Weierstrass Points and Torsion Points on Tropical Curves

by

David Harry Richman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2020

Doctoral Committee:

Professor David Speyer, Chair
Professor Alexander Barvinok
Professor Mattias Jonsson
Professor Jeffrey Lagarias
Professor Mark Newman

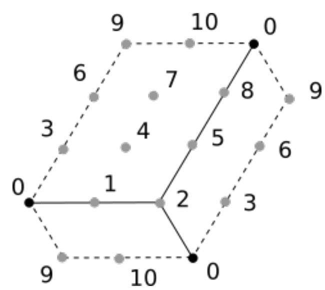

9
10
0
6
7
8
9
3
4
5
6
0
1
2
3
9
10
0

David Harry Richman

hrichman@umich.edu

ORCID iD: 0000-0002-0101-0521

This thesis is dedicated to my father, David Ross Richman, and to my grandfather, Alexander Richman.

ACKNOWLEDGEMENTS

> As a small child, I already wanted to be a writer. I am now doing a job that a child told me to do.
>
> Elif Batuman

I am grateful to my advisor, David Speyer, for generously sharing his time and enthusiasm in talking about math with me; in particular, about the math that appears in this thesis. I thank Jeffrey Lagarias for giving lots of good advice; also to Michael Filaseta and Wei-Tian Li. Thank you to the members of my thesis committee, Alexander Barvinok, Mattias Jonsson, and Mark Newman, for teaching me interesting and useful mathematics.

The results in this thesis benefitted from conversations with Matthew Baker, Melody Chan, David Jensen, Eric Katz, Sam Payne, Nathan Pflueger, Farbod Shokrieh and Dmitry Zakharov. Thank you to Sachi Hashimoto for telling me about the Manin–Mumford conjecture, which led to the work in Chapter 5. The idea to consider the higher-degree case in Chapter 5 was suggested by David Speyer. This work was supported by NSF grants DMS-1600223 and DMS-1701576, and a Rackham Predoctoral Fellowship.

Thank you to my classmates for their friendship and providing an intellectually stimulating environment, including: Will Dana, Lara Du, Jonathan Gerhard, Trevor Hyde, Gracie Ingermanson, Hyung Kyu Jun, Luby Lu, Visu Makam, Takumi Murayama, Andrew O'Desky, Samantha Pinella, Emanuel Reinecke, Salman Siddiqi, Matt Stevenson, Phil Tosteson, Farrah Yhee and Feng Zhu.

Thank you to my family for their support. Thank you to my mom for getting me interested in math from a young age. I did not want to be a writer as a small child, but I did already like thinking about math and this was entirely due to her encouragement.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF SYMBOLS

Symbol

G	a finite, connected graph with vertex set $V(G)$ and edge set $E(G)$
(G,ℓ)	a combinatorial model for a metric graph, where $\ell : E(G) \to \mathbb{R}_{>0}$ is an edge-length function
Γ	a compact, connected metric graph
$\mathrm{val}(x)$	valence of a point x on a metric graph
$\mathrm{PL}_{\mathbb{R}}(\Gamma)$	continuous, piecewise linear functions on Γ
$\mathrm{PL}_{\mathbb{Z}}(\Gamma)$	continuous, piecewise $\mathbb{Z}$-linear functions on Γ
$\Delta(f)$	the principal divisor associated to a piecewise ($\mathbb{Z}$-)linear function f, value of the metric graph Laplacian $\Delta : \mathrm{PL}_{\mathbb{Z}}(\Gamma) \to \mathrm{Div}(\Gamma)$
D	a divisor on a metric graph or an algebraic curve
$\deg(D)$	degree of divisor D
$r(D)$	the Baker–Norine rank of D
$[D]$	a divisor class; the set of divisors linearly equivalent to D
$\|D\|$	the space of effective divisors linearly equivalent to D
$\mathrm{red}_x[D]$	the x-reduced divisor equivalent to D, where $x \in \Gamma$
$\mathrm{br}[D]$	the break divisor equivalent to a divisor D of degree g
$\mathrm{Div}(\Gamma)$	divisors on Γ (with $\mathbb{Z}$-coefficients)
$\mathrm{Div}_{\mathbb{R}}(\Gamma)$	divisors on Γ with $\mathbb{R}$-coefficients, i.e. $\mathrm{Div}(\Gamma) \otimes_{\mathbb{Z}} \mathbb{R}$
$\mathrm{Div}^d(\Gamma)$	divisors of degree d on Γ
$\mathrm{Pic}(\Gamma)$	Picard group of Γ, cokernel of $\Delta : \mathrm{PL}_{\mathbb{Z}}(\Gamma) \to \mathrm{Div}(\Gamma)$
$\mathrm{Pic}^d(\Gamma)$	divisor classes of degree d on Γ
$\mathrm{Sym}^d(\Gamma)$	effective divisors of degree d on Γ
$\mathrm{Eff}^d(\Gamma)$	effective divisor classes of degree d on Γ
$\mathrm{Br}^g(\Gamma)$	the space of break divisors on Γ
$W(D)$	Weierstrass locus of divisor D
$W^{\mathrm{st}}(D)$	stable Weierstrass locus of D
j_z^y	voltage function $\Gamma \to \mathbb{R}$ when unit current sent from y to z
$r(x,y) = r(x,y;\Gamma)$	effective resistance between points x and y on Γ
$\mu = \mu_\Gamma$	the Zhang canonical measure on Γ
$\mathrm{Jac}(\Gamma)$	Jacobian group of Γ, consisting of degree 0 divisor classes
AJ_q	Abel–Jacobi map $\Gamma \to \mathrm{Jac}(\Gamma)$, with basepoint $q \in \Gamma$
$AJ_Q^{(d)}$	higher-degree Abel–Jacobi map $\Gamma^d \to \mathrm{Jac}(\Gamma)$, with base-divisor $Q \in \mathrm{Eff}^d(\Gamma)$
$\mathrm{Jac}(\Gamma)_{\mathrm{tors}}$	torsion subgroup of $\mathrm{Jac}(\Gamma)$
$\mathrm{Jac}(G)$	critical group of G, consisