that edge is called a *bridge*, and a graph with no bridges is called *bridgeless*. By definition, a graph is 2-connected if and only if it is bridgeless.

A *metric graph* is a pair (Γ, ℓ) , where Γ is a graph and ℓ is a length function $\ell : E(\Gamma) \rightarrow \mathbb{R}^+$. A *abstract tropical curve* is a metric graph with unbounded ends allowed; in this case, ℓ maps to $\mathbb{R}^+ \cup \{\infty\}$. For instance, any tropical curve as defined in Section 1.2 is an abstract tropical curve, with the metric ℓ coming from the lattice length of the edges. See Figure 1.5 for some examples of abstract tropical curves.




Figure 1.5: Three abstract tropical curves, of genus 1, 3, and 0

Abstract tropical curves behave analogously to classical algebraic curves, which are Riemann surfaces, in many ways. Of particular interest to us is a *divisor theory* that works on these metric graphs, which we will now describe; see [GK08, MZ08] for more details. A *divisor* D on a metric graph Γ is an element of the free abelian group on Γ , i.e., a finite formal sum of points

$$D = a_1 \cdot p_1 + \dots + a_n \cdot p_n,$$

where all p_i 's are distinct. We say the *degree of D at p_i* is a_i , and the *degree of D* is $a_1 + \dots + a_n$. The divisor D is said to be *effective* if all the a_i are non-negative.

A *rational function* on Γ is a continuous piecewise linear function (with finitely many “pieces”) with integer slopes. To a rational function f we associate a divisor $\text{div}(f)$ whose degree at each point p is the sum of the outgoing slopes of f at p . A divisor of this form is called *principal*. Note that the degree of a principal divisor is always zero. Sometimes we will write f^{trop} instead of f to remind the reader that the function is on an abstract tropical curve.




Figure 1.6: The graph of a rational function f^{trop} on an abstract tropical curve Γ .

As an example, consider Figure 1.6. Here Γ consists of four vertices and three edges arranged in a Y-shape, and the image of Γ under a rational function f is illustrated lying above it. The leftmost vertex is a zero of order 2, since there is an outgoing slope of -2 and no other outgoing slopes. The next kink in the graph is a pole of order 1, since the outgoing slopes are 2 and -1 and $2 + (-1) = 1$. Moving along in this direction we have a pole of order 4, a zero of order 4, at one endpoint a pole of order 1, and at the other endpoint no zeros or poles. Note that, counting multiplicity, there are six zeros and six poles. The numbers agree, as in the classical case.

Two divisors D and D' are said to be *linearly equivalent* if $D - D' = \text{div}(f)$ for some rational function f . A divisor D has *rank r* if for every effective divisor E of degree r , $D - E$ is linearly equivalent to an effective divisor (and r is the largest integer with this property).

A *canonical divisor* on Γ is any divisor linearly equivalent to

$$K_\Gamma := \sum_{v \in \Gamma} (\text{valence}(v) - 2) \cdot v.$$

Its degree is $2g - 2$, where g is the genus of Γ . A helpful result in Section 2.3 will be the *tropical Riemann-Roch theorem* [BN07], which asserts that if Γ is a metric graph and D is any divisor on Γ , then

$$\text{rank}(D) + \text{rank}(K_\Gamma - D) = \deg(D) + 1 - g.$$

Divisor theory on an abstract tropical curve Γ is sometimes phrased in the language of *chip-firing* [HMY12]. An effective divisor $D = \sum a_i \cdot p_i$ is referred to as a *chip configuration*, and we say that there are a_i chips at the point $p_i \in \Gamma$. Let ℓ be a positive real number and $\Gamma' \subset \Gamma$ a (not necessarily connected) subgraph. The corresponding *chip firing move* $CF(\Gamma', \ell)$ is the tropical rational function $-\min(\ell, d(p, \Gamma'))$, where d measures the distance from p to Γ' . This function is equal to 0 on Γ' , has slope -1 away from Γ' within ℓ of Γ' , and is equal to $-\ell$ on the rest of the graph. The boundary points of Γ' are the zeros of the tropical rational function, and the other bend points are the poles. This means that $CF(\Gamma', \ell)$ witnesses the linear equivalence of these two sets of points. We can think of $CF(\Gamma', \ell)$ as “firing chips” from the boundary points of Γ' , which end up at the other set of points. An example is illustrated in Figure 1.7. Here, Γ is a line segment of length 6, Γ' is the union of the two endpoints, ℓ is 1, and the graph of $CF(\Gamma', \ell)$ is illustrated as a dashed line. The two endpoints “fire” to the points labelled with crosses, and the corresponding divisors are linearly equivalent.




Figure 1.7: A graph with endpoints “firing” to a linearly equivalent pair of points

Many definitions made for abstract tropical curves are analogous to classical definitions. For instance, a metric graph Γ is *hyperelliptic* if it has a divisor of degree 2 and rank 1 [BN09]. This is equivalent to a graph having a degree 2 harmonic morphism to a tree [Cha13], just as classical hyperelliptic curves have a degree 2 morphism to the projective line. An embedded version of hyperelliptic graphs will be discussed in Sections 3.5 and 3.6.

We will now assume that our graphs have genus $g \geq 2$, and will only consider trivalent, leafless, connected graphs. Such a graph has $2g - 2$ vertices and $3g - 3$ edges, and the number of such trivalent graphs (up to homeomorphism) for $g = 2, 3, \dots, 10$ is 2, 5, 17, 71, 388, 2592, 21096, 204638, 2317172; see [Bal76] and [Cha12, Prop. 2.1]. Any trivalent abstract tropical curve will admit a deformation retract onto such a graph, and there is a unique such minimal graph; this is the *skeleton* described in Section 1.2.

We will let \mathbb{M}_g denote the *moduli space of metric graphs* of genus g . This moduli space is obtained by gluing together finitely many orthants $\mathbb{R}_{\geq 0}^m$, where $m \leq 3g - 3$, one for each combinatorial type of graph, modulo the identifications corresponding to graph automorphisms. These automorphisms endow the moduli space \mathbb{M}_g with the structure of a *stacky fan*. We refer to [BMV11, Cha12] for the complete definition of \mathbb{M}_g , combinatorial details, and applications in algebraic geometry. The maximal cones of \mathbb{M}_g correspond to trivalent graphs of genus g . Since each of these graphs has $3g - 3$ edges, \mathbb{M}_g is pure of dimension $3g - 3$. It is the tropical analogue of the classical *moduli space of curves* of genus g , a connection that will be fleshed out in the following section.

Inside \mathbb{M}_g is a stacky subfan $\mathbb{M}_{g, \text{hyp}}$, which consists of all graphs in \mathbb{M}_g that are hyperelliptic. Unlike \mathbb{M}_g , it is *not* pure dimensional; for instance, $\mathbb{M}_{3, \text{hyp}}$ has maximal cells of

dimension 6 and dimension 5. However, the subfan consisting of 2-connected graphs inside $\mathbb{M}_{g,\text{hyp}}$ is pure of dimension $2g - 1$, and each 2-connected hyperelliptic graph can be constructed as follows [Cha13]. Let T be a metric tree with at most trivalent nodes. Let T' be a copy of T , and connect corresponding vertices $v \in T$ and $v' \in T'$ by $3 - \text{valence}(v)$ edges. The resulting graph G will be 2-connected and hyperelliptic. A graph that arises from this construction is called a *ladder*, and all 2-connected hyperelliptic graphs are ladders. This fact will be crucial in Section 3.5.

1.4 Curves over non-Archimedean fields

Let K be an algebraically closed field that is complete with respect to a surjective non-Archimedean valuation $\text{val} : K^* \rightarrow \mathbb{R}$. Let C be a smooth curve of genus g over K . Associated to the algebraic curve C is the *Berkovich analytification* C^{an} [Ber90]. As a set, C^{an} is the collection of multiplicative seminorms $|\cdot|$ on the coordinate ring of C extending the usual norm on K , and it is endowed with the weakest topology such that the map $f \mapsto |f|$ is continuous. There is a natural metric on $C^{\text{an}} \setminus C$, with the points of C thought of as being infinitely far away. Moreover, the space C^{an} is path-connected (in contrast to C , which is totally disconnected in the non-Archimedean topology). The space C^{an} contains a finite metric graph Σ called the *minimal Berkovich skeleton* of C , onto which there is a deformation retract of C^{an} . The genus of Σ is at most g . If the genus of Σ is equal to g , we call C a *Mumford curve* [Mum72b]. Such curves are the subject of Chapter 4, which will also describe the Berkovich analytic line $(\mathbb{P}^1)^{\text{an}}$ in detail.

More generally, one can define the analytification of any variety over K . Payne showed that the Berkovich analytification of an affine variety X is the inverse limit of all tropicalizations of X [Pay09]. More precisely, given an embedding $\iota : X \rightarrow \mathbb{A}^m$, write $\text{Trop}(X, \iota)$ for the tropicalization of X under this embedding. Given another embedding $\iota' : X \rightarrow \mathbb{A}^n$ and an equivariant morphism $\varphi : \mathbb{A}^m \rightarrow \mathbb{A}^n$ such that $\iota' = \varphi \circ \iota$, we have that $\text{Trop}(\varphi)$ maps $\text{Trop}(X, \iota)$ into $\text{Trop}(X, \iota')$. Then

$$X^{\text{an}} \cong \varprojlim \text{Trop}(X, \iota),$$

where the inverse limit is in the category of topological spaces and is taken over all affine embeddings ι of X and all such maps $\text{Trop}(\varphi)$. The map from X^{an} to a particular $\text{Trop}(X, \iota)$ is given by mapping a multiplicative seminorm $|\cdot|$ to $(-\log |f_1|, \dots, -\log |f_m|)$, where f_1, \dots, f_m are the rational functions giving the embedding $\iota : X \rightarrow \mathbb{A}^m$.

This means we can view tropical geometry as studying snapshots of Berkovich analytic spaces. In general, the more information a tropical curve preserves about the Berkovich analytic curve and the Berkovich skeleton Σ , the more helpful that tropicalization is. In particular, we say a tropicalization is *faithful* if Σ is mapped homeomorphically and isometrically onto its image. In the case of faithful tropicalizations, all uses of the term *skeleton* in this chapter agree. Baker, Payne, and Rabinoff gave criteria in [BPR12] for when a tropicalization is faithful; for instance, for tropical plane curves it suffices that the curve is

smooth. They also proved that every curve C has a faithful tropicalization in *some* dimension, which was later improved to 3 dimensions in [BR13] (at the cost of possibly allowing non-embeddings). The question of which curves admit faithful tropicalizations in the plane is a major focus of Chapter 3.

Let \mathcal{M}_g denote the moduli space of smooth curves of genus g , where $g \geq 2$. By our hypotheses on the field K , every metric graph G of genus g arises as the Berkovich skeleton from some curve C over K . This defines a surjective tropicalization map from (the K -valued points in) the moduli space of smooth curves of genus g to the moduli space of metric graphs of genus g :

$$\text{trop} : \mathcal{M}_g \rightarrow \mathbb{M}_g. \quad (1.4.1)$$

Both spaces have dimension $3g - 3$ for $g \geq 2$. It was shown by Abramovich, Caporaso, and Payne [ACP14] that this tropicalization agrees with “naive set-theoretic tropicalization”, and is faithful.

Consider a plane curve defined by a Laurent polynomial $f = \sum_{(i,j) \in \mathbb{Z}^2} c_{ij} x^i y^j \in K[x^\pm, y^\pm]$ with Newton polygon P . For τ a face of P we let $f|_\tau = \sum_{(i,j) \in \tau} c_{ij} x^i y^j$, and say that the curve cut out by f is *non-degenerate* if for all faces τ of P , the curve defined by $f|_\tau$ has no singularities in $(K^*)^2$. Non-degenerate curves are useful for studying many subjects in algebraic geometry, including singularity theory [Kou76], the theory of sparse resultants [GKZ08], and real algebraic curves in maximal condition [Mik00].

Let P be any lattice polygon in \mathbb{R}^2 with g interior lattice points. We write \mathcal{M}_P for the Zariski closure (inside \mathcal{M}_g) of the set of curves that appear as non-degenerate plane curves over K with Newton polygon P . This space was introduced by Koelman [Koe91]. Consider the union over all relevant polygons:

$$\mathcal{M}_g^{\text{nd}} := \bigcup_P \mathcal{M}_P. \quad (1.4.2)$$

This moduli space was introduced and studied by Castryck and Voight in [CV09]. In Chapter 3 we introduce and study the tropical analogues of \mathcal{M}_P and $\mathcal{M}_g^{\text{nd}}$.

Chapter 2

Non-proper tropical intersections

In this chapter we consider tropical varieties that intersect in components that are higher dimensional than expected. In Section 2.1 we consider pairs of plane curves that intersect in one-dimensional components, and prove constraints on where classical points can map to within these intersections. Section 2.2 presents explicit examples of these non-proper intersections. Section 2.3 studies the non-proper intersections of smooth plane tropical quartic curves with their bitangents, and provides a natural segue for the following chapter on moduli of smooth plane tropical curves.

Sections 2.1 and 2.2 are based on the single-author paper “Tropical images of intersection points” [Mor15], appearing in *Collectanea Mathematica*. Section 2.3 is based on the paper “Bitangents of tropical plane quartic curves” [BLM⁺14], coauthored with Matthew Baker, Yoav Len, Nathan Pflueger, and Qingchun Ren.

2.1 Intersections of plane curves

Let K be an algebraically closed non-Archimedean field with a nontrivial valuation $\text{val} : K^* \rightarrow \mathbb{R}$, such as the Puiseux series $\mathbb{C}\{\{t\}\}$. Consider two plane curves $X, Y \subset (K^*)^2$ intersecting in a finite number of points. We are interested in the image of the intersection points under tropicalization; that is, in $\text{Trop}(X \cap Y)$ inside of $\text{Trop}(X) \cap \text{Trop}(Y) \subset \mathbb{R}^2$. It was shown in [OP13, Theorem 1.1] that if $\text{Trop}(X) \cap \text{Trop}(Y)$ is zero-dimensional in a neighborhood of a point in the intersection, then that point is in $\text{Trop}(X \cap Y)$. More generally, they showed this for varieties X and Y under the assumption that $\text{Trop}(X) \cap \text{Trop}(Y)$ has codimension $\text{codim } X + \text{codim } Y$ in a neighborhood of the point. It follows that if $\text{Trop}(X) \cap \text{Trop}(Y)$ is a finite set, then $\text{Trop}(X \cap Y) = \text{Trop}(X) \cap \text{Trop}(Y)$.

It is possible for $\text{Trop}(X) \cap \text{Trop}(Y)$ to have higher dimensional components, namely finite unions of line segments and rays. It was shown in [OR11] that if $\text{Trop}(X) \cap \text{Trop}(Y)$ is bounded, then each connected component of $\text{Trop}(X) \cap \text{Trop}(Y)$ has the “right” number of images of points in $X \cap Y$, counted with multiplicity. They further showed that the theorem holds for components of $\text{Trop}(X) \cap \text{Trop}(Y)$ that are unbounded, after a suitable

compactification. In this context, the “right” number is the number of points in the *stable tropical intersection* of that connected component, which is defined as $\lim_{\varepsilon \rightarrow 0} (\text{Trop}(X) + \varepsilon \cdot v) \cap \text{Trop}(Y)$, where v is a generic vector and ε is a real number [OR11, 4]. For instance, in Figure 2.1, two tropical plane curves intersect in a line segment, but have stable tropical intersection equal to the union of the two endpoints of the line segment; this is the limit of the transverse intersection of the curves under generic translation. The *stable intersection divisor* of such a pair of intersecting curves is the formal sum of points in the stable tropical intersection, with coefficients recording intersection multiplicity.

We offer the following example to illustrate this higher-dimensional component phenomenon. This will motivate the main question of this section: as we vary X and Y over curves with the same tropicalizations, how does the set $\text{Trop}(X \cap Y)$ vary inside of the fixed set $\text{Trop}(X) \cap \text{Trop}(Y)$?

Example 2.1.1. Let $K = \mathbb{C}\{\{t\}\}$ and let $f, g \in K[x, y]$ be $f(x, y) = c_1 + c_2x + c_3y$ and $g(x, y) = c_4x + c_5xy + tc_6y$, where $c_i \in K$ and $\text{val}(c_i) = 0$ for all i . Let $X, Y \subset (K^*)^2$ be the curves defined by f and g , respectively.




Figure 2.1: $\text{Trop}(X)$ and $\text{Trop}(Y)$, before and after a small shift of $\text{Trop}(Y)$

Regardless of our choice of c_i , $\text{Trop}(X)$ and $\text{Trop}(Y)$ will be as pictured in Figure 2.1, with $\text{Trop}(X)$ and $\text{Trop}(Y)$ intersecting in the line segment L from $(0, 0)$ to $(1, 0)$. However, X and Y only intersect in two points (or one point with multiplicity 2). The natural question is: as we vary the coefficients while keeping valuations (and thus tropicalizations) fixed, what are the possible images of the two intersection points within L ?

A reasonable guess is that the intersection points map to the stable tropical intersection $\{(0, 0), (1, 0)\}$, and indeed this does happen for a generic choice of coefficients c_1, \dots, c_6 . However, as shown in Example 2.2.1, one can choose coefficients such that the intersection points map to any pair of points in L of the form $(r, 0)$ and $(1 - r, 0)$, where $0 \leq r \leq \frac{1}{2}$. These possible configurations are illustrated in Figure 2.2.

The main result of this section is that the points $\text{Trop}(X \cap Y)$ inside of $\text{Trop}(X) \cap \text{Trop}(Y)$ must be linearly equivalent to the stable tropical intersection via particular tropical rational functions. To distinguish tropical rational functions from classical rational functions in this section, they will be written as f^{trop} , g^{trop} , or h^{trop} instead of f , g , or h . At times we will refer to the support of a tropical rational function f^{trop} , which we will denote $\text{supp}[f^{\text{trop}}]$. In




Figure 2.2: Possible images of $X \cap Y$ in $\text{Trop}(X) \cap \text{Trop}(Y)$.

all the examples discussed in Section 2.2, essentially every such configuration is achievable. Conjecture 2.1.8 expresses our hope that this always holds.

Theorem 2.1.2. *Let $X, Y \subset (K^*)^2$ where $X \cap Y$ is equal to the multiset $\{p_1, \dots, p_n\}$. Let E be the stable intersection divisor of $\text{Trop}(X)$ and $\text{Trop}(Y)$, and let D be*

$$D = \sum_i \text{trop}(p_i).$$

Then there exists a tropical rational function h^{trop} on $\text{Trop}(X)$ such that $(h^{\text{trop}}) = D - E$ and $\text{supp}[h^{\text{trop}}] \subset \text{Trop}(X) \cap \text{Trop}(Y)$.

This result is an instance of understanding the difference between a tropical object and a tropicalization. The tropicalization of the intersection is a zero-dimensional, and the tropical intersection is one-dimensional; however, the finite set sits inside the larger set in a very particular way.

We will present two proofs of Theorem 2.1.2, the first using some additional assumptions on $\text{Trop}(X)$. Although the result in that case is weaker, it illustrates connections to Berkovich theory, which seems to be a better context for proving a reverse direction. We also present an alternate argument using tropical modifications, which proves the full scope of the theorem.

Example 2.1.3. Let X and Y be as in Example 2.1.1. We will consider tropical rational functions on $\text{Trop}(X) \cap \text{Trop}(Y)$ such that

- (i) the stable intersection points are the poles (possibly canceling with zeros), and
- (ii) the tropical rational function takes on the same value at every boundary point of $\text{Trop}(X) \cap \text{Trop}(Y)$.

If we insist that the “same value” in condition (ii) is 0, we may extend these tropical rational functions to all of $\text{Trop}(X)$ by setting them equal to 0 on $\text{Trop}(X) \setminus \text{Trop}(Y)$. This yields tropical rational functions on $\text{Trop}(X)$ with $\text{supp}[h^{\text{trop}}] \subset \text{Trop}(X) \cap \text{Trop}(Y)$, as in Theorem 2.1.2. Instances of the types of such tropical rational functions on $L = \text{Trop}(X) \cap \text{Trop}(Y)$ from our example are illustrated in Figure 2.3.

As asserted by Theorem 2.1.2, all possible image intersection sets in $\text{Trop}(X) \cap \text{Trop}(Y)$ arise as the zero set of such a tropical rational function. Equivalently, the stable intersection divisor and the image of intersection divisor are linearly equivalent via one of these functions.




Figure 2.3: Graphs of tropical rational functions on $\text{Trop}(X) \cap \text{Trop}(Y)$

Remark 2.1.4. It is not quite the case that the zero set of every such tropical rational function (from Example 2.1.3) is attainable as the image of the intersections of X and Y (with changed coefficients). For instance, such a tropical rational function could have zeros at $(\frac{\sqrt{2}}{2}, 0)$ and $(1 - \frac{\sqrt{2}}{2}, 0)$, which cannot be the coordinate-wise valuations of *any* points on X and Y since they have irrational coordinates. However, if we insist that the tropical rational functions have zeros at points with rational coefficients (since $\mathbb{Q} = \text{val}(K^*)$), all zero sets can be achieved as the images of intersections. This is the content of Conjecture 2.1.8. Alternatively, we could base-change to a larger field whose value group is \mathbb{R} , and then all zero sets are indeed achieved as images of intersections in our examples.

Since we can tropicalize a curve to obtain a tropical curve, we would like to tropicalize a rational function on a curve and obtain a tropical rational function on a tropical curve. Let h be a rational function on a curve X . Naïvely, we might try to define the “tropicalization of h ”, denoted $\text{trop}(h)$, as follows. For every point w in the image of $X \setminus \{\text{zeros and poles of } h\}$ under tropicalization, lift that point to $p \in X$, and define

$$\text{trop}(h)(w) = \text{val}(h(p)).$$

Extend this function to all of $\text{Trop}(X)$ by continuity.

Unfortunately this is not quite well-defined, because $\text{val}(h(p))$ depends on which lift p of w we choose. However, as suggested by Matt Baker, this definition can be made rigorous if at least one of the two tropicalizations is suitably faithful in a Berkovich sense. Let h be a rational function on X , and assume that there is a canonical section s to the map $X^{\text{an}} \rightarrow \text{Trop}(X)$, where X^{an} is the analytification of X . For $w \in \text{Trop}(X)$, define

$$\text{trop}(h)(w) = \log |h|_{s(w)},$$

where $|\cdot|_{s(w)}$ is the seminorm corresponding to the point $s(w)$ in X^{an} . This rational function has the desired properties.

Remark 2.1.5. In [BPR12, §6] one can find conditions to guarantee that there exists a canonical section s to the map $X^{\text{an}} \rightarrow \text{Trop}(X)$. For instance, if $\text{Trop}(X)$ is smooth in the sense that it comes from a unimodular triangulation of its Newton polygon, such a section will exist.

We are ready to prove Theorem 2.1.2 with the additional assumption that $\text{Trop}(X)$ has a canonical section s to X^{an} .

Proof of Theorem 2.1.2 when there is a canonical section $s : \text{Trop}(X) \rightarrow X^{\text{an}}$. Let f and g be the defining equations of X and Y , respectively. Let $g' \in K[x, y]$ have the same tropical polynomial as g , and let Y' be the curve defined by g' . We have that $\text{Trop}(Y) = \text{Trop}(Y')$, and for generic g' we have that $\text{Trop}(X \cap Y')$ is the stable tropical intersection of $\text{Trop}(X)$ and $\text{Trop}(Y)$. (Here “generic” means chosen with coefficients so that no leading terms of coefficients cancel in the resultant of f and g' .)

Recall that p_1, \dots, p_n denote the intersection points of X and Y , possibly with repeats. Let p'_1, \dots, p'_m denote the intersection points of X and Y' , with duplicates in the case of multiplicity.

Consider the rational function $h = \frac{g}{g'}$ on X , which has zeros at the intersection points of X and Y and poles at the intersection points of X and Y' . Since the section $s : \text{Trop}(X) \rightarrow X^{\text{an}}$ exists, we may tropicalize h . This gives a tropical rational function $\text{trop}(h)$ on $\text{Trop}(X)$ with divisor

$$(\text{trop}(h)) = \text{trop}(p_1) + \dots + \text{trop}(p_n) - \text{trop}(p'_1) - \dots - \text{trop}(p'_m) = D - E.$$

We claim that $\text{trop}(h)$ is the desired h^{trop} from the statement of the theorem. All that remains to show is that $\text{supp}(\text{trop}(h)) \subset \text{Trop}(X) \cap \text{Trop}(Y)$. If $w \in \text{Trop}(X) \setminus \text{Trop}(Y)$, then $|g|_{s(w)} = |g'|_{s(w)}$ because g and g' both have bend locus $\text{Trop}(Y)$, and w is away from $\text{Trop}(Y)$. This means that $\text{trop}(h)(w) = |h|_{s(w)} = |g|_{s(w)} - |g'|_{s(w)} = 0$ on $\text{Trop}(X) \setminus \text{Trop}(Y)$. This completes the proof. \square

Remark 2.1.6. Since we have our result in terms of linear equivalence, we get as a corollary that the configurations of points differ by a sequence of chip-firing moves on the graph $\text{Trop}(X) \cap \text{Trop}(Y)$ by [HMY12]. Any chip-firing must take place within this graph, so boundary points cannot “fire” into $\text{Trop}(X) \setminus \text{Trop}(Y)$ or $\text{Trop}(Y) \setminus \text{Trop}(X)$.

Remark 2.1.7. If $\text{Trop}(X) \cap \text{Trop}(Y)$ is unbounded (for instance, if $\text{Trop}(X) = \text{Trop}(Y)$), then it is possible to have zeros of the rational function “at infinity.” This is OK, and can be made sense of using a compactifying fan as in [OR11, S3]. See Example 2.2.3 for an instance of this phenomenon.

An alternate approach to understanding the images of intersection points is by using *tropical modifications* [Mik06], [BLdM12, §4], which can be thought of as a tropical version of the classical algebro-geometric process of “blowing up” varieties in codimension 1. For a general hypersurface, the process works as follows. Let $F(z)$ be a polynomial in n variables over the field K , and let $Y = (K^*)^n \setminus V(F)$. We can embed Y into $(K^*)^{n+1}$ by the map $\Phi : z \mapsto (z, F(z))$. The image $\Phi(Y)$ is cut out by $z_{n+1} - F(z_1, \dots, z_n) = 0$. The tropical variety $W = \text{Trop}(\Phi(Y))$ is then called the *tropical modification* of \mathbb{R}^n defined by $F(z)$. It is defined by the tropical polynomial $z_{n+1} \oplus \text{trop}(F)(z_1, \dots, z_n)$. Projecting onto the first n -coordinates, we obtain a surjective map $\pi : W \rightarrow \mathbb{R}^n$, which is one-to-one above $\mathbb{R}^n \setminus \text{Trop}(V(F))$ with $\pi^{-1}(p)$ a half ray in the direction $(0, \dots, 0, -1)$ for $p \in \text{Trop}(V(F))$.

We now give a proof for the full version of Theorem 2.1.2.

Proof of Theorem 2.1.2 using tropical modifications. Let X, Y, f, g, g', D , and E be as in the previous proof. Let g^{trop} and $(g')^{\text{trop}}$ be the tropical polynomials defined by g and g' , respectively.

Let $g(X) \subset (K^*)^2 \times K$ be the curve that is the closure of $\{(p, g(p)) \mid p \in X\}$. Its tropicalization $\text{Trop}(g(X))$ is contained in the tropical hypersurface in \mathbb{R}^3 determined by the polynomial $z = g^{\text{trop}}$, and projects onto $\text{Trop}(X)$. Call this projection π . Note that outside of $\text{Trop}(Y)$, π is one-to-one, and $\text{Trop}(g(X))$ agrees with $\text{Trop}(g'(X))$.

By [BLdM12, Lemma 4.4], the infinite vertical rays in $\pi^{-1}(\text{Trop}(X) \cap \text{Trop}(Y))$ correspond to the intersection points of X and Y , and so lie above the support of the divisor D on $\text{Trop}(X)$. Delete the vertical rays from $\pi^{-1}(\text{Trop}(X) \cap \text{Trop}(Y))$, and decompose the remaining line segments into one or more layers, where each layer gives the graph of a piecewise linear function on $\text{Trop}(X) \cap \text{Trop}(Y)$. (If deleting the vertical rays makes π a bijection, there will be only one layer.) Call these piecewise linear functions ℓ_1, \dots, ℓ_k . The tropical rational function

$$h^{\text{trop}} := \sum_{i=1}^k (\ell_i - (g')^{\text{trop}})$$

has value 0 outside of $\text{Trop}(X) \cap \text{Trop}(Y)$ because of the agreement of $\text{Trop}(g(X))$ and $\text{Trop}(g'(X))$, and has divisor $D - E$. This is our desired tropical rational function. \square

Theorem 2.1.2 places a constraint on the configurations of images of intersection points mapping into tropicalizations. The following conjecture posits that essentially all these configurations are attainable.

Conjecture 2.1.8. *Assume we are given $\text{Trop}(X)$ and $\text{Trop}(Y)$ and a tropical rational function h^{trop} on $\text{Trop}(X)$ with simple poles precisely at the stable tropical intersection points and zeros in some configuration (possibly canceling some of the poles) with coordinates in the value group of K , such that $\text{supp}[h^{\text{trop}}] \subset \text{Trop}(X) \cap \text{Trop}(Y)$. Then there exist classical curves X and Y with the given tropicalizations such that $\text{trop}(p_1), \dots, \text{trop}(p_n)$ are the zeros of h^{trop} .*

Proof Strategy. We consider the space of all configurations of zeros of rational functions on $\text{Trop}(X) \cap \text{Trop}(Y)$ satisfying the given properties. This forms a polyhedral complex.

- First, we will prove that we can achieve the configurations corresponding to the vertices of this complex.
- Next, let E be an edge connecting V and V' , where the configuration given by V is achieved by X and Y and the configuration given by V' is achieved by X' and Y' . We will prove that we can achieve any configuration along the edge by deforming (X, Y) to (X', Y') . This deformation will consist of varying higher-order terms in the coefficients of the defining polynomials of X and Y . This will show that all points on edges of the complex correspond to achievable configurations.

- We will continue this process (vertices giving edges, edges giving faces, etc.) to show that all points in the complex correspond to achievable configurations. □

For an illustration of this process, see Example 2.2.2 and Figure 2.6.

2.2 Examples of non-proper intersections

In these examples we consider curves X and Y over the field of Puiseux series $K = \mathbb{C}\{\{t\}\}$.

Example 2.2.1. Let f and g be as in Example 2.1.1. Treating them as elements of $(K[x])[y]$, their resultant is

$$-c_2c_5x^2 + (c_3c_4 - c_1c_5 - tc_2c_6)x - tc_1c_6.$$

The two roots of this quadratic polynomial in x , which are the x -coordinates of the two points in $X \cap Y$, have valuations equal to the slopes of the Newton polygon of the resultant. Assuming no cancellation in the coefficient of x , the valuations of the coefficients are 0, 0, and 1, giving slopes 0 and 1. For any rational number $r > 0$ we may choose $c_1 = 1 - t^r - t$ and all other $c_i = 1$, giving $\text{val}(c_3c_4 - c_1c_5 - tc_2c_6) = \text{val}(t^r) = r$. If $r \leq \frac{1}{2}$, then the two line segments of the Newton polygon have slopes of r and $1 - r$; and if $r \geq \frac{1}{2}$, then the slopes are both $\frac{1}{2}$. These cases are illustrated in Figure 2.2 and correspond to rational functions illustrated in Figure 2.3. This means all possible images of intersections allowed by Theorem 2.1.2 with rational coordinates are achievable, so Conjecture 2.1.8 holds for this example.

Example 2.2.2. Consider conic curves X and Y cut out the polynomials $f(x, y) = c_1x + c_2y + c_3xy$ and $g(x, y) = c_4x + c_5y + c_6xy + t(c_7x^2 + c_8y^2 + c_9)$ respectively, where $\text{val}(c_i) = 0$ for all i . The tropicalizations of X and Y are shown in Figure 2.4, and intersect in three line segments joined at a point.




Figure 2.4: $\text{Trop}(X)$ and $\text{Trop}(Y)$, before and after a small shift of $\text{Trop}(Y)$

The stable tropical intersection consists of four points: $(-1, 0)$, $(0, -1)$, $(1, 1)$, and $(0, 0)$. The possible images of $\text{Trop}(X \cap Y)$ must be linearly equivalent to these via a rational function equal to 0 on the three exterior points. (If viewed as a rational function on $\text{Trop}(X)$, this function it must be 0 on all of $\text{Trop}(X) \setminus \text{Trop}(Y)$ in addition to the three extreme points.) This gives us intersection configurations of three possible types:

- (i) $\{(-(p-r), 0), (0, -p), (p, p), (-r, 0)\}$ where $0 \leq r \leq p/2$;
- (ii) $\{(-p, 0), (0, -(p-r)), (p, p), (0, -r)\}$ where $0 \leq r \leq p/2$; and
- (iii) $\{(-p, 0), (0, -p), (p-r, p-r), (r, r)\}$ where $0 \leq r \leq p/2$.

To achieve a type (i) configuration, set $f(x, y) = x + y + xy$ and $g(x, y) = (1 + 2t^{1-p+r})x + (1 + t^{1-p})y + xy + t(x^2 + y^2 + 1)$; if $r > 0$, the 2 can be omitted from the coefficient of x in g . The Newton polygons of two polynomials, namely the resultants of f and g with respect to x and with respect to y , show that $\text{Trop}(X \cap Y) = \{(-(p-r), 0), (0, -p), (p, p), (-r, 0)\}$. Type (ii) and (iii) are achieved similarly, so Conjecture 2.1.8 holds for this example.

For instance, if $f(x, y) = x + y + xy$ and $g(x, y) = (1 + t^{1/2})x + (1 + t^{1/3})y + xy + t(x^2 + y^2 + 1)$, then $\text{Trop}(X \cap Y) = \{(2/3, 2/3), (0, -2/3), (-1/2, 0), (-1/6, 0)\}$. The formal sum of these points is linearly equivalent to the stable intersection divisor, as illustrated by the rational function in Figure 2.5. This is the tropicalization of the rational function $h(x, y) = \frac{(1+t^{1/2})x+(1+t^{1/3})y+xy+t(x^2+y^2+1)}{2x+4y+xy+t(x^2+y^2+1)}$, where $g'(x, y) := 2x + 4y + xy + t(x^2 + y^2 + 1)$ was chosen so that $\text{Trop}(X) \cap \text{Trop}(V(g'))$ is the stable tropical intersection of $\text{Trop}(X)$ and $\text{Trop}(Y)$.




Figure 2.5: The graph of $\text{trop}(h)$ on $\text{Trop}(X) \cap \text{Trop}(Y)$, with zeros at the dots and poles at the x's

We can also consider this example in view of the outlined method of proof for Conjecture 2.1.8. Considering each intersection configuration as a point in \mathbb{R}^8 (since we have four points in \mathbb{R}^2), we obtain a moduli space M for the possible tropical images of $X \cap Y$. The structure of this space is related to the notion of tropical convexity, as discussed in [Luo13]. As illustrated in Figure 2.6, M consists of three triangles glued along one edge. The hope is that if vertices like A and E can be achieved, then it is possible to slide along the edge and




Figure 2.6: A moduli space of intersection configurations, with seven examples

achieve points like D . For instance, if we set

$$\begin{aligned}
 f_A(x, y) &= f_E(x, y) = f_{AE,r} = x + y + xy \\
 g_A &= (1 + t^0)x + 4y + xy + t(x^2 + y^2 + 1) \\
 g_E &= (1 + t^{1/2})x + 4y + xy + t(x^2 + y^2 + 1) \\
 g_{AE,r} &= (1 + t^r)x + 4y + xy + t(x^2 + y^2 + 1),
 \end{aligned}$$

then f_A and g_A give configuration A , f_E and g_E give configuration E , and $f_{AE,r}$ and $g_{AE,r}$ give all configurations along the edge AE as r varies from 0 to $\frac{1}{2}$.

Example 2.2.3. Let X and Y be distinct lines defined by $f(x, y) = c_1 + c_2x + c_3y$ and $g(x, y) = c_6 + c_4x + c_5y$ with $\text{val}(c_i) = 0$ for all i . These lines tropicalize to the same tropical line centered at the origin, with stable tropical intersection equal to the single point $(0, 0)$. Any point on $\text{Trop}(X) = \text{Trop}(X) \cap \text{Trop}(Y)$ is linearly equivalent to $(0, 0)$ via a tropical rational function on $\text{Trop}(X)$, so Theorem 2.1.2 puts no restrictions on the image of $p = X \cap Y$ under tropicalization. As predicted by Conjecture 2.1.8, all possibilities can be achieved:

- (i) For $\text{trop}(p) = (r, 0)$, let $f(x, y) = 1 + x + y$, $g(x, y) = (1 + t^r) + x + y$.
- (ii) For $\text{trop}(p) = (0, r)$, let $f(x, y) = 1 + x + y$, $g(x, y) = 1 + (1 + t^r)x + y$.
- (iii) For $\text{trop}(p) = (-r, -r)$, let $f(x, y) = 1 + x + y$, $g(x, y) = 1 + x + (1 + t^r)y$.

The point $(0, 0)$ is also linearly equivalent to points at infinity, as witnessed by rational functions with constant slope 1 on an entire infinite ray. Mapping p “to infinity” means that X and Y cannot intersect in $(K^*)^2$, so we can choose equations for X and Y that give p a coordinate equal to 0, such as $x + y + 1 = 0$ and $x + 2y + 1 = 0$.

Example 2.2.4. Let X and Y be the curves defined by

$$f(x, y) = xy + t(c_1x + c_2y^2 + c_3x^2y)$$

$$g(x, y) = xy + t(d_1x + d_2y^2 + d_3x^2y)$$

respectively, where $\text{val}(c_i) = \text{val}(d_i) = 0$ for all i . This means $\text{Trop}(X)$ and $\text{Trop}(Y)$ are the same, and are as pictured in Figure 2.7.




Figure 2.7: $\text{Trop}(X) = \text{Trop}(Y) = \text{Trop}(X \cap Y)$, with vertices at $(-1, 1)$, $(2, 1)$, and $(-1, -2)$

The resultant of f and g with respect to the variable y is

$$\begin{aligned} & t^4(c_2^2d_1^2 - 2c_1c_2d_1d_2 + c_1^2d_2^2)x^2 + t^2(c_1c_2 - c_2d_1 - c_1d_2 + d_1d_2)x^3 \\ & + t^3(-c_2c_3d_1 - c_1c_3d_2 + 2c_3d_1d_2 + 2c_1c_2d_3 - c_2d_1d_3 - c_1d_2d_3)x^4 \\ & + t^4(c_3^2d_1d_2 - c_2c_3d_1d_3 - c_1c_3d_2d_3 + c_1c_2d_3^2)x^5, \end{aligned}$$

and the resultant of f and g with respect to the variable x is

$$\begin{aligned} & t^4(c_2c_3d_1^2 - c_1c_3d_1d_2 - c_1c_2d_1d_3 + c_1^2d_2d_3)y^3 \\ & + t^3(2c_2c_3d_1 - c_1c_3d_2 - c_3d_1d_2 - c_1c_2d_3 - c_2d_1d_3 + 2c_1d_2d_3)y^4 \\ & + t^2(c_2c_3 - c_3d_2 - c_2d_3 + d_2d_3)y^5 + t^4(c_3^2d_2^2 - 2c_2c_3d_2d_3 + c_2^2d_3^2)y^6. \end{aligned}$$

The stable tropical intersection consists of the three vertices of the triangle. Let us consider possible configurations of the three intersection points that have all three intersection points lying on the triangle, rather than on the unbounded rays. These are the configurations of zeros of rational functions with poles precisely at the three vertices; let h^{trop} be such a function. Label the vertices clockwise starting with $(-1, 1)$ as v_1, v_2, v_3 . Starting from v_1 and going clockwise, label the poles of h^{trop} as w_1, w_2, w_3 . Let δ_i denote the signed lattice distance between v_i and w_i , with counterclockwise distance negative. Then a necessary condition for the w_i 's to be the poles of h^{trop} is $\delta_1 + \delta_2 + \delta_3 = 0$: otherwise the value of h^{trop} would have a nonzero net change in one rotation around the cycle. In fact this condition is sufficient to guarantee the existence of such an h^{trop} . It follows that the w_i 's cannot be in all different or all the same line segment of triangle, as all different would have $\delta_1 + \delta_2 + \delta_3 > 0$ and all

the same would have $\delta_1 + \delta_2 + \delta_3 \neq 0$. Hence we need only show that each configuration with exactly two w_i 's on the same edge satisfying $\delta_1 + \delta_2 + \delta_3 = 0$ is achievable.

There are six cases to handle, since there are three choices for the edge with a pair of points and then two choices for the edge with the remaining point. We will focus on the case where w_1 and w_2 are on the edge connecting v_1 and v_2 , and w_3 is on the edge connecting v_2 and v_3 , as shown in Figure 2.8. Let $\delta_1 = r$ and $\delta_2 = -s$, where $r, s > 0$, and $2 - s \geq -1 + r$. It follows that $\delta_3 = -(r - s)$, and that $r > s$ by the position of w_3 .




Figure 2.8: The desired configuration of intersection points, where $\delta_1 = r > 0$, $\delta_2 = -s < 0$, and $\delta_3 = -(r - s) < 0$.

To achieve the configuration specified by r and s , set

$$c_1 = 3 + t^r, c_2 = 3, c_3 = 1, d_1 = 3, d_2 = 3 + 2t^{r-s}, d_3 = 2.$$

The valuations of the coefficients of the resultant polynomial with x terms are $4 + 2(r - s)$ for x^2 , $2 + 2r - 3$ for x^3 , $3 + r - s$ for x^4 , and 4 for x^5 . It follows that the valuations of the x -coordinates are $2 - s$, $-1 + r$, and $-1 - s + r$. When coupled with rational function restrictions, this implies that the intersection points of X and Y tropicalize to $(-1 + r, 1)$, $(2 - s, 1)$, and $(-1 - s + r, -2 - s + r)$, which are indeed the points w_1 , w_2 , and w_3 we desired.

The five other cases with all three intersection points in the triangle are handled similarly, and the cases with one or more intersection point on an infinite ray are even simpler.

These examples provide not only a helpful check of Theorem 2.1.2, but also evidence that all possible intersection configurations can in fact be achieved. Future work towards proving this might be of a Berkovich flavor, or may have more to do with tropical modifications, as highlighted in the two proofs of Theorem 2.1.2. Regardless of the approach, future investigations should not only look towards proving Conjecture 2.1.8, but also towards algorithmically lifting tropical intersection configurations to curves yielding them.

2.3 Bitangents of tropical plane quartic curves

In classical algebraic geometry, smooth plane quartics are the simplest examples of algebraic curves that are not hyperelliptic, as well as the simplest examples of canonically embedded

curves. The study of the enumerative geometry of these curves dates back to at least 1834, when Plücker [Plü34] showed that a smooth plane quartic curve over \mathbb{C} has 28 bitangent lines. The main results in this section are the following two theorems, which show that smooth tropical plane quartics behave similarly to their algebraic counterparts:

Theorem 2.3.1. *Every smooth tropical plane quartic curve admits exactly 7 equivalence classes of bitangent lines.*

(See Definitions 2.3.4 and 2.3.11 for the precise meaning of bitangent lines and equivalence classes thereof.)

Theorem 2.3.2. *Every smooth tropical plane quartic curve is nonhyperelliptic.*

Many classical arguments regarding smooth algebraic plane quartic curves break down in the tropical setting because smooth tropical quartics are not canonically embedded in the strongest possible sense: while the stable intersection with any tropical line gives a canonical divisor (cf. Lemma 2.3.6), it is not the case that every canonical divisor arises in this way. For example, the standard representative of the canonical divisor of a smooth tropical plane quartic (when viewed as a metric graph) does not generally lie on a tropical line. Thus many standard algebro-geometric arguments are invalid tropically. However, a combination of geometric results and case-by-case analysis will allow us to combine the embedded and abstract theories of tropical curves and prove Theorems 2.3.1 and 2.3.14.

A smooth tropical plane quartic curve is obtained as a dual graph of a unimodular triangulation of the lattice points of the triangle $T_4 = \text{conv}\{(0, 0), (0, 4), (4, 0)\}$. Since T_4 has three interior lattice points, the genus of the tropical curve (i.e., the genus of its skeleton considered as a metric graph) is three. The skeleton of such a curve is then a trivalent, connected, leafless graph of genus 3. There are exactly five such graphs up to homeomorphism, illustrated in Figure 2.9. We name each graph (ℓbc) , where ℓ is the number of loops, b is the number of bi-edges, and c is the number of *cut-edges* (also known as *bridges*). This labelling uniquely identifies the five graphs.




Figure 2.9: The five trivalent graphs of genus 3

Of these five combinatorial types of graph, only (000), (020), (111), and (212) occur as the skeleton of a smooth plane quartic curve. The graph (303) is ruled out by Proposition 3.5.4, and also by an enumeration of all unimodular triangulations of T_4 as performed in




Figure 2.10: Smooth tropical plane quartic curves and the dual triangulations. The skeletons are highlighted.

Section 3.4. Examples of smooth tropical plane quartic curves with skeletons of the other four types are illustrated in Figure 2.10.

As usual, the skeleton of a smooth plane tropical curve can be viewed as a metric graph by letting the length of an edge be its lattice length. In the same way, we can (more generally) consider *any* connected compact subset of a smooth tropical plane curve as a metric graph.

The definition of the metric on a smooth tropical plane curve is justified in part by the following result:

Lemma 2.3.3. *Let C be a smooth tropical plane curve and let L_1, L_2 be tropical lines. Suppose that Σ is a connected subgraph of C that contains the skeleton and the stable intersections $L_1 \cdot C$ and $L_2 \cdot C$. Then $L_1 \cdot C$ and $L_2 \cdot C$ are linearly equivalent as divisors on Σ .*

Proof. For $i = 1, 2$, let f_i be the pullback to C of the defining piecewise-linear equation of the tropical line L_i , and let φ_i be the rational function on Σ obtained by restricting f_i to Σ .

By [AR10], the divisor $\text{div}(f_i)$ on C is exactly the stable intersection divisor of L_i and C . For a point x in the interior of Σ , we see from the definitions that $\text{div}(\varphi_i)(x)$ coincides with $\text{div}(f_i)(x)$. For x on the boundary of Σ , the difference between $\text{div}(f_i)(x)$ and $\text{div}(\varphi_i)(x)$ equals the outgoing slope of f_i on the infinite edge of C emanating from x . If x is on an edge pointing in the $(1, 1)$ direction, then since we are assuming that the intersection of L_i and C is contained in Σ , it follows that x is not in the region spanned by the two rays of L_i pointing in the $(-1, 0)$ and $(0, -1)$ directions. Therefore, the slope of f_i on the infinite edge emanating from x is 1. Similarly, if x is on an infinite edge pointing in a different direction, the slope of f_i is zero on that edge.

Therefore, $\text{div}(\varphi_i)$ is the divisor obtained from $\text{div}(f_i)(x)$ by subtracting a chip from every boundary point of Σ on an infinite edge pointing in the $(1, 1)$ direction. It follows that

$\varphi_1 - \varphi_2$ is a rational function on Σ whose associated principal divisor is exactly $C \cdot L_1 - C \cdot L_2$, as required. \square

Definition 2.3.4. A tropical line L is said to be *bitangent* to a smooth tropical plane quartic curve C if $L \cap C$ consists of two components each with stable intersection multiplicity 2, or one component with stable intersection multiplicity 4. A tropical bitangent line is called *skeletal* if its intersection with the tropical curve is contained in the skeleton.

An outline of a classical proof that there are 28 bitangent lines is as follows. Every smooth algebraic plane quartic curve X is canonically embedded, so any line section is a canonical divisor. It follows that bitangent lines correspond bijectively to linear equivalence classes of effective divisors D of degree two such that $2D$ is a canonical divisor (since X is not hyperelliptic, one can in fact remove the phrase “linear equivalence classes”). Such divisors are called *effective theta characteristics*. The set of all theta characteristics (effective or not) is canonically a torsor for $\text{Jac}(X)[2]$, which has order 64. Via a non-trivial analysis of bilinear forms in characteristic 2 as in [DO88, Ch. VIII.2], one proves that there are exactly 28 effective theta characteristics, obtaining the desired result.

One can define theta characteristics for a metric graph in the same way:

Definition 2.3.5. A *theta characteristic* on a metric graph Γ is a linear equivalence class of divisors D such that $2D$ is linearly equivalent to the canonical divisor on Γ . A theta characteristic is called *effective* if one can choose D to be an effective divisor.

Zharkov [Zha10] proves that a metric graph Γ of genus g has 2^g theta characteristics, of which exactly $2^g - 1$ are effective. The fact that there are 2^g theta characteristics (instead of 2^{2g} as in the classical case) comes from the fact that $\text{Jac}(\Gamma)$ is a *real* (rather than complex) torus of dimension g . In both the classical and tropical situations, the set of theta characteristics is naturally a torsor for the 2-torsion in the Jacobian. Zharkov’s proof that all but one of the theta characteristics on Γ is effective does not seem to have an algebraic analogue; for the reader’s convenience we summarize some of the ideas behind Zharkov’s proof in Lemma 2.3.7 below.

Zharkov’s theorem implies that an abstract metric graph Γ of genus 3 has exactly 7 effective theta characteristics. However, unlike the classical case, it is not obvious (and in fact not true) that bitangent lines to a smooth tropical plane quartic Γ are in bijection with effective theta characteristics. In fact, it is not even completely obvious *a priori* how to define a tropical bitangent line. With our definition (which seems to be the only reasonable one), there are examples of smooth tropical plane quartics with infinitely many bitangent lines; see Example 2.3.8. (This is reminiscent of Vigeland’s example of cubic surfaces with infinitely many lines rather than 27 [Vig10]; a possible relationship between tropical plane quartic curves and tropical cubic surfaces warrants further investigation.) Moreover, unlike the classical case, it appears to be subtle to prove that for every effective theta characteristic $[D]$, there is a tropical line L and a divisor $D' \sim D$ such that $L \cdot C = 2D'$. We will prove this via a case-by-case analysis of the different possible smooth tropical plane quartic curves.

First, however, we show that while tropical plane quartics are not canonically embedded, it is still true that every line section belongs to the class of the canonical divisor. We thank Yang An for suggesting the following proof:

Lemma 2.3.6. *For a smooth tropical plane quartic curve C and a tropical line L , the stable intersection divisor $L \cdot C$ is canonical.*

Proof. By Lemma 2.3.3, any two divisors obtained as line sections of C are linearly equivalent. Now let L be a tropical line, and let D be the divisor $L \cdot C$. By Bézout's Theorem [RGST05], the degree of D is 4, and by the tropical Riemann-Roch theorem, it is canonical if and only if its rank is 2. Let Γ be a compact subset of C containing the support of D , considered as a metric graph. Choose a loopless model G for Γ , and let E be any effective divisor of degree 2 supported on the vertices of G .

If E consists of two distinct vertices u and v , let L' be a tropical line passing through them. Since u and v are vertices, they are contained in the stable intersection $L' \cdot C$ (even when the intersection is not transversal). By the argument above, $L \cdot C$ is linearly equivalent to $L' \cdot C$, and so D is equivalent to a divisor containing E .

Otherwise, E consists of a double point at some vertex w of G . Let L be a tropical line with a vertex at w . Suppose first that, locally at w , Γ is parallel to L , i.e. Γ is trivalent at w with edges pointing towards $(-1, 0)$, $(0, -1)$, $(1, 1)$. Since C is not a line, at least one of these edges is finite and the stable intersection contains a point at w and at the vertices at the end of those edges. By chip-firing the complement of a neighborhood of w , we obtain a divisor containing at least two chips at w . If, on the other hand, L is not parallel to Γ at w , the stable intersection is easily seen to have multiplicity 2 at w .

Since the vertices of G are a rank-determining set by [Luo11], the rank of D is at least 2, and therefore D is canonical. \square

We now describe Zharkov's algorithm for finding the theta characteristics of a metric graph Γ (see 2.3.8 for an example). Place a nonzero $(\mathbb{Z}/2\mathbb{Z})$ -flow on the graph; equivalently, choose a subset of the edges so that around each vertex, the number of selected edges equals 0 modulo 2. Let S be the support of the flow. Orient the edges of $\Gamma - S$ indicating movement away from S . At the points where the orientations conflict, place a number of chips equal to $\#\{\text{incoming edges}\} - 1$. For generic lengths on the graph, these points are the local minimizers of distance to the set S .

Lemma 2.3.7 (Zharkov). *Let Γ be a metric graph of genus g . Then:*

1. *Each of the $2^g - 1$ divisors constructed as above is an effective theta characteristic.*
2. *Different choices of $(\mathbb{Z}/2\mathbb{Z})$ -flows yield non-linearly equivalent divisors.*
3. *The remaining theta characteristic on Γ is non-effective.*

We refer the reader to [Zha10] for an explicit description of the non-effective theta characteristic, which will not be needed for this work.

Example 2.3.8. Consider the genus 3 metric graph Σ in Figure 2.11. (We have not included lengths, since it will not affect the combinatorics of this example.) Place a nonzero $(\mathbb{Z}/2\mathbb{Z})$ -flow on it, and let $P + Q$ be the corresponding effective theta characteristic for Σ .

We then consider a tropical curve C which has Σ as its minimal skeleton. Drawing a line through P and Q does indeed give a bitangent to C , as shown in Figure 2.11 (showing the tropical curve without its exterior branches). One component of the intersection, containing P , is a line segment with stable tropical intersection number 2. The other component, the point Q , is a tropical intersection of multiplicity 2. Hence $L \cdot C = 2P + 2Q$.

Note that the bitangent line in Figure 2.11 can be translated horizontally while remaining bitangent, leaving P alone but moving Q . This gives an infinite family of bitangents whose intersection divisors with C are linearly equivalent. This behavior can be seen already in the metric graph Σ : there are effective divisors linearly equivalent to $P + Q$, namely those leaving P fixed and moving Q along the edge containing it.




Figure 2.11: The construction of a theta characteristic, and a corresponding bitangent line

We used Zharkov’s algorithm to compute the 7 effective theta characteristics for each of the first four graphs in Figure 2.9, and summarize the output in Figure 2.16. The computation depends slightly on the relative lengths of certain edges, but is made easier by the fact that smooth plane quartic curves are not hyperelliptic (Theorem 2.3.14).

We will now show that each effective theta characteristic on a smooth tropical plane quartic curve is represented by a bitangent line:

Proposition 2.3.9. *Let $P + Q$ be an effective theta characteristic for a smooth tropical plane quartic curve C , and let L be a tropical line connecting P and Q . Then L is a bitangent line to C .*

In the course of proving this, we will split into three possible cases:

- (i) $P + Q$ is *rigid*, i.e., not linearly equivalent to any other effective divisor.

- (ii) One or both of P and Q is *flexible*, i.e., can be moved independently of the other while maintaining linear equivalence.
- (iii) $P + Q$ is linearly equivalent to effective divisors obtainable only by moving the two points in tandem.

The most difficult case ends up being (iii). Before proving Proposition 2.3.9, we will prove the following lemma to help with this case.

Lemma 2.3.10. *Let P and Q be two distinct points of Γ such that $P + Q$ is not rigid, and neither P nor Q is flexible (as defined in Proposition 2.3.9). Then $\Gamma \setminus \{P, Q\}$ is disconnected.*

Proof. Let φ be a non-constant rational function such that $P + Q + \text{div}(\varphi)$ is effective, and let M be the global maximum of φ . Then $\text{div}(\varphi) < 0$ on any boundary point of $\varphi^{-1}(M)$ of valency greater than 1. Therefore, the global maximum of φ is obtained at P and Q , and any other point where the global maximum is obtained is an interior point of $\varphi^{-1}(M)$. Moreover, since $\text{div}(\varphi)$ is -1 at P and Q , the slope of φ is -1 on a single edge leaving P, Q , and 0 on the others.

Let x be a point on an edge leaving P where the slope is 0 , and y a point on an edge leaving Q where the slope is -1 . Then x is in $\varphi^{-1}(M)$, and y is not, so any path between x and y passes through the boundary of M , namely through either P or Q . Therefore, by removing P and Q , the graph becomes disconnected. \square

We are now ready to prove that effective theta characteristics give rise to bitangent lines.

Proof of Proposition 2.3.9. First, if the stable tropical intersection $L \cdot C$ does not contain P or Q , we can move the divisor $L \cdot C$ within $L \cap C$ to a linearly equivalent divisor $P + Q + R + S$, with all four points still in $L \cap C$. We wish to show that $P + Q + R + S = 2P + 2Q$.

By the definition of a theta divisor, we know that $2P + 2Q$ is canonical, and by Lemma 2.3.6, $L \cdot C$ is canonical for any line L . Thus $2P + 2Q \sim P + Q + R + S$, and so $P + Q \sim R + S$. There are three cases to deal with. See Figure 2.16 and the preceding discussion for examples.

- (i) $P + Q$ is rigid. In this case, it follows immediately that $P + Q = R + S$.
- (ii) One or both of P and Q is *flexible*. In this case, we claim that it suffices to show that the intersection multiplicity of L with C at each flexible point is 2. Indeed, if the multiplicity at P is 2, and Q cannot move independently, then $L \cdot C = 2P + Q + S \sim K$ together with $2P + 2Q \sim K$ implies that $Q \sim S$. Since we assumed that Q cannot move independently, we obtain $Q = S$, namely $L \cdot C = 2P + 2Q$.

To see that the multiplicity at each flexible point is 2, we notice that, by [BN07, Corollary 4.7], a chip can only move independently when it is on a bridge edge. This can only occur for a curve with skeleton (111) or (212). The slopes of the bridge edges in these curves are always either -2 or $-\frac{1}{2}$. By the exhaustive search illustrated in

Figure 2.16, the (212) graph has at most four theta characteristics with a chip on a bridge, and the (111) graph has at most three such theta characteristics (depending on the edge lengths). Moreover, the chip on a bridge with slope -2 is always obtained by intersection with the edge of a tropical line pointing in the $(-1, 0)$ direction, and the chip on an edge with slope $-\frac{1}{2}$ is obtained by intersection with an edge pointing in the $(0, -1)$ direction. Either way, the intersection is of multiplicity 2.

- (iii) $P + Q$ is linearly equivalent to effective divisors obtainable only by moving the two points in tandem. By Lemma 2.3.10, this can only occur when P and Q form a 2-edge-cut. By the classification in Figure 2.16, this happens only for the graphs (020) and (212).




Figure 2.12: The divisor D_1 on (020), and the divisor D_2 on (212).

We begin with the graph (020). From the combinatorial type of C , we know that the corresponding Newton subdivision has exactly two edges between the interior points $(1, 1)$, $(2, 1)$, and $(1, 2)$. By symmetry, we may assume that $(1, 1)$ is connected to the two other points, giving a Newton subdivision of the form shown on the left side of Figure 2.13. The simple cycle $\gamma \subset C$ containing the support of the divisor D_1 is determined by which other points connect to $(1, 1)$ in the triangulation, with all possibilities for the edges and a corresponding cycle shown in Figure 2.13. The relevant points are labelled by i with $1 \leq i \leq 9$, and the edges are labelled by e_i , where e_i is dual to the edge connecting point i to $(1, 1)$. Although not all triangulations will have all of the edges, each triangulation will have some subset of them making up the cycle γ .

Let L_E be a Euclidean line with slope 1 passing through the midpoint of the edge of C dual to the segment connecting $(1, 1)$ and $(2, 2)$. The line L_E must pass through exactly one more point of the cycle γ , and this point must in fact be contained in $e_4 \cup e_5 \cup e_6$; this is because the edges e_1 , e_2 , and e_3 have slopes ≥ 0 and ≤ 1 , and the edges e_7 , e_8 , and e_9 have slopes ≥ 1 , making it impossible for L_E to pass through them due to the position of the edges.

We claim that the tropical line L defined by placing the vertex at the point $p_E = L_E \cap (e_4 \cup e_5 \cup e_6)$ is bitangent to C . It will certainly pass through the edge of C dual




Figure 2.13: Triangulations for the (020) combinatorial type, and the middle cycle.

to the segment connecting $(1, 1)$ and $(2, 2)$ with multiplicity 2. We must show that the other component of $C \cap L$ also has multiplicity 2. We can handle this in several cases:

1. If p is the interior of e_4 , then $C \cap L$ contains the vertical line segment below p on e_4 , which will have multiplicity 2. A similar argument holds if p is in the interior of e_6 .
2. If p is in the interior of e_5 , then p is an isolated intersection point of C and L , and will have multiplicity 2 due to slopes.
3. If $p = e_4 \cap e_5$, then that component of $C \cap L$ is a horizontal ray in the direction $(-1, 0)$. A small perturbation will put us in one of the previous two cases (or give two intersection points), meaning it is an intersection of multiplicity 2. A similar argument holds if $p = e_5 \cap e_6$.
4. If $p = e_4 \cap e_6$ (meaning there is no e_5), a small perturbation will give either a vertical line segment contained in e_4 , a horizontal line segment contained in e_6 , a diagonal line segment dual to the edge connecting points 4 and 6, or a pair of points. This too will be an intersection of multiplicity 2.

Hence L is indeed a bitangent line to C . To verify that it corresponds to the theta characteristic D_1 , note that it must correspond to *some* theta characteristic, and D_1 is the only one with both points on the middle cycle.

The proof for the (212) case is nearly identical. Due to the combinatorial type, we know there are no edges between the three interior points in the triangulation, and by symmetry we may assume the point $(1, 1)$ corresponds to the middle cycle. This means there is an edge connecting $(1, 1)$ to $(2, 2)$. At first glance, the other points $(1, 1)$ may connect to are the points $1, \dots, 9$ as in the previous proof; in fact, constraints on the triangulation narrow the points down to 4, 5, 6. If we consider a Euclidean line through the midpoint of the edge corresponding to the edge connecting $(1, 1)$ to $(2, 2)$, it must (as in the previous case) pass through $e_4 \cup e_5 \cup e_6$. Letting L be the tropical line having a vertex at that point, we once again argue that it is bitangent to C , and find that it corresponds to the theta characteristic D_2 .

□

We have shown that effective theta characteristics give rise to bitangent lines. However, we have seen (cf. Example 2.3.8) that a smooth tropical plane quartic curve can have infinitely many bitangent lines, and so we wish to define a notion of equivalence for tropical bitangents so that if one bitangent can be moved to another while preserving the bitangency condition, then the two are equivalent.

Definition 2.3.11. Two bitangent lines are *equivalent* if they correspond to linearly equivalent theta characteristics.

We are now ready to prove Theorem 2.3.1, that every smooth tropical plane quartic curve admits exactly 7 equivalence classes of bitangent lines.

Proof of Theorem 2.3.1. First we note that any bitangent line is equivalent to a skeletal bitangent line. To see this, suppose that a tropical line intersects C at the points P and Q with multiplicity 2, and that either of them is on the infinite branches and not on the skeleton Σ of the curve. Then, since the infinite branches are leaf edges, the theta divisor $P + Q$ is linearly equivalent to a divisor $P' + Q'$ on the skeleton. By Proposition 2.3.9, the line passing through $P' + Q'$ is a bitangent line.

By Lemma 2.3.7, there are exactly seven classes of effective theta characteristic on the skeleton of a tropical quartic curve C . By Proposition 2.3.9 each theta characteristic gives rise to a skeletal bitangent line, and by definition the classes of theta characteristics correspond to classes of bitangent lines. Hence there are *at least* seven bitangent lines up to equivalence. For a bitangent line L to C , the stable intersection $L \cdot C$ is linearly equivalent to a theta characteristic by Lemma 2.3.6, meaning that there are *at most* seven bitangent lines up to equivalence. This completes the proof. □

If X is a smooth plane quartic for which $\text{Trop}(X)$ is tropically smooth, it is natural to guess that the 28 bitangents on X specialize in groups of 4. This has been proven in the case that $\text{Trop}(X)$ is a smooth tropical curve with skeleton (000) by Melody Chan and Pakawut Jiradilok [CJ15].

Conjecture 2.3.12. *Let X be a smooth plane quartic and assume that $\text{Trop}(X)$ is tropically smooth. Then each odd theta characteristic of Γ is the specialization of four effective theta characteristics of X (counted with multiplicity).*

As is well known, a smooth algebraic plane quartic curve is never hyperelliptic. We now prove a tropical analogue of this statement. This is a useful result in computing theta characteristics for genus 3 graphs, as it put restrictions on possible edge lengths (cf. Figure 2.16).

We will use the following necessary condition for hyperellipticity.

Lemma 2.3.13. *Let Γ be a metric graph with a 2-edge-cut. If Γ is hyperelliptic, then the two edges of the cut are of equal length.*

Proof. Let e_1 and e_2 be the two edges of the cut, and let Γ' and Γ'' be the two connected components of $\Gamma \setminus \{e_1 \cup e_2\}$. By the metric graph analogue of [BN09, Corollary 5.10], we may assume that Γ is leafless with no bridge edges. Let $D = p + q$ be a divisor of degree 2 and rank 1 on Γ . Since D has positive rank, we may assume that p is on the intersection between e_1 and Γ' .

We will begin by showing that q is on $e_1 \cup e_2$. Assume for the sake of contradiction that it is not, and choose some point x in the interior of e_2 . Then one easily checks that D is x -reduced, contradicting the fact that D has rank 1. Moreover, q has to be on the intersection between Γ' and e_2 . Otherwise, choose a point y in Γ' . Then again, D is y -reduced, and we arrive at a contradiction.

Now, move p and q along the cut while preserving linear equivalence, until one of them reaches Γ'' . By the same argument as before (switching the roles of Γ' and Γ'' , and of p and q if necessary), both p and q must reach Γ'' together. It follows that e_1 and e_2 have the same length. \square

Let C be a smooth tropical plane quartic, so that its skeleton must have one of the first four combinatorial types depicted in Figure 2.9. According to [Cha13], out of these four, only the ones shown in Figure 2.14 have hyperelliptic metric realizations. The following theorem shows that none of these realizations occur for the skeleton of a smooth tropical plane quartic curve.




Figure 2.14: The three combinatorial types of graphs realizable as trivalent hyperelliptic metric graphs of genus 3, and the same graphs with lengths making them nonhyperelliptic

Theorem 2.3.14. *Every smooth tropical plane quartic curve is nonhyperelliptic.*

Proof. By Lemma 2.3.13, whenever a hyperelliptic metric graph contains a 2-edge-cut, that cut consists of two edges of equal length. As we will see, this never occurs for smooth tropical plane quartic curves.

Consider a unimodular triangulation of the standard triangle with vertices $(0, 0)$, $(4, 0)$, $(0, 4)$, corresponding to a tropical curve C . There are four possible combinatorial types of the skeleton Γ of C , namely the first four pictured in Figure 2.9. The graph (000) is never hyperelliptic [Cha13], regardless of the length of the edges. The other three can be hyperelliptic if given the appropriate lengths. We will show that such lengths are never achieved by Γ .




Figure 2.15: Edges coming from $(1, 1)$, and the corresponding cycle.

Without loss of generality, we may assume that the point $(1, 1)$ corresponds to the large cycle γ in the middle as in the second row of Figure 2.14. Each edge emanating from $(1, 1)$ in the triangulation corresponds to an edge of the cycle. Figure 2.15 shows all possibilities for such edges and the corresponding cycle.

In the three combinatorial types in question, the 2-edge-cut consists of certain edges of the cycle γ . The first edge, which we will denote e_1 , is dual to the edge connecting $(1, 1)$ with $(2, 2)$ in the triangulation. The other edge of the cut, denoted e_2 , comes from the edges in the triangulation emanating from $(1, 1)$ with angles strictly between $\frac{\pi}{4}$ and 2π . While a triangulation does not uniquely determine the lattice lengths of the edges in a tropical curve, we claim that the data of the slopes of the edges suffices to conclude that e_2 must be longer than e_1 . We will show this for each of the three combinatorial types separately. Let ℓ_1 and ℓ_2 be the lengths of the edges e_1 and e_2 .

1. (020). In this case, the cycle γ decomposes as $e_1-e_h-e_2-e_v$, where e_h and e_v are horizontal and vertical edges with lengths ℓ_h and ℓ_v , respectively. Consider the line segments in e_2 that are dual to edges in the triangulation connecting $(1, 1)$ with points on the horizontal edge of the Newton polygon. The sum of the horizontal widths of these segments must be at least the sum of the horizontal widths of e_1 and e_h : otherwise the cycle γ would not be closed. Since these line segments in e_2 have integer slopes, each of them has lattice length equal to horizontal width. The same holds for e_1 and e_h , implying $\ell_2 \geq \ell_1 + \ell_h > \ell_1$.
2. (111). In this case, without loss of generality the cycle γ decomposes as $e_1-e_h-e_2$, where e_h is a horizontal edge of length ℓ_h . The same argument from the (020) case holds, giving $\ell_2 \geq \ell_1 + \ell_h > \ell_1$.
3. (212). In this case, e_1 and e_2 form the whole cycle γ . Since edges must separate $(1, 1)$ from $(1, 2)$ and $(2, 1)$ in the triangulation, there are very few possibilities for the form of e_2 . Either e_2 consists of a vertical segment and a horizontal segment, which implies $\ell_2 = 2\ell_1$, or e_2 consists of a vertical, a diagonal, and a horizontal segment with lengths ℓ_v , ℓ_d , and ℓ_h . In this latter case, $\ell_v = \ell_h$ and $\ell_d = \ell_1 - \ell_v$, so $\ell_2 = \ell_v + \ell_d + \ell_h = \ell_v + (\ell_1 - \ell_v) + \ell_v = \ell_v + \ell_1 > \ell_1$.

This proves the claim, and hence the theorem. \square

We close this section by filling in some details in the proof of Proposition 2.3.9. In particular, we made use of knowledge about the 7 theta characteristics of each of the four graphs arising as the skeleton of a smooth plane quartic curve. In Figure 2.16, we illustrate all the theta characteristics for the four types of genus 3 graphs relevant to us. Each graph has the support of a nonzero $(\mathbb{Z}/2\mathbb{Z})$ -flow in bold, and in most cases the corresponding theta characteristic is illustrated as a pair of circular points. In the cases where edge lengths might change the combinatorial position of the points of the theta characteristic, the other possibility is illustrated by a pair of crosses. There are degenerate cases where one of these moves to a vertex, but this will not affect our arguments. We have also taken advantage of Theorem 2.3.14, which allows us to assume asymmetry for the middle cycle in the second, third, and fourth columns.

Labeling the columns 1, 2, 3, and 4 and the rows A, B, C, D, E, F, and G, we have the following classification of these 28 theta characteristics by cases (i), (ii), and (iii) of Proposition 2.3.9:

- (i) The 20 not falling into cases (ii) or (iii).
- (ii) The 6 theta characteristics 3E, 3F, 3G, 4D, 4F, 4G.
- (iii) The 2 theta characteristics 2E, 4E.




Figure 2.16: The 7 theta characteristics for each type of graph

Chapter 3

Smooth tropical plane curves

This chapter seeks to answer the question: when is a metric graph the skeleton of a plane tropical curve? In Section 3.1 we define $\mathbb{M}_g^{\text{planar}}$, the moduli space of graphs of genus g that arise in this way. We prove the dimension of this space is $2g + 1$ for $g > 3$ and $g \neq 7$ using honeycomb curves in Section 3.2. In Section 3.3 we discuss methodology for algorithmically computing $\mathbb{M}_g^{\text{planar}}$ for a fixed g , which we then carry out up to genus 5 in Section 3.4 for nonhyperelliptic tropical plane curves. We close the chapter with two sections on hyperelliptic curves, with Section 3.5 focusing on theory and Section 3.6 focusing on computations.

Except for the majority of Section 3.5, this chapter’s content comes from the paper “Moduli of tropical plane curves” [BJMS15], coauthored with Michael Joswig, Sarah Brodsky, and Bernd Sturmfels and appearing in *Research in the Mathematical Sciences*.

3.1 The moduli space of tropical plane curves

Let K be an algebraically closed field that is complete with respect to a surjective non-Archimedean valuation $\text{val} : K^* \rightarrow \mathbb{R}$, and let $g \geq 2$. As discussed in Section 1.4, there is a surjective tropicalization map from (the K -valued points in) the moduli space of smooth curves of genus g to the moduli space of metric graphs of genus g :

$$\text{trop} : \mathcal{M}_g \rightarrow \mathbb{M}_g. \quad (3.1.1)$$

Both spaces have dimension $3g - 3$ for $g \geq 2$. Inside of \mathcal{M}_g is $\mathcal{M}_g^{\text{nd}}$, the locus of non-degenerate plane curves over K of genus g . This moduli space was introduced by Castryck and Voight in [CV09]. It has dimension $2g + 1$, and admits a decomposition

$$\mathcal{M}_g^{\text{nd}} := \bigcup_P \mathcal{M}_P \quad (3.1.2)$$

by Newton polygons P , where each P has g interior lattice points.

We now define tropical spaces analogous to \mathcal{M}_P and $\mathcal{M}_g^{\text{nd}}$, and throughout this chapter perform a computational study of these new spaces. Fix a (convex) lattice polygon P with $g = \#(\text{int}(P) \cap \mathbb{Z}^2)$. Let \mathbb{M}_P be the closure in \mathbb{M}_g of the set of metric graphs that are realized as the skeletons of smooth tropical plane curves with Newton polygon P . For a fixed regular unimodular triangulation Δ of P , let \mathbb{M}_Δ be the closure of the cone of metric graphs that are skeletons of smooth tropical curves dual to Δ . These curves all have the same skeleton G , and \mathbb{M}_Δ is a convex polyhedral cone in the orthant $\mathbb{R}_{\geq 0}^{3g-3}$ of metrics on G . Working modulo automorphisms of G , we identify \mathbb{M}_Δ with its image in the stacky fan \mathbb{M}_g .

Now fix the graph of the skeleton G but vary the triangulation. The resulting subset of $\mathbb{R}_{\geq 0}^{3g-3}$ is a finite union of closed convex polyhedral cones, so it can be given the structure of a polyhedral fan. Moreover, by appropriate subdivisions, we can choose a fan structure that is invariant under the symmetries of G , and hence the image in the moduli space \mathbb{M}_g is a stacky fan:

$$\mathbb{M}_{P,G} := \bigcup_{\substack{\Delta \text{ triangulation of } P \\ \text{with skeleton } G}} \mathbb{M}_\Delta. \quad (3.1.3)$$

We note that \mathbb{M}_P is represented inside \mathbb{M}_g by finite unions of convex polyhedral cones:

$$\mathbb{M}_P = \bigcup_{\substack{G \text{ trivalent graph} \\ \text{of genus } g}} \mathbb{M}_{P,G} = \bigcup_{\substack{\Delta \text{ regular unimodular} \\ \text{triangulation of } P}} \mathbb{M}_\Delta. \quad (3.1.4)$$

The *moduli space of tropical plane curves of genus g* is the following stacky fan inside \mathbb{M}_g :

$$\mathbb{M}_g^{\text{planar}} := \bigcup_P \mathbb{M}_P. \quad (3.1.5)$$

Here P runs over isomorphism classes of lattice polygons with g interior lattice points. The number of such classes is finite by Proposition 3.3.2.

Write P_{int} for the convex hull of the g interior lattice points of P . This is called the *interior hull* of P . If the dimension of P_{int} is 1, we say P is a *hyperelliptic polygon*, and if the dimension of P_{int} is 2, we say P is a *nonhyperelliptic polygon*. It is sometimes helpful to decompose $\mathbb{M}_g^{\text{planar}}$ as

$$\mathbb{M}_g^{\text{planar}} = \mathbb{M}_{g,\text{hyp}}^{\text{planar}} \cup \mathbb{M}_{g,\text{nonhyp}}^{\text{planar}},$$

where

$$\mathbb{M}_{g,\text{hyp}}^{\text{planar}} := \bigcup_{P \text{ hyperelliptic}} \mathbb{M}_P, \quad \mathbb{M}_{g,\text{nonhyp}}^{\text{planar}} := \bigcup_{P \text{ nonhyperelliptic}} \mathbb{M}_P$$

and both unions only include polygons with g interior lattice points. These two spaces are *not* disjoint, although it will be shown in Section 3.5 that the overlap only arises from taking closures of the dense subsets of achievable graphs.

We summarize the objects discussed so far in a diagram of surjections and inclusions:

$$\begin{array}{ccccccc}
 \mathcal{M}_P & \subseteq & \mathcal{M}_g^{\text{nd}} & \subseteq & \mathcal{M}_g & & \\
 \downarrow & & \downarrow & & \downarrow & & \\
 \text{trop}(\mathcal{M}_P) & \subseteq & \text{trop}(\mathcal{M}_g^{\text{nd}}) & \subseteq & \text{trop}(\mathcal{M}_g) & & (3.1.6) \\
 \cup & & \cup & & \parallel & & \\
 \mathbb{M}_\Delta & \subseteq & \mathbb{M}_{P,G} & \subseteq & \mathbb{M}_P & \subseteq & \mathbb{M}_g^{\text{planar}} & \subseteq & \mathbb{M}_g
 \end{array}$$

By the Structure Theorem for Tropical Varieties [MS15, section 3.3], the dimensions of \mathcal{M}_P and $\mathcal{M}_g^{\text{nd}}$ are preserved under the tropicalization map (3.1.1). The images $\text{trop}(\mathcal{M}_P)$ and $\text{trop}(\mathcal{M}_g^{\text{nd}})$ are stacky fans that live inside $\mathbb{M}_g = \text{trop}(\mathcal{M}_g)$ and have the expected dimension. Furthermore, all maximal cones in $\text{trop}(\mathcal{M}_P)$ have the same dimension since \mathcal{M}_P is irreducible (in fact, unirational).

Our first result reveals that $\mathbb{M}_g^{\text{planar}}$ is full dimensional inside of $\text{trop}(\mathcal{M}_g^{\text{nd}})$:

Theorem 3.1.1. *For all $g \geq 2$ there exists a lattice polygon P with g interior lattice points such that \mathbb{M}_P has the dimension expected from classical algebraic geometry, namely,*

$$\dim(\mathbb{M}_g^{\text{planar}}) = \dim(\mathbb{M}_P) = \begin{cases} 3 & \text{if } g = 2, \\ 6 & \text{if } g = 3, \\ 16 & \text{if } g = 7, \\ 2g + 1 & \text{otherwise.} \end{cases} \quad (3.1.7)$$

In each case, the cone \mathbb{M}_Δ of honeycomb curves supported on P attains this dimension.

Honeycomb curves are introduced in Section 3.2, which also gives the proof of Theorem 3.1.1. Whenever we speak about the “dimension expected from classical algebraic geometry”, as we do in Theorem 3.1.1, this refers to the formulas for $\dim(\mathcal{M}_P)$ and $\dim(\mathcal{M}_g^{\text{nd}})$ that were derived by Castryck and Voight [CV09].

For $g \geq 3$, the inclusions between the second row and the third row in (3.1.6) are strict, by a wide margin. Once again, we see a difference between a tropicalization and a tropical analog, namely between tropicalizations of plane curves and tropical plane curves. One main objective of this chapter is to understand how the latter sit inside the former.

For example, consider $g = 3$ and $T_4 = \text{conv}\{(0, 0), (0, 4), (4, 0)\}$. Disregarding the hyper-elliptic locus, equality holds in the second row:

$$\text{trop}(\mathcal{M}_{T_4}) = \text{trop}(\mathcal{M}_3^{\text{nd}}) = \text{trop}(\mathcal{M}_3) = \mathbb{M}_3. \quad (3.1.8)$$

This is the stacky fan in [Cha12, Figure 1]. The space $\mathbb{M}_{T_4} = \mathbb{M}_{3, \text{nonhyp}}^{\text{planar}}$ of tropical plane quartics is also six-dimensional, but it is smaller. It fills up less than 30% of the curves in \mathbb{M}_3 ; see Corollary 3.4.2. Most metric graphs of genus 3 do *not* come from plane quartics.

For $g = 4$, the canonical curve is a complete intersection of a quadric surface with a cubic surface. If the quadric is smooth then we get a curve of bidegree $(3, 3)$ in $\mathbb{P}^1 \times \mathbb{P}^1$. This

leads to the Newton polygon $R_{3,3} = \text{conv}\{(0, 0), (3, 0), (0, 3), (3, 3)\}$. Singular surfaces lead to families of genus 4 curves of codimension 1 and 2 that are supported on two other polygons [CV09, section 6]. As we shall see in Theorem 3.4.3, \mathbb{M}_P has the expected dimension for each of the three polygons P . Furthermore, $\mathbb{M}_4^{\text{planar}}$ is strictly contained in $\text{trop}(\mathcal{M}_4^{\text{nd}})$. Detailed computations that reveal our spaces for $g = 3, 4, 5$ are presented in Section 3.4.

We close this section with a consideration of classical algebraic geometry. Let \mathcal{T}_g denote the *trigonal locus* in the moduli space \mathcal{M}_g . It is well known that \mathcal{T}_g is an irreducible subvariety of dimension $2g + 1$ when $g \geq 5$. For a proof see [FL08, Proposition 2.3]. A recent theorem of Ma [Ma14] states that \mathcal{T}_g is a rational variety for all g .

We note that Ma's work, as well as the classical approaches to trigonal curves, are based on the fact that canonical trigonal curves of genus g are realized by a certain special polygon P . This is either the rectangle in (3.2.3) or the trapezoid in (3.2.4). These polygons appear in [CV09, Section 12], where they are used to argue that \mathcal{T}_g defines one of the irreducible components of $\mathcal{M}_g^{\text{nd}}$, namely, \mathcal{M}_P . The same P appear in the next section, where they serve to prove one inequality on the dimension in Theorem 3.1.1. The combinatorial moduli space \mathbb{M}_P is full-dimensional in the tropicalization of the trigonal locus. The latter space, denoted $\text{trop}(\mathcal{T}_g)$, is contained in the space of trigonal metric graphs, by Baker's Specialization Lemma [Bak08, section 2].

In general, $\mathcal{M}_g^{\text{nd}}$ has many irreducible components other than the trigonal locus \mathcal{T}_g . As a consequence, there are many skeletons in $\mathbb{M}_g^{\text{planar}}$ that are not trigonal in the sense of metric graph theory. This is seen clearly in the top dimension for $g = 7$, where $\dim(\mathcal{T}_7) = 15$ but $\dim(\mathcal{M}_7^{\text{nd}}) = 16$. The number 16 comes from the family of trinodal sextics in [CV09, §12].

3.2 Honeycomb curves

In this section we prove Theorem 3.1.1. This will be done using the special family of *honeycomb curves*, which are tropical curves dual to *honeycomb triangulations*. The polygons admitting such a triangulation are given by four integer parameters a, b, c and d that satisfy the constraints

$$0 \leq c \leq a, b \leq d \leq a + b. \tag{3.2.1}$$

To such a quadruple (a, b, c, d) , we associate the polygon

$$H_{a,b,c,d} = \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq a \text{ and } 0 \leq y \leq b \text{ and } c \leq x + y \leq d\}.$$

If all six inequalities in (3.2.1) are non-redundant then $H_{a,b,c,d}$ is a hexagon. Otherwise it can be a pentagon, quadrangle, triangle, segment, or just a point. The number of lattice points is

$$\#(H_{a,b,c,d} \cap \mathbb{Z}^2) = ad + bd - \frac{1}{2}(a^2 + b^2 + c^2 + d^2) + \frac{1}{2}(a + b - c + d) + 1,$$

and, by Pick's Theorem [DLRS10, Theorem 9.3.5], the number of interior lattice points is

$$g = \#((H_{a,b,c,d})_{\text{int}} \cap \mathbb{Z}^2) = ad + bd - \frac{1}{2}(a^2 + b^2 + c^2 + d^2) - \frac{1}{2}(a + b - c + d) + 1.$$

The *honeycomb triangulation* Δ subdivides $H_{a,b,c,d}$ into $2ad + 2bd - (a^2 + b^2 + c^2 + d^2)$ unit triangles. It is obtained by slicing $H_{a,b,c,d}$ with the vertical lines $\{x = i\}$ for $0 < i < a$, the horizontal lines $\{y = j\}$ for $0 < j < b$, and the diagonal lines $\{x + y = k\}$ for $c < k < d$. The tropical curves dual to Δ look like honeycombs, as in Figure 3.1. The corresponding skeletons G are called *honeycomb graphs*.




Figure 3.1: The honeycomb triangulation of $H_{5,4,2,5}$, the tropical curve, and its skeleton

If $P = H_{a,b,c,d}$, then its interior hull P_{int} is a honeycomb polygon as well: a translate of P_{int} can be obtained from P by decreasing the values of a, b, c, d by an appropriate amount. For instance, if $P = H_{5,4,2,5}$ is the polygon in Figure 3.1 then its interior hull is $P_{\text{int}} = H_{3,3,1,2} + (1, 1)$.

Lemma 3.2.1. *Let Δ be the honeycomb triangulation of $P = H_{a,b,c,d}$. Then*

$$\dim(\mathbb{M}_\Delta) = \#(P_{\text{int}} \cap \mathbb{Z}^2) + \#(\partial P_{\text{int}} \cap \mathbb{Z}^2) + \#\text{vertices}(P_{\text{int}}) - 3.$$

Proof. The honeycomb graph G consists of $g = \#(P_{\text{int}} \cap \mathbb{Z}^2)$ hexagons. The hexagons associated with lattice points on the boundary of P_{int} have vertices that are 2-valent in G . Such 2-valent vertices get removed, so these boundary hexagons become cycles with fewer than six edges. In the orthant $\mathbb{R}_{\geq 0}^{3g-3}$ of all metrics on G , we consider the subcone of metrics that arise from P . In addition to the non-negativity constraints, this convex cone is defined by

- (a) one linear inequality for each vertex of P_{int} ;
- (b) one linear equation for each lattice point in the relative interior of an edge of P_{int} ;
- (c) two linear equations for each lattice point in the interior of P_{int} .

These inequalities and equations can be seen as follows. Let $\ell_1, \ell_2, \ell_3, \ell_4, \ell_5, \ell_6$ denote the lengths of the edges (labeled cyclically) of a hexagon in a honeycomb curve. Then

$$\ell_1 + \ell_2 = \ell_4 + \ell_5, \quad \ell_2 + \ell_3 = \ell_5 + \ell_6, \quad \text{and} \quad \ell_3 + \ell_4 = \ell_6 + \ell_1.$$

These three equations are linearly dependent, and they give rise to the inequalities in (a) and to the equations in (b) and (c). The linear equations (b) and (c), when taken over all

hexagons, have a triangular structure. These linear equations are thus linearly independent. This implies that the codimension of the cone \mathbb{M}_Δ inside the orthant $\mathbb{R}_{\geq 0}^{3g-3}$ equals

$$\text{codim}(\mathbb{M}_\Delta) = (\#(\partial P_{\text{int}} \cap \mathbb{Z}^2) - \#\text{vertices}(P_{\text{int}})) + 2 \cdot \#(\text{int}(P_{\text{int}}) \cap \mathbb{Z}^2). \quad (3.2.2)$$

This expression can be rewritten as

$$g + \#(\text{int}(P_{\text{int}}) \cap \mathbb{Z}^2) - \#\text{vertices}(P_{\text{int}}) = 2g - \#(\partial P_{\text{int}} \cap \mathbb{Z}^2) - \#\text{vertices}(P_{\text{int}}).$$

Subtracting this codimension from $3g - 3$, we obtain the desired formula. \square

Proof of Theorem 3.1.1. For the classical moduli space $\mathcal{M}_g^{\text{nd}}$, formula (3.1.7) was proved in [CV09]. That dimension is preserved under tropicalization. The inclusion $\mathbb{M}_g^{\text{planar}} \subseteq \text{trop}(\mathcal{M}_g^{\text{nd}})$ in (3.1.6) shows that the right-hand side in (3.1.7) is an upper bound on the dimension of $\mathbb{M}_g^{\text{planar}}$.

To prove the lower bound, we choose P to be a specific honeycomb polygon with honeycomb triangulation Δ . The form of the polygon depends on the parity of the genus g . If $g = 2h$ is even then we take the rectangle

$$R_{3,h+1} = H_{3,h+1,0,h+4} = \text{conv}\{(0,0), (0,h+1), (3,0), (3,h+1)\}. \quad (3.2.3)$$

The interior hull of $R_{3,h+1}$ is the rectangle

$$(R_{3,h+1})_{\text{int}} = \text{conv}\{(1,1), (1,h), (2,1), (2,h)\} \cong R_{1,h-1}.$$

All $g = 2h$ lattice points of this polygon lie on the boundary. From Lemma 3.2.1, we see that $\dim(\mathbb{M}_\Delta) = g + g + 4 - 3 = 2g + 1$. If $g = 2h + 1$ is odd then we take the trapezoid

$$H_{3,h+3,0,h+3} = \text{conv}\{(0,0), (0,h+3), (3,0), (3,h)\}. \quad (3.2.4)$$

The convex hull of the interior lattice points in $H_{3,h+3,0,h+3}$ is the trapezoid

$$(H_{3,h+3,0,h+3})_{\text{int}} = \text{conv}\{(1,1), (1,h+1), (2,1), (2,h)\}.$$

All $g = 2h + 1$ lattice points of this polygon lie on its boundary, and again $\dim(\mathbb{M}_\Delta) = 2g + 1$.

For all $g \geq 4$ with $g \neq 7$, this matches the upper bound obtained from [CV09]. We conclude that $\dim(\mathbb{M}_P) = \dim(\mathbb{M}_g) = 2g + 1$ holds in all of these cases. For $g = 7$ we take $P = H_{4,4,2,6}$. Then P_{int} is a hexagon with $g = 7$ lattice points. From Lemma 3.2.1, we find $\dim(\mathbb{M}_\Delta) = 7 + 6 + 6 - 3 = 16$, so this matches the upper bound. Finally, for $g = 3$, we will see $\dim(\mathbb{M}_{T_4}) = 6$ in Section 3.4, and also follows from Proposition 3.2.2. The case $g = 2$ follows from the discussion in Example 3.3.4. \square

The honeycomb polygons and triangulations used in the above proof are illustrated in Figure 3.2. Shown in order are polygons for genus 2, 3, 7, 5, and 6, with the polygons for genus 5 and 6 generalizing to higher genus.




Figure 3.2: Honeycomb triangulations from the proof of Theorem 3.1.1

We close this section by considering two special families of honeycomb curves: those arising from the triangles T_d for $d \geq 4$ and rectangles $R_{d,e}$ for $d, e \geq 3$. The triangle T_d corresponds to curves of degree d in the projective plane \mathbb{P}^2 . Their genus is $g = (d - 1)(d - 2)/2$. The case $d = 4, g = 3$ of smooth plane quartics will be the first topic in Section 3.4. The rectangle $R_{d,e}$ corresponds to curves of bidegree (d, e) in $\mathbb{P}^1 \times \mathbb{P}^1$. Their genus is $g = (d - 1)(e - 1)$. The case $d = e = 3$ appears in our computation of $\mathbb{M}_4^{\text{planar}}$ in Section 3.4.

Proposition 3.2.2. *Let P be the triangle T_d with $d \geq 4$ or the rectangle $R_{d,e}$ with $d, e \geq 3$. The moduli space \mathbb{M}_P of tropical plane curves has the expected dimension inside \mathbb{M}_g , namely,*

$$\dim(\mathbb{M}_{T_d}) = \frac{1}{2}d^2 + \frac{3}{2}d - 8 \quad \text{and} \quad \text{codim}(\mathbb{M}_{T_d}) = (d - 2)(d - 4), \quad \text{whereas}$$

$$\dim(\mathbb{M}_{R_{d,e}}) = de + d + e - 6 \quad \text{and} \quad \text{codim}(\mathbb{M}_{R_{d,e}}) = 2(de - 2d - 2e + 3).$$

In particular, the honeycomb triangulation defines a cone \mathbb{M}_Δ of this maximal dimension.

Proof. For our standard triangles and rectangles, the formula (3.2.2) implies

$$\begin{aligned} \text{codim}(\mathbb{M}_{T_d}) &= 3(d - 3) - 3 + 2 \cdot \frac{1}{2}(d - 4)(d - 5), \\ \text{codim}(\mathbb{M}_{R_{d,e}}) &= 2((d - 2) + (e - 2)) - 4 + 2 \cdot (d - 3)(e - 3). \end{aligned}$$

Subtracting from $3g - 3 = \dim(\mathbb{M}_g)$, we get the desired formulas for $\dim(\mathbb{M}_P)$. □

The above dimensions are those expected from algebraic geometry. Plane curves with Newton polygon T_d form a projective space of dimension $\frac{1}{2}(d + 2)(d + 1) - 1$ on which the 8-dimensional group $\text{PGL}(3)$ acts effectively, while those with $R_{d,e}$ form a space of dimension $(d + 1)(e + 1) - 1$ on which the 6-dimensional group $\text{PGL}(2)^2$ acts effectively. In each case, $\dim(\mathcal{M}_P)$ equals the dimension of the family of all curves minus the dimension of the group.

3.3 Computational methods

In this section we present algorithms and initial examples for computing the space $\mathbb{M}_g^{\text{planar}}$. For additional background on geometric combinatorics, the reader is referred to the book by De Loera, Rambau, and Santos [DLRS10].

Let P be a lattice polygon, and let $A = P \cap \mathbb{Z}^2$ be the set of lattice points in P . Recall from Section 1.2 that height functions $h : A \rightarrow \mathbb{R}$ are identified with tropical polynomials,

and that the tropical plane curve defined by such a polynomial is dual to the subdivision Δ of P induced by h . The *secondary cone* $\Sigma(\Delta)$ of Δ is the closure of the relatively open polyhedral cone in \mathbb{R}^A consisting of the set of all height functions inducing Δ , and the *secondary fan* of A is the complete polyhedral fan in \mathbb{R}^A consisting of all secondary cones $\Sigma(\Delta)$.

We now derive an inequality representation for the secondary cone $\Sigma(\Delta)$ as follows. Consider any four points $a = (a_1, a_2)$, $b = (b_1, b_2)$, $c = (c_1, c_2)$ and $d = (d_1, d_2)$ in A such that the triples (c, b, a) and (b, c, d) are clockwise-oriented triangles of Δ . Then we require

$$\det \begin{pmatrix} 1 & 1 & 1 & 1 \\ a_1 & b_1 & c_1 & d_1 \\ a_2 & b_2 & c_2 & d_2 \\ h(a) & h(b) & h(c) & h(d) \end{pmatrix} \geq 0. \tag{3.3.1}$$

This is a linear inequality for $h \in \mathbb{R}^A$. It can be viewed as a “flip condition”, determining which of the two diagonals of a quadrilateral are in the subdivision. We have one such inequality for each interior edge bc of Δ . The set of solutions to these linear inequalities is the secondary cone $\Sigma(\Delta)$. From this it follows that the lineality space $\Sigma(\Delta) \cap -\Sigma(\Delta)$ of the secondary cone is 3-dimensional. It is the space $\text{Lin}(A)$ of functions $h \in \mathbb{R}^A$ that are restrictions of affine-linear functions on \mathbb{R}^2 . We usually identify $\Sigma(A)$ with its image in $\mathbb{R}^A/\text{Lin}(A)$, which is a pointed cone of dimension $\#A - 3$. That pointed cone has finitely many rays and we represent these by vectors in \mathbb{R}^A .

Suppose that Δ has E interior edges and g interior vertices. We consider two linear maps

$$\mathbb{R}^A \xrightarrow{\lambda} \mathbb{R}^E \xrightarrow{\kappa} \mathbb{R}^{3g-3}. \tag{3.3.2}$$

The map λ takes h and outputs the vector whose bc -coordinate equals (3.3.1). This determinant is nonnegative: it is precisely the (lattice) length of the edge of the tropical curve C that is dual to bc . Hence $\lambda(h)$ is the vector whose $\#E$ coordinates are the lengths of the bounded edges of C .

Each edge e of the skeleton G is a concatenation of edges of C . The second map κ adds up the corresponding lengths. Thus the composition (3.3.2) is the linear map with e^{th} coordinate

$$(\kappa \circ \lambda)(h)_e = \sum_{\substack{bc: \text{the dual of } bc \\ \text{contributes to } e}} \lambda(h)_{bc} \quad \text{for all edges } e \text{ of } G.$$

By definition, the secondary cone is mapped into the nonnegative orthant under λ . Hence

$$\Sigma(\Delta) \xrightarrow{\lambda} \mathbb{R}_{\geq 0}^E \xrightarrow{\kappa} \mathbb{R}_{\geq 0}^{3g-3}. \tag{3.3.3}$$

Note that $\kappa(\lambda(h))$ is the vector whose $\#E$ coordinates are the lengths of the edges of G . Our discussion implies the following result on the cone of metric graphs arising from Δ :

Proposition 3.3.1. *The cone \mathbb{M}_Δ is the image of the secondary cone $\Sigma(\Delta)$ under $\kappa \circ \lambda$.*

Given any lattice polygon P , we seek to compute the moduli space \mathbb{M}_P via the decompositions in (3.1.4). Our line of attack towards that goal can now be summarized as follows:

1. compute all regular unimodular triangulations of $A = P \cap \mathbb{Z}^2$ up to symmetry;
2. sort the triangulations into buckets, one for each trivalent graph G of genus g ;
3. for each triangulation Δ with skeleton G , compute its secondary cone $\Sigma(\Delta) \subset \mathbb{R}^A$;
4. for each secondary cone $\Sigma(\Delta)$, compute its image \mathbb{M}_Δ in the moduli space \mathbb{M}_g via (3.3.3);
5. merge the results to get the fans $\mathbb{M}_{P,G} \subset \mathbb{R}^{3g-3}$ in (3.1.3) and the moduli space \mathbb{M}_P in (3.1.4).

Step 1 is based on computing the secondary fan of A . There are two different approaches to doing this. The first, more direct, method is implemented in `Gfan` [Jen]. It starts out with one regular triangulation of Δ , e.g. a placing triangulation arising from a fixed ordering of A . This comes with an inequality description for $\Sigma(\Delta)$, as in (3.3.1). From this, `Gfan` computes the rays and the facets of $\Sigma(\Delta)$. Then `Gfan` proceeds to an adjacent secondary cone $\Sigma(\Delta')$ by producing a new height function from traversing a facet of $\Sigma(\Delta)$. Iterating this process results in a breadth-first-search through the edge graph of the *secondary polytope* of A .

The second method starts out the same. But it passes from Δ to a neighboring triangulation Δ' that need not be regular. It simply performs a purely combinatorial restructuring known as a *bistellar flip*. The resulting breadth-first search is implemented in `TOPCOM` [Ram02].

Neither algorithm is generally superior to the other, and sometimes it is difficult to predict which one will perform better. The flip-algorithm may suffer from wasting time by also computing non-regular triangulations, while the polyhedral algorithm is genuinely costly since it employs exact rational arithmetic. The flip-algorithm also uses exact coordinates, but only in a preprocessing step which encodes the point configuration as an oriented matroid. Both algorithms can be modified to enumerate all regular *unimodular* triangulations up to symmetry only. For our particular planar instances, we found `TOPCOM` to be more powerful.

We start Step 2 by computing the dual graph of a given Δ . The nodes are the triangles and the edges record incidence. Hence each node has degree 1, 2 or 3. We then recursively delete the nodes of degree 1. Next, we recursively contract edges which are incident with a node of degree 2. The resulting trivalent graph G is the skeleton of Δ . It often has loops and multiple edges. In this process we keep track of the history of all deletions and contractions.

Steps 3 and 4 are carried out using `polymake` [GJ00]. Here the buckets or even the individual triangulations can be treated in parallel. The secondary cone $\Sigma(\Delta)$ is defined in \mathbb{R}^A by the linear inequalities $\lambda(h) \geq 0$ in (3.3.1). From this we compute the facets and rays of $\Sigma(\Delta)$. This is essentially a convex hull computation. In order to get unique rays modulo $\text{Lin}(A)$, we fix $h = 0$ on the three vertices of one particular triangle. Since the cones are rather small, the choice of the convex hull algorithm does not matter much. For

details on state-of-the-art convex hull computations and an up-to-date description of the `polymake` system see [AGH⁺14].

For Step 4, we apply the linear map $\kappa \circ \lambda$ to all rays of the secondary cone $\Sigma(\Delta)$. Their images are vectors in \mathbb{R}^{3g-3} that span the moduli cone $\mathbb{M}_\Delta = (\kappa \circ \lambda)(\Sigma(\Delta))$. Via a convex hull computation as above, we compute all the rays and facets of \mathbb{M}_Δ .

The cones \mathbb{M}_Δ are generally not full-dimensional in \mathbb{R}^{3g-3} . The points in the relative interior are images of interior points of $\Sigma(\Delta)$. Only these represent smooth tropical curves. However, it can happen that another cone $\mathbb{M}_{\Delta'}$ is a face of \mathbb{M}_Δ . In that case, the metric graphs in the relative interior of that face are also realizable by smooth tropical curves.

Step 5 has not been fully automatized yet, but we carry it out in a case-by-case manner. This will be described in detail for curves of genus 3 in Sections 3.4 and 3.6.

We now come to the question of what lattice polygons P should be the input for Step 1. Our point of departure towards answering that question is the following finiteness result.

Proposition 3.3.2. *For every fixed genus $g \geq 1$, there are only finitely many lattice polygons P with g interior lattice points, up to integer affine isomorphisms in \mathbb{Z}^2 .*

Proof and Discussion. Scott [Sco76] proved that $\#(\partial P \cap \mathbb{Z}^2) \leq 2g + 7$, and this bound is sharp. This means that the number of interior lattice points yields a bound on the total number of lattice points in P . This result was generalized to arbitrary dimensions by Hensley [Hen83]. Lagarias and Ziegler [LZ91] improved Hensley's bound and further observed that there are only finitely many lattice polytopes with a given total number of lattice points, up to unimodular equivalence [LZ91, Theorem 2]. Castryck [Cas12] gave an algorithm for finding all lattice polygons of a given genus, along with the number of lattice polygons for each genus up to 30. We remark that the assumption $g \geq 1$ is essential, as there are lattice triangles of arbitrarily large area and without any interior lattice point. \square

Proposition 3.3.2 ensures that the union in (3.1.5) is finite. However, from the full list of polygons P with g interior lattice points, only very few will be needed to construct $\mathbb{M}_g^{\text{planar}}$. To show this, and to illustrate the concepts seen so far, we now discuss our spaces for $g \leq 2$.

Example 3.3.3. For $g = 1$, only one polygon P is needed in (3.1.5), and only one triangulation Δ is needed in (3.1.4). We take $P = \text{conv}\{(0, 0), (0, 3), (3, 0)\}$, since every smooth genus 1 curve is a plane cubic, and we let Δ be the honeycomb triangulation from Section 3.2. The skeleton G is a cycle whose length is the tropical j -invariant [BPR12, section 7.1]. We can summarize this as follows:

$$\mathbb{M}_\Delta = \mathbb{M}_{P,G} = \mathbb{M}_P = \mathbb{M}_1^{\text{planar}} = \mathbb{M}_1 = \mathbb{R}_{\geq 0}. \tag{3.3.4}$$

All inclusions in (3.1.6) are equalities for this particular choice of (P, Δ) .

Example 3.3.4. In classical algebraic geometry, all smooth curves of genus $g = 2$ are hyperelliptic, and they can be realized with the Newton polygon $P = \text{conv}\{(0, 0), (0, 2), (6, 0)\}$. There are two trivalent graphs of genus 2, namely, the *theta graph* $G_1 = \Theta$ and the *dumbbell*




Figure 3.3: The triangulations Δ_1 , Δ'_1 , and Δ_2

graph $G_2 = \circ-\circ$. The moduli space \mathbb{M}_2 consists of two quotients of the orthant $\mathbb{R}_{\geq 0}^3$, one for each graph, glued together. For nice drawings see Figures 3 and 4 in [Cha12]. Figure 3.3 shows three unimodular triangulations Δ_1 , Δ'_1 , and Δ_2 of P such that almost all metric graphs in \mathbb{M}_2 are realized by a smooth tropical curve C dual to Δ_1 , Δ'_1 , or Δ_2 . We say “almost all” because here the three edges of G_1 cannot have all the same length [CDMY14, Proposition 4.7]. The triangulations Δ_1 and Δ'_1 both give G_1 as a skeleton. If $a \geq b \geq c$ denote the edge lengths on G_1 , then the curves dual to Δ_1 realize all metrics with $a \geq b > c$, and the curves dual to Δ'_1 realize all metrics with $a > b = c$. The triangulation Δ_2 gives G_2 as a skeleton, and the curves dual to it achieve all possible metrics. Since our 3-dimensional cones are closed by definition,

$$(\mathbb{M}_{\Delta_1} \cup \mathbb{M}_{\Delta'_1}) \cup \mathbb{M}_{\Delta_2} = \mathbb{M}_{P,G_1} \cup \mathbb{M}_{P,G_2} = \mathbb{M}_P = \mathbb{M}_2^{\text{planar}} = \mathbb{M}_2 = [\text{Cha12, Figure 3}]. \tag{3.3.5}$$

In Section 3.6 we extend this analysis to hyperelliptic curves of $g \geq 3$. The graphs G_1 and G_2 represent the *chains* for $g = 2$. For more information on hyperelliptic metric graphs see [Cha13].

With $g = 1, 2$ out of the way, we now assume $g \geq 3$. We follow the approach of Castryck and Voight [CV09] in constructing polygons P that suffice for the union (3.1.5).

Lemma 3.3.5. *Let $P \subseteq Q$ be lattice polygons with $P_{\text{int}} = Q_{\text{int}}$. Then \mathbb{M}_P is contained in \mathbb{M}_Q .*

Proof. Every regular unimodular triangulation Δ of P can be extended to a regular unimodular triangulation Δ' of Q . (This is a special property of planar triangulations: it does not hold in higher dimensions.) This means that every tropical curve C dual to Δ is contained in a curve C' dual to Δ' , except for unbounded edges of C . The skeleton and its possible metrics remain unchanged, since $P_{\text{int}} = Q_{\text{int}}$. We conclude that $\mathbb{M}_\Delta = \mathbb{M}_{\Delta'}$. The unions for P and Q in (3.1.4) show that $\mathbb{M}_P \subseteq \mathbb{M}_Q$. \square

This lemma shows that we only need to consider *maximal polygons*, i.e. those P that are maximal with respect to inclusion for fixed P_{int} . If P_{int} is 2-dimensional then this determines P uniquely. Namely, suppose that $P_{\text{int}} = \{(x, y) \in \mathbb{R}^2 : a_i x + b_i y \leq c_i \text{ for } i = 1, 2, \dots, s\}$, where $\gcd(a_i, b_i, c_i) = 1$ for all i . Let P be the polygon $\{(x, y) \in \mathbb{R}^2 : a_i x + b_i y \leq c_i + 1 \text{ for } i = 1, 2, \dots, s\}$. If P is a lattice polygon then it is a maximal lattice polygon. However, it can happen that P has non-integral vertices. In that case, the given P_{int} is not the interior of any lattice polygon.

The maximal polygon P is not uniquely determined by P_{int} when P_{int} is a line segment. For each $g \geq 2$ there are $g + 2$ distinct *hyperelliptic trapezoids* to be considered. We shall see in Theorem 3.6.1 that for our computational purposes it suffices to use the triangle $\text{conv}\{(0, 0), (0, 2), (2g + 2, 0)\}$.

Here is the list of all maximal nonhyperelliptic polygons we use as input for the pipeline described above.

Proposition 3.3.6. *Up to isomorphism there are precisely 12 maximal polygons P such that P_{int} is 2-dimensional and $3 \leq g = \#(P_{\text{int}} \cap \mathbb{Z}^2) \leq 6$. For $g = 3$, there is a unique type, namely, $T_4 = \text{conv}\{(0, 0), (0, 4), (4, 0)\}$. For $g = 4$ there are three types:*

$$\begin{aligned} Q_1^{(4)} = R_{3,3} &= \text{conv}\{(0, 0), (0, 3), (3, 0), (3, 3)\}, & Q_2^{(4)} &= \text{conv}\{(0, 0), (0, 3), (6, 0)\}, \\ Q_3^{(4)} &= \text{conv}\{(0, 2), (2, 4), (4, 0)\}. \end{aligned}$$

For $g = 5$ there are four types of maximal polygons:

$$\begin{aligned} Q_1^{(5)} &= \text{conv}\{(0, 0), (0, 4), (4, 2)\}, & Q_2^{(5)} &= \text{conv}\{(2, 0), (5, 0), (0, 5), (0, 2)\}, \\ Q_3^{(5)} &= \text{conv}\{(2, 0), (4, 2), (2, 4), (0, 2)\}, & Q_4^{(5)} &= \text{conv}\{(0, 0), (0, 2), (2, 0), (4, 4)\}. \end{aligned}$$

For $g = 6$ there are four types of maximal polygons:

$$\begin{aligned} Q_1^{(6)} = T_5 &= \text{conv}\{(0, 0), (0, 5), (5, 0)\}, & Q_2^{(6)} &= \text{conv}\{(0, 0), (0, 7), (3, 0), (3, 1)\}, \\ Q_3^{(6)} = R_{3,4} &= \text{conv}\{(0, 0), (0, 4), (3, 0), (3, 4)\}, & Q_4^{(6)} &= \text{conv}\{(0, 0), (0, 4), (2, 0), (4, 2)\}. \end{aligned}$$

The notation we use for polygons is as follows. We write $Q_i^{(g)}$ for maximal polygons of genus g , but we also use a systematic notation for families of polygons, including the *triangles* $T_d = \text{conv}\{(0, 0), (0, d), (d, 0)\}$ and the *rectangles* $R_{d,e} = \text{conv}\{(0, 0), (d, 0), (0, e), (d, e)\}$.

Proposition 3.3.6 is found by exhaustive search, using Castryck's method in [Cas12]. We started by classifying all types of lattice polygons with precisely g lattice points. These are our candidates for P_{int} . For instance, for $g = 5$, there are six such polygons. Four of them are the interior hulls of the polygons $Q_i^{(5)}$ with $i = 1, 2, 3, 4$. The other two are the triangles

$$\text{conv}\{(1, 1), (1, 4), (2, 1)\} \quad \text{and} \quad \text{conv}\{(1, 1), (2, 4), (3, 2)\}.$$

However, neither of these two triangles arises as P_{int} for any lattice polygon P .

For each genus g , we construct the stacky fans $\mathbb{M}_g^{\text{planar}}$ by computing each of the spaces $\mathbb{M}_{Q_i^{(g)}}$ and then subdividing their union appropriately. This is then augmented in Section 3.6 by the spaces \mathbb{M}_P where P_{int} is not two-dimensional, but is instead a line segment.

3.4 Computations of $\mathbb{M}_g^{\text{planar}}$

In this section we explicitly compute the nonhyperelliptic part of $\mathbb{M}_3^{\text{planar}}$, and perform similar computations for genus 4 and 5.

In classical algebraic geometry, all nonhyperelliptic smooth curves of genus 3 are plane quartics. Their Newton polygon $T_4 = \text{conv}\{(0, 0), (0, 4), (4, 0)\}$ is the unique maximal non-hyperelliptic polygon with $g = 3$ in Proposition 3.3.6. To compute $\mathbb{M}_3^{\text{planar}}$, we must compute \mathbb{M}_{T_4} , and along the way we will characterize the dense subset of metric graphs that are realized as the skeleton of a smooth tropical quartic. The full moduli space for tropical plane curves of genus 3 is then obtained as

$$\mathbb{M}_3^{\text{planar}} = \mathbb{M}_{T_4} \cup \mathbb{M}_{3,\text{hyp}}^{\text{planar}}. \tag{3.4.1}$$

The space $\mathbb{M}_{3,\text{hyp}}^{\text{planar}}$ is computed in Section 3.6. As expected from classical algebraic geometry, $\dim(\mathbb{M}_{T_4}) = 6$ and $\dim(\mathbb{M}_{3,\text{hyp}}^{\text{planar}}) = 5$.




Figure 3.4: The five trivalent graphs of genus 3, with letters labeling each graph’s six edges

The stacky fan \mathbb{M}_3 of *all* metric graphs of genus 3 has five maximal cones, as shown in [Cha12, Figure 4]. These correspond to the five (leafless) trivalent graphs of genus 3, pictured in Figure 3.4. As in Section 2.3, each graph is labeled by the triple (ℓbc) , where ℓ is the number of loops, b is the number of bi-edges, and c is the number of cut-edges. Here, ℓ , b , and c are single digit numbers, so there is no ambiguity to this notation. Our labeling and ordering is largely consistent with [Bal76].

Although \mathbb{M}_{T_4} has dimension 6, it is not pure due to the realizable metrics on (111). It also misses one of the five cones in \mathbb{M}_3 : the graph (303) cannot be realized in \mathbb{R}^2 by Proposition 3.5.4. The restriction of \mathbb{M}_{T_4} to each of the other four cones is given by a finite union of convex polyhedral subcones, characterized by the following piecewise-linear formulas:

Theorem 3.4.1. *A graph in \mathbb{M}_3 arises from a smooth tropical quartic if and only if it is one of the first four graphs in Figure 3.4, with edge lengths satisfying the following, up to symmetry:*

- (000) is realizable if and only if $\max\{x, y\} \leq u$, $\max\{x, z\} \leq v$ and $\max\{y, z\} \leq w$, where
 - at most two of the inequalities can be equalities, and
 - if two are equalities, then either x, y, z are distinct and the edge (among u, v, w) that connects the shortest two of x, y, z attains equality, or $\max\{x, y, z\}$ is attained exactly twice, and the edge connecting those two longest does not attain equality.

- (020) is realizable if and only if $v \leq u$, $y \leq z$, and $w + \max\{v, y\} \leq x$, and if the last inequality is an equality, then: $v = u$ implies $v < y < z$, and $y = z$ implies $y < v < u$.
- (111) is realizable if and only if $w < x$ and

$$(v + w = x \text{ and } v < u) \text{ or } (v + w < x \leq v + 3w \text{ and } v \leq u) \text{ or}$$

$$(v + 3w < x \leq v + 4w \text{ and } v \leq u \leq 3v/2) \text{ or} \tag{3.4.2}$$

$$(v + 3w < x \leq v + 4w \text{ and } 2v = u) \text{ or } (v + 4w < x \leq v + 5w \text{ and } v = u).$$
- (212) is realizable if and only if $w < x \leq 2w$.

To understand the qualifier “up to symmetry” in Theorem 3.4.1, it is worthwhile to read off the automorphisms from the graphs in Figure 3.4. The graph (000) is the complete graph on four vertices. Its automorphism group is the symmetric group of order 24. The automorphism group of the graph (020) is generated by the three transpositions $(u v)$, $(y z)$, $(w x)$ and the double transposition $(u y)(v z)$. Its order is 16. The automorphism group of the graph (111) has order 4, and it is generated by $(u v)$ and $(w x)$. The automorphism group of the graph (212) is generated by $(u z)(v y)$ and $(w x)$, and has order 4. The automorphism group of the graph (303) is the symmetric group of order 6. Each of the five graphs contributes an orthant $\mathbb{R}_{\geq 0}^6$ modulo the action of that symmetry group to the stacky fan \mathbb{M}_3 .

Table 3.1: Dimensions of the 1278 moduli cones \mathbb{M}_Δ within \mathbb{M}_{T_4}

$G \setminus \dim$	3	4	5	6	$\#\Delta$'s
(000)	18	142	269	144	573
(020)		59	216	175	450
(111)		10	120	95	225
(212)			15	15	30
total	18	211	620	429	1278

Proof of Theorem 3.4.1. This is based on explicit computations as described in Section 3.3. The symmetric group S_3 acts on the triangle T_4 . We enumerated all unimodular triangulations of T_4 up to that symmetry. There are 1279 (classes of) such triangulations, and of these precisely 1278 are regular. The unique non-regular triangulation is a refinement of Figure 1.3. For each regular triangulation we computed the graph G and the polyhedral cone \mathbb{M}_Δ . Each \mathbb{M}_Δ is the image of the 12-dimensional secondary cone of Δ . We found that \mathbb{M}_Δ has dimension 3, 4, 5 or 6, depending on the structure of the triangulation Δ . A census is given by Table 3.1. For instance, 450 of the 1278 triangulations Δ have the skeleton $G = (020)$. Among these 450, we found that 59 have $\dim(\mathbb{M}_\Delta) = 4$, 216 have $\dim(\mathbb{M}_\Delta) = 5$, and 175 have $\dim(\mathbb{M}_\Delta) = 6$.

For each of the 1278 regular triangulations Δ we checked that the inequalities stated in Theorem 3.4.1 are valid on the cone $\mathbb{M}_\Delta = (\kappa \circ \lambda)(\Sigma(\Delta))$. This proves that the dense realizable part of \mathbb{M}_{T_4} is contained in the polyhedral space described by our constraints.

For the converse direction, we need to go through the four cases and construct a planar tropical realization of each metric graph that satisfies our constraints. We shall now do this.




Figure 3.5: A triangulation that realizes almost all realizable graphs of type (000)

All realizable graphs of **type (000)**, except for lower-dimensional families, arise from a single triangulation Δ , shown in Figure 3.5 with its skeleton. The cone \mathbb{M}_Δ is six-dimensional. Its interior is defined by $x < \min\{u, v\}$, $y < \min\{u, w\}$, and $z < \min\{v, w\}$. Indeed, the parallel segments in the outer edges can be arbitrarily long, and each outer edge be as close as desired to the maximum of the two adjacent inner edges. This is accomplished by putting as much length as possible into a particular edge and pulling extraneous parts back.




Figure 3.6: Triangulations giving all metrics in the cases (i) through (v) for the graph (000)

There are several lower dimensional collections of graphs we must show are achievable:

- (i) $y < x = u$, $\max\{x, z\} < v$, $\max\{y, z\} < w$; (dim = 5)
- (ii) $y = x = u$, $\max\{x, z\} < v$, $\max\{y, z\} < w$; (dim = 4)
- (iii) $z < y < x < v$, $u = x$, $w = y$; (dim = 4)
- (iv) $z < y < x < u$, $v = x$, $w = y$; (dim = 4)
- (v) $z < y = x = v = w < u$. (dim = 3)

In Figure 3.6 we show triangulations realizing these five special families. Dual edges are labeled

$$(1, 1) \overset{x}{-} (1, 2) \overset{y}{-} (2, 1) \overset{z}{-} (1, 1).$$




Figure 3.7: A triangulation that realizes almost all realizable graphs of type (020)

Next, we consider **type (020)**. Again, except for some lower-dimensional cases, all graphs arise from single triangulation, pictured in Figure 3.7. The interior of \mathbb{M}_Δ is given by $v < u$, $y < z$, and $w + \max\{v, y\} < x$. There are several remaining boundary cases, all of whose graphs are realized by the triangulations in Figure 3.8:

- (i) $v < u, y < z, w + \max\{v, y\} = x$; (dim = 5)
- (ii) $u = v, y < z, w + \max\{v, y\} < x$; (dim = 5)
- (iii) $u = v, y = z, w + \max\{v, y\} < x$; (dim = 4)
- (iv) $u = v, v < y < z, w + \max\{v, y\} = x$. (dim = 4)




Figure 3.8: Triangulations giving all metrics in the cases (i) through (iv) for the graph (020)

Type (111) is the most complicated. We begin by realizing the metric graphs that lie in $\text{int}(\mathbb{M}_{T_4, (111)})$. These arise from the second and third cases in the disjunction (3.4.2).

We assume $w < x$. The triangulation to the left in Figure 3.9 realizes all metrics on (111) satisfying $v + w < x < v + 3w$ and $v < u$. The dilation freedom of u, y , and z is clear. To see that the edge x can have length arbitrarily close to $v + 3w$, simply dilate the double-headed segment to be as long as possible, with some very small length given to the next two segments counterclockwise. Shrinking the double-headed segment as well as the vertical segment of x brings the length close to $v + w$. The triangulation to the right in Figure 3.9 realizes all metrics satisfying $v + 3w < x < v + 4w$ and $v < u < 3v/2$. Dilation of x is more free due to the double-headed segment of slope 1/2, while dilation of u is more restricted.

Many triangulations are needed in order to deal with low-dimensional case. In Figure 3.10 we show triangulations that realize each of the following families of type (111) graphs:

- (i) $v + w < x < v + 5w, v = u$; (dim = 5)
- (ii) $v + w < x < v + 4w, 2v = u$; (dim = 5)




Figure 3.9: Triangulations of type (111) realizing $v + w < x < v + 2x$ and $v < u$ (on the left) and $v + 3w < x < v + 4w$ and $v < u < 3v/2$ (on the right)

- (iii) $v + w = x, v < u;$ (dim = 5)
- (iv) $x = v + 3w, v < u;$ (dim = 5)
- (v) $x = v + 4w, v < u \leq 3v/2;$ (dim = 5)
- (vi) $x = v + 5w, v = u;$ (dim = 4)
- (vii) $x = v + 4w, 2v = u.$ (dim = 4)




Figure 3.10: Triangulations of type (111) that realize the boundary cases (i) through (vii)

All graphs of **type (212)** can be achieved with the two triangulations in Figure 3.11. The left gives all possibilities with $w < x < 2w$, and the right realizes $x = 2w$. The edges u, v, y, z are completely free to dilate. This completes the proof of Theorem 3.4.1. \square




Figure 3.11: Triangulations giving graphs of type (212) giving $w < x < 2w$ and $x = 2w$

The space \mathbb{M}_{T_4} is not pure-dimensional because of the graphs (111) with $u = v$ and $v + 4w < x < v + 5w$. These appear in the five-dimensional \mathbb{M}_Δ where Δ is the leftmost triangulation in Figure 3.10, but \mathbb{M}_Δ is not contained in the boundary of any six-dimensional $\mathbb{M}_{\Delta'}$.

We close this section by suggesting an answer to the following question: *What is the probability that a random metric graph of genus 3 can be realized by a tropical plane quartic?*

To examine this question, we need to endow the moduli space \mathbb{M}_3 with a probability measure. Here we fix this measure as follows. We assume that the five trivalent graphs G are equally likely, and all non-trivalent graphs have probability 0. The lengths on each trivalent graph G specify an orthant $\mathbb{R}_{\geq 0}^6$. We fix a probability measure on $\mathbb{R}_{\geq 0}^6$ by normalizing so that $u + v + w + x + y + z = 1$, and we take the uniform distribution on the resulting 5-simplex. With this probability measure on the moduli space \mathbb{M}_3 we are asking for the ratio of volumes

$$\text{vol}(\mathbb{M}_3^{\text{planar}})/\text{vol}(\mathbb{M}_3). \tag{3.4.3}$$

This ratio is a rational number, which we computed from our data in Theorem 3.4.1.

Corollary 3.4.2. *The rational number in (3.4.3) is 31/105. This means that, in the measure specified above, about 29.5% of all metric graphs of genus 3 come from tropical plane quartics.*

Proof and Explanation. The graph (303) is not realizable, since none of the 1278 regular unimodular triangulations of the triangle T_4 has gives rise to it. So, its probability is zero. For the other four trivalent graphs in Figure 3.4 we compute the volume of the realizable edge lengths, using the inequalities in Theorem 3.4.1. The result of our computations is the table

Graph	(000)	(020)	(111)	(212)	(303)
Probability	4/15	8/15	12/35	1/3	0

A non-trivial point in verifying these numbers is that Theorem 3.4.1 gives the constraints only up to symmetry. We must apply the automorphism group of each graph in order to obtain the realizable region in its 5-simplex $\{(u, v, w, x, y, z) \in \mathbb{R}_{\geq 0}^6 : u+v+w+x+y+z = 1\}$. Since we are measuring volumes, we are here allowed to replace the regions described in Theorem 3.4.1 by their closures. For instance, consider type (020). After taking the closure, and after applying the automorphism group of order 16, the realizability condition becomes

$$\max(\min(u, v), \min(y, z)) \leq |x - w|. \tag{3.4.4}$$

The probability that a uniformly sampled random point in the 5-simplex satisfies (3.4.4) is equal to 8/15. The desired probability (3.4.3) is the average of the five numbers in the table. □

Notice that asking for those probabilities only makes sense since the dimension of the moduli space agrees with the number of skeleton edges. In view of (3.1.7) this occurs for the three genera $g = 2, 3, 4$. For $g \geq 5$ the number of skeleton edges exceeds the dimension of the moduli space. Hence, in these cases, the probability that a random metric graph can be realized by a tropical plane curve must be 0. For $g = 2$ that probability is one; see Example 3.3.4. For $g = 4$ that probability is less than 0.5% by Corollary 3.4.4 below.

We will now compute the moduli space of tropical plane curves of genus 4. This is

$$\mathbb{M}_4^{\text{planar}} = \mathbb{M}_{Q_1^{(4)}} \cup \mathbb{M}_{Q_2^{(4)}} \cup \mathbb{M}_{Q_3^{(4)}} \cup \mathbb{M}_{4,\text{hyp}}^{\text{planar}},$$

where $Q_i^{(4)}$ are the three genus 4 polygons in Proposition 3.3.6. They are shown in Figure 3.12.




Figure 3.12: The three nonhyperelliptic genus 4 polygons, and a triangulation

There are 17 trivalent genus 4 graphs, of which 16 are planar. These were first enumerated in [Bal76], and are shown in Figure 3.14. All have 6 vertices and 9 edges. The labels (ℓbc) are as before: ℓ is the number of loops, b the number of bi-edges, and c the number of cut-edges. This information is enough to uniquely determine the graph with the exception of (000) , where “A” indicates the honeycomb graph and “B” the complete bipartite graph $K_{3,3}$. Although some of the names are the same as those for graphs of genus 3, this will not create a conflict as for the remainder of the section the names will only refer to graphs of genus 4 unless we specify otherwise.

Up to their respective symmetries, the square $Q_1^{(4)} = R_{3,3}$ has 5941 unimodular triangulations, the triangle $Q_2^{(4)}$ has 1268 unimodular triangulations, and the triangle $Q_3^{(4)}$ has 20 unimodular triangulations. We computed the cone \mathbb{M}_Δ for each triangulation Δ , and we ran the pipeline of Section 3.3. We summarize our findings as the following result:

Theorem 3.4.3. *Of the 17 trivalent graphs of genus 4, precisely 13 are realizable by tropical plane curves. The moduli space $\mathbb{M}_4^{\text{planar}}$ is 9-dimensional, but it is not pure: the left decomposition in (3.1.4) has components (3.1.3) of dimensions 7, 8 and 9. That decomposition is explained in Table 3.2.*




Figure 3.13: Triangulations Δ of $Q_1^{(4)}$ with $\dim(\mathbb{M}_\Delta) = 9$




Figure 3.14: The 17 trivalent graphs of genus 4. All are planar except for (000)B.

The four non-realizable graphs are (000)B, (213), (314) and (405). This is obvious for (000)B, because $K_{3,3}$ is not planar. The other three are similar to the genus 3 graph (303), and are ruled out by Proposition 3.5.4. The 13 realizable graphs G appear in the rows in Table 3.2. The first three columns correspond to the polygons $Q_1^{(4)}$, $Q_2^{(4)}$ and $Q_3^{(4)}$. Each entry is the number of regular unimodular triangulations Δ of $Q_i^{(4)}$ with skeleton G . The entry is blank if no such triangulation exists. Six of the graphs are realized by all three polygons, five are realized by two polygons, and two are realized by only one polygon. For instance, the graph (303) comes from a unique triangulation of the triangle $Q_3^{(4)}$, shown on the right in Figure 3.12. Neither $Q_1^{(4)}$ nor $Q_2^{(4)}$ can realize this graph.

Our moduli space $\text{M}_4^{\text{planar}}$ has dimension 9. We know this already from Proposition 3.2.2, where the square $Q_1^{(4)}$ appeared as $R_{3,3}$. In classical algebraic geometry, that square serves as the Newton polygon for canonical curves of genus 4 lying on a smooth quadric surface. In Table 3.2, we see that all realizable graphs except for (303) arise from triangulations of $R_{3,3}$. However, only five graphs allow for the maximal number of degrees of freedom. Corresponding triangulations are depicted in Figure 3.13.

Table 3.2: The number of triangulations for graphs of genus 4 and their moduli dimensions

G	$\#\Delta_{Q_1^{(4)},G}$	$\#\Delta_{Q_2^{(4)},G}$	$\#\Delta_{Q_3^{(4)},G}$	$\dim(\mathbb{M}_{Q_1^{(4)},G})$	$\dim(\mathbb{M}_{Q_2^{(4)},G})$	$\dim(\mathbb{M}_{Q_3^{(4)},G})$
(000)A	1823	127	12	9	8	7
(010)	2192	329	2	9	8	7
(020)	351	194		9	8	
(021)	351	3		9	7	
(030)	334	23	1	9	8	7
(101)	440	299	2	8	8	7
(111)	130	221		8	8	
(121)	130	40	1	8	8	7
(122)	130	11		8	7	
(202)	15	25		7	7	
(212)	30	6	1	7	7	7
(223)	15			7		
(303)			1			7
total	5941	1278	20			

The last three columns in Table 3.2 list the dimensions of the moduli space $\mathbb{M}_{Q_i^{(4)},G}$, which is the maximal dimension of any cone \mathbb{M}_Δ where Δ triangulates $Q_i^{(4)}$ and has skeleton G . More detailed information is furnished in Table 3.3. The three subtables (one each for $i = 1, 2, 3$) explain the decomposition (3.1.3) of each stacky fan $\mathbb{M}_{Q_i^{(4)},G}$. The row sums in Table 3.3 are the first three columns in Table 3.2. For instance, the graph (030) arises in precisely 23 of the 1268 triangulations Δ of the triangle $Q_2^{(4)}$. Among the corresponding cones \mathbb{M}_Δ , three have dimension six, twelve have dimension seven, and eight have dimension eight.

Equipped with these data, we can now extend the probabilistic analysis of Corollary 3.4.2 from genus 3 to genus 4. As before, we assume that all 17 trivalent graphs are equally likely and we fix the uniform distribution on each 8-simplex that corresponds to one of the 17 maximal cones in the 9-dimensional moduli space \mathbb{M}_4 . The five graphs that occur with positive probability are those with $\dim(\mathbb{M}_{Q_1^{(4)},G}) = 9$. Full-dimensional realizations were seen in Figure 3.13. The result of our volume computations is the following table:

Graph	(000)A	(010)	(020)	(021)	(030)
Probability	0.0101	0.0129	0.0084	0.0164	0.0336

In contrast to the exact computation in Corollary 3.4.2, our probability computations for genus 4 rely on a Monte-Carlo simulation, with one million random samples for each graph.

Corollary 3.4.4. *Less than 0.5% of all metric graphs of genus 4 come from plane tropical curves. More precisely, the fraction is approximately $\text{vol}(\mathbb{M}_4^{\text{planar}})/\text{vol}(\mathbb{M}_4) = 0.004788$.*

Table 3.3: All cones \mathbb{M}_Δ from triangulations Δ of the three polygons in Figure 3.12

$G \setminus \dim$	$Q_1^{(4)}$					$Q_2^{(4)}$				$Q_3^{(4)}$			
	5	6	7	8	9	5	6	7	8	4	5	6	7
(000)	103	480	764	400	76	5	52	60	10	1	6	3	2
(010)	38	423	951	652	128	7	113	155	54			1	1
(020)	3	32	152	128	36		53	100	41				
(021)	3	32	152	128	36		1	2					
(030)		45	131	122	36		3	12	8				1
(101)	15	155	210	60		19	122	128	30		1		1
(111)		10	80	40			52	126	43				
(121)		35	65	30			8	20	12				1
(122)		10	80	40					1				
(202)			15					25					
(212)		15	15				4	2					1
(223)			15										
(303)													1

By Theorem 3.6.1, $\mathbb{M}_{4,\text{hyp}}^{\text{planar}} = \mathbb{M}_{E_{g+2}^{(g)}}$. This space is 7-dimensional, with 6 maximal cones corresponding to the chains (020), (021), (111), (122), (202), and (223). The graphs (213), (314), and (405) are hyperelliptic if given the right metric, but beyond not being chain graphs, these are not realizable in the plane even as combinatorial types by Proposition 3.5.4.

We will now consider the moduli space of tropical plane curves of genus 5. That space is

$$\mathbb{M}_5^{\text{planar}} = \mathbb{M}_{Q_1^{(5)}} \cup \mathbb{M}_{Q_2^{(5)}} \cup \mathbb{M}_{Q_3^{(5)}} \cup \mathbb{M}_{Q_4^{(5)}} \cup \mathbb{M}_{5,\text{hyp}}^{\text{planar}},$$

where $Q_1^{(5)}$, $Q_2^{(5)}$, $Q_3^{(5)}$, $Q_4^{(5)}$ are the four genus 5 polygons in Proposition 3.3.6. They are shown in Figure 3.15. Modulo their respective symmetries, the numbers of unimodular triangulations of these polygons are: 508 for $Q_1^{(5)}$, 147908 for $Q_2^{(5)}$, 162 for $Q_3^{(5)}$, and 968 for $Q_4^{(5)}$.




Figure 3.15: The genus 5 polygons $Q_1^{(5)}$, $Q_2^{(5)}$, $Q_3^{(5)}$ and $Q_4^{(5)}$

We applied the pipeline of Section 3.3 to all these triangulations. The outcome of our computations is the following result which is the genus 5 analogue to Theorem 3.4.3.

Theorem 3.4.5. *Of the 71 trivalent graphs of genus 5, precisely 38 are realizable by smooth tropical plane curves. The four polygons satisfy $\dim(\mathbb{M}_{Q_i^{(5)}}) = 9, 11, 10, 10$ for $i = 1, 2, 3, 4$.*

These 38 realizable graphs are illustrated in Figure 3.16. All but one of them arises from $Q_1^{(5)}$ or $Q_2^{(5)}$. The remaining graph, realized only by a single triangulation of $Q_4^{(5)}$, is illustrated in Figure 3.17. This is reminiscent of the genus 4 graph (303), which was realized only by the triangulation of $Q_3^{(4)}$ in Figure 3.12. The other 37 graphs are realized by at least two of the polygons $Q_1^{(5)}, \dots, Q_4^{(5)}, E_7^{(5)}$.




Figure 3.16: The 38 trivalent graphs of genus 5 appearing in $\mathbb{M}_5^{\text{planar}}$




Figure 3.17: A genus 5 graph, and the unique triangulation that realizes it

The other 33 trivalent graphs of genus 5 are pictured in Figure 3.18. Of these, 26 can be quickly seen not to appear in $\mathbb{M}_5^{\text{planar}}$. Four of the graphs are non-planar, fifteen are

sprawling, and seven are *crowded* (these latter two properties will be defined in the next section, which will also prove that any genus g graph with either property cannot appear in $\mathbb{M}_g^{\text{planar}}$). However, the remaining seven graphs are not ruled out by any yet-known concise criteria, and at present have only been shown to be unrealizable in a tropical plane curve by our massive computation. This contrasts the cases of genus 3 and 4, where all graphs not appearing were either sprawling or nonplanar.




Figure 3.18: The 33 trivalent graphs of genus 5 not appearing in $\mathbb{M}_5^{\text{planar}}$. The first four are nonplanar, the next fifteen are sprawling, and the seven after that are crowded.

The process we have carried out for genus $g = 3, 4$, and 5 can be continued for $g \geq 6$. As the genus increases so does computing time, so it may be prudent to limit the computations to special cases of interest. For $g = 6$ we might focus on the triangle $Q_1^{(6)} = T_5$. This is of particular interest as it is the Newton polygon of a smooth plane quintic curve. This triangle has 561885 regular unimodular triangulations up to symmetry.

Although T_5 is interesting as the Newton polygon of plane quintics, it has the downside that \mathbb{M}_{T_5} is not full-dimensional inside $\mathbb{M}_6^{\text{planar}}$. Proposition 3.2.2 implies that $\dim(\mathbb{M}_{T_5}) = 12$, while $\dim(\mathbb{M}_6^{\text{planar}}) = 13$, and this dimension is attained by the rectangle $R_{3,4}$ as in (3.2.3).

This might lead us to focus on *full-dimensional polygons* of genus g . By this we mean polygons P whose moduli space \mathbb{M}_P has the dimension in (3.1.7). For each genus from 3 to 5, our results show that there is a unique full-dimensional polygon, namely, T_4 , $R_{3,3}$, and $Q_2^{(5)}$. The proof of Theorem 3.1.1 furnishes an explicit example for each genus $g \geq 6$: take the rectangle in (3.2.3) or the trapezoid in (3.2.4) if $g \neq 7$, or the hexagon $H_{4,4,2,6}$ if $g = 7$.

Preliminary calculations show that there are exactly two full-dimensional maximal polygons for $g = 6$, namely, $Q_3^{(6)} = R_{3,4}$ and $Q_4^{(6)}$ from Proposition 3.3.6.

3.5 Tropical hyperelliptic curves in the plane: theory

Recall that a polygon P of genus g is *hyperelliptic* if P_{int} is a line segment of length $g - 1$. The moduli space of hyperelliptic tropical plane curves of genus g is

$$\mathbb{M}_{g,\text{hyp}}^{\text{planar}} := \bigcup_P \mathbb{M}_P,$$

where the union is over all hyperelliptic polygons P of genus g .

It is reasonable to ask about the relationship between $\mathbb{M}_{g,\text{hyp}}^{\text{planar}}$ and $\mathbb{M}_{g,\text{hyp}}$, the locus of all hyperelliptic graphs in \mathbb{M}_g . The easier direction is that the first is contained in the second: assuming P_{int} is a horizontal line segment, a 2-to-1 map from a tropical curve (dual to a subdivision of P) to a line is given by vertical projection and bridge-dilating. This section will be spent proving the following theorem, which shows that the relationship is as nice as can be hoped for.

Theorem 3.5.1. *If a smooth tropical plane curve with Newton polygon P has a hyperelliptic skeleton, then P is a hyperelliptic polygon.*

This means that, at least before taking closures,

$$\mathbb{M}_{g,\text{hyp}}^{\text{planar}} = \mathbb{M}_{g,\text{hyp}} \cap \mathbb{M}_g^{\text{planar}}.$$

This is a generalization of Theorem 2.3.14, which proved the result for genus 3.

The space $\mathbb{M}_{g,\text{hyp}}^{\text{planar}}$ is contained in *the moduli space of chains* $\mathbb{M}_g^{\text{chain}}$, which we now define. Start with a line segment on $g - 1$ nodes where the $g - 2$ edges have arbitrary non-negative lengths. Double each edge so that the resulting parallel edges have the same length, and attach two loops of arbitrary lengths at the endpoints. Now, each of the $g - 1$ nodes is 4-valent. There are two possible ways to split each node into two nodes connected by an edge of arbitrary length. Any metric graph arising from this procedure is called a *chain of genus g* . Although there are 2^{g-1} possible choices in this procedure, some give isomorphic graphs. There are $2^{g-2} + 2^{\lfloor (g-2)/2 \rfloor}$ combinatorial types of chains of genus g . In genus 3 the chains are (020), (111), and (212) in Figure 3.4, and in genus 4 they are (020), (021), (111), (122), (202), and (223) in Figure 3.14.

By construction, there are $2g - 1$ degrees of freedom for the edge lengths in a chain of genus g , so each such chain defines an orthant $\mathbb{R}_{\geq 0}^{2g-1}$. We write $\mathbb{M}_g^{\text{chain}}$ for the stacky subfan of \mathbb{M}_g consisting of all chains. Note that $\mathbb{M}_g^{\text{chain}}$ is strictly contained in the space $\mathbb{M}_{g,\text{hyp}}$ of all hyperelliptic metric graphs. Hyperelliptic graphs arise by the same construction from any tree with $g - 1$ nodes, whereas for chains that tree must be a line segment.

Our strategy for proving Theorem 3.5.1 is as follows. We will define two types of graphs, *sprawling* and *crowded*, which we will show are never the skeletons of smooth tropical plane curves. In Proposition 3.5.11 we will use these criteria to show that a hyperelliptic graph that is a smooth tropical plane curve’s skeleton must be a chain. It will then suffice to show that if a polygon gives rise to a hyperelliptic chain, it must be a hyperelliptic polygon.

Definition 3.5.2. A connected, trivalent, leafless graph G is called *sprawling* if there exists a vertex s of G such that $G \setminus \{s\}$ consists of three distinct components.

Remark 3.5.3. Each component of $G \setminus \{s\}$ must have genus at least one; otherwise G would not have been leafless. The vertex s need not be unique. The genus 3 graph (303) in Figure 3.4 is sprawling, as are the genus 4 graphs (213), (314), and (405) in Figure 3.14.

Lemma 3.5.4. *Sprawling graphs are never the skeletons of smooth tropical plane curves.*

This was originally proven in [CDMY14, Prop. 4.1]. We present our own proof for completeness.

Proof. Suppose for the sake of contradiction that the skeleton of a smooth tropical plane curve C is a sprawling graph G with separating vertex s . After a change of coordinates, we may assume that the directions emanating from s are $(1, 1)$, $(0, -1)$, and $(-1, 0)$. The curve C is dual to a unimodular triangulation Δ of a polygon $P \subset \mathbb{R}^2$. Let $T \in \Delta$ be the triangle dual to s . We may take $T = \text{conv}\{(0, 0), (0, 1), (1, 0)\}$ after an appropriate translation of P . Let P_1, P_2, P_3 be the subpolygons of P corresponding to the components of $G \setminus \{s\}$. After relabeling, we have $P_1 \cap P_2 = \{(0, 1)\}$, $P_1 \cap P_3 = \{(0, 0)\}$, and $P_2 \cap P_3 = \{(1, 0)\}$. Each P_i has at least one interior lattice point, since each component of $G \setminus \{s\}$ must have genus at least 1.




Figure 3.19: The triangle T with angles formed between it and the boundary edges of P

Let $\alpha, \beta, \gamma, \delta$ be angles between the triangle T and the boundary edges of P emanating from the vertices of T , as pictured in Figure 3.19. Since P is convex, we know $\alpha + \beta \leq 3\pi/4$, $\gamma < \pi/2$, and $\delta < 3\pi/4$. As P_1 contains at least one interior lattice point, and $\gamma < \pi/2$, we

must also have that $\alpha > \pi/2$; otherwise $P_1 \subset (\infty, 0] \times [0, 1]$, which has no interior lattice points. Similarly, as P_2 has at least one interior lattice point and $\delta < 3\pi/4$, we must have that $\beta > \pi/4$. But we now have that $\alpha + \beta > \pi/2 + \pi/4 = 3\pi/4$, a contradiction. Thus, the skeleton of C cannot be a sprawling graph, as originally assumed. \square

Definition 3.5.5. A planar embedding of a connected, trivalent, leafless planar graph G is called *crowded* if either: there exist two bounded faces sharing at least two edges; or, there exists a bounded face sharing an edge with itself. If all planar embeddings of such a G are crowded, we say that G is *crowded*.

Lemma 3.5.6. *Crowded graphs are never the skeletons of smooth plane tropical curves.*

Proof. Suppose the skeleton of a smooth plane tropical curve is a planar embedding of a crowded graph G . If the embedding has two bounded faces F and F' sharing at least two edges, let p and p' be the corresponding interior lattice points of the tropical curve's Newton polygon. Since F and F' share at least two edges, p and p' must be connected by at least two edges in the corresponding triangulation of the Newton polygon. This is impossible, since the only possible edge between p and p' is the unique line segment connecting them. A similar argument holds if the embedding of G has a bounded face sharing an edge with itself. These contradictions prove the claim. \square

It is relatively simple to check that a graph is sprawling, but checking that a graph is crowded is *a priori* difficult, since a condition must hold over *all* planar embeddings of a graph. The following criterion lets us check for crowdedness from a single planar embedding. We will use the fact that since G is trivalent, the dual graph G^* on $g + 1$ vertices is well-defined, *independent of the embedding of G* .

Proposition 3.5.7. *Let G be a connected, planar, leafless, trivalent graph, and let $M \in \mathbb{N}^{(g+1) \times (g+1)}$ be the adjacency matrix of the dual graph G^* . Let $M_{[i]}$ denote the submatrix of M obtained by deleting the i^{th} row and the i^{th} column. Then G is crowded if and only if for all i , $M_{[i]}$ has either an entry ≥ 2 , or a diagonal entry ≥ 1 .*

Proof. Assume G is crowded. Embed G on a sphere, and project to obtain a planar embedding with the i^{th} face unbounded. Then $M_{[i]}$ is the adjacency matrix of the subgraph of G^* obtained by deleting the vertex corresponding to the i^{th} face and all adjacent edges. Either some pair of the bounded faces must share at least two edges, and so $M_{[i]}$ must have some entry ≥ 2 ; or a bounded face shares an edge with itself, and so $M_{[i]}$ must have some diagonal entry ≥ 1 .

Now assume that for all i , either $M_{[i]}$ has an entry ≥ 2 , or a diagonal entry ≥ 1 . Let k be the index corresponding to the unbounded face. If ℓ is the index of the unbounded face of G , we have that either $(M_{[\ell]})_{j,k} \geq 2$ for some j, k or $(M_{[\ell]})_{i,i} \geq 1$ for some i . In the former case, the j^{th} and k^{th} faces are bounded and share at least two edges. By the invariance of M up to relabeling rows and columns, this will hold for all planar embeddings of G , so G is crowded. A similar argument holds for the latter case. \square

In the following results we will refer to a *2-connected component* of a graph G . This is a subgraph of G that is a connected component of the graph obtained by deleting all bridges from G and smoothing over 2-valent vertices. Note that a 2-connected component will be a 2-edge connected graph. If G is sprawling, then the vertex that witnesses this property will be a 2-connected component. A *nontrivial* 2-connected component is one that is not a vertex.

Lemma 3.5.8. *Let G' be a 2-connected component of a planar graph G . If G' is crowded then so is G .*

Proof. Embed G , then delete everything that is not part of G' , smoothing over the resulting 2-valent vertices. Label the (bounded and unbounded) faces F_1, \dots, F_k . Now add back in the rest of G . Since G' is a 2-connected component of G , the faces F_1, \dots, F_k are preserved as faces (without being split), and the number of edges shared by pairs (not necessarily distinct) amongst F_1, \dots, F_k have either remained the same or increased. Applying Proposition 3.5.7 shows that G is crowded. \square

Remark 3.5.9. It is possible for a graph to have a crowded subgraph without being crowded, as long as that subgraph is not a 2-connected component. See Figure 3.20 for an example.




Figure 3.20: A graph that isn't crowded, but has a crowded subgraph

Proposition 3.5.10. *If a hyperelliptic graph is 2-connected, either it is a chain or it is crowded.*

Proof. Let G be a 2-connected hyperelliptic graph of genus g . By [Cha13, Theorem 4.9], G is a ladder over tree with $g - 1$ nodes, each with valency at most three. Note that G is the 2-connected chain of genus g if and only if the tree is a line segment. Assume G is not a chain. Then the tree must contain a trivalent vertex v , which we will label v' in the second copy of T used to obtain G . This means G is of the form shown in Figure 3.21, where each unknown box contains at least one bi-edge. This means that each pair of faces amongst the faces F_1, F_2, F_3 shares two edges. It follows that the adjacency matrix of the dual graph begins

$$\begin{pmatrix} 0 & 2 & 2 & \cdots \\ 2 & 0 & 2 & \cdots \\ 0 & 2 & 2 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Proposition 3.5.7 implies that G is crowded. □




Figure 3.21: A non-chain 2-connected hyperelliptic graph with three faces labeled.

Proposition 3.5.11. *If G is the hyperelliptic skeleton of a smooth tropical plane curve, then G is a chain.*

Our proof of this proposition will show something even stronger: that if a graph is the skeleton of a smooth tropical plane curve, and there exists *some* metric on that graph that makes it hyperelliptic (not necessarily the metric given by that embedding), then that graph must be a chain under that new metric.

Proof. Let G be such a hyperelliptic skeleton, meaning it comes with a given embedding into the plane which cannot be a crowded embedding. Each 2-connected component of G must be hyperelliptic and not crowded by Lemma 3.5.8, and so by Proposition 3.5.10 each 2-connected component of G must be either a chain or just a vertex.

Let G' be a nontrivial 2-connected component of G , from which G' inherits a non-crowded embedding into the plane. Since G' is a chain, it must be embedded in the standard chain embedding in Figure 3.22: any other embedding is crowded, as can be read off from the adjacency matrix of a dual graph. The only edges that could possibly connect G' to the rest of G are an edge from the middle of e_0 and an edge from the middle of e_g . This is because a 2-connected component connecting to G from any other e_i (or with multiple edges from e_0 or from e_g) would make a bounded face of G share an edge with itself, meaning the embedding was crowded.

It follows that each 2-connected component of G has at most one incoming and one outgoing edge. As with any graph, shrinking the 2-connected components down to nodes yields a tree T . Lemma 3.5.4 implies that T must be a line segment: if T had any trivalent nodes, the corresponding 2-connected component would have to be a vertex, and so G would be sprawling. Considering the structure of each nontrivial 2-connected component, we conclude that G must be a chain. □




Figure 3.22: The standard embedding of a chain, with vertical edges labelled e_0 to e_g

This proposition implies that $\mathbb{M}_g^{\text{planar}} \cap \mathbb{M}_{g,\text{hyp}} \subset \mathbb{M}_g^{\text{planar}} \cap \mathbb{M}_g^{\text{chain}}$. However, it is not immediately clear that there is no contribution from nonhyperelliptic polygons, which could in principle give rise to hyperelliptic chains. The following proposition rules this out, and is the last ingredient we need in order to prove Theorem 3.5.1.

Proposition 3.5.12. *Let G be the skeleton of a smooth tropical plane curve with nonhyperelliptic Newton polygon P . If G is combinatorially a chain, then the metric on G is not hyperelliptic.*

In the proof of this proposition, when we say “ G is a chain”, we mean that G is a chain combinatorially, possibly without a hyperelliptic metric.

Proof. Let Δ be the unimodular triangulation of P dual to the smooth plane tropical curve with skeleton G . The order on the distinguished cycles c_1, \dots, c_g of G induces a natural ordering on the interior lattice points of P , which we will call p_1, \dots, p_g . Since P is nonhyperelliptic, there exists some triple p_i, p_{i+1}, p_{i+2} of these interior lattice points that are not collinear. We will assume that the cycle c_{i+1} shares an edge with c_i and an edge with c_{i+2} ; the other cases with at least one bridge coming from c_{i+1} are handled similarly. Dually, this means that Δ contains the line segments $p_i p_{i+1}$ and $p_{i+1} p_{i+2}$. After a change of coordinates we may assume $p_i = (0, 1)$, $p_{i+1} = (0, 0)$, and $p_{i+2} = (1, 0)$.

Since G is a chain, the triangulation Δ does not contain the line segment $p_i p_{i+2}$. This means that some line segment in Δ containing p_{i+1} must separate p_i and p_{i+2} . By the convexity of P , this means that the point $q = (1, 1)$ is contained in P , and in fact $p_{i+1} q$ is a line segment in Δ . Since G is a chain, it follows that q is a boundary point of P . Since P is convex, there is no segment in Δ containing p_{i+1} that separates p_i from q , or p_{i+2} from q . It follows that $p_i q$ and $p_{i+2} q$ are both segments in the triangulation Δ . In the dual tropical curve, let e_h be dual to $p_i p_{i+1}$; e_v be dual to $p_{i+1} p_{i+2}$; e_1 be dual to $p_{i+1} q$; and e_2 be the remainder of the cycle c_{i+1} . This is illustrated in Figure 3.23. Let ℓ_h, ℓ_v, ℓ_1 , and ℓ_2 denote the lengths of these edges, respectively.

Let q_1, q_2, \dots, q_n denote the lattice points of P that Δ connects to p_{i+1} , ordered counterclockwise starting with q (so that $q_1 = q$, $q_2 = p_i$ and $q_n = p_{i+2}$). Write these points in coordinates as $q_j = (a_j, b_j)$. For $3 \leq j \leq n - 1$, at least one of a_j and b_j must be negative due to the placement of p_i and p_{i+2} . In fact, for all j , at least one of a_j and b_j is -1 : if not, convexity of P would imply the existence of a new interior lattice point that is connected to p_{i+1} by a segment Δ , which is impossible as G is a chain.




Figure 3.23: A portion of the triangulation Δ of P , and part of the dual tropical curve

In the tropical embedding of the graph G , the edge e_2 is made up of line segments that are dual to q_3, q_4, \dots, q_{n-1} . Consider the line segments in e_2 dual to q_i 's of the form $(a_i, -1)$. The sum of the horizontal widths of these segments must be at least the sum of the horizontal widths of e_1 and e_h : otherwise the cycle c_{i+1} would not be closed. Since these line segments in e_2 have slopes in \mathbb{Z} , each of them has lattice length equal to horizontal width. The same holds for e_1 and e_h , implying $\ell_2 \geq \ell_1 + \ell_h > \ell_1$. This means G has edges of a two-cut with different lengths, namely e_1 and e_2 with lengths $\ell_1 \neq \ell_2$. By Lemma 2.3.13, the metric on G cannot be hyperelliptic. This completes the proof. \square

It is worth remarking that it is not immediately obvious from Figure 3.23 that e_2 is longer than e_1 , since we are considering lattice length rather than Euclidean length. For instance, if $\ell_1 = \ell_v = 1$, $\ell_h = 2$, and e_2 consists of a single line segment with slope $2/3$, then $\ell_1 = \ell_2$. This is ruled out by constraints on the lattice polygon P , but the result does require more work than it might initially seem.

The results of this section now allow us to prove that if hyperelliptic smooth tropical plane curves only come from hyperelliptic Newton polygons.

Proof of Theorem 3.5.1. Let C be a smooth plane tropical curve with Newton polygon P and a hyperelliptic skeleton G . By Proposition 3.5.11, the graph G must be a chain. If P were not a hyperelliptic polygon, then by Proposition 3.5.12 the chain G could not be hyperelliptic as assumed. We conclude that P must be a hyperelliptic polygon. \square

3.6 Tropical hyperelliptic curves in the plane: computation

We will now consider how to compute $\mathbb{M}_{g,\text{hyp}}^{\text{planar}}$, which is the union of all \mathbb{M}_P 's with P hyperelliptic with g interior lattice points. Unlike when the interior hull P_{int} is two-dimensional, there does not exist a unique maximal hyperelliptic polygon P with given P_{int} . However,

there are only finitely many such polygons up to isomorphism. These are

$$E_k^{(g)} := \text{conv}\{(0, 0), (0, 2), (g + k, 0), (g + 2 - k, 2)\} \quad \text{for } 1 \leq k \leq g + 2.$$

These hyperelliptic polygons interpolate between the rectangle $E_1^{(g)} = R_{g+1,2}$ and the triangle $E_{g+2}^{(g)}$. The five maximal hyperelliptic polygons for genus $g = 3$ are pictured in Figure 3.24.




Figure 3.24: The five maximal hyperelliptic polygons of genus 3

At first, it seems that to compute $\mathbb{M}_{g,\text{hyp}}^{\text{planar}}$ one must compute $\mathbb{M}_{E_k^{(g)}}$ for all k , and take the union. By [KZ03, Proposition 3.4], all triangulations of hyperelliptic polygons are regular, so we need not worry about non-regular triangulations arising in the TOPCOM computations described in Section 3.3. We next show that it suffices to consider the triangle:

Theorem 3.6.1. *For each genus $g \geq 2$, the hyperelliptic triangle $E_{g+2}^{(g)}$ satisfies*

$$\mathbb{M}_{E_{g+2}^{(g)}} = \mathbb{M}_{g,\text{hyp}}^{\text{planar}}. \tag{3.6.1}$$

The equality holds even before taking closures of the spaces of realizable graphs.

To prove Theorem 3.6.1 we must show that any metric graph arising from a maximal hyperelliptic polygon $E_k^{(g)}$ also arises from the hyperelliptic triangle $E_{g+2}^{(g)}$. Given a triangulation Δ of $E_k^{(g)}$, our proof constructs a triangulation Δ' of $E_{g+2}^{(g)}$ that gives rise to the same collection of metric graphs, so that $\mathbb{M}_\Delta = \mathbb{M}_{\Delta'}$, with equality holding even before taking closures of the spaces of realizable graphs. Before our proof, we illustrate this construction with the following example.

Example 3.6.2. Let Δ be the triangulation of $R_{4,2}$ pictured on the left in Figure 3.25 along with a metric graph G arising from it. We claim there exists a triangulation Δ' of the hyperelliptic triangle $E_5^{(3)}$ that gives rise to the exact same set of metric graphs. This triangulation is also pictured in Figure 3.25. We will explain why these triangulations give the same metric graphs, and then explain how to construct Δ' algorithmically from Δ .

The possible metrics on G are determined by the slopes of the line segments emanating from the vertical edges. For instance, consider the constraints on v and y imposed by the width w (which equals x). If most of the w and x edges are made up of the segments emanating from v , we find y close to $v + 2w$. If instead most of the w and x edges are made up of the segments emanating from y , we find y close to $v - 2w$. Interpolating gives graphs achieving $v - 2w < y < v + 2w$. This only depends on the *difference* of the slopes emanating either left or right from the edges v and y : the same constraints would be imposed if the slopes emanating from v to the right were 2 and 0 rather than 1 and -1 . Boundary behavior determines constraints on u and z , namely $v < u$ and $y < z$.




Figure 3.25: Triangulations of $R_{4,2}$ and $E_5^{(3)}$, giving rise to skeletons with the same possible metrics.

The skeleton G' arising from Δ' has the same combinatorial type as G . Moreover, each vertical edge in G' has outgoing edges to the left and to the right, whose slope differences are the same as in G . For instance, from the vertical edge v , both graphs have the first segment of w and the first segment in x emanating to the right. Although the slopes of these segments are different in the two graphs, the difference in the slopes is 2 for both. This slope similarity, combined with similar boundary behavior, this shows that G and G' have the exact same achievable metrics. In other words, $\mathbb{M}_\Delta = \mathbb{M}_{\Delta'}$, with equality even before taking closures of the realizable graphs.

We now explain how to construct Δ' from Δ , an algorithm spelled out explicitly in the proof of Theorem 3.6.1. We start by adding edges from $(0, 2)$ to the interior lattice points (since any unimodular triangulation of $E_5^{(3)}$ must include these edges), and then add additional edges based on the combinatorial type of Δ , as pictured in Figure 3.26.




Figure 3.26: The start of the construction of the triangulation Δ'

Next we add edges connecting the interior lattice points to the lower edge of the triangle. We will ensure that the outgoing slopes from the vertical edges in the Γ' have the same difference as in Γ . For $i = 1, 2, 3$, we will connect $(i, 1)$ to all points between $(2i + a_i, 0)$ and $(2i + b_i, 0)$ where a_i is the difference between the reciprocals of the slopes of the leftmost edges from $(i, 1)$ to the upper and lower edges of $R_{4,2}$ in Δ , and b_i is defined similarly but with the rightmost edges. Here we take the reciprocal of ∞ to be 0. In the dual tropical curve, this translates to slopes emanating from vertical edges in the tropical curve having the same difference as from Δ .

We compute $a_1 = \frac{1}{-1} - \frac{1}{1} = -2$ and $b_1 = \frac{1}{\infty} - \frac{1}{\infty} = 0$. Since $2 \cdot 1 + a_1 = 0$ and $2 \cdot 1 + b_1 = 2$, we add edges from $(1, 1)$ to $(0, 0)$, to $(0, 2)$, and to all points in between, in this case just $(0, 1)$. We do similarly for the other two interior lattice points, as pictured in the first three triangles in Figure 3.27. The fourth triangle includes the edges $(0, 1) - (1, 1)$

and $(3, 1) - (4, 1)$, which ensures the same constraints as from Δ on the first and third loops of the corresponding metric graph.




Figure 3.27: Several steps leading up to Δ' , on the right

Proof of Theorem 3.6.1. We need to prove $\mathbb{M}_{E_{g+2}^{(g)}} = \mathbb{M}_{g, \text{hyp}}^{\text{planar}}$. Given any triangulation Δ of a hyperelliptic polygon $E_k^{(g)}$, we shall construct a triangulation Δ' of $T := E_{g+2}^{(g)}$ such that $\mathbb{M}_\Delta = \mathbb{M}_{\Delta'}$. Our construction will show that the equality even holds before taking the closure of the sets of achievable metrics on the skeletons.

Before we define Δ' , we will need to define several numbers attached to a triangulation of a maximal hyperelliptic polygon, and describe the corresponding realizable dual graphs.

Let Δ be a triangulation of $E_k^{(g)} = \text{conv}\{(0, 0), (0, 2), (g + k, 0), (g + 2 - k, 2)\}$, where $1 \leq k \leq g + 2$. For each i , let

- σ_i^{NW} be the slope of the leftmost edge from $(i, 1)$ to the top edge of $E_k^{(g)}$,
- σ_i^{NE} be the slope of the rightmost edge from $(i, 1)$ to the top edge of $E_k^{(g)}$,
- σ_i^{SW} be the slope of the leftmost edge from $(i, 1)$ to the bottom edge of $E_k^{(g)}$, and
- σ_i^{SE} be the slope of the rightmost edge from $(i, 1)$ to the bottom edge of $E_k^{(g)}$.

Each (possibly infinite) number is well-defined as no edge of Δ can separate $(i, 1)$ from the top or the bottom edge of $E_k^{(g)}$. It is possible for $\sigma_i^{\text{NW}} = \sigma_i^{\text{NE}}$ or $\sigma_i^{\text{SW}} = \sigma_i^{\text{SE}}$. See Figure 3.28 for an example with $\sigma_2^{\text{NW}} = -\frac{1}{2}$, $\sigma_2^{\text{NE}} = \infty$, $\sigma_2^{\text{SW}} = 1$, and $\sigma_2^{\text{SE}} = -\frac{1}{3}$.




Figure 3.28: A triangulation Δ highlighting the point $(2, 1)$, with the corresponding σ 's in Δ and s 's in the dual tropical cycle.

Let $s_i^{\text{NW}} = -1/\sigma_i^{\text{NW}}$, so that in a tropical curve dual to Δ , s_i^{NW} is the slope of the edge dual to the edge that determines σ_i^{NW} . We define s_i^{NE} , s_i^{SW} , and s_i^{SE} similarly. See Figure 3.28 for an example with $s_2^{\text{NW}} = 2$, $s_2^{\text{NE}} = 0$, $s_2^{\text{SW}} = -1$, and $s_2^{\text{SE}} = 3$. Based on the vertices of $E_k^{(g)}$, we have the bounds

$$i + k - (g + 2) \leq s_i^{\text{NE}} \leq s_i^{\text{NW}} \leq i$$

and

$$-i \leq s_i^{\text{SW}} \leq s_i^{\text{SE}} \leq g + k - i.$$

Define $d_i^{\text{W}} = s_i^{\text{NW}} - s_i^{\text{SW}}$ and $d_i^{\text{E}} = s_i^{\text{NE}} - s_i^{\text{SE}}$. We'll call $d(\Delta) = (d_1^{\text{W}}, d_1^{\text{E}}, d_2^{\text{W}}, d_2^{\text{E}}, \dots, d_g^{\text{W}}, d_g^{\text{E}})$ the d -vector of Δ . The bounds on the slopes imply

$$2i - (2g + 2) \leq d_i^{\text{E}} \leq d_i^{\text{W}} \leq 2i.$$

We also have $s_{i+1}^{\text{NW}} = s_i^{\text{NE}} + 1$ and $s_i^{\text{SW}} = s_i^{\text{SE}} - 1$, so $d_{i+1}^{\text{W}} = d_i^{\text{E}} + 2$. If there is ever any ambiguity as to which triangulation an s_i or a d_i is coming from, we will use such notation as $s_i^{\text{SW}}(\Delta)$.

To describe the lengths on a planar chain graph dual to Δ , we will use the following notation for edge lengths:

- ℓ_1 and ℓ_g for leftmost and rightmost edges.
- w_i with $2 \leq i \leq g - 1$ for the width of the i^{th} loop.
- $h_{i,i+1}$ for the length of an edge shared by the i^{th} and $(i + 1)^{\text{th}}$ loops (taken to be 0 if these cycles do not share an edge).
- $b_{i,i+1}$ for the length of a cut-edge connecting the i^{th} and $(i + 1)^{\text{th}}$ loops (taken to be 0 if these cycles are not connected by a cut-edge).

Figure 3.29 illustrates these labels for a particular chain graph with $g = 6$.




Figure 3.29: A labelled skeleton of a chain tropical curve with $g = 6$.

Let us now describe the metric chain graphs achieved in the plane dual to Δ with a given d -vector, *regardless of which polygon E_k^g is triangulated by Δ* . By the convexity of each loop, any tropical curve dual to Δ must have edge lengths satisfying

$$d_i^{\text{E}} \leq (h_{i-1,i} - h_{i,i+1})/w_i \leq d_i^{\text{W}}.$$

Moreover, these inequalities fail to be strict if and only if $d_i^{\text{E}} = d_i^{\text{W}}$ (which by convexity means $s_i^{\text{NW}} = s_i^{\text{NE}}$ and $s_i^{\text{SW}} = s_i^{\text{SE}}$). If $d_i^{\text{E}} \neq d_i^{\text{W}}$, then all lengths between d_i^{E} and d_i^{W} are

achievable. Since any nonzero $b_{i,i+1}$ can be made an arbitrary length, we conclude that the d -vector $d(\Delta)$ determines the set of all graphs achievable by Δ , except for possible constraints on ℓ_1 and ℓ_g ; namely, we can have exactly the choices of w_i 's and $h_{i,i+1}$'s compatible with

$$\begin{cases} d_i^E < (h_{i-1,i} - h_{i,i+1})/w_i < d_i^W \text{ if } d_i^E \neq d_i^W, \\ d_i^E = (h_{i-1,i} - h_{i,i+1})/w_i = d_i^W \text{ otherwise.} \end{cases}$$

We now have the necessary machinery to define the unimodular triangulation Δ' of T and prove it has the desired properties. It may be helpful to refer to Example 3.6.2, where an example of Δ' is constructed explicitly. First, Δ' contains the edges connecting $(0, 2)$ to $(i, 1)$ for $1 \leq i \leq g$, as these edges must be present in any unimodular triangulation of T . Next, if Δ includes the edge from $(i, 1)$ to $(i + 1, 1)$, then this edge is in Δ' ; if not then Δ' includes the edge from $(0, 2)$ to $(2i + 1, 0)$. This ensures that the combinatorial type of the dual graph of Δ' will be the same as that of Δ .

For each i from 1 to g , Δ' contains the edges connecting $(i, 1)$ to $(2i - d_i^W(\Delta), 0)$, to $(2i - d_i^E(\Delta), 0)$, and to any points in between. These points are in T since

$$0 = 2i - 2i \leq 2i - d_i^W(\Delta) \leq 2i - d_i^E(\Delta) \leq 2i - (2i - (2g + 2)) = 2g + 2.$$

Suppose there is a conflict in these edges preventing a coherent triangulation, say coming from the points $(i, 1)$ and $(i + 1, 1)$. Then we know that $2i - d_i^E(\Delta) < 2(i + 1) - d_{i+1}^W(\Delta)$, which can be rewritten as $d_{i+1}^W(\Delta) < d_i^E(\Delta) + 2$. But we know that $d_{i+1}^W(\Delta) = d_i^E + 2(\Delta)$, a contradiction. Hence there can be no conflict in the edges.

To complete the definition of Δ' , we must define what is happening in the vicinity of $(1, 1)$ and $(g, 1)$. By symmetry, it will suffice to describe the case for $(1, 1)$. Let $(n, 0)$ be the leftmost point of the bottom edge of T connected to $(1, 1)$.

- (i) If $n = 0$, let Δ' include the edge $(0, 1) - (1, 1)$.
- (ii) If $n \geq 2$, let Δ' include the edge $(0, 1) - (1, 1)$ and all edges $(0, 1) - (0, m)$ with $0 \leq m \leq n$.
- (iii) If $n = 1$ and Δ contains the edge $(0, 1) - (1, 1)$, let Δ' include the edges $(0, 1) - (1, 1)$ and $(0, 1) - (1, 0)$.
- (iv) If $n = 1$ and Δ contains the edge $(0, 1) - (1, 1)$, let Δ' contain $(0, 2) - (1, 0)$ and $(0, 1) - (1, 0)$.

It's worth noting that in cases (i) and (ii), we completed Δ' in the vicinity of $(1, 1)$ in the only way possible. A symmetric definition in the vicinity of $(g, 1)$ completes the definition of Δ' .

To show that Δ' and Δ give rise to the same set of metric graphs, we first show that $d(\Delta) = d(\Delta')$. We have $s_i^{\text{NW}}(\Delta') = s_i^{\text{NE}}(\Delta') = i$, $s_i^{\text{SW}}(\Delta') = (2i - d_i^W(\Delta)) - i$, and $s_i^{\text{SE}} = (2i - d_i^E(\Delta)) - i$, so

$$d_i^W(\Delta') = i - (2i - d_i^W(\Delta)) - i = d_i^W(\Delta)$$

and

$$d_i^E(\Delta') = i - (2i - d_i^E(\Delta)) - i = d_i^E(\Delta),$$

as desired.

Since $d(\Delta) = d(\Delta')$, we know that Δ and Δ' give rise to the same set of graphs up to perhaps having different constraints on ℓ_1 and ℓ_g . It is enough to show the constraints on ℓ_1 are the same for the two triangulations, as symmetry will show the same holds for ℓ_g . We will focus on case (iii), wherein $n = 1$ and Δ does not contain $(0, 1) - (1, 1)$. The other three cases are similar.

Considering the edges emanating from $(1, 1)$ in Δ' , the constraints on ℓ_1 are $h_{12} < \ell_1$ if Δ' contains $(1, 1) - (2, 0)$, and $h_{12} < \ell_1 < 2h_{12}$ otherwise. The exact same is true of Δ : if Δ' contains $(1, 1) - (2, 0)$, then $d_1^E(\Delta) = d_1^E(\Delta') \leq 1 - 1 = 0$, which allows for arbitrary scaling of ℓ_1 that can be truncated by the edge $(1, 0) - (1, 1)$. And if Δ' does not contain $(1, 1) - (2, 0)$, then $d_1^E(\Delta) = d_1^E(\Delta') = -1$. The total length of ℓ_1 is then twice the width w of ℓ_1 plus the height of the edge dual to $(1, 0) - (1, 1)$, which will be at most $h_{12} - w$. As all lengths must be positive, we know $0 < w < h_{12}$, which leads to the bound of $h_{12} < \ell_1 < 2h_{12}$. As all such graphs are realizable, the constraints on ℓ_1 are the same for both triangulations.

We have shown that the metric graphs arising from Δ' are exactly the same as the metric graphs arising from Δ . This gives a containment $\mathbb{M}_{E_k^{(g)}} \subset \mathbb{M}_T$, true even before taking closures. This completes the proof. \square

We will now compute $\mathbb{M}_{3,\text{hyp}}^{\text{planar}}$. By (3.4.1) and Theorem 3.6.1, it suffices to compute the 5-dimensional space $\mathbb{M}_{E_{g+2}^{(g)}}$. As in Section 3.4 this will require exhibiting particular triangulations giving rise to families of metric graphs. For the purposes of presenting such details, the hyperelliptic triangle is not especially user-friendly, as the triangulations become hard to parse visually. An explicit computation as in Section 3.3 reveals that the rectangle $E_1^{(3)} = R_{4,2}$ realizes precisely the same metric graphs as the triangle $E_5^{(3)}$. With this, Theorem 3.6.1 implies $\mathbb{M}_{3,\text{hyp}}^{\text{planar}} = \mathbb{M}_{R_{4,2}}$. To complete the computation of $\mathbb{M}_3^{\text{planar}}$ begun in Section 3.4, it thus suffices to analyze the more user-friendly rectangle $R_{4,2}$.

Table 3.4: Dimensions of the moduli cones \mathbb{M}_Δ for $R_{4,2}$ and $E_5^{(3)}$

$G \setminus \dim$	$R_{4,2}$				$E_5^{(3)}$			
	3	4	5	$\#\Delta$'s	3	4	5	$\#\Delta$'s
(020)	42	734	1296	2072	42	352	369	763
(111)		211	695	906		90	170	260
(212)			127	127			25	25
total	42	945	2118	3105	42	442	564	1048

It follows from Theorem 2.3.14 that $\mathbb{M}_{R_{4,2}}$ and \mathbb{M}_{T_4} have disjoint interiors. Moreover, $\mathbb{M}_{R_{4,2}}$ is not contained in \mathbb{M}_{T_4} . This highlights a crucial difference between (3.1.8) and (3.4.1). The former concerns the tropicalization of classical moduli spaces, so the hyperelliptic locus

lies in the closure of the nonhyperelliptic locus. The analogous statement is false for tropical plane curves. To see that \mathbb{M}_{T_4} does not contain $\mathbb{M}_{R_{4,2}}$ consider the (020) graph with all edge lengths equal to 1. By Theorems 3.4.1 and 3.6.3, this metric graph is in $\mathbb{M}_{R_{4,2}}$ but not in \mathbb{M}_{T_4} . What follows is the hyperelliptic analogue to the nonhyperelliptic Theorem 3.4.1.

Theorem 3.6.3. *A graph in \mathbb{M}_3 arises from $R_{4,2}$ if and only if it is one of the graphs (020), (111), or (212) in Figure 3.4, with edge lengths satisfying the following, up to symmetry:*

- (020) is realizable if and only if $w = x$, $v \leq u$, $v \leq y \leq z$, and

$$(y < v + 2w) \text{ or } (y = v + 2w \text{ and } y < z)$$

$$\text{or } (y < v + 3w \text{ and } u \leq 2v) \text{ or } (y = v + 3w \text{ and } u \leq 2v \text{ and } y < z)$$

$$\text{or } (y < v + 4w \text{ and } u = v) \text{ or } (y = v + 4w \text{ and } u = v \text{ and } y < z).$$
(3.6.2)
- (111) is realizable if and only if $w = x$ and $\min\{u, v\} \leq w$.
- (212) is realizable if and only if $w = x$.

Proof. This is based on an explicit computation as described in Section 3.3. The hyperelliptic rectangle $R_{4,2}$ has 3105 unimodular triangulations up to symmetry. All triangulations are regular. For each such triangulation we computed the graph G and the polyhedral cone \mathbb{M}_Δ . Each \mathbb{M}_Δ has dimension 3, 4, or 5, with census given on the left in Table 3.4. For each cone \mathbb{M}_Δ we then checked that the inequalities stated in Theorem 3.6.3 are satisfied. This proves that the dense realizable part of $\mathbb{M}_{R_{4,2}}$ is contained in the polyhedral space described by our constraints.




Figure 3.30: Triangulations giving all realizable hyperelliptic metrics for the graph (020)

For the converse direction, we construct a planar tropical realization of each metric graph that satisfies our constraints. For the graph (020), we consider eleven cases:

- | | |
|--|-----------|
| (i) $y < v + 2w$, $u \neq v$, $y \neq z$; | (dim = 5) |
| (ii) $y = v + 2w$, $u \neq v$, $y \neq z$; | (dim = 5) |
| (iii) $(y < v + 3w, v < u < 2v, y \neq z)$ or $(y < v + 2w, u \neq v, y < z < 2y)$; | (dim = 5) |
| (iv) $(y < v + 3w, u = 2v, y \neq z)$ or $(y < v + 2w, u \neq v, z = 2y)$; | (dim = 4) |
| (v) $(y < v + 3w, v < u < 2v, y = z)$ or $(y < v + 4w, u = v, y < z < 2z)$; | (dim = 4) |
| (vi) $(y < v + 3w, u = 2v, y = z)$ or $(y < v + 4w, u = v, z = 2y)$; | (dim = 3) |
| (vii) $y = v + 3w, v < u < 2v, y \neq z$; | (dim = 4) |
| (viii) $y = v + 3w, u = 2v, y \neq z$; | (dim = 3) |
| (ix) $(y < v + 4w, u = v, y \neq z)$ or $(y < v + 2w, y = z, u \neq v)$; | (dim = 3) |

- (x) $y < v + 4w, u = v, y = z;$ (dim = 3)
- (xi) $y = v + 4w, u = v, y \neq z.$ (dim = 3)

The disjunction of (i),(ii),...,(xi) is equivalent to (3.6.2). Triangulations giving all metric graphs satisfying each case are pictured in Figure 3.30. Next to the first triangulation is a metric graph arising from it.

Next we deal with graph **(111)**. Here we need two triangulations, one for $u \neq v$ and one for $u = v$. They are pictured in Figure 3.31. The left gives $u \neq v$, and the middle gives $u = v$.




Figure 3.31: Triangulations realizing hyperelliptic metrics for the graphs (111) and (212)

Finally, for the graph **(212)**, the single triangulation on the right in Figure 3.31 suffices. □

Chapter 4

Algorithms for Mumford curves

Since their initial appearance in the 1970s, a rich theory behind Mumford curves has been developed, largely in the 1980s in such works as [GvdP80]. However, prior to the work presented here, there have been few numerical algorithms for working with them (an exception being a treatment of hyperelliptic Mumford curves, mostly genus 2, in [Kad07] from 2007). We have designed and implemented algorithms that accomplish Mumford curve-based tasks over \mathbb{Q}_p previously absent from the realm of computation, and have made many seemingly theoretical and opaque objects hands-on and tractable.

We begin in Section 4.1 with a brief review of the theory of Mumford curves, and a summary of the goals of our algorithms. After discussing in Section 4.2 a technical hypothesis (“good position”) for the input for our algorithms, we present our main algorithms in Section 4.3. In Section 4.4, we present an algorithm to achieve the “good position” hypothesis that allows the other algorithms to run efficiently, which in doing so verifies that the input group is Schottky (or proves that the group is not Schottky). This is in some ways the most important result of this chapter, as the algorithms in Section 4.3 rely heavily upon it.

We made extensive use of the software package `sage` [S⁺13]. Supplementary files can be found at http://math.berkeley.edu/~ralph42/mumford_curves_supp.html. There are minor changes in the `sage` implementation from the description of the algorithms in this chapter. The changes are made only for convenience in implementation, and they do not affect the behavior of the algorithms.

This chapter’s content comes from the paper “Algorithms for Mumford curves” [MR15], coauthored with Qingchun Ren and appearing in the *Journal of Symbolic Computation*.

4.1 The theory of Mumford curves

Let K be an algebraically closed field complete with respect to a nontrivial non-Archimedean valuation. Unless otherwise stated, $|\cdot|$ will denote a choice of norm on K coming from this valuation. Let $R = \{x \in K \mid \text{val}(x) \geq 0\}$ be the valuation ring of K . This is a local ring with unique maximal ideal $M = \{x \in K \mid \text{val}(x) > 0\}$. Let $k = R/M$ denote the residue field

of K . We are most interested in the field of p -adic numbers \mathbb{Q}_p , which unfortunately is not algebraically closed. (For this case, $R = \mathbb{Z}_p$, the ring of p -adic integers, and $k = \mathbb{F}_p$, the field with p elements.) Therefore for theoretical purposes we will often consider $K = \mathbb{C}_p$, the complete algebraic closure of \mathbb{Q}_p . (In this case R is much larger, and k is the algebraic closure of \mathbb{F}_p .) In most of this chapter, choosing elements of \mathbb{C}_p that happen to be elements of \mathbb{Q}_p as inputs for algorithms yields an output once again in \mathbb{Q}_p . This “ \mathbb{Q}_p in, \mathbb{Q}_p out” property means we may take K to be \mathbb{Q}_p for our algorithmic purposes, while still considering $K = \mathbb{C}_p$ when more convenient for the purposes of theory. Much of the theory presented here works for other non-Archimedean fields, such as the field of Puiseux series $\mathbb{C}\{\{t\}\}$.

We recall some standard definitions and notation for p -adic numbers; for further background on the p -adics, see [Hol01]. For a prime p , the p -adic valuation $\text{val}_p : \mathbb{Q}^* \rightarrow \mathbb{Z}$ is defined by $\text{val}_p\left(p^v \frac{m}{n}\right) = v$, where m and n are not divisible by p . The usual p -adic norm $|\cdot|_p$ on \mathbb{Q} is defined for $a \in \mathbb{Q}^*$ by $|a|_p = \frac{1}{p^{\text{val}_p(a)}}$ and for 0 by $|0|_p = 0$. This means that large powers of p are small in absolute value, and small powers of p are large in absolute value. We will usually omit the subscript p from both $|\cdot|_p$ and val_p .

The completion of \mathbb{Q} with respect to the p -adic norm is denoted \mathbb{Q}_p , and is called the field of p -adic numbers. Each nonzero element b of \mathbb{Q}_p can be written uniquely as

$$b = \sum_{n=v}^{\infty} a_n p^n,$$

where $v \in \mathbb{Z}$, $a_v \neq 0$ and $a_n \in \{0, 1, \dots, p-1\}$ for all n . The p -adic valuation and norm extend to this field, and such a sum will have $\text{val}(b) = v$ and $|b| = \frac{1}{p^v}$. In analog to decimal expansions, we will sometimes write

$$b = \dots a_N a_{N-1} \dots a_3 a_2 a_1 a_0 . a_{-1} a_{-2} \dots a_v,$$

where the expression trails to the left since higher powers of p are smaller in p -adic absolute value. We may approximate $b \in \mathbb{Q}_p$ by a finite sum

$$b \approx \sum_{n=v}^N a_n p^n,$$

which will give an error of size at most $\frac{1}{p^{N+1}}$.

Consider the group $\text{PGL}(2, K)$, which acts on $\mathbb{P}^1(K)$ by treating elements as column vectors. That is, a matrix acts on the point $(a : b) \in \mathbb{P}^1(K)$ by multiplication with the vector $\begin{pmatrix} a \\ b \end{pmatrix}$. Viewed on an affine patch, the elements of this group act as fractional linear transformations. We are interested in the action of certain subgroups of $\text{PGL}(2, K)$ called *Schottky groups*, because a Schottky group minimally generated by $g \geq 2$ elements will give rise to a curve of genus g .

Definition 4.1.1. A 2×2 matrix over K is *hyperbolic* if it has two eigenvalues with different valuations. A *Schottky group* $\Gamma \leq \text{PGL}(2, K)$ is a finitely generated subgroup such that every non-identity element is hyperbolic.

There are many equivalent definitions of Schottky groups, including the following useful characterization.

Proposition 4.1.2. *A subgroup of $PGL(2, K)$ is Schottky if and only if it is free, discrete, and finitely generated.*

As remarked in the introduction of [Mum72b], if the matrices have entries in a locally compact field (like \mathbb{Q}_p), the definition of Schottky is equivalent to asking that Γ has no elements of finite order. So if we are working with generators in $\mathbb{Q}_p^{2 \times 2}$, we may replace “free” with “torsion free.” For an algorithm to test if a group is Schottky, see Section 4.4.

Let Γ be a Schottky group minimally generated by $\gamma_1, \dots, \gamma_g$. The above proposition implies that each element $\gamma \in \Gamma$ can be written as a unique shortest product $h_1 h_2 \cdots h_k$, where each $h_i \in \{\gamma_1, \dots, \gamma_g, \gamma_1^{-1}, \dots, \gamma_g^{-1}\}$. This product is called the *reduced word* for γ .

Let Σ be the set of points in $\mathbb{P}^1(K)$ that are fixed points of elements of Γ or limit points of the fixed points. The group Γ acts nicely on $\Omega := \mathbb{P}^1(K) \setminus \Sigma$; for this reason we will sometimes refer to Σ as *the set of bad points* for Γ .

Theorem 4.1.3 (Mumford, [Mum72b]). *Let $\Gamma = \langle \gamma_1, \dots, \gamma_g \rangle$ and Ω be as above. Then Ω/Γ is analytically isomorphic to a curve of genus g . We call such a curve a Mumford curve.*

Mumford also showed that a curve is Mumford if and only if it is smooth with totally degenerate reduction over the residue field. It follows that Mumford curves are precisely those smooth curves of genus g whose Berkovich skeleton is a graph with genus g .

In a companion paper to [Mum72b] (see [Mum72a]), Mumford also considered abelian varieties over non-Archimedean fields. He showed that these could be represented as $(K^*)^g/Q$, where $Q \in (K^*)^{g \times g}$ is called a *period matrix* for the abelian variety, and represents the multiplicative subgroup generated by its columns.

The algorithms in Section 4.3 accomplish the following tasks, where we denote Ω/Γ by C . The first and third algorithms are numerical in nature, and so will only give an approximation their outputs.

- Algorithm 4.3.3: Given a Schottky group Γ , find a period matrix Q for the abelian variety $\text{Jac}(C)$.
- Algorithm 4.3.9: Given a Schottky group Γ , find a triple (G, ℓ, h) , where
 - G is a graph,
 - ℓ is a length function on G such that the metric graph (G, ℓ) is the skeleton of C^{an} (the analytification of C), and
 - h is a natural equivalence $h : R^g \rightarrow G$ from the rose graph on g petals;

this data specifies a point in the tropical Teichmüller space described in [CMV13].

- Algorithm 4.3.13: Given a Schottky group Γ , find points in a canonical embedding of the curve C into \mathbb{P}^{g-1} .

Our algorithms take advantage of a property that makes non-Archimedean valued fields like \mathbb{Q}_p special: the ultrametric inequality $|x+y| \leq \max\{|x|, |y|\}$. As a result of this property, errors do not accumulate in the computation. Thus we avoid a dangerous hazard present in doing numerical computation over \mathbb{R} or \mathbb{C} . Nonetheless, our computational problems are hard in nature. Efficient computation for similar problems is not common in the literature even for genus 2 case. Our algorithms are capable of solving genus 2 and some genus 3 examples on a laptop in reasonable time (several minutes). However, they are less efficient for larger genera. The reason is that the running time grows exponentially as the requirement on the precision of the output (in terms of the number of digits) grows. A future goal is to find a way to reduce the running time for the algorithms.

4.2 Good fundamental domains

This section introduces *good fundamental domains* and the notion of *good position* for generators of Schottky groups, both of which will play key roles in our algorithms for Mumford curves. Our main algorithms in Section 4.3 require as input Schottky generators in good position, without which the rate of convergence of approximations will drop drastically. For our method of putting generators into good position, see Section 4.4.

We start with the usual projective line \mathbb{P}^1 , then discuss the analytic projective line $(\mathbb{P}^1)^{\text{an}}$. Our treatment of good fundamental domains follows Gerritzen and van der Put [GvdP80]. The notion is also discussed by Kadziela [Kad07]. The introduction to the analytic projective line follows Baker, Payne and Rabinoff [BPR12].

Definition 4.2.1. An *open ball* in \mathbb{P}^1 is either a usual open ball $B(a, r) = \{x \in K : |x - a| < r\}$ or the complement of a usual closed ball $B(a, r)^+ = \{x \in K : |x - a| \leq r\}$. A *closed ball* is either a usual closed ball or the complement of a usual open ball.

The open balls generate a topology on \mathbb{P}^1 . Both open balls and closed balls are simultaneously open and closed in this topology, as is the case for any non-Archimedean field due to the ultrametric inequality $|x + y| \leq \max\{|x|, |y|\}$. Let $|K^*|$ denote the image of K^* under $|\cdot|$. If $r \in |K^*|$, the open ball and the closed ball of radius r are distinguished by whether there exist two points x, y in the ball such that $|x - y|$ equals the diameter. The complement of an open ball is a closed ball, and vice versa.

Definition 4.2.2. A *good fundamental domain* $F \subset \mathbb{P}^1$ corresponding to the generators $\gamma_1, \dots, \gamma_g$ is the complement of $2g$ open balls $B_1, \dots, B_g, B'_1, \dots, B'_g$, such that corresponding closed balls $B_1^+, \dots, B_g^+, B'_1, \dots, B'_g$ are disjoint, and $\gamma_i(\mathbb{P}^1 \setminus B'_i) = B_i^+$ and $\gamma_i^{-1}(\mathbb{P}^1 \setminus B_i) = B_i'^+$ for all i . The *interior* of F is $F^\circ = \mathbb{P}^1 \setminus (B_1^+ \cup \dots \cup B_g^+ \cup B_1'^+ \cup \dots \cup B_g'^+)$. The *boundary* of F is $F \setminus F^\circ$.

The definition above implies that $\gamma_i(\mathbb{P}^1 \setminus B_i'^+) = B_i$ and $\gamma_i^{-1}(\mathbb{P}^1 \setminus B_i^+) = B_i'$ for all i .

Example 4.2.3. All the groups in this example are Schottky. This follows from the existence of a good fundamental domain [GvdP80].

(1) Let $K = \mathbb{C}_3$ and Γ be the group generated by

$$\gamma_1 = \begin{bmatrix} -5 & 32 \\ -8 & 35 \end{bmatrix}, \gamma_2 = \begin{bmatrix} -13 & 80 \\ -8 & 43 \end{bmatrix}.$$

Both matrices have eigenvalues 27 and 3. The matrix γ_1 has left eigenvectors $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 4 \\ 1 \end{pmatrix}$, and γ_2 has left eigenvectors $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 5 \\ 1 \end{pmatrix}$. We use the convention that $(z_1 : z_2) = z_1/z_2$. Then, $F = \mathbb{P}^1 \setminus (B_1 \cup B'_1 \cup B_2 \cup B'_2)$ where $B_1 = B(4, 1/9)$, $B'_1 = B(1, 1/9)$, $B_2 = B(5, 1/9)$, $B'_2 = B(2, 1/9)$ is a good fundamental domain relative to the generators γ_1 and γ_2 . This can be verified as follows. First rewrite

$$\gamma_1 z = \frac{-5z + 32}{-8z + 35} = 4 + \frac{27(z-1) - 81}{-8(z-1) + 27}.$$

Suppose that $z \in B'_1 = B(1, 1/9)$. Then, $\text{val}(27(z-1)) = 3 + \text{val}(z-1) \geq 3 + 2 = 5$, and $\text{val}(81) = 4$. So $\text{val}(27(z-1) - 81) = 4$. Also, $\text{val}(-8(z-1) + 27) \geq \min(\text{val}(8(z-1)), \text{val}(27)) > \min(2, 3) = 2$. So,

$$|\gamma_1 z - 4| = \left| \frac{27(z-1) - 81}{-8(z-1) + 27} \right| > \frac{3^{-4}}{3^{-2}} = 1/9.$$

So $\gamma_1(B'_1) \subset \mathbb{P}^1 \setminus B_1^+$. The other three conditions can be verified similarly.

(2) Let $K = \mathbb{C}_3$ and Γ be the group generated by

$$\gamma_1 = \begin{bmatrix} -79 & 160 \\ -80 & 161 \end{bmatrix}, \gamma_2 = \begin{bmatrix} -319 & 1600 \\ -80 & 401 \end{bmatrix}.$$

Both matrices have eigenvalues 81 and 1. The matrix γ_1 has left eigenvectors $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$, and the matrix γ_2 has left eigenvectors $\begin{pmatrix} 4 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 5 \\ 1 \end{pmatrix}$. Then, $F = \mathbb{P}^1 \setminus (B_1 \cup B'_1 \cup B_2 \cup B'_2)$ where $B_1 = B(2, 1/9)$, $B'_1 = B(1, 1/9)$, $B_2 = B(5, 1/9)$, $B'_2 = B(4, 1/9)$ is a good fundamental domain relative to the generators γ_1 and γ_2 .

(3) Let $K = \mathbb{C}_3$, and let Γ be the group generated by

$$\gamma_1 = \begin{bmatrix} 121 & -120 \\ 40 & -39 \end{bmatrix}, \gamma_2 = \begin{bmatrix} 121 & -240 \\ 20 & -39 \end{bmatrix}, \gamma_3 = \begin{bmatrix} 401 & -1600 \\ 80 & -319 \end{bmatrix}.$$

All three generators have eigenvalues 1 and 3^4 . The element γ_1 has eigenvectors $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 3 \\ 1 \end{pmatrix}$. The element γ_2 has eigenvectors $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 6 \\ 1 \end{pmatrix}$. The element γ_3 has eigenvectors $\begin{pmatrix} 4 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 5 \\ 1 \end{pmatrix}$. Then, $F = \mathbb{P}^1 \setminus (B_1 \cup B'_1 \cup B_2 \cup B'_2 \cup B_3 \cup B'_3)$ where $B_1 = B(1, 1/9)$, $B'_1 = B(3, 1/9)$, $B_2 = B(2, 1/9)$, $B'_2 = B(6, 1/9)$, $B_3 = B(4, 1/9)$, $B'_3 = B(5, 1/9)$ is a good fundamental domain relative to the generators γ_1 , γ_2 and γ_3 .

The following lemma follows from Definition 4.2.2 by induction (see [Kad07, Theorem 6.2]).

Lemma 4.2.4. *Let F and $\gamma_1, \dots, \gamma_g$ be as in Definition 4.2.2, and let $\gamma \in \Gamma \setminus \left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right\}$ and $b \in \mathbb{P}^1(K)$. Write the reduced word for γ as $h_1 h_2 \cdots h_k$, where $k \geq 1$ and $h_i \in \{\gamma_1^\pm, \dots, \gamma_g^\pm\}$ for all i . Assume that $b \notin B'_j$ if $h_k = \gamma_j$ and $b \notin B_j$ if $h_k = \gamma_j^{-1}$. Then we have*

$$\gamma b \in \begin{cases} B_i^+, & \text{if } h_1 = \gamma_i, \\ B_i'^+, & \text{if } h_1 = \gamma_i^{-1}. \end{cases}$$

Proof. To simplify notation we will outline the proof for the case where $h_i \in \{\gamma_1, \dots, \gamma_g\}$ for all i , and then describe how to generalize to the case of $h_i \in \{\gamma_1^\pm, \dots, \gamma_g^\pm\}$.

Write $h_i = \gamma_{a_i}$ for each i . Since $h_k = \gamma_{a_k}$, we know by assumption that $b \notin B'_{a_k}$. By Definition 4.2.2 we have $\gamma_{a_k}(\mathbb{P}^1 \setminus B'_{a_k}) = B_{a_k}^+$, so $h_k b \in B_{a_k}^+$. By the disjointness of the $2g$ closed balls, we know that $h_k b \notin B'_{a_{k-1}}$, and since $\gamma_{a_{k-1}}(\mathbb{P}^1 \setminus B'_{a_{k-1}}) = B_{a_{k-1}}^+$, we have $h_{k-1} h_k b \in B_{a_{k-1}}^+$. We may continue in this fashion until we find that $h_1 h_2 \cdots h_k b \in B_{a_1}^+$.

The only possible obstruction to the above argument in the case of $h_i \in \{\gamma_1^\pm, \dots, \gamma_g^\pm\}$ occurs if $h_i \cdots h_k b \in B_{a_i}^+$ and $h_{i-1} = \gamma_{a_i}$ (or, similarly, if $h_i \cdots h_k b \in B_{a_i}^+$ and $h_{i-1} = \gamma_{a_i}^{-1}$), since the above argument needs γ_{a_i} to act on $\mathbb{P}^1 \setminus B_{a_i}^+$. However, this situation arises precisely when $h_i = \gamma_{a_i}^{-1} = h_{i-1}^{-1}$, meaning that the word is not reduced. Since we've assumed $h_1 \cdots h_k$ is reduced, we have the desired result. \square

For a fixed set of generators of Γ , there need not exist a good fundamental domain. If there exists a good fundamental domain for some set of free generators of Γ , we say that the generators are *in good position*. Gerritzen and van der Put [GvdP80, §I.4] proved that there always exists a set of generators in good position. They also proved the following desirable properties for good fundamental domains.

Theorem 4.2.5. *Let Γ be a Schottky group, Σ its set of bad points, and $\Omega = \mathbb{P}^1 \setminus \Sigma$.*

- (1) *There exists a good fundamental domain F for some set of generators $\gamma_1, \dots, \gamma_g$ of Γ .*
- (2) *If $\gamma \neq \text{id}$, then $\gamma F^\circ \cap F = \emptyset$.*
- (3) *If $\gamma \notin \{\text{id}, \gamma_1, \dots, \gamma_g, \gamma_1^{-1}, \dots, \gamma_g^{-1}\}$, then $\gamma F \cap F = \emptyset$.*
- (4) $\cup_{\gamma \in \Gamma} \gamma F = \Omega$.

The statements (2), (3), and (4) imply that Ω/Γ can be obtained from F by gluing the boundary of F . More specifically, $B_i^+ \setminus B_i$ is glued with $B_i'^+ \setminus B_i'$ via the action of γ_i . We have designed the following subroutine, which takes any point p in Ω and finds a point q in F such that they are equivalent modulo the action of Γ . This subroutine is useful in developing the algorithms in Section 3 and 4.

Subroutine 4.2.6 (Reducing a point into a good fundamental domain).

Require: Matrices $\gamma_1, \dots, \gamma_g$ generating a Schottky group Γ , a good fundamental domain $F = \mathbb{P}^1 \setminus (B_1 \cup \cdots \cup B_g \cup B_1' \cup \cdots \cup B_g')$ associated to these generators, and a point $p \in \Omega$.

Ensure: A point $q \in F$ and an element $\gamma \in \Gamma$ such that $q = \gamma p$.

- 1: Let $q \leftarrow p$ and $\gamma \leftarrow \text{id}$.
- 2: **while** $q \notin F$ **do**
- 3: If $q \in B'_i$, let $q \leftarrow \gamma_i q$ and $\gamma \leftarrow \gamma_i \gamma$.
- 4: Otherwise, if $q \in B_i$, let $q \leftarrow \gamma_i^{-1} q$ and $\gamma \leftarrow \gamma_i^{-1} \gamma$.
- 5: **end while**
- 6: **return** q and γ .

Proof. The correctness of this subroutine is clear. It suffices to prove that the algorithm always terminates. Given $p \in \Omega$, if $p \notin F$, by Theorem 4.2.5, there exists $\gamma^\circ = h_1 h_2 \cdots h_k \in \Gamma$ (where each h_j is γ_i or γ_i^{-1} for some i) such that $\gamma^\circ p \in F$. Without loss of generality, we may assume that γ° is chosen such that k is the smallest. Steps 3,4 and Lemma 4.2.4 make sure that we always choose $q \leftarrow h_k q$ and $\gamma \leftarrow h_k \gamma$. Therefore, this subroutine terminates with $\gamma = \gamma^\circ$. \square

We can extend the definition of good fundamental domains to the analytic projective line $(\mathbb{P}^1)^{\text{an}}$. As discussed in Section 1.4, analytification of an algebraic variety is defined in terms of multiplicative seminorms. For our special case $(\mathbb{P}^1)^{\text{an}}$, there is a simpler description. As detailed in [Bak07], the Berkovich projective line $(\mathbb{P}^1)^{\text{an}}$ consists of four types of points:

- Type 1 points are just the usual points of \mathbb{P}^1 .
- Type 2 points correspond to closed balls $B(a, r)^+$ where $r \in |K^\times|$.
- Type 3 points correspond to closed balls $B(a, r)^+$ where $r \notin |K^\times|$.
- Type 4 points correspond to equivalence classes of sequences of nested closed balls $B_1^+ \supset B_2^+ \supset \cdots$ such that their intersection is empty.

There is a metric on the set of Type 2 and Type 3 points, defined as follows: let P_1 and P_2 be two such points and let $B(a_1, r_1)^+$ and $B(a_2, r_2)^+$ be the corresponding closed balls.

- (1) If one of them is contained in the other, say $B(a_1, r_1)^+$ is contained in $B(a_2, r_2)^+$, then the distance $d(P_1, P_2)$ is $\log_p(r_2/r_1)$.
- (2) In general, there is a unique smallest closed ball $B(a_3, r_3)^+$ containing both of them. Let P_3 be the corresponding point. Then, $d(P_1, P_2)$ is defined to be $d(P_1, P_3) + d(P_3, P_2)$.

This metric makes $(\mathbb{P}^1)^{\text{an}}$ a tree with infinite branching, as we now describe. There is a unique path connecting any two points P_1 and P_2 . In case (1) above, the path is defined by the isometry $t \mapsto B(a_1, p^t)^+$, $t \in [\log(r_1), \log(r_2)]$. It is straightforward to check that $B(a_1, r_2)^+ = B(a_2, r_2)^+$. In case (2) above, the path is the concatenation of the paths from P_1 to P_3 and from P_3 to P_2 . Then, Type 1 points become limits of Type 2 and Type 3 points with respect to this metric. More precisely, if $x \neq \infty$, then it lies at the limit of the path $t \mapsto B(x, p^{-t})^+$, $t \in [0, +\infty)$. Type 1 points behave like leaves of the tree at infinity. For any

two Type 1 points x, y , there is a unique path in $(\mathbb{P}^1)^{\text{an}}$ connecting them, which has infinite length.

Definition 4.2.7. Let Σ be a discrete subset in \mathbb{P}^1 . The subtree of $(\mathbb{P}^1)^{\text{an}}$ *spanned by* Σ , denoted $T(\Sigma)$, is the union of all paths connecting all pairs of points in Σ .

An *analytic open ball* $B(a, r)^{\text{an}}$ is a subset of $(\mathbb{P}^1)^{\text{an}}$ whose set of Type 1 points is just $B(a, r)$ and whose Type 2, 3, and 4 points correspond to closed balls $B(a', r')^+ \subset B(a, r)$ and the limits of sequences of such closed balls. An *analytic closed ball* is similar, with $B(a, r)$ replaced with $B(a, r)^+$. Just as in the case of balls in \mathbb{P}^1 , the analytic closed ball $(B^+)^{\text{an}}$ is not the closure of B^{an} in the metric topology of $(\mathbb{P}^1)^{\text{an}}$. The complement of an analytic open ball is an analytic closed ball, and vice versa. In an analytic closed ball $(B(a, r)^+)^{\text{an}}$ such that $r \in |K^\times|$, the *Gaussian point* is the Type 2 point corresponding to $B(a, r)^+$. An *analytic annulus* is $B \setminus B'$, where B and B' are analytic balls such that $B' \subsetneq B$. If B is an analytic open (resp. closed) ball and B' is an analytic closed (resp. open) ball, then $B \setminus B'$ is an *analytic open annulus* (resp. *analytic closed annulus*). A special case of analytic open annulus is the complement of a point in an analytic open ball.

Any element of $\text{PGL}(2, K)$ sends open balls to open balls and closed balls to closed balls. Thus, there is a well defined action of $\text{PGL}(2, K)$ on $(\mathbb{P}^1)^{\text{an}}$.

Definition 4.2.8. A *good fundamental domain* $F \subset (\mathbb{P}^1)^{\text{an}}$ corresponding to the generators $\gamma_1, \dots, \gamma_g$ is the complement of $2g$ analytic open balls $B_1^{\text{an}}, \dots, B_g^{\text{an}}, B_1'^{\text{an}}, \dots, B_g'^{\text{an}}$, such that the corresponding analytic closed balls $(B_1^+)^{\text{an}}, \dots, (B_g^+)^{\text{an}}, (B_1'^+)^{\text{an}}, \dots, (B_g'^+)^{\text{an}}$ are disjoint, and that $\gamma_i((\mathbb{P}^1)^{\text{an}} \setminus B_i'^{\text{an}}) = (B_i^+)^{\text{an}}$ and $\gamma_i^{-1}((\mathbb{P}^1)^{\text{an}} \setminus B_i^{\text{an}}) = (B_i'^+)^{\text{an}}$. The *interior* of F is $F^\circ = (\mathbb{P}^1)^{\text{an}} \setminus ((B_1^+)^{\text{an}} \cup \dots \cup (B_g^+)^{\text{an}} \cup (B_1'^+)^{\text{an}} \cup \dots \cup (B_g'^+)^{\text{an}})$. The *boundary* of F is $F \setminus F^\circ$.

Definition 4.2.8 implies that $\gamma_i((\mathbb{P}^1)^{\text{an}} \setminus (B_i'^+)^{\text{an}}) = B_i^{\text{an}}$ and $\gamma_i^{-1}((\mathbb{P}^1)^{\text{an}} \setminus (B_i^+)^{\text{an}}) = B_i'^{\text{an}}$.

We now argue that there is a one-to-one correspondence between good fundamental domains in \mathbb{P}^1 and good fundamental domains in $(\mathbb{P}^1)^{\text{an}}$. (This fact is well-known, though seldom explicitly stated in the literature; for instance, it's taken for granted in the later chapters of [GvdP80].) If $\mathbb{P}^1 \setminus (B_1 \cup \dots \cup B_g \cup B_1' \cup \dots \cup B_g')$ is a good fundamental domain in \mathbb{P}^1 , then $(\mathbb{P}^1)^{\text{an}} \setminus (B_1^{\text{an}} \cup \dots \cup B_g^{\text{an}} \cup B_1'^{\text{an}} \cup \dots \cup B_g'^{\text{an}})$ is a good fundamental domain in $(\mathbb{P}^1)^{\text{an}}$. Indeed, since the closed balls $B_1^+, \dots, B_g^+, B_1'^+, \dots, B_g'^+$ are disjoint, and the corresponding analytic closed balls consist of points corresponding to closed balls contained in $B_1^+, \dots, B_g^+, B_1'^+, \dots, B_g'^+$ and their limits, the analytic closed balls are also disjoint. Conversely, if $(\mathbb{P}^1)^{\text{an}} \setminus (B_1^{\text{an}} \cup \dots \cup B_g^{\text{an}} \cup B_1'^{\text{an}} \cup \dots \cup B_g'^{\text{an}})$ is a good fundamental domain in $(\mathbb{P}^1)^{\text{an}}$, then $\mathbb{P}^1 \setminus (B_1 \cup \dots \cup B_g \cup B_1' \cup \dots \cup B_g')$ is a good fundamental domain in \mathbb{P}^1 , because the classical statement can be obtained from the analytic statement by considering only Type 1 points. This correspondence allows us to abuse notation by not distinguishing the classical case and the analytic case. Theorem 4.2.5 is also true for analytic good fundamental domains.

Another analytic object of interest to us is the (Berkovich) skeleton of the analytification of the genus g curve Ω/Γ (a task that is part of Algorithm 4.3.9), so we close this section

with background information on this object. The following definitions are taken from Baker, Payne and Rabinoff [BPR12], with appropriate simplification.

Definition 4.2.9. (1) The *skeleton* of an open annulus $B \setminus B'$ is the straight path between the Gaussian point of B^+ and the Gaussian point of B'^+ .

- (2) Let C be a complete smooth curve over K . A *semistable vertex set* V is a finite set of Type 2 points in C^{an} such that $C^{\text{an}} \setminus V$ is the disjoint union of open balls and open annuli. The *skeleton corresponding to V* is the union of V with all skeletons of these open annuli.
- (3) If $\text{genus}(C) \geq 2$, then C^{an} has a unique *minimal skeleton*. The minimal skeleton is the intersection of all skeletons. If $\text{genus}(C) \geq 2$ and C is complete, then the minimal skeleton is a finite metric graph. We call this minimal skeleton the *Berkovich skeleton* of C^{an} .

Definition 4.2.10. An *algebraic semistable model* of a smooth curve C over K is a flat and proper scheme X over R whose generic fiber X_K is isomorphic to C and whose special fiber X_k satisfies

- X_k is a connected and reduced curve, and
- all singularities of X_k are ordinary double points.

Work towards algorithmic computation of semistable models is discussed in [AW12, §1.2] and [BW14, §3.1], though such a computation is in general hard to carry out.

Semistable models are related to skeletons in the following way: take a semistable model X of C . Associate a vertex for each irreducible component of X_k . For each ordinary intersection of two irreducible components in X_k , connect an edge between the two corresponding vertices. The resulting graph is combinatorially a skeleton of C^{an} .

4.3 Computations from good starting data

If we have a Schottky group $\Gamma = \langle \gamma_1, \dots, \gamma_g \rangle$ in terms of its generators, there are many objects we wish to compute for the corresponding curve Ω/Γ , such as the Jacobian of the curve, the Berkovich skeleton of the curve, and a canonical embedding for the curve. In this section we present algorithms for numerically computing these three objects, given the input of a Schottky group with generators in good position. For an algorithm that puts arbitrary generators of a Schottky group into good position, see Section 4.4.

Remark 4.3.1. Several results in this section are concerned with the accuracy of numerical approximations. Most of our results will be of the form

$$\left| \frac{\text{estimate}}{\text{actual}} - 1 \right| = \text{size of error term} \leq \text{a small real number of the form } p^{-N},$$

where we think of $N \gg 0$. This is equivalent to

$$\frac{\text{estimate}}{\text{actual}} - 1 = \text{error term} = \text{a } p\text{-adic number of the form } bp^N,$$

where $|b| \leq 1$. So, since $|p^N| = p^{-N}$, the *size* of the error term is a *small* power of p , while the *error term itself* is a *large* power of p (possibly with a constant that doesn't matter much).

Rearranging the second equation gives

$$\text{estimate} = \text{actual} + \text{actual} \cdot bp^N,$$

meaning that we are considering not the *absolute* precision of our estimate, but rather the *relative* precision. In this case we would say that our estimate is of *relative precision* $O(p^N)$. So if we desire relative precision $O(p^N)$, we want the *actual* error term to be p^N (possibly with a constant term with nonnegative valuation), and the *size* of the error term to be at most p^{-N} .

Given a Schottky group $\Gamma = \langle \gamma_1, \dots, \gamma_g \rangle$, we wish to find a period matrix Q so that $\text{Jac}(\Omega/\Gamma) \cong (K^*)^g/Q$. For previous work on this computation in the genus-2 case, see [Tei88].

First we will set some notation. For any parameters $a, b \in \Omega$, we introduce the following analytic function in the unknown z , called a *theta function*:

$$\Theta(a, b; z) := \prod_{\gamma \in \Gamma} \frac{z - \gamma a}{z - \gamma b}.$$

Note that if Γ is defined over \mathbb{Q}_p and $a, b, z \in \mathbb{Q}_p$, then $\Theta(a, b; z) \in \mathbb{Q}_p \cup \{\infty\}$. (This is an instance of “ \mathbb{Q}_p in, \mathbb{Q}_p out.”) For any $\alpha \in \Gamma$ and $a \in \Omega$, we can specialize to

$$u_\alpha(z) := \Theta(a, \alpha a; z).$$

It is shown in [GvdP80, II.3] that the function $u_\alpha(z)$ is in fact independent of the choice of a . This is because for any choice of $a, b \in \Omega$ we have

$$\begin{aligned} \frac{\Theta(a, \alpha a; z)}{\Theta(b, \alpha b; z)} &= \prod_{\gamma \in \Gamma} \left(\frac{z - \gamma a}{z - \gamma \alpha a} \frac{z - \gamma \alpha b}{z - \gamma b} \right) = \prod_{\gamma \in \Gamma} \left(\frac{z - \gamma a}{z - \gamma b} \frac{z - \gamma \alpha b}{z - \gamma \alpha a} \right) \\ &= \prod_{\gamma \in \Gamma} \frac{z - \gamma a}{z - \gamma b} \cdot \prod_{\gamma \in \Gamma} \frac{z - \gamma \alpha b}{z - \gamma \alpha a} = \prod_{\gamma \in \Gamma} \frac{z - \gamma a}{z - \gamma b} \cdot \prod_{\gamma \in \Gamma} \frac{z - \gamma b}{z - \gamma a} \\ &= \Theta(a, b; z) \cdot \Theta(b, a; z) = 1. \end{aligned}$$

From [GvdP80, VI.2] we have a formula for the period matrix Q of $\text{Jac}(\Omega/\Gamma)$:

Theorem 4.3.2. *The period matrix Q for $\text{Jac}(\Omega/\Gamma)$ is given by*

$$Q_{ij} = \frac{u_{\gamma_i}(z)}{u_{\gamma_i}(\gamma_j z)},$$

where z is any point in Ω .

As shown in [GvdP80, II.3], the choice of z does not affect the value of Q_{ij} .

Theorem 4.3.2 implies that in order to compute each Q_{ij} , it suffices to find a way to compute $\Theta(a, b; z)$. Since a theta function is defined as a product indexed by the infinite group Γ , approximation will be necessary. Recall that each element γ in the free group generated by $\gamma_1, \dots, \gamma_g$ can be written in a unique shortest product $h_1 h_2 \dots h_k$ called the reduced word, where each $h_i \in \{\gamma_1, \dots, \gamma_g, \gamma_1^{-1}, \dots, \gamma_g^{-1}\}$. We can approximate $\Theta(a, b; z)$ by replacing the product over Γ with a product over Γ_m , the set of elements of Γ whose reduced words have length $\leq m$. More precisely, we approximate $\Theta(a, b; z)$ with

$$\Theta_m(a, b; z) := \prod_{\gamma \in \Gamma_m} \frac{z - \gamma a}{z - \gamma b},$$

where

$$\Gamma_m = \{h_1 h_2 \dots h_k \mid 0 \leq k \leq m, h_i \in \{\gamma_1^\pm, \dots, \gamma_g^\pm\}, h_i \neq h_{i+1}^{-1} \text{ for any } i\}.$$

With this approximation method, we are ready to describe an algorithm for computing Q .

Algorithm 4.3.3 (Period Matrix Approximation).

Require: Matrices $\gamma_1, \dots, \gamma_g \in \mathbb{Q}_p^{2 \times 2}$ generating a Schottky group Γ in good position, and an integer n to specify the desired relative precision.

Ensure: An approximation for a period matrix Q for $\text{Jac}(\Omega/\Gamma)$ up to relative precision $O(p^n)$.

- 1: Choose suitable p -adic numbers a and z as described in Theorem 4.3.6.
- 2: Based on n , choose a suitable positive integer m as described in Remark 4.3.7.
- 3: **for** $1 \leq i, j \leq g$ **do**
- 4: Compute $Q_{ij} = \Theta_m(a, \gamma_i(a); z) / \Theta_m(a, \gamma_j(a); z)$
- 5: **end for**
- 6: **return** Q .

The complexity of this algorithm is in the order of the number of elements in Γ_m , which is exponential in m . The next issue is that to achieve certain precision in the final result, we need to know how large m needs to be. Given a good fundamental domain F for the generators $\gamma_1, \dots, \gamma_g$, we are able to give an upper bound on the error in our estimation of Θ by Θ_m . (Algorithm 4.3.3 would work even if the given generators were not in good position, but would in general require a very large m to give the desired convergence. See Example 4.3.8(4).)

To analyze the convergence of the infinite product

$$\Theta(a, \gamma_i(a); z) = \prod_{\gamma \in \Gamma} \frac{z - \gamma a}{z - \gamma \gamma_i a},$$

we need to know where γa and $\gamma\gamma_i a$ lie. We can determine this by taking the metric of $(\mathbb{P}^1)^{\text{an}}$ into consideration. Assume that ∞ lies in the interior of F . Let $S = \{P_1, \dots, P_g, P'_1, \dots, P'_g\}$ be the set of points corresponding to the set of closed balls $\{B_1^+, \dots, B_g^+, B_1'^+, \dots, B_g'^+\}$ from the characterization of the good fundamental domain. Let c be the smallest pairwise distance between these points. This distance c will be key for determining our choice of m in the algorithm.

Proposition 4.3.4. *Let F , S , and c be as above. Suppose the reduced word for γ is $h_1 h_2 \cdots h_k$, where $k \geq 0$. Then $d(\gamma P_i, S) \geq kc$ for all i unless $h_k = \gamma_i^{-1}$, and $d(\gamma P'_i, S) \geq kc$ unless $h_k = \gamma_i$.*

Proof. We will prove this proposition by induction. If $k = 0$, there is nothing to prove. Let $k > 0$, and assume that the claim holds for all integers n with $0 \leq n < k$. Without loss of generality, we may assume $h_1 = \gamma_1$. Let B^+ be the closed disk corresponding to P_i . By Lemma 4.2.4, we have $\gamma(B^+) \subset B_1$. This means P_1 lies on the unique path from γP_i to ∞ . Since we assumed $\infty \in F$, p_1 lies on the unique path from γP_i to any point in S . Thus,

$$\begin{aligned} d(\gamma P_i, S) &= d(\gamma P_i, P_1) \\ &= d(\gamma_1^{-1} \gamma P_i, \gamma_1^{-1} P_1) \\ &= d(h_2 h_3 \cdots h_k P_i, P'_1). \end{aligned}$$

Let $P = P_j$ if $h_2 = \gamma_j$ and $P = P'_j$ if $h_2 = \gamma_j^{-1}$. By the same argument as above, P lies on the unique path from $h_2 h_3 \cdots h_k P_i$ to P'_1 . The reducedness of the word $h_1 h_2 \cdots h_k$ guarantees that $P \neq P'_1$. So

$$\begin{aligned} d(\gamma P_i, S) &= d(h_2 h_3 \cdots h_k P_i, P'_1) \\ &= d(h_2 h_3 \cdots h_k P_i, P) + d(P, P'_1) \\ &\geq (k-1)c + c = kc. \end{aligned}$$

The last step follows from the inductive hypothesis. The proof of the second part of this proposition is similar. \square

Proposition 4.3.5. *Let F , S , and c be as above. Let $z \in F$ and $a \in B_i^+ \setminus B_i'$ such that a , z , and ∞ are distinct modulo the action of Γ . Suppose the reduced word for $\gamma \in \Gamma$ is $h_1 h_2 \cdots h_k$. If $k \geq 2$ and $h_k \neq \gamma_i^{-1}$, then*

$$\left| \frac{z - \gamma a}{z - \gamma\gamma_i a} - 1 \right| \leq p^{-c(k-1)}.$$

Proof. Our choice of a guarantees that both a and $\gamma_i a$ are in F . Without loss of generality, we may assume that $h_k = \gamma_1$. Then, both $h_k a$ and $h_k \gamma_i a$ are in B_1^+ . So both γa and $\gamma\gamma_i a$ lie in $h_1 h_2 \cdots h_{k-1} B_1^+$, which is contained in some $B = B_j$ or B'_j . By Proposition 4.3.4, the points in $(\mathbb{P}^1)^{\text{an}}$ corresponding to the disks $h_1 h_2 \cdots h_{k-1} B_1^+$ and B^+ have distance at least

$c(k-1)$. This implies $\text{diam}(h_1 h_2 \cdots h_{k-1} B_1^+) \leq p^{-c(k-1)} \text{diam}(B^+)$. Therefore, $|\gamma a - \gamma \gamma_i a| \leq p^{-c(k-1)} \text{diam}(B^+)$. On the other hand, since $z \notin B$ and $\gamma \gamma_i a \in B$, we have $|z - \gamma \gamma_i a| \geq \text{diam}(B^+)$. This means that

$$\left| \frac{z - \gamma a}{z - \gamma \gamma_i a} - 1 \right| = \left| \frac{\gamma \gamma_i a - \gamma a}{z - \gamma \gamma_i a} \right| \leq \frac{p^{-c(k-1)} \text{diam}(B^+)}{\text{diam}(B^+)} = p^{-c(k-1)},$$

as claimed. \square

We are now ready to prove our approximation theorem, which is a new result that allows one to determine the accuracy of an approximation of a ratio of theta functions. It is similar in spirit to [Kad07, Theorem 6.10], which is an approximation result for a particular subclass of Schottky groups called *Whittaker groups*, corresponding to hyperelliptic Mumford curves. Our result is more general, as there are many Schottky groups that are not Whittaker.

Theorem 4.3.6. *Suppose that the given generators $\gamma_1, \dots, \gamma_g$ of Γ are in good position, with corresponding good fundamental domain F and disks $B_1, \dots, B_g, B'_1, \dots, B'_g$. Let $m \geq 1$. In Algorithm 4.3.3, if we choose $a \in B_i^+ \setminus B'_i$ and $z \in B_j^+ \setminus B'_j$ such that $a \neq z$, then*

$$\left| \frac{\Theta_m(a, \gamma_i(a); z) / \Theta_m(a, \gamma_i(a); \gamma_j(z))}{\Theta(a, \gamma_i(a); z) / \Theta(a, \gamma_i(a); \gamma_j(z))} - 1 \right| \leq p^{-cm},$$

where c is the constant defined above.

Proof. Our choice of z guarantees that both z and $\gamma_j z$ are in F . Thus, if ∞ lies in the interior of F , then this theorem follows directly from Proposition 4.3.5. The last obstacle is to remove the assumption on ∞ . We observe that Q_{ij} is a product of cross ratios:

$$\frac{\Theta(a, \gamma_i a; z)}{\Theta(a, \gamma_i a; \gamma_j z)} = \prod_{\gamma \in \Gamma} \frac{(z - \gamma a)(\gamma_j z - \gamma \gamma_i a)}{(z - \gamma \gamma_i a)(\gamma_j z - \gamma a)}.$$

Therefore, each term is invariant under any projective automorphism of \mathbb{P}^1 . Under such an automorphism, any point in the interior of F can be sent to ∞ . \square

Remark 4.3.7. If we wish to use Algorithm 4.3.3 to compute a period matrix Q with relative precision $O(p^n)$ (meaning that we want $p^{-cm} \leq p^{-n}$ in Theorem 4.3.6), we must first compute c . As above, c is defined to be the minimum distance between pairs of the points $P_1, \dots, P_g, P'_1, \dots, P'_g \in (\mathbb{P}^1)^{\text{an}}$ corresponding to the balls $B_1, \dots, B_g, B'_1, \dots, B'_g$ that characterize our good fundamental domain. Once we have computed c (perhaps by finding a good fundamental domain using the methods of Section 4.4), then by Theorem 4.3.6 we must choose m such that $cm \geq n$, so $m = \lceil n/c \rceil$ will suffice.

As a special case of the approximation theorem, suppose that we want to compute the period matrix for the tropical Jacobian of C , which is the matrix $(\text{val}(Q_{ij}))_{g \times g}$. We need only to compute Q_{ij} up to relative precision $O(1)$. Thus, setting $m = 0$ suffices. In this case, each of the products $\Theta_m(a, \gamma_i(a); z)$, $\Theta_m(a, \gamma_i(a); \gamma_j(z))$ has only one term.

Example 4.3.8. (1) Let Γ be the Schottky group in Example 4.2.3(1). Choose the same good fundamental domain, with $B_1 = B(4, 1/9)$, $B'_1 = B(1, 1/9)$, $B_2 = B(5, 1/9)$, and $B'_2 = B(2, 1/9)$. The four balls correspond to four points in the tree $(\mathbb{P}^1)^{\text{an}}$. We need to find the pairwise distances between the points P_1, P'_1, P_2 , and P'_2 in $(\mathbb{P}^1)^{\text{an}}$. Since the smallest ball containing both B_1^+ and B'_1^+ is $B^+(1, 1/3)$, both P_1 and P'_1 are distance $\text{val}((1/3)/(1/9)) = \text{val}(3) = 1$ from the point corresponding to $B^+(1, 1/3)$, so P_1 and P'_1 are distance 2 from one another. Similar calculations give distances of 2 between P_2 and P'_2 , and of 4 between P_1 or P'_1 and P_2 or P'_2 . In fact, the distance between P_i and P'_i equals the difference in the valuations of the two eigenvalues of γ_i . This allows us to construct the subtree of $(\mathbb{P}^1)^{\text{an}}$ spanned by P_1, P_2, P'_1, P'_2 as illustrated in Figure 4.1. The minimum distance between them is $c = 2$. To approximate Q_{11} , we take $a = 10$ and $z = 19$. To compute Q up to relative precision $O(p^{10})$, we need $2m \geq 10$ (this is the equation $cm \geq n$ from Remark 4.3.7), so choosing $m = 5$ works. The output of the algorithm is $Q_{11} = (\dots 220200000100)_3$. Similarly, we can get the other entries in the matrix Q :

$$Q = \begin{bmatrix} (\dots 220200000100)_3 & (\dots 0101010101)_3 \\ (\dots 0101010101)_3 & (\dots 220200000100)_3 \end{bmatrix}.$$

(2) Let Γ be the Schottky group in Example 4.2.3(2). Choose the same good fundamental domain. Again, we need $m = 5$ for relative precision $O(p^{10})$. The algorithm outputs

$$Q = \begin{bmatrix} (\dots 12010021010000)_3 & (\dots 002000212200)_3 \\ (\dots 002000212200)_3 & (\dots 12010021010000)_3 \end{bmatrix}.$$

(3) Let Γ be the Schottky group in Example 4.2.3(3). Choose the same good fundamental domain. The minimum distance between the corresponding points in $(\mathbb{P}^1)^{\text{an}}$ is 2, so we may take $m = 10/2 = 5$ to have relative precision up to $O(p^{10})$. Our algorithm outputs

$$Q = \begin{bmatrix} (\dots 11201000010000)_3 & (\dots 12020022210)_3 & (\dots 20020002120)_3 \\ (\dots 12020022210)_3 & (\dots 10101010010000)_3 & (\dots 020201120.1)_3 \\ (\dots 20020002120)_3 & (\dots 020201120.1)_3 & (\dots 21010100010000)_3 \end{bmatrix}.$$

(4) Let $K = \mathbb{C}_3$ and Γ be the group generated by

$$\gamma_1 = \begin{bmatrix} -5 & 32 \\ -8 & 35 \end{bmatrix}, \gamma_2 = \gamma_1^{100} \begin{bmatrix} -13 & 80 \\ -8 & 43 \end{bmatrix}$$

The group is the same as in part (1) of this set of examples, but the generators are not in good position. To achieve the same precision, m needs to be up to 100 times greater than in part (1), because the γ_2 in part (1) now has a reduced word of length 101. Since the running time grows exponentially in m , it is not feasible to approximate Q using Algorithm 4.3.3 with these generators as input.

We will now change gears and discuss finding the Berkovich skeleton of a Mumford curve coming from a given Schottky group over K , together with some data on its homotopy

group. This is a relatively easy task, assuming that the given generators $\gamma_1, \dots, \gamma_g$ are in good position, and that we are also given a fundamental domain $F = \mathbb{P}^1 \setminus (B_1 \cap \dots \cap B_g \cap B'_1 \cap \dots \cap B'_g)$. Without loss of generality, we may assume that $\infty \in F^\circ$. Let $P_1, \dots, P_g, P'_1, \dots, P'_g \in (\mathbb{P}^1)^{\text{an}}$ be the Gaussian points of the disks $B_1^+, \dots, B_g^+, B'_1, \dots, B'_g$.

Let R_g be the *rose graph* on g leaves (with one vertex and g loops), and let r_1, \dots, r_g be the loops. A homotopy equivalence $h : R^g \rightarrow G$ must map r_1, \dots, r_g to g loops of G that generate $\pi_1(G)$, so to specify h it will suffice to label g such loops of G with $\{s_1, \dots, s_g\}$ and orientations. It is for this reason that we call h a *marking* of G .

Algorithm 4.3.9 (Berkovich Skeleton Construction).

Require: Matrices $\gamma_1, \dots, \gamma_g \in \mathbb{Q}_p^{2 \times 2}$ generating a Schottky group Γ , together with a good fundamental domain $F = \mathbb{P}^1 \setminus (B_1 \cap \dots \cap B_g \cap B'_1 \cap \dots \cap B'_g)$.

Ensure: The triple (G, ℓ, h) with (G, ℓ) the Berkovich skeleton as a metric graph with h a marking presented as g labelled oriented loops of G .

- 1: Construct the subtree in $(\mathbb{P}^1)^{\text{an}}$ spanned by $P_1, \dots, P_g, P'_1, \dots, P'_g$, including lengths.
- 2: Label the unique shortest path from P_i to P'_i as s_i , remembering orientation.
- 3: Identify each P_i with P'_i , and declare the length of the new edge containing $P_i = P'_i$ to be the sum of the lengths of the edges that were joined to form it.
- 4: Define h by the labels s_i , with each s_i now an oriented loop.
- 5: **return** the resulting labeled metric graph (G, ℓ, h) .

Proof. The proof is essentially given in [GvdP80, I 4.3]. □

Remark 4.3.10. It's worth noting that this algorithm can be done by hand if a good fundamental domain is known. If $P_1, P_2 \in (\mathbb{P}^1)^{\text{an}}$ are the points corresponding to the disjoint closed balls $B(a_1, r_1)^+$ and $B(a_2, r_2)^+$, then the distance between P_1 and P_2 is just the sum of their distances from P_3 corresponding to $B(a_3, r_3)^+$, where $B(a_3, r_3)^+$ is the smallest closed ball containing both a_1 and a_2 . The distance between P_i and P_3 is just $\text{val}(r_3/r_i)$ for $i = 1, 2$. Once all pairwise distances are known, constructing (G, ℓ) is simple. Finding h is simply a matter of drawing the orientation on the loops formed by each pair (P_i, P'_i) and labeling that loop s_i . This process is illustrated three times in Example 4.3.12.

Remark 4.3.11. The space parameterizing labelled metric graphs (G, ℓ, h) (identifying those with markings that are homotopy equivalent) is called *Outer space*, and is denoted X_g . It is shown in [CMV13] that X_g sits inside tropical Teichmüller space as a dense open set, so Algorithm 4.3.9 can be viewed as computing a point in tropical Teichmüller space.

Example 4.3.12. (1) Let Γ be the Schottky group in Example 4.2.3(1). Choose the same good fundamental domain, with $B_1 = B(4, 1/9)$, $B'_1 = B(1, 1/9)$, $B_2 = B(5, 1/9)$, and $B'_2 = B(2, 1/9)$. We have constructed the subtree of $(\mathbb{P}^1)^{\text{an}}$ spanned by P_1, P_2, P'_1, P'_2 as illustrated in Figure 4.1 in Example 4.3.8(1). After identifying P_1 with P'_1 and P_2 with P'_2 , we get the “dumbbell” graph shown in Figure 4.1, with both loops having length 2 and the connecting edge having length 2.

(2) Let Γ be the Schottky group in Example 4.2.3(2). Choose the same good fundamental domain. The subtree of $(\mathbb{P}^1)^{\text{an}}$ spanned by P_1, P_2, P'_1, P'_2 is illustrated in Figure 4.2. After

Figure 4.1: The tree in Example 4.3.12(1), and the Berkovich skeleton.





Figure 4.2: The tree in Example 4.3.12(2), and the Berkovich skeleton.



identifying P_1 with P'_1 and P_2 with P'_2 , we get the “theta” graph shown in Figure 4.2, with two edges of length 2 and one edge of length 2.


(3) Let Γ be the Schottky group in Example 4.2.3(3). The subtree of $(\mathbb{P}^1)^{\text{an}}$ spanned by $P_1, P_2, P_3, P'_1, P'_2, P'_3$ is illustrated in Figure 4.3. After identifying P_1 with P'_1 , P_2 with P'_2 and P_3 with P'_3 , we get the “honeycomb” graph shown in Figure 4.3, with interior edges of length 1 and exterior edges of length 2.

Our final algorithm of this section seeks to study the canonical embedding of a Mumford curve coming from a given Schottky group Γ . From [GvdP80, VI.4], we have that

$$\omega_i(z) := w_i(z)dz = \frac{u'_{\gamma_i}(z)}{u_{\gamma_i}(z)}dz$$

are g linearly independent analytic differentials on Ω that are invariant under the action of Γ . Therefore, they define g linearly independent differentials on $C = \Omega/\Gamma$. Gerritzen and van der Put [GvdP80, VI.4] also state that these form a basis of the space of Γ -invariant analytic differentials. Since the space of algebraic differentials on C has dimension g , it must be generated by these g differentials. Therefore, the canonical embedding has the following

Figure 4.3: The tree in Example 4.3.12(3), and the Berkovich skeleton.



form:

$$C \rightarrow \mathbb{P}^{g-1},$$

$$z \mapsto \left(\frac{u'_{\gamma_1}(z)}{u_{\gamma_1}(z)} : \dots : \frac{u'_{\gamma_g}(z)}{u_{\gamma_g}(z)} \right).$$

It therefore suffices to approximate the derivative $u'_\alpha(z)$. A naïve approach is to consider the approximation

$$u'_\alpha(z) \approx \frac{u_\alpha(z+h) - u_\alpha(z)}{h}.$$

We can do better by taking advantage of the product form of $u_\alpha(z)$:

$$\begin{aligned} u'_\alpha(z) &= \frac{d}{dz} \prod_{\gamma \in \Gamma} \frac{z - \gamma a}{z - \gamma \alpha a} \\ &= \sum_{\gamma \in \Gamma} \left(\frac{d}{dz} \left(\frac{z - \gamma a}{z - \gamma \alpha a} \right) \prod_{\gamma' \in \Gamma, \gamma' \neq \gamma} \frac{z - \gamma' a}{z - \gamma' \alpha a} \right) \\ &= u_\alpha(z) \sum_{\gamma \in \Gamma} \frac{d}{dz} \left(\frac{z - \gamma a}{z - \gamma \alpha a} \right) \left(\frac{z - \gamma a}{z - \gamma \alpha a} \right)^{-1} \\ &= u_\alpha(z) \sum_{\gamma \in \Gamma} \frac{\gamma a - \gamma \alpha a}{(z - \gamma a)(z - \gamma \alpha a)}. \end{aligned}$$

Algorithm 4.3.13 (Canonical Embedding).

Require: Matrices $\gamma_1, \dots, \gamma_g \in \mathbb{Q}_p^{2 \times 2}$ generating a Schottky group Γ in good position, an element $z \in K$, and an integer n to determine precision.

Ensure: An approximation for the image of z under *the* canonical embedding $\Omega/\Gamma \rightarrow \mathbb{P}^{g-1}$ determined by the choice of generators.

- 1: Based on n , choose a suitable positive integer m as described in Remark 4.3.15.
- 2: **for** $i = 1$ to g **do**
- 3: Choose a suitable element $a \in K$ as described in Proposition 4.3.14.
- 4: Compute

$$w_i = \sum_{\gamma \in \Gamma_m} \frac{\gamma a - \gamma \gamma_i a}{(z - \gamma a)(z - \gamma \gamma_i a)}.$$

- 5: **end for**
- 6: **return** $(w_1 : \cdots : w_g)$.

With appropriate choice of a , we can provide a lower bound on the precision of the result in terms of m . Fortunately, we can choose different values of a to approximate

$$\sum_{\gamma \in \Gamma} \frac{\gamma a - \gamma \gamma_i a}{(z - \gamma a)(z - \gamma \gamma_i a)}$$

for different γ_i . As in Proposition 4.3.5, we choose $a \in B_i'^+ \setminus B_i'$ to ensure that both a and $\gamma_i a$ are in F .

Proposition 4.3.14. *If we choose $a \in B_i'^+ \setminus B_i'$ in Algorithm 4.3.13, and assuming $z \in F$, then*

$$\left| \sum_{\gamma \in \Gamma_m} \frac{\gamma a - \gamma \gamma_i a}{(z - \gamma a)(z - \gamma \gamma_i a)} - \frac{u'_{\gamma_i}(z)}{u_{\gamma_i}(z)} \right| \leq p^{-mc - \log_p(d)},$$

where c is the minimum pairwise distance between $P_1, \dots, P_g, P'_1, \dots, P'_g$, and d is the minimum diameter of $B_1, \dots, B_g, B'_1, \dots, B'_g$.

Proof. Let $\gamma \in \Gamma$ have reduced word $\gamma = h_1 h_2 \cdots h_k$. We have seen in the proof of Proposition 4.3.5 that $|\gamma a - \gamma \gamma_i a| \leq p^{-(k-1)c} \text{diam}(B^+)$, $|z - \gamma a| \geq \text{diam}(B^+)$ and $|z - \gamma \gamma_i a| \geq \text{diam}(B^+)$, where B is one of $B_1, \dots, B_g, B'_1, \dots, B'_g$. Thus,

$$\begin{aligned} \left| \frac{\gamma a - \gamma \gamma_i a}{(z - \gamma a)(z - \gamma \gamma_i a)} \right| &\leq \frac{p^{-(k-1)c} \text{diam}(B^+)}{\text{diam}(B^+)^2} \\ &\leq p^{-(k-1)c} d^{-1} \\ &= p^{-(k-1)c - \log_p(d)}. \end{aligned}$$

Since the difference between our approximation and the true value is the sum over terms where γ has reduced words of length $\geq m + 1$, we conclude that the error has absolute value at most $p^{-mc - \log_p(d)}$. \square

In the last proposition, we assumed $z \in F$. If $z \notin F$, we can do an extra step and replace z by some γz such that $\gamma z \in F$, with the help of Subroutine 4.2.6. This step does not change the end result because the theta functions are invariant under the action of Γ .

Remark 4.3.15. If we wish to use Algorithm 4.3.13 to compute the image of a point under the canonical embedding with accuracy up to the n^{th} p -adic digit, we must first compute c and d . Recall that c is defined to be the minimum distance between pairs of the points $P_1, \dots, P_g, P'_1, \dots, P'_g \in (\mathbb{P}^1)^{\text{an}}$ corresponding to the balls $B_1^+, \dots, B_g^+, B_1'^+, \dots, B_g'^+$ that characterize our good fundamental domain, and d is the minimum diameter of $B_1, \dots, B_g, B'_1, \dots, B'_g$. Once we have computed c and d , then by Proposition 4.3.14 we must choose m such that $p^{-mc}d^{-1} \leq p^{-n}$. We could also think of it as choosing m such that $mc + \log_p(d) \geq n$.

Remark 4.3.16. As was the case with Algorithm 4.3.3, we may run Algorithm 4.3.13 even if the input generators are not in good position, and it will approximate images of points in the canonical embedding. However, we will not have control over the rate of convergence, which will in general be very slow.

Example 4.3.17. Let Γ be the Schottky group in Example 4.2.3(3). Choose the same good fundamental domain. We will compute the image of the field element 17 under the canonical embedding (we have chosen 17 as it is in Ω for this particular Γ). The minimum diameter is $d = 1/9$, and the minimum distance is $c = 2$. To get absolute precision to the order of p^{-10} , we need $p^{-mc}d^{-1} \leq p^{-10}$, i.e. $m \geq 6$. Applying Algorithm 4.3.13 with $m = 6$ gives us the following point in \mathbb{P}^2 :

$$((\dots 2100012121)_3 : (\dots 2211022001.1)_3 : (\dots 2221222111.1)_3).$$

This point lies on the canonical embedding of the genus 3 Mumford curve Ω/Γ . Any genus 3 curve is either a hyperelliptic curve or a smooth plane quartic curve. However, it is impossible for a hyperelliptic curve to have the skeleton in Figure 4.3 (see [Cha12, Theorem 4.15]), so Ω/Γ must be a smooth plane quartic curve. Its equation has the form

$$C_1x^4 + C_2x^3y + C_3x^3z + C_4x^2y^2 + C_5x^2yz + C_6x^2z^2 \\ + C_7xy^3 + C_8xy^2z + C_9xyz^2 + C_{10}xz^3 + C_{11}y^4 + C_{12}y^3z + C_{13}y^2z^2 + C_{14}yz^3 + C_{15}z^4.$$

Using linear algebra over \mathbb{Q}_3 , we can solve for its 15 coefficients by computing 14 points on the curve and plugging them into the equation. The result is

$$\begin{aligned} C_1 &= 1, & C_2 &= (\dots 11101)_3, & C_3 &= (\dots 00211)_3, \\ C_4 &= (\dots 1020.2)_3, & C_5 &= (\dots 110.21)_3, & C_6 &= (\dots 1002.1)_3, \\ C_7 &= (\dots 122)_3, & C_8 &= (\dots 222.02)_3, & C_9 &= (\dots 222.02)_3, \\ C_{10} &= (\dots 21101)_3, & C_{11} &= (\dots 2122)_3, & C_{12} &= (\dots 2201)_3, \\ C_{13} &= (\dots 0202.2)_3, & C_{14} &= (\dots 10102)_3, & C_{15} &= (\dots 01221)_3. \end{aligned}$$

For the Newton subdivision and tropicalization of this plane quartic, see the following discussion, in which we consider the interactions of the three algorithms of Section 4.3.

We close this section by checking that the three algorithms give results consistent with one another and with some mathematical theory. We will use our running example of a genus 3 Mumford curve from Examples 4.3.8(3), 4.3.12(3), and 4.3.17, for which we have computed a period matrix of the Jacobian, the Berkovich skeleton, and a canonical embedding.

First we will look at the period matrix and the Berkovich skeleton, and verify that these outputs are consistent. Recall that for the period matrix Q of $\text{Jac}(\Omega/\Gamma)$, we have

$$Q_{ij} = \frac{u_{\gamma_i}(z)}{u_{\gamma_i}(\gamma_j z)}.$$

Motivated by this, we define

$$Q : \Gamma \times \Gamma \rightarrow K^*$$

$$(\alpha, \beta) \mapsto \frac{u_\alpha(z)}{u_\alpha(\beta z)},$$

where our choice of $z \in \Omega$ does not affect the value of $Q(\alpha, \beta)$. (Note that $Q_{ij} = Q(\gamma_i, \gamma_j)$.) As shown in [GvdP80, VI, 2], the kernel of Q is the commutator subgroup $[\Gamma, \Gamma]$ of Γ , and Q is symmetric and positive definite (meaning $|Q(\alpha, \alpha)| < 1$ for any $\alpha \notin \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \pmod{[\Gamma, \Gamma]}$). Moreover, the following theorem holds (see [vdP92, Theorem 6.4]).

Theorem 4.3.18. *Let G be the Berkovich skeleton of Ω/Γ , and let $\pi_1(G)$ be its homotopy group, treating G as a topological space. There is a canonical isomorphism $\varphi: \Gamma \rightarrow \pi_1(G)$ such that $\text{val}(Q(\gamma, \gamma')) = \langle \varphi^{ab}(\gamma), \varphi^{ab}(\gamma') \rangle$, where $\langle p_1, p_2 \rangle$ denotes the shared edge length of the oriented paths p_1 and p_2 .*

The map φ is made very intuitive by considering the construction of G in Algorithm 4.3.9: a generator γ_i of Γ yields two points $P_i, P'_i \in (\mathbb{P}^1)^{\text{an}}$ (corresponding to balls containing the eigenvalues of γ_i), and these points are glued together in constructing G . So γ_i corresponds to a loop around the cycle resulting from this gluing; after abelianization, this intuition is made rigorous.

Consider the matrix Q computed in Example 4.3.8(3). Worrying only about valuations, we have

$$\text{val}(Q) = \begin{bmatrix} 4 & 1 & 1 \\ 1 & 4 & -1 \\ 1 & -1 & 4 \end{bmatrix}.$$

For $i = 1, 2, 3$, let s_i be the oriented loop in G arising from gluing P_i and P'_i . In light of Theorem 4.3.18, we expect to find shared edge lengths

$$\langle s_1, s_1 \rangle = \langle s_2, s_2 \rangle = \langle s_3, s_3 \rangle = 4,$$

$$\langle s_1, s_2 \rangle = \langle s_1, s_3 \rangle = 1,$$

and

$$\langle s_2, s_3 \rangle = -1.$$

That is, each cycle length should be 4, and the common edge of each distinct pair of cycles should have length 1, with the orientation of s_1 agreeing with the orientation of s_2 (respectively, s_3) on the shared edge and the orientation of s_2 disagreeing with the orientation




Figure 4.4: The Newton polygon of the plane quartic curve in Example 4.3.17, and the corresponding tropical curve in \mathbb{R}^2 (drawn using the max convention). Each edge of infinite length has weight 2, and all other edges have weight 1.

of s_3 on the shared edge. This is indeed what we found in Example 4.3.12(3), with edge lengths and orientations shown in Figure 4.3. This example has shown how the outputs of Algorithms 4.3.3 and 4.3.9 can be checked against one another.

We will now consider the relationship between the Berkovich skeleton and the canonical embedding for this example. In particular, we will compute a tropicalization of the curve from the canonical embedding and see how well this represents the Berkovich skeleton.

We compute the tropicalization of the curve as outlined in Section 1.2, starting with the tropical polynomial obtained by taking valuations of the coefficients in the quartic planar equation computed in Example 4.3.17. The Newton polygon of this quartic is a triangle with side length 4, and the coefficients give a subdivision of the Newton polygon as shown in Figure 4.4. The tropicalization of the curve is combinatorially the dual graph of this subdivision, and using the max convention of tropical geometry it sits in \mathbb{R}^2 as shown in Figure 4.4, with the common point of the three cycles at $(0, 0)$.

Let us compare the cycles in the tropicalization with the cycles in the Berkovich skeleton. We know *a priori* from [BPR12, §6.23] that this tropicalization is faithful since all vertices are trivalent and are adjacent to at least one edge of weight one. This means that lattice lengths on the tropicalization must agree with lengths of the Berkovich skeleton, which we now verify. Each cycle in the tropicalization is made up of five line segments, and for each cycle two segments are length $\frac{1}{2}$ and three are length 1. This gives a length of 4, as we'd expect based on Example 4.3.12(3). Moreover, each shared edge has lattice length 1, as was the case in the Berkovich skeleton. Thus we have checked the outputs of Algorithms 4.3.9 and 4.3.13 against one another.

4.4 Finding good starting data

The previous section describes several algorithms that compute various objects from a set of free generators of a Schottky group, assuming that the generators are in good position, and (in Algorithm 4.3.9) that a good fundamental domain is given together with the generators. This content of this section is what allows us to make this assumption. Algorithm 4.4.8 takes an arbitrary set of free generators of a Schottky group and outputs a set of free generators that are in good position, together with a good fundamental domain. This algorithm can be modified as described in Remark 4.4.10 to perform a ‘‘Schottky test’’; in particular, given a set of g invertible matrices generating a group Γ , the modified algorithm will either

- return a set of g free generators of Γ in good position together with a good fundamental domain, which is a certificate that Γ is Schottky;
- return a relation satisfied by the input matrices, which is a certificate that the generators do not freely generate the group; or
- return a non-hyperbolic, non-identity matrix $\gamma \in \Gamma$, which is a certificate that Γ is not Schottky.

Before presenting Algorithm 4.4.8, we will first develop some theory for trees, and then define useful subroutines. Our starting point is a remark in Gerritzen and van der Put’s book:

Proposition 4.4.1. *[GvdP80, III 2.12.3] Let Γ be Schottky and Σ and Ω be as usual. Let $T(\Sigma)$ be the subtree of $(\mathbb{P}^1)^{\text{an}}$ spanned by Σ . Then the minimal skeleton of Ω/Γ is isomorphic to $T(\Sigma)/\Gamma$.*

This statement is essential for our algorithm, because it helps reducing problems involving $(\mathbb{P}^1)^{\text{an}}$ to problems involving the much simpler tree $T(\Sigma)$. Though $T(\Sigma)$ is not finite, it is a finitely branching tree: it consists of vertices and edges such that each vertex is connected with finitely many edges. A good fundamental domain in $(\mathbb{P}^1)^{\text{an}}$ can be obtained from a *good fundamental domain* in $T(\Sigma)$, defined as follows:

Definition 4.4.2. A *principal subtree* T of $T(\Sigma)$ is a connected component of $T(\Sigma) \setminus \{e\}$ for some edge e of $T(\Sigma)$. An *extended principal subtree* is $T^+ = T \cup \{e\}$.

Definition 4.4.3. A *good fundamental domain* S in $T(\Sigma)$ for a set of free generators $\gamma_1, \dots, \gamma_g$ of Γ is the complement of $2g$ principal subtrees $T_1, \dots, T_g, T'_1, \dots, T'_g$, such that $T_1^+, \dots, T_g^+, T'_1^+, \dots, T'_g^+$ are disjoint, and that $\gamma_i(T(\Sigma) \setminus T_i^+) = T_i$ and $\gamma_i^{-1}(T(\Sigma) \setminus T_i^+) = T'_i$. The *interior* of S is $S^\circ = T(\Sigma) \setminus (T_1^+ \cup \dots \cup T_g^+ \cup T'_1^+ \cup \dots \cup T'_g^+)$. The *boundary* of S is $S \setminus S^\circ$.

In other words, S is a connected finite subtree of $T(\Sigma)$ with $2g$ boundary edges (R_i, Q_i) and (R'_i, Q'_i) , where $Q_i, Q'_i \notin S$, such that $\gamma_i(R'_i, Q'_i) = (Q_i, R_i)$. Given this data, the principal subtree T_i (resp. T'_i) is the connected component of $T(\Sigma) \setminus (R_i, Q_i)$ (resp. $T(\Sigma) \setminus (R'_i, Q'_i)$) that

is disjoint from S . Given a good fundamental domain S in $T(\Sigma)$, one can find a good fundamental domain in $(\mathbb{P}^1)^{\text{an}}$ as follows. Without loss of generality, we may assume that the retraction of ∞ to $T(\Sigma)$ is in the interior of S . Then, Q_i and R_i correspond to two nested balls $B(a_i, r_i)^+ \subset B(a_i, R_i)^+$. Define $B_i = B(a_i, \sqrt{r_i R_i})$. Define B'_i similarly.

Proposition 4.4.4. *Let B_i, B'_i be as above. Then $F = (\mathbb{P}^1)^{\text{an}} \setminus (B_1 \cup \dots \cup B_g \cup B'_1 \cup \dots \cup B'_g)$ is a good fundamental domain.*

Proof. Let Q_i, R_i be as above. Let P_i be the midpoint of the segment of the boundary edge (Q_i, R_i) . Then, P_i corresponds to the ball B_i^+ . Let π denote the retraction from $(\mathbb{P}^1)^{\text{an}}$ to $T(\Sigma)$. Again, we may assume that $\pi(\infty)$ is in the interior of S . For any $P \in B_i^+$, the unique path from P to ∞ passes through P_i . Therefore, $\pi(P)$ lies on the union of T_i with the segment (P_i, Q_i) , which is a subset of T_i^+ . Hence, the condition that T_i^+ and T'_i^+ are disjoint implies that the retraction of the B_i^+ and the B'_i^+ are disjoint. Thus, the B_i^+ and B'_i^+ are disjoint.

Let (Q'_i, R'_i) be the boundary edge of T'_i , and let P'_i be its midpoint. Since $\gamma_i(Q'_i, R'_i) = (R_i, Q_i)$, it sends the midpoint P'_i to P_i . Since B'_i is a connected component in $(\mathbb{P}^1)^{\text{an}} \setminus \{P'_i\}$, the element γ_i must send B'_i to a connected component of $(\mathbb{P}^1)^{\text{an}} \setminus \{P_i\}$. One of the connected components in $(\mathbb{P}^1)^{\text{an}} \setminus \{P_i\}$ is $(\mathbb{P}^1)^{\text{an}} \setminus B_i^+$. Since γ_i sends $Q'_i \in B'_i$ to $R_i \in (\mathbb{P}^1)^{\text{an}} \setminus B_i^+$, it must send B'_i to $(\mathbb{P}^1)^{\text{an}} \setminus B_i^+$. Similarly, γ_i^{-1} sends B_i to $(\mathbb{P}^1)^{\text{an}} \setminus B'_i^+$. Thus F is a good fundamental domain in $(\mathbb{P}^1)^{\text{an}}$. \square

One can establish properties of S similar to Theorem 4.2.5. They can be derived either combinatorially or from Proposition 4.4.4.

The following algorithm constructs a good fundamental domain S in $T(\Sigma)$.

Subroutine 4.4.5 (Good Fundamental Domain Construction).

Require: An “agent” knowing all vertices and edges of $T = T(\Sigma)$, and the map $T(\Sigma) \rightarrow T(\Sigma)/\Gamma$, where Γ is defined over \mathbb{Q}_p .

Ensure: A good fundamental domain S in $T(\Sigma)$.

- 1: Choose a vertex P of T . Let P_1, \dots, P_k be all neighbors of P in T .
- 2: Let $V \leftarrow \{P\}$, $E \leftarrow \emptyset$, $O \leftarrow \{(P, P_1), \dots, (P, P_k)\}$, $I \leftarrow \emptyset$, $A \leftarrow \emptyset$.
- 3: **while** $O \neq \emptyset$ **do**
- 4: Choose $(Q, Q') \in O$, remove it from O and add it to E .
- 5: Let Q, Q_1, \dots, Q_k be all neighbors of Q' in T .
- 6: Add Q' to V .
- 7: **for** each Q_k **do**
- 8: With the help of the “agent” in the input, determine if (Q_k, Q') is conjugate to some edge $(R, R') \in O$, i.e. $\gamma Q_k = R$ and $\gamma Q' = R'$ for some $\gamma \in \Gamma$.
- 9: **if** so **then**
- 10: Remove (R, R') from O .
- 11: Add $(Q', Q_k), (R, R')$ to I .
- 12: Add γ to A .

```

13:     else
14:         Add  $(Q_k, Q')$  to  $O$ .
15:     end if
16: end for
17: end while
18: return  $S = V \cup E \cup I$ . (The edges in  $I$  are the boundary edges, and  $A$  is a set of free
    generators of  $\Gamma$  in good position.)

```

Proof. Consider the map from $T(\Sigma)$ to $G = T(\Sigma)/\Gamma$. Let P be as in Step (1). Suppose that a “fire” starts at $P \in T(\Sigma)$ and the image of P in G . In each step, when we choose the edge (Q, Q') in Step (4) and add a vertex Q' to V in Step (5), we “propagate” the fire from Q to Q' , and “burn” Q' together with halves of all edges connecting to Q' . Also, we “burn” the corresponding part in G . Suppose two fires meet each other in G . In this case, both halves of an edge in G are burned, but it corresponds to two half burned edges in $T(\Sigma)$. If so, we stop the fire by removing the edges from O and adding them to I (Step (9)). The algorithm terminates when the whole graph G is burned. The burned part S' of $T(\Sigma)$ is a lifting of G . Then, V is the set of vertices of S' , E is the set of whole edges in S' , and I is the set of half edges in S' . The fact that they form a good fundamental domain follows from the method in the proof of [GvdP80, I (4.3)]. \square

This subroutine requires an “agent” knowing everything about $T(\Sigma)$. It is hard to construct such an “agent” because $T(\Sigma)$ is infinite. Therefore, we approximate $T(\Sigma)$ by a finite subtree. One candidate is $T(\Sigma_m)$, where Σ_m is the set of fixed points of elements of Γ_m . Recall that Γ_m is the set of elements of Γ whose reduced words in terms of the given generators have lengths at most m . We go one step further: we approximate $T(\Sigma)$ by $T(\Gamma_m a)$, where a is any point in Σ .

Lemma 4.4.6. *For any $a \in K$, we have $T(\Gamma a) \supset T(\Sigma)$. Furthermore, if $a \in \Sigma$, then $T(\Gamma a) = T(\Sigma)$.*

Proof. For any $g \in \Gamma$, the fixed point corresponding to the eigenvalue with larger absolute value is the limit of the sequence a, ga, g^2a, \dots . The other fixed point is the limit of the sequence $a, g^{-1}a, g^{-2}a, \dots$. Therefore, every point in Σ is either in Γa or a limit point of Γa . Therefore, $T(\Gamma a) \supset T(\Sigma)$. The second statement is clear. \square

We can construct a complete list of vertices and edges in $T(\Gamma_m a)$. Then, the map from $T(\Gamma_m a)$ to $T(\Sigma)/\Gamma$ can be approximated in the following way: for each pair of vertices P, Q (resp. edges e, f in $T(\Gamma_m a)$ and each given generator γ_i , check if $\gamma_i P = Q$ (resp. $\gamma_i e = f$). If so, then we identify them. Note that this method may not give the correct map, because two vertices P and Q in $T(\Gamma_m a)$ may be conjugate via the action of some $h_1 h_2 \cdots h_k \in \Gamma$, where some intermediate step $h_l h_{l+1} \cdots h_k P \notin T(\Gamma_m a)$. Due to this flaw, we need a way to certify the correctness of the output.

Subroutine 4.4.7 (Good Fundamental Domain Certification).

Require: Generators $\gamma_1, \dots, \gamma_g \in \mathbb{Q}_p^{2 \times 2}$ of a Schottky group Γ , and a quadruple (V, E, I, A) , where V is a set of vertices in $T(\Sigma)$, E and I are sets of edges of $T(\Sigma)$, I contains k pairs of edges $(P_i, Q_i), (P'_i, Q'_i)$, where $P_i, P'_i \in V$, $Q_i, Q'_i \notin V$, and A contains k elements a_i in Γ .

Ensure: TRUE if $S = V \cup E \cup I$ is a good fundamental domain in $T(\Sigma)$ for the set of generators A , and I is the set of boundary edges. FALSE otherwise.

- 1: If $k \neq g$, **return** FALSE.
- 2: If S is not connected, **return** FALSE.
- 3: If any element of I is not a terminal edge of S , **return** FALSE.
- 4: If any $(P_i, Q_i) \neq a_i(Q'_i, P'_i)$, **return** FALSE.
- 5: Choose P in the interior of S .
- 6: **for** $h \in \{\gamma_1, \dots, \gamma_g, \gamma_1^{-1}, \dots, \gamma_g^{-1}\}$ **do**
- 7: Using a variant of Subroutine 4.2.6, find point $P' \in S$ and group element $\gamma \in \langle a_1, \dots, a_k \rangle$ such that $P' = \gamma(hP)$.
- 8: If $P \neq P'$, **return** FALSE.
- 9: **end for**
- 10: **return** TRUE.

Proof. Steps 1–4 verify that S satisfies the definition of a good fundamental domain in $T(\Sigma)$ for the set of generators a_1, \dots, a_g . In addition, we need to verify that a_1, \dots, a_g generate the same group as the given generators $\gamma_1, \dots, \gamma_g$. This is done by Steps 5–9. If $P = P'$ in Step 8, then there exists $\gamma \in \langle a_1, \dots, a_k \rangle$ such that $\gamma hP = P$. We are assuming $a_i \in \Gamma$ in the input, so $\gamma h \in \Gamma$. Since the action of Γ on $(\mathbb{P}^1)^{\text{an}} \setminus \Sigma$ is free, we have $\gamma h = \text{id}$. Thus, $h \in \langle a_1, \dots, a_k \rangle$. If $P = P'$ for all h , then $\Gamma = \langle a_1, \dots, a_k \rangle$.

Otherwise, if $P \neq P'$ in Step 8 for some h , then there exists $\gamma' \in \Gamma$ such that $\gamma'P = P'$. For any $\gamma \in \langle a_1, \dots, a_k \rangle$ other than identity, we have $\gamma P' \notin s^\circ$ by a variant of Lemma 4.2.4. Therefore, $\Gamma \neq \langle a_1, \dots, a_k \rangle$. \square

If the certification fails, we choose a larger m and try again, until it succeeds. We are ready to state our main algorithm for this section:

Algorithm 4.4.8 (Turning Arbitrary Generators into Good Generators).

Require: Free generators $\gamma_1, \dots, \gamma_g \in \mathbb{Q}_p^{2 \times 2}$ of a Schottky group Γ .

Ensure: Free generators a_1, \dots, a_g of Γ , together with a good fundamental domain $F = (\mathbb{P}^1)^{\text{an}} \setminus (B_1 \cup \dots \cup B_g \cup B'_1 \cup \dots \cup B'_g)$ for this set of generators.

- 1: Let $m = 1$.
- 2: Let a be a fixed point of some γ_i .
- 3: Compute all elements in $\Gamma_m a$.
- 4: Find all vertices and edges of $T(\Gamma_m a)$.
- 5: Approximate the map $T(\Gamma_m a) \rightarrow T(\Sigma)/\Gamma$.
- 6: Use Subroutine 4.4.5 to construct a subgraph $S = V \cup E \cup I$ of $T(\Gamma_m a)$ and a subset $A \subset \Gamma$.

- 7: Use Subroutine 4.4.7 to determine if $S = V \cup E \cup I$ is a good fundamental domain in $T(\Sigma)$.
- 8: If not, increment m and go back to Step 2.
- 9: Compute B_i and B'_i from S using the method in Proposition 4.4.4.
- 10: **return** generators A and good fundamental domain $F = (\mathbb{P}^1)^{\text{an}} \setminus (B_1 \cup \dots \cup B_g \cup B'_1 \cup \dots \cup B'_g)$.

Proof. The correctness of the algorithm follows from the proof of Subroutine 4.4.7. It suffices to prove that the algorithm eventually terminates. Assume that we have the “agent” in Subroutine 4.4.5. Since Subroutine 4.4.5 terminates in a finite number of steps, the computation involves only finitely many vertices and edges in $T(\Sigma)$. If m is sufficiently large, $T(\Sigma_m)$ will contain all vertices and edges involved in the computation. Moreover, for any pair of vertices or edges in $T(\Sigma_m)$ that are identified in $T(\Sigma)/\Gamma$, there exists a sequence of actions by the given generators of Γ that sends one of them to the other, so there are finitely many intermediate steps. If we make m even larger so that $T(\Sigma_m)$ contains all these intermediate steps, we get the correct approximation of the map $T(\Sigma) \rightarrow T(\Sigma)/\Gamma$. This data is indistinguishable from the “agent” in the computation of Subroutine 4.4.5. Thus, it will output the correct good fundamental domain. \square

Remark 4.4.9. The performance of the algorithm depends on how “far” the given generator is from a set of generators in good position, measured by the lengths of the reduced words of the good generators in terms of the given generators. If the given generators is close to a set of generators in good position, then a relatively small m is sufficient for $T(\Gamma_m a)$ to contain all relevant vertices. Otherwise, a larger m is needed. For example, in the genus 2 case, this algorithm terminates in a few minutes for our test cases where each given generator has a reduced word of length ≤ 4 in a set of good generators. However, the algorithm is not efficient on Example 4.3.8 (4), where one of the given generators has a reduced word of length 101. One possible way of speeding up the algorithm is to run the non-Euclidean Euclidean algorithm developed by Gilman [Gil14] on the given generators.

Remark 4.4.10. We may relax the requirement that the input matrices freely generate a Schottky group by checking that every element in Γ_m not coming from the empty word is hyperbolic before Step 3 in Algorithm 4.4.8. If the group is Schottky and freely generated by the input matrices, the algorithm will terminate with a good fundamental domain. Otherwise, Step 7 will never certify a correct good fundamental domain, but the hyperbolic test will eventually fail when a non-hyperbolic matrix is generated. In particular, if the identity matrix is generated by a nonempty word, the generators are not free (though they may or may not generate a Schottky group); and if a non-identity non-hyperbolic matrix is generated, the group is not Schottky. Thus, Algorithm 4.4.8 is turned into a Schottky test algorithm. Again, the non-Euclidean Euclidean algorithm in [Gil14] is a possible ingredient for a more efficient Schottky test algorithm.

Chapter 5

The tropical commuting variety

There are various ways to study the pairs of $n \times n$ matrices X and Y over a field K that commute under matrix multiplication. Linear algebraically, one can ask for criteria to determine commutativity. Algebro-geometrically, one can study the *commuting variety*, which is cut out by the n^2 polynomials $(XY)_{ij} - (YX)_{ij}$. These perspectives and many other variants have been studied in the classical setting [OCV11, §5]. This chapter considers similar questions for tropical and tropicalized matrices.

In Section 5.1, we present the necessary background on linear algebra over tropical semirings, and give criteria for the commutativity of matrices over the tropical semiring. Section 5.2 presents three tropical spaces related to commuting spaces: the tropical commuting set, the tropical commuting prevariety, and the tropical commuting variety. We determine the containment relations for these three spaces, which are all distinct for $n \times n$ matrices when $n \geq 3$.

This chapter’s content comes from the paper “The tropical commuting variety” [MT15], coauthored with Ngoc M. Tran.

5.1 Tropical linear algebra

As in Section 1.1, the tropical min-plus algebra $(\overline{\mathbb{R}}, \oplus, \odot)$ is defined by $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$, $a \oplus b = \min(a, b)$, and $a \odot b = a + b$. A pair of $n \times n$ tropical matrices $A = (a_{ij})$, $B = (b_{ij})$ commute if $A \odot B = B \odot A$, where matrix multiplication takes place in the min-plus algebra. Explicitly, this means that for all $1 \leq i, j \leq n$,

$$\min_{s=1, \dots, n} a_{is} + b_{sj} = \min_{s=1, \dots, n} b_{is} + a_{sj}.$$

Tropical linear algebra has extensive applications to discrete events systems [BCOQ92], scheduling [But10], and pairwise ranking [Tra13], amongst others. However, the tropical analogues of many fundamental results in classical linear algebra remain open. Commutativity of tropical matrices is one such example. Classically, if $A, B \in \mathbb{C}^{n \times n}$ where A has n distinct eigenvalues, then $AB = BA$ if and only if B can be written as a polynomial in A

[OCV11, §5]. Moreover, if B has n distinct eigenvalues, then $AB = BA$ if and only if A and B are simultaneously diagonalizable. In a similar spirit, we have the following criterion for a special class of matrices called *polytropes*, defined in Definition 5.1.3, to commute tropically. Here the Kleene star A^* of an $n \times n$ polytrope A is the finite geometric sum $A \oplus A^{\odot 2} \oplus \dots \oplus A^{\odot n}$. Our result states:

Theorem 5.1.1. *Suppose $A, B \in \mathbb{R}^{n \times n}$ are polytropes. If $A \oplus B = (A \oplus B)^*$, then $A \odot B = B \odot A$. If $A \odot B = B \odot A$, then $(A \oplus B)^{\odot 2} = (A \oplus B)^*$. In particular, for $n = 2, 3$, $A \odot B = B \odot A$ if and only if $A \oplus B = A \odot B$.*

Previous works on commuting tropical matrices have also focused on polytropes [KSS12, LP12], due to their special role as the projection to the tropical eigenspace [KSS12, SSB09]. To the best of our knowledge, this is the first necessary and sufficient characterization of commutativity for polytropes for $n < 4$.

Before proving Theorem 5.1.1, we will present some notation and basic facts in tropical linear algebra. See [But10, §1-3] for more details.

If n is a positive integer, let $[n] = \{1, 2, \dots, n\}$. We will write tropical matrix multiplication as $A \odot B$ to remind the reader of the min-linear nature of the algebra. Let I denote the tropical identity matrix, with 0 on the diagonal and ∞ elsewhere. Let $\mathbb{TP}^{n-1} := \mathbb{R}^n / \mathbb{R}(1, \dots, 1)$ be tropical projective space. If $C \subset \mathbb{R}^n$ is closed under scalar tropical multiplication, we shall identify it with the set in \mathbb{TP}^{n-1} obtained by normalizing the first coordinate to be 0. The image of a matrix A , denoted $\text{im}(A)$, is an example of such a set. The *tropical convex hull* between two points $x, y \in \mathbb{R}^n$ is

$$[x, y] = \{a \odot x \oplus b \odot y : a, b \in \mathbb{R}\}.$$

As a set in \mathbb{TP}^{n-1} , this is called the tropical line segment between x and y . A *tropical polytope*, also known as tropical semi-module, is the tropical convex hull of finitely many points. The image of an $n \times n$ matrix is a tropical polytope with at most n distinct vertices in \mathbb{TP}^{n-1} . For an $n \times n$ matrix A with tropical eigenvalue 0, the *Kleene star* of A is the matrix $A^* = I \oplus \bigoplus_{i=1}^{\infty} A^{\odot i}$. This in fact equals to the finite sum $I \oplus \bigoplus_{i=1}^n A^{\odot i}$.

Geometrically, we can view a matrix $A \in \mathbb{R}^{n^2}$ as a map $A : \mathbb{TP}^{n-1} \rightarrow \mathbb{TP}^{n-1}$. Each of the columns of A defines a point in \mathbb{TP}^{n-1} , and the image of A is the tropical convex hull of these points.

Definition 5.1.2. A matrix $A \in \mathbb{R}^{n \times n}$ is a *premetric* if $A_{ii} = 0$, $A_{ij} > 0$ for all $i \neq j \in [n]$.

In this case, A has eigenvalue 0, and its image in \mathbb{TP}^{n-1} is a full-dimensional tropical simplex whose main cell has type $(0, 1, \dots, n-1)$ in the sense of [DS04].

Definition 5.1.3. A matrix $A \in \mathbb{R}^{n \times n}$ is a *polytrope* if A is a premetric, and for all $i, j, k \in [n]$, $A_{ij} \leq A_{ik} + A_{kj}$.

There are many equivalent characterizations of polytrope, e.g., that it is a premetric and $A = A^{\odot 2}$, or that it is a Kleene star of some matrix [But10, §4]. A polytrope A has eigenvalue 0, and the n columns of A are its n eigenvectors. If A is a polytrope, then the image of A in \mathbb{TP}^{n-1} is a full-dimensional tropical polytope that is also convex in the usual Euclidean sense.

For any matrix $A \in \mathbb{R}^{n^2}$ and $b \in \text{im}(A)$, we can consider its preimage under A , i.e. the set of $x \in \mathbb{R}^n$ such that $A \odot x = b$ [But10, §3.1-3.2]. If A is a polytrope, this preimage has a simple and explicit form. We note that the following theorem is a special case of Theorem 3.1.1 in [But10], attributed to Cunninghame-Green (1960) and Zimmerman (1976). This result was also independently re-discovered by Krivulin [Kri12].

Theorem 5.1.4. *Let $A \in \mathbb{R}^{n^2}$ be a polytrope. Define $I = \{i_1, \dots, i_k\}$ for some $1 \leq k \leq n$. Suppose that $b \in \text{im}(A)$ has the form*

$$b = \bigoplus_{i \in I} a_i \odot A_i = a_1 \odot A_{i_1} \oplus \dots \oplus a_k \odot A_{i_k}. \tag{5.1.1}$$

Then $A \odot x = b$ if and only if

$$x = b + \sum_{j \in [n] \setminus I} t_j \mathbf{e}_j, \tag{5.1.2}$$

where $t_j \geq 0$, and \mathbf{e}_j is the j -th standard basis vector.




Figure 5.1: The action of a polytrope on \mathbb{TP}^2 .

The above theorem implies that A is a projection of \mathbb{TP}^{n-1} onto its image, which is the tropical convex hull of the images of the vectors. This is illustrated in Figure 5.1 for a 3×3 polytrope with the columns in \mathbb{TP}^2 as dots and their tropical convex hull in grey. As the matrix acts on \mathbb{TP}^2 , the three columns of the matrix are fixed, as is their tropical convex hull. The remainder of the plane except for three rays is divided into three regions that are

mapped in the directions $(0, -1)$, $(-1, 0)$, or $(1, 1)$. This maps each point to an upside-down tropical line with center at one of the three columns. The rays of these lines that are not in the tropical convex hull are mapped to the point at the center of the tropical line.

Note that $A \odot B = B \odot A$ means for each $i = 1, \dots, n$, the projection of the i -th column of B onto the image of A equals the projection of the i -th column of A onto the image of B . Thus, Theorem 5.1.4 gives an easy geometric check if two polytropes commute. We now give explicit examples for $n = 3$, using three polytropes whose images are illustrated in the first picture in Figure 5.2. Let A , B , and C have the images of their columns labelled by circles, boxes, and crosses, respectively.




Figure 5.2: The images of three polytropes for $n = 3$, some commuting and some not

Example 5.1.5. The matrices A and B (with circles and boxes) commute. Consider $\text{im}(A) \cap \text{im}(B)$, which is a hexagon. The vertices of the hexagon are the vertices of $\text{im}(A \odot B)$; to see this, simply map the columns of B to $\text{im}(A)$ in the natural way. Similarly, the vertices of the hexagon are the vertices of $\text{im}(B \odot A)$. These vertices are illustrated as dots in the second picture in Figure 5.2. It follows that $A \odot B = B \odot A$.

Example 5.1.6. The matrices A and C (with circles and crosses) do not commute. The pentagon $\text{im}(A) \cap \text{im}(C)$ is *not* $\text{im}(A \odot C)$ (or $\text{im}(C \odot A)$). For instance, the upper-right cross vertex is not mapped to this intersection by the action of A . This is illustrated by dots in the third picture in Figure 5.2. This means that A and C do not commute.

We now collect some useful facts about premetrics. Only the last two statements are new, and they are needed for the proof of Theorem 5.1.1. Therefore, we only prove those statements.

Lemma 5.1.7. *If $A, B \in \mathbb{R}^{n \times n}$ are premetrics, then the following hold:*

1. $A \odot B \leq A \oplus B$.
2. $A^{\odot(n-1)} = A^*$.

3. $A^{\odot 2} = A$ if and only if $A = A^*$.
4. $A \odot x = x$ if and only if x is in the image of A^* .
5. $\text{im}((A \odot B)^*) = \text{im}((A \oplus B)^*) = \text{im}(A^*) \cap \text{im}(B^*)$.
6. $(A \odot B)^* = (B \odot A)^* = (A \oplus B)^*$.

Proof. The last statement is the matrix multiplication version of the preceding statement, so let us prove the latter. By one characterization of the Kleene star, [JK10]

$$\begin{aligned}\text{im}(A^*) &= \{x \in \mathbb{TP}^{n-1} : x_i - x_j \leq A_{ij}\} \\ \text{im}(B^*) &= \{x \in \mathbb{TP}^{n-1} : x_i - x_j \leq B_{ij}\}.\end{aligned}$$

This implies

$$\text{im}(A^*) \cap \text{im}(B^*) = \{x \in \mathbb{TP}^{n-1} : x_i - x_j \leq \min(A_{ij}, B_{ij})\} = \text{im}((A \oplus B)^*).$$

Consider the first equality, that is, the claim that $\text{im}((A \odot B)^*) = \text{im}((A \oplus B)^*)$. As before,

$$\text{im}((A \odot B)^*) = \{x \in \mathbb{TP}^{n-1} : x_i - x_j \leq (A \odot B)_{ij}\}$$

Now, $(A \odot B)_{ij} = \min_k A_{ik} + B_{kj} = \min\{A_{ij}, B_{ij}, \min_{k \neq i, j} A_{ik} + B_{kj}\}$. Thus

$$\text{im}((A \odot B)^*) \subseteq \text{im}((A \oplus B)^*).$$

Conversely, suppose that $x \in \text{im}(A^*) \cap \text{im}(B^*)$. By the fourth statement of the lemma,

$$A \odot B \odot x = B \odot A \odot x = x,$$

therefore $x \in \text{im}((A \odot B)^*)$. So $\text{im}(A^*) \cap \text{im}(B^*) \subseteq \text{im}((A \odot B)^*)$. This proves the desired equality. \square

With these results, we are now ready to prove the main result of this section.

Proof of Theorem 5.1.1.

Suppose that $A \oplus B = (A \oplus B)^*$. By Lemma 5.1.7, $A \oplus B = (A \oplus B)^2$. We have

$$A \oplus B = A^{\odot 2} \oplus B^{\odot 2} \oplus A \odot B \oplus B \odot A = A \oplus B \oplus A \odot B \oplus B \odot A.$$

This implies $A \oplus B \leq A \odot B, B \odot A$. By Lemma 5.1.7, $A \odot B, B \odot A \leq A \oplus B$. So we must have

$$A \odot B, B \odot A = A \oplus B,$$

which then implies $A \odot B = B \odot A$. Now, suppose that $A \odot B = B \odot A$. For any $m \geq 2$,

$$(A \oplus B)^{\odot m} = \bigoplus_{k=1}^m A^{\odot k} \odot B^{m-k} = \bigoplus_{k=1}^m A \odot B = A \oplus B \oplus A \odot B = A \odot B$$

Therefore, $(A \oplus B)^{\odot 2} = (A \oplus B)^*$. \square

Corollary 5.1.8. *For $n = 3$, $A \odot B = B \odot A$ if and only if $A \oplus B = (A \oplus B)^*$.*

Proof. The theorem supplies the “if” direction. For the converse, note that $A \odot B = B \odot A$ implies

$$(A \oplus B)^2 = A \oplus B \oplus A \odot B = A \odot B.$$

Now, suppose for the sake of contradiction that $A \odot B$ is strictly smaller than $A \oplus B$ at some coordinate, say, $(1, 2)$. That is,

$$(A \odot B)_{12} = \min\{A_{11} + B_{12}, A_{12} + B_{22}, A_{13} + B_{32}\}$$

But A and B have zero diagonals, and so

$$(A \odot B)_{12} = \min\{B_{12}, A_{12}, A_{13} + B_{32}\}.$$

For strict inequality to occur, we necessarily have $(A \odot B)_{12} = A_{13} + B_{32}$. But A and B are polytropes, so

$$A_{12} \leq A_{13} + A_{32}, B_{12} \leq B_{13} + B_{32}.$$

Therefore,

$$A_{32} > B_{32}, \quad A_{13} < B_{13}.$$

On the other hand, $(A \odot B)_{12} = (B \odot A)_{12}$, and by the same argument, we necessarily have

$$(B \odot A)_{12} = B_{13} + A_{32} < B_{12} < B_{13} + B_{32},$$

which implies $A_{32} < B_{32}$, a contradiction. Hence there is no coordinate $(i, j) \in [3] \times [3]$ such that $A \odot B_{ij} < A_{ij} \oplus B_{ij}$. In other words, $A \odot B = A \oplus B$, which then implies $A \oplus B = (A \oplus B)^2$. \square

Theorem 5.1.1 implies the set inclusion

$$\{(A, B) : A \oplus B = (A \oplus B)^*\} \subseteq \{(A, B) : A \odot B = B \odot A\} \subseteq \{(A, B) : (A \oplus B)^2 = (A \oplus B)^*\}.$$

For $n = 3$, the corollary implies

$$\begin{aligned} \{(A, B) : A \oplus B = (A \oplus B)^*\} &= \{(A, B) : A \odot B = B \odot A\} \\ &\subset \{(A, B) : (A \oplus B)^2 = (A \oplus B)^*\} = \mathbb{R}^{2n^2} \end{aligned}$$

These inclusions are strict for $n \geq 4$. Consider the following two examples for $n = 4$.

Example 5.1.9. $[A \odot B = B \odot A \text{ but } A \oplus B > (A \oplus B)^*]$

Let

$$A = \begin{bmatrix} 0.00 & 4.10 & 3.43 & 0.95 \\ 4.94 & 0.00 & 1.20 & 5.89 \\ 3.74 & 4.44 & 0.00 & 4.69 \\ 3.39 & 6.92 & 2.48 & 0.00 \end{bmatrix}, \quad B = \begin{bmatrix} 0.00 & 1.11 & 8.21 & 9.02 \\ 6.74 & 0.00 & 7.61 & 9.82 \\ 9.96 & 9.56 & 0.00 & 9.77 \\ 1.03 & 2.14 & 1.36 & 0.00 \end{bmatrix}.$$

One can check that $A \odot B = B \odot A$, but $A \oplus B$ differs from $(A \oplus B)^2$ in the $(1, 2)$ entry:

$$(A \oplus B)_{13} = 3.43 > (A \oplus B)_{13}^2 = 2.31.$$

Example 5.1.10 $((A \oplus B)^2 = (A \oplus B)^*$ but $A \odot B \neq B \odot A$). Let

$$A = \begin{bmatrix} 0.00 & 1.09 & 4.02 & 3.33 \\ 6.77 & 0.00 & 2.93 & 3.47 \\ 7.77 & 8.00 & 0.00 & 6.20 \\ 3.30 & 1.85 & 1.39 & 0.00 \end{bmatrix}, \quad B = \begin{bmatrix} 0.00 & 5.02 & 1.45 & 2.58 \\ 3.53 & 0.00 & 2.01 & 2.12 \\ 7.10 & 3.57 & 0.00 & 1.13 \\ 7.71 & 6.04 & 2.47 & 0.00 \end{bmatrix}.$$

The following is an example for $n = 3$ that shows that it is not sufficient to have $A \odot B = A \oplus B$: one needs $A \oplus B = (A \odot B) \oplus (B \odot A)$ for A and B to commute.

Example 5.1.11.

$$A = \begin{bmatrix} 0.00 & 6.4 & 6.10 \\ 3.01 & 0.0 & 0.54 \\ 5.41 & 2.4 & 0.00 \end{bmatrix}, \quad B = \begin{bmatrix} 0.00 & 2.25 & 5.04 \\ 6.81 & 0.00 & 2.79 \\ 4.02 & 6.27 & 0.00 \end{bmatrix}.$$

In this case, $A \odot B = A \oplus B$, but $B \odot A \neq A \oplus B$. These two matrices differ in the $(1, 3)$ coordinate

$$(B \odot A)_{13} = 2.79 < (A \oplus B)_{13} = 5.04,$$

so in particular, $A \odot B \neq B \odot A$.

5.2 Three tropical commuting spaces

Let K be an algebraically closed non-Archimedean field with non-trivial valuation, such as the Puiseux series over \mathbb{C} , and fix an integer $n \geq 2$. Let $S_n = K[\{x_{ij}, y_{ij}\}_{i,j \in \{1, \dots, n\}}]$ and let $I_n \subset S$ be the ideal generated by the n^2 elements of the form

$$\sum_{k=1}^n x_{ik}y_{kj} - \sum_{\ell=1}^n x_{\ell j}y_{i\ell} \tag{5.2.1}$$

where $i, j \in \{1, \dots, n\}$. We call the variety $V(I_n)$ the $n \times n$ commuting variety over K . It is irreducible and has dimension $n^2 + n$ [GS00, MT55]. Its classical points correspond to pairs of matrices $X, Y \in K^{n \times n}$ that commute. The situation is more subtle tropically. In this section, we consider three tropical spaces:

- The *tropical commuting set* \mathcal{TS}_n , which is the collection of all pairs of $n \times n$ tropical commuting matrices in \mathbb{R}^{2n^2} .
- The *tropical commuting variety* \mathcal{TC}_n , which is the tropicalization of the commuting variety.
- The *tropical commuting prevariety* $\mathcal{T}_{\text{pre},n}$, which is the tropical prevariety (as defined in Section 1.1) cut out by the n^2 equations in (5.2.1).

The classical analogs of these three spaces are all identical. We will see that tropically, all of them are different. The difference between the first two illustrates the discrepancy between *tropical commuting matrices* and *tropicalizations of commuting matrices*.

As usual, the tropicalization of a variety over such a field can be defined as the Euclidean closure of the image of the variety under coordinate-wise valuation. For completeness we recall an alternate definition of a tropical variety: for $\omega = (\omega^x, \omega^y) \in \mathbb{R}^{n^2} \times \mathbb{R}^{n^2}$, $f \in \mathbb{R}[x_{ij}, y_{ij}]$, let $in_\omega(f)$ denote the initial form of f , and $in_\omega(I_n) := \langle in_\omega(f) : f \in I_n \rangle$ the initial ideal of I_n . The tropical variety $\mathcal{T}(I_n)$ is the subcomplex of the Gröbner fan of I_n consisting of cones C_ω where $in_\omega(I_n)$ does not contain a monomial.

Our first result concerns the lineality space of $\mathcal{T}(I_n)$, denoted $\text{lin}(I_n)$. This is the set of $\omega \in \mathbb{R}^{2n^2}$ such that $in_\omega(I_n) = I_n$. In our case, this set is a subspace of dimension $n + 1$, which coincides with the lineality space of the Gröbner fan of I_n .

Proposition 5.2.1. *Suppose $n \geq 3$. For $\omega = (\omega^x, \omega^y) \in \mathbb{R}^{n^2} \times \mathbb{R}^{n^2}$, $\omega \in \text{lin}(I_n)$ if and only if there exists $a, b \in \mathbb{R}$ and $c \in \mathbb{R}^n$ such that for all $i, j \in [n]$*

$$\omega_{ii}^x = a, \omega_{jj}^y = b, \omega_{ij}^y = \omega_{ij}^x - a + b, \text{ and } \omega_{ij}^x = c_i - c_j + a. \quad (5.2.2)$$

In particular, $\text{lin}(I_n)$ has dimension $n + 1$ for $n \geq 3$. For $n = 2$, $\omega \in \text{lin}(I_2)$ if and only if there exist $a, b \in \mathbb{R}$ such that

$$\omega_{11}^x = \omega_{22}^x = a, \omega_{11}^y = \omega_{22}^y = b, \omega_{12}^y = \omega_{12}^x - a + b, \text{ and } \omega_{21}^y = \omega_{21}^x - a + b. \quad (5.2.3)$$

In particular, $\text{lin}(I_2)$ has dimension 4.

Proof. We shall prove that ω satisfies (5.2.2) if and only if $in_\omega(g_{ij}) = g_{ij}$ for all $i, j \in [n]$. Since the g_{ij} 's generate I_n , this then implies $in_\omega(I_n) = I_n$.

Suppose ω is such that $in_\omega(g_{ij}) = g_{ij}$. For each fixed $i, j \in [n]$, the monomials $x_{ii}y_{ij}$ and $y_{ij}x_{jj}$ have equal weights. Thus $\omega_{ii}^x = \omega_{jj}^x = a$ for all $i, j \in [n]$. Similarly, $\omega_{ii}^y = \omega_{jj}^y = b$. Now, $x_{ii}y_{ij}$ and $x_{ij}y_{jj}$ have equal weights. Thus

$$\omega_{ij}^y = \omega_{ij}^x - \omega_{ii}^x + \omega_{jj}^y = \omega_{ij}^x - a + b \quad (5.2.4)$$

Consider a triple $i, j, k \in [n]$ of distinct indices. The monomials $x_{ik}y_{kj}$ and $x_{ij}y_{jj}$ have equal weights. Thus

$$0 = \omega_{ik}^x + \omega_{kj}^y - (\omega_{ij}^x + b) = \omega_{ik}^x + (\omega_{kj}^x - a + b) - (\omega_{ij}^x + b) = \omega_{ik}^x + \omega_{kj}^x - a - \omega_{ij}^x. \quad (5.2.5)$$

Since this holds for all triples $i, j, k \in [n]$, we necessarily have

$$\omega_{ij}^x = c_i - c_j + a \quad (5.2.6)$$

for some $c \in \mathbb{R}^n$. Thus, ω is of the form given in (5.2.2).

Finally, for $n = 2$, (5.2.4) still holds. So we have (5.2.3). Define $\omega_{12}^x = c$, $\omega_{21}^x = d$, we see that $\text{lin}(I_2)$ is a linear subspace of \mathbb{R}^8 of dimension 4, parametrized by four parameters a, b, c, d . \square

We now consider the three tropical spaces for $n = 2$. The tropical variety \mathcal{TC}_2 lives in an 8-dimensional ambient space, corresponding to the four x_{ij} and the four y_{ij} coordinates. It is 6-dimensional, with a 4-dimensional lineality space. Modding out by this lineality space gives a 2-dimensional fan with f-vector $(1 \ 4 \ 6)$, meaning there are four rays and six 2-dimensional cones, as illustrated in Figure 5.3, with the rays meeting at the origin. The tropical variety is simplicial and pure.




Figure 5.3: The tropical variety \mathcal{TC}_2 modulo its lineality space

Computation with `gfan` [Jen] shows that the tropical prevariety equals the tropical variety. In other words, the three polynomials

$$\begin{aligned} g_{11} &= x_{12}y_{21} - y_{12}x_{21}, \\ g_{12} &= x_{11}y_{12} + x_{12}y_{22} - y_{11}x_{12} - y_{12}x_{22}, \\ g_{21} &= x_{21}y_{11} + x_{22}y_{21} - y_{21}x_{11} - y_{22}x_{21} \end{aligned}$$

form a tropical basis for $\mathcal{T}(I_2)$. We summarize this and slightly more in the following proposition.

Proposition 5.2.2. *We have $\mathcal{T}_{pre,2} = \mathcal{TC}_2 = \mathcal{TS}_2 \cap \{a_{12} + b_{21} = a_{21} + b_{12}\}$. The homogeneity space is*

$$\omega_{11}^x = \omega_{22}^x = a, \omega_{11}^y = \omega_{22}^y = b, \omega_{12}^y = \omega_{12}^x - a + b, \text{ and } \omega_{21}^y = \omega_{21}^x - a + b.$$

Proof. We have proven everything except the relationship between \mathcal{TS}_2 and the other spaces. If $(A, B) \in \mathcal{TS}_2$, then two of the generators of our tropical basis, namely g_{12} and g_{21} , are tropically satisfied. The final generator g_{11} is tropically satisfied if and only if $a_{12} + b_{21} = a_{21} + b_{12}$, giving the claimed equality. \square

Example 5.2.3. Let $K = \mathbb{C}\{\{t\}\}$ be the field of Puiseux series over \mathbb{C} with the usual valuation. Proposition 5.2.2 tells us when commuting 2×2 tropical matrices with entries in $\text{val}(K)$ can

be lifted to commuting tropical matrices in K . Since the pair of matrices

$$\left(\begin{pmatrix} 0 & 4 \\ 2 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 3 \\ 1 & -1 \end{pmatrix} \right)$$

satisfies $4 + 1 = 2 + 3$, so they can be lifted, for instance to the pair of matrices

$$\left(\begin{pmatrix} 1+t & t^4 \\ t^2 & 2 \end{pmatrix}, \begin{pmatrix} 1 & t^3 \\ t & t^{-1} \end{pmatrix} \right).$$




Figure 5.4: The three spaces for $n = 2$

The relationship between the three spaces for $n = 2$ is illustrated in Figure 5.4. We now give an example to demonstrate that the containment $\mathcal{T}_{\text{pre},2} \subset \mathcal{TS}_2$ really is proper.

Example 5.2.4. Consider the pair of matrices $\left(\begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right)$. These commute under tropical matrix multiplication, but do not tropically satisfy the polynomial $g_{11} = x_{12}y_{21} - y_{12}x_{21}$. Thus, this pair of matrices is in \mathcal{TS}_2 , but not in $\mathcal{T}_{\text{pre},2} = \mathcal{TC}_2$.

For higher dimensions, the containment relation between the three sets is as pictured in Figure 5.5. We will state and prove the result for $n = 3$. The proofs for cases with $n > 3$ are similar.




Figure 5.5: The three spaces for $n > 2$

Proposition 5.2.5. *We have $\mathcal{TC}_3 \subsetneq \mathcal{T}_{\text{pre},3} \cap \mathcal{TS}_3$, and neither $\mathcal{T}_{\text{pre},3}$ nor \mathcal{TS}_3 are contained in one another.*

Proof. To see that each region in Figure 5.5 is really nonempty, consider the following examples.

- (a) The pair of matrices $A = \begin{bmatrix} 0 & 2 & 0 \\ 2 & 0 & 8 \\ 0 & 4 & 0 \end{bmatrix}$, $B = \begin{bmatrix} 12 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 6 \end{bmatrix}$ is in $(\mathcal{T}_{\text{pre},3} \cap \mathcal{TS}_3) \setminus \mathcal{TC}_3$.

Indeed, direct computation shows that $(A, B) \in \mathcal{T}_{\text{pre},3} \cap \mathcal{TS}_3$. Computations with **gfan** show that the initial monomial ideal with this weight vector contains the monomial $x_{31}y_{12}y_{31}y_{21}$. Thus, (A, B) does not lie in the tropical variety. The polynomial with this leading term is given by

$$(XY - YX)_{31}y_{32}y_{21} - (XY - YX)_{32}y_{31}y_{21} - (XY - YX)_{21}y_{31}y_{32}. \quad (5.2.7)$$

Each of the three terms $(XY - YX)_{31}$, $(XY - YX)_{32}$ and $(XY - YX)_{21}$ is a sum of six monomials, two of which are initial monomials. This gives 18 monomials in total with 6 initial monomials. However, the six initial monomials come in three pairs, which are cancelled out by the signs. So (5.2.7) has 12 monomials, and the weights are such that there is a unique leading term.

- (b) The pair of matrices $C = \begin{bmatrix} 0 & 1 & 4 \\ 1 & 0 & 4 \\ 4 & 4 & 0 \end{bmatrix}$, $D = \begin{bmatrix} 0 & 2 & 4 \\ 1 & 0 & 4 \\ 4 & 4 & 0 \end{bmatrix}$ is in $\mathcal{TS}_3 \setminus \mathcal{T}_{\text{pre},3}$.

Indeed, direct computation shows that these matrices commute, and that containment in $\mathcal{T}_{\text{pre},3}$ fails on the (1, 1) and the (2, 2) entries of the products.

- (c) The pair of matrices $E = \begin{bmatrix} 0 & 1 & 0 \\ 3 & 0 & 1 \\ 0 & 3 & 0 \end{bmatrix}$, $F = \begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 0 \\ 1 & 0 & 3 \end{bmatrix}$ is in $\mathcal{T}_{\text{pre},3} \setminus \mathcal{TS}_3$.

Indeed, direct computation shows that these matrices fail to commute in the (3, 3) entry of the products, and that $(E, F) \in \mathcal{T}_{\text{pre},3}$.

In summary, we have $(A, B) \in (\mathcal{T}_{\text{pre},3} \cap \mathcal{TS}_3) \setminus \mathcal{TC}_3$, $(C, D) \in \mathcal{TS}_3 \setminus \mathcal{T}_{\text{pre},3}$, and $(E, F) \in \mathcal{T}_{\text{pre},3} \setminus \mathcal{TS}_3$. \square

We close this section with a study of the geometry of \mathcal{TC}_3 and $\mathcal{T}_{\text{pre},3}$. The tropical variety \mathcal{TC}_3 lives in an 18-dimensional ambient space, corresponding to the nine x_{ij} and the nine y_{ij} coordinates. It is 12-dimensional, with a 4-dimensional lineality space. Modding out gives us an 8-dimensional space. The f-vector is

$$(1 \ 1658 \ 23755 \ 143852 \ 481835 \ 972387 \ 1186489 \ 808218 \ 235038),$$

which ranges from the 1658 rays to the 235,038 8-dimensional cones. The tropical variety is pure, but not simplicial.

The tropical prevariety higher in dimension than the tropical variety. The prevariety is neither pure nor simplicial. Modulo the lineality space, its largest cones are of dimension 10. Its f -vector is

$$(1 \ 146 \ 2290 \ 16322 \ 66193 \ 162886 \ 241476 \ 199030 \ 71766 \ 2397 \ 58).$$

As shown in the proof of Proposition 5.2.5, apart from the generators of the prevariety, the tropical basis for \mathcal{TC}_3 necessarily contains the polynomial

$$(XY - YX)_{31}Y_{32}Y_{21} - (XY - YX)_{32}Y_{31}Y_{21} - (XY - YX)_{21}Y_{31}Y_{32}$$

and all of its permutations under $\mathbb{S}_3 \times \mathbb{S}_2$, by permuting the rows and columns of the matrices simultaneously, and swapping X and Y . By a similar argument, another set of polynomials in the tropical basis are all permutations of

$$(XY - YX)_{12}Y_{21} - (XY - YX)_{21}Y_{12}.$$

However, these two sets of polynomials alone cannot account for the gap in the dimension of the maximal cones between \mathcal{TC}_3 and $\mathcal{T}_{\text{pre},3}$. We suspect that the full tropical basis of \mathcal{TC}_3 contains many more polynomials. Computing this basis explicitly is an interesting computational challenge.

As a first step to computing the tropical basis of \mathcal{TC}_3 , we study the analogue of $\mathcal{T}_{\text{pre},3}$ and \mathcal{TC}_3 for pairs of commuting *symmetric matrices*, so that $X = X^T$ and $Y = Y^T$. These live in a 12-dimensional ambient space, corresponding to the six x_{ij} and the six y_{ij} coordinates. The ideal I_3^{sym} is generated by the following three polynomials:

$$\begin{aligned} (XY)_{12} - (YX)_{12} &= x_{11}y_{12} - y_{11}x_{12} + x_{12}y_{22} - y_{12}x_{22} + x_{13}y_{23} - y_{13}x_{23} \\ (XY)_{13} - (YX)_{13} &= x_{11}y_{13} - y_{11}x_{13} + x_{12}y_{23} - y_{12}x_{23} + x_{13}y_{33} - y_{13}x_{33} \\ (XY)_{23} - (YX)_{23} &= x_{12}y_{13} - y_{12}x_{13} + x_{22}y_{23} - y_{22}x_{23} + x_{23}y_{33} - y_{23}x_{33}. \end{aligned}$$

The symmetric tropical commuting variety is 9-dimensional, with a 2-dimensional lineality space. Its f -vector is

$$(1 \ 66 \ 705 \ 3246 \ 7932 \ 10878 \ 8184 \ 2745).$$

The symmetric tropical commuting prevariety is only one dimension bigger. It has dimension 10, also with a 2-dimensional lineality space. Its f -vector is

$$(1 \ 39 \ 375 \ 1716 \ 4359 \ 6366 \ 5136 \ 1869 \ 6).$$

Under the action of $\mathbb{S}_3 \times \mathbb{S}_2$, the six cones of dimension ten form three orbits. We name them type I, II and III. Type I has orbit size 1, with initial monomials

$$x_{13}y_{23} - x_{23}y_{13}, \quad x_{12}y_{23} - x_{23}y_{12}, \quad x_{12}y_{13} - x_{13}y_{12}.$$

Type II has orbit size 2, with initial monomials

$$x_{12}y_{11} - x_{12}y_{22}, \quad x_{13}y_{11} - x_{13}y_{33}, \quad x_{23}y_{22} - x_{23}y_{33}.$$

Type III has orbit size 3, with initial monomials

$$x_{11}y_{12} - x_{12}y_{11}, \quad x_{11}y_{13} - x_{13}y_{11}, \quad x_{12}y_{13} - x_{13}y_{12}.$$

In theory, since there are three generators with six terms, there can be at most $\binom{6}{2}^3 = 15^3$ possible cones of the symmetric tropical commuting prevariety with maximal dimension. It remains to be understood why only the above six cones are full-dimensional.

As has occurred throughout this work, we find a discrepancy between a tropicalization and a natural tropical analog: the tropicalization of the commuting variety does not contain all pairs of matrices that commute tropically. This is not to say that either object is the *wrong* object to study: the tropicalization of the commuting variety has nicer geometric properties and is more relevant for lifting to a non-Archimedean field, while the tropical commuting set is more natural when working with min-plus linear algebra. They are both relevant spaces in their own right, and an important goal is to understand the difference between them, as with many objects in tropical geometry.

Bibliography

- [ACP14] D. Abramovich, L. Caporaso, and S. Payne. The tropicalization of the moduli space of curves, 2014. <http://arxiv.org/abs/1212.0373>.
- [AGH⁺14] B. Assarf, E. Gawrilow, K. Herr, M. Joswig, B. Lorenz, A. Paffenholz, and T. Rehn. polymake in linear and integer programming, 2014. <http://arxiv.org/abs/1408.4653>.
- [AR10] L. Allermann and J. Rau. First steps in tropical intersection theory. *Mathematische Zeitschrift*, 264:633–670, 2010.
- [AW12] K. Arzdorf and S. Wewers. Another proof of the semistable reduction theorem, 2012. <http://arxiv.org/abs/1211.4624>.
- [Bak07] M. Baker. An introduction to berkovich analytic spaces and non-archimedean potential theory on curves, 2007. notes from the 2007 Arizona Winter School on p-adic geometry, http://people.math.gatech.edu/~mbaker/pdf/aws07mb_v4.pdf.
- [Bak08] M. Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008. With an appendix by Brian Conrad.
- [Bal76] A. T. Balaban, editor. *Chemical applications of graph theory*. Academic Press [Harcourt Brace Jovanovich, Publishers], London-New York, 1976.
- [BCOQ92] F. L. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and linearity*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Ltd., Chichester, 1992. An algebra for discrete event systems.
- [Ber90] V. G. Berkovich. *Spectral theory and analytic geometry over non-Archimedean fields*, volume 33 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1990.
- [BJMS15] S. Brodsky, M. Joswig, R. Morrison, and B. Sturmfels. Moduli of tropical plane curves. *Res. Math. Sci.*, 2:2:4, 2015.

- [BLdM12] E. A. Brugallé and L. M. López de Medrano. Inflection points of real and tropical plane curves. *J. Singul.*, 4:74–103, 2012.
- [BLM⁺14] M. Baker, Y. Len, R. Morrison, N. Pflueger, and Q. Ren. Bitangents of tropical plane quartic curves, 2014. <http://arxiv.org/abs/1404.7568>.
- [BMV11] S. Brannetti, M. Melo, and F. Viviani. On the tropical Torelli map. *Adv. Math.*, 226(3):2546–2586, 2011.
- [BN07] M. Baker and S. Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215(2):766–788, 2007.
- [BN09] M. Baker and S. Norine. Harmonic morphisms and hyperelliptic graphs. *Int. Math. Res. Not. IMRN*, (15):2914–2955, 2009.
- [BPR12] M. Baker, S. Payne, and J. Rabinoff. Nonarchimedean geometry, tropicalization, and metrics on curves, 2012. <http://arxiv.org/abs/1104.0320>.
- [BR13] M. Baker and J. Rabinoff. The skeleton of the jacobian, the jacobian of the skeleton, and lifting meromorphic functions from tropical to algebraic curves, 2013. <http://arxiv.org/abs/1308.3864>.
- [But10] P. Butkovič. *Max-linear systems: theory and algorithms*. Springer Monographs in Mathematics. Springer-Verlag London, Ltd., London, 2010.
- [BW14] I. I. Bouw and S. Wewers. Computing l -functions and semistable reduction of superelliptic curves, 2014. <http://arxiv.org/abs/1211.4459>.
- [Cas12] W. Castryck. Moving out the edges of a lattice polygon. *Discrete Comput. Geom.*, 47(3):496–518, 2012.
- [CDMY14] D. Cartwright, A. Dudzik, M. Manjunath, and Y. Yao. Embeddings and immersions of tropical curves, 2014. <http://arxiv.org/abs/1409.7372>.
- [Cha12] M. Chan. Combinatorics of the tropical Torelli map. *Algebra Number Theory*, 6(6):1133–1169, 2012.
- [Cha13] M. Chan. Tropical hyperelliptic curves. *J. Algebraic Combin.*, 37(2):331–359, 2013.
- [CJ15] M. Chan and P. Jiradilok. Theta characteristics of tropical k_4 curves, 2015. <http://arxiv.org/abs/1503.05776>.
- [CMV13] M. Chan, M. Melo, and F. Viviani. Tropical Teichmüller and Siegel spaces. In *Algebraic and combinatorial aspects of tropical geometry*, volume 589 of *Contemp. Math.*, pages 45–85. Amer. Math. Soc., Providence, RI, 2013.

- [CV09] W. Castryck and J. Voight. On nondegeneracy of curves. *Algebra Number Theory*, 3(3):255–281, 2009.
- [DLRS10] J. A. De Loera, J. Rambau, and F. Santos. *Triangulations*, volume 25 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2010. Structures for algorithms and applications.
- [DO88] I. Dolgachev and D. Ortland. Point sets in projective spaces and theta functions. *Astérisque*, (165):210 pp. (1989), 1988.
- [DS04] M. Develin and B. Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27 (electronic), 2004.
- [FL08] C. Fontanari and E. Looijenga. A perfect stratification of \mathcal{M}_g for $g \leq 5$. *Geom. Dedicata*, 136:133–143, 2008.
- [Gil14] J. Gilman. The non-Euclidean Euclidean algorithm. *Adv. Math.*, 250:227–241, 2014.
- [GJ00] E. Gawrilow and M. Joswig. polymake: a framework for analyzing convex polytopes. In G. Kalai and G. M. Ziegler, editors, *Polytopes — Combinatorics and Computation*, pages 43–74. Birkhäuser, 2000.
- [GK08] A. Gathmann and M. Kerber. A Riemann-Roch theorem in tropical geometry. *Math. Z.*, 259(1):217–230, 2008.
- [GKZ08] I. M. Gelfand, M. Kapranov, and A. Zelevinsky. *Discriminants, resultants and multidimensional determinants*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, 2008. Reprint of the 1994 edition.
- [GS00] R. M. Guralnick and B. A. Sethuraman. Commuting pairs and triples of matrices and related varieties. *Linear Algebra Appl.*, 310(1-3):139–148, 2000.
- [GvdP80] L. Gerritzen and M. van der Put. *Schottky groups and Mumford curves*, volume 817 of *Lecture Notes in Mathematics*. Springer, Berlin, 1980.
- [Hen83] D. Hensley. Lattice vertex polytopes with interior lattice points. *Pacific J. Math.*, 105(1):183–191, 1983.
- [HMY12] C. Haase, G. Musiker, and J. Yu. Linear systems on tropical curves. *Math. Z.*, 270(3-4):1111–1140, 2012.
- [Hol01] J. E. Holly. Pictures of ultrametric spaces, the p -adic numbers, and valued fields. *Amer. Math. Monthly*, 108(8):721–728, 2001.
- [Jen] A. N. Jensen. Gfan, a software system for Gröbner fans and tropical varieties. Available at <http://home.imf.au.dk/jensen/software/gfan/gfan.html>.

- [JK10] M. Joswig and K. Kulas. Tropical and ordinary convexity combined. *Adv. Geom.*, 10(2):333–352, 2010.
- [Kad07] S. Kadziela. *Rigid analytic uniformization of hyperelliptic curves*. ProQuest LLC, Ann Arbor, MI, 2007. Thesis (Ph.D.)—University of Illinois at Urbana-Champaign.
- [Koe91] R. Koelman. *The number of moduli of families of curves on toric surfaces*. 1991. Thesis (Ph.D.)—Katholieke Universiteit te Nijmegen.
- [Kou76] A. G. Kouchnirenko. Polyèdres de Newton et nombres de Milnor. *Invent. Math.*, 32(1):1–31, 1976.
- [Kri12] N. K. Krivulin. A solution of a tropical linear vector equation. *Recent Advances in Computer Engineering Series*, 5:244–249, 2012.
- [KSS12] R. D. Katz, H. Schneider, and S. Sergeev. On commuting matrices in max algebra and in classical nonnegative algebra. *Linear Algebra Appl.*, 436(2):276–292, 2012.
- [KZ03] V. Kaibel and G. M. Ziegler. Counting lattice triangulations. In *Surveys in combinatorics, 2003 (Bangor)*, volume 307 of *London Math. Soc. Lecture Note Ser.*, pages 277–307. Cambridge Univ. Press, Cambridge, 2003.
- [LP12] J. Linde and M. Puente. Commuting normal idempotent tropical matrices: an algebraic-geometric approach, 2012. <http://arxiv.org/abs/1209.0660>.
- [Luo11] Y. Luo. Rank-determining sets of metric graphs. *J. Combin. Theory Ser. A*, 118(6):1775–1793, 2011.
- [Luo13] Y. Luo. Tropical convexity and canonical projections, 2013. <http://arxiv.org/abs/1304.7963>.
- [LZ91] J. C. Lagarias and G. M. Ziegler. Bounds for lattice polytopes containing a fixed number of interior points in a sublattice. *Canad. J. Math.*, 43(5):1022–1035, 1991.
- [Ma14] S. Ma. The rationality of the moduli spaces of trigonal curves, 2014. <http://arxiv.org/abs/1207.0184>.
- [Mik00] G. Mikhalkin. Real algebraic curves, the moment map and amoebas. *Ann. of Math. (2)*, 151(1):309–326, 2000.
- [Mik06] G. Mikhalkin. Tropical geometry and its applications. In *International Congress of Mathematicians. Vol. II*, pages 827–852. Eur. Math. Soc., Zürich, 2006.
- [Mor15] R. Morrison. Tropical images of intersection points. *Collect. Math.*, 66(2):273–283, 2015.

- [MR15] R. Morrison and Q. Ren. Algorithms for Mumford curves. *J. Symbolic Comput.*, 68(part 2):259–284, 2015.
- [MS15] D. Maclagan and B. Sturmfels, editors. *Introduction to Tropical Geometry*. Graduate Texts in Math., Vol. 161, American Math. Soc, 2015.
- [MT55] T. S. Motzkin and O. Taussky. Pairs of matrices with property L . II. *Trans. Amer. Math. Soc.*, 80:387–401, 1955.
- [MT15] R. Morrison and N. M. Tran. The tropical commuting variety, 2015. <http://arxiv.org/abs/1501.03070>.
- [Mum72a] D. Mumford. An analytic construction of degenerating abelian varieties over complete rings. *Compositio Math.*, 24:239–272, 1972.
- [Mum72b] D. Mumford. An analytic construction of degenerating curves over complete local rings. *Compositio Math.*, 24:129–174, 1972.
- [MZ08] G. Mikhalkin and I. Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and abelian varieties, Contemp. Math., vol. 465*,. Amer. Math. Soc., Providence, RI, 2008.
- [OCV11] K. C. O’Meara, J. Clark, and C. I. Vinsonhaler. *Advanced topics in linear algebra*. Oxford University Press, Oxford, 2011. Weaving matrix problems through the Weyr form.
- [OP13] B. Osserman and S. Payne. Lifting tropical intersections. *Doc. Math.*, 18:121–175, 2013.
- [OR11] B. Osserman and J. Rabinoff. Lifting non-proper tropical intersections, 2011. <http://arxiv.org/abs/1109.5733>.
- [Pay09] S. Payne. Analytification is the limit of all tropicalizations. *Math. Res. Lett.*, 16(3):543–556, 2009.
- [Plü34] J. Plücker. Solution d’une question fondamentale concernant la théorie générale des courbes. *J. Reine Angew. Math.*, 12:105–108, 1834.
- [Ram02] J. Rambau. TOPCOM: Triangulations of point configurations and oriented matroids. In A. M. Cohen, X.-S. Gao, and N. Takayama, editors, *Mathematical Software—ICMS 2002*, pages 330–340. World Scientific, 2002.
- [RGST05] J. Richter-Gebert, B. Sturmfels, and T. Theobald. First steps in tropical geometry. In *Idempotent mathematics and mathematical physics*, volume 377 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 2005.

- [S⁺13] W. Stein et al. *Sage Mathematics Software (Version x.y.z)*. The Sage Development Team, 2013. <http://www.sagemath.org>.
- [Sco76] P. R. Scott. On convex lattice polygons. *Bull. Austral. Math. Soc.*, 15(3):395–399, 1976.
- [SSB09] S. Sergeev, H. Schneider, and P. Butkovič. On visualization scaling, subeigenvectors and Kleene stars in max algebra. *Linear Algebra Appl.*, 431(12):2395–2406, 2009.
- [Tei88] J. Teitelbaum. p -adic periods of genus two Mumford-Schottky curves. *J. Reine Angew. Math.*, 385:117–151, 1988.
- [Tra13] N. M. Tran. Pairwise ranking: choice of method can produce arbitrarily different rank order. *Linear Algebra Appl.*, 438(3):1012–1024, 2013.
- [vdP92] M. van der Put. Discrete groups, Mumford curves and theta functions. *Ann. Fac. Sci. Toulouse Math. (6)*, 1(3):399–438, 1992.
- [Vig10] M. D. Vigeland. Smooth tropical surfaces with infinitely many tropical lines. *Ark. Mat.*, 48(1):177–206, 2010.
- [Zha10] I. Zharkov. Tropical theta characteristics. In *Mirror symmetry and tropical geometry*, volume 527 of *Contemp. Math.*, pages 165–168. Amer. Math. Soc., Providence, RI, 2010.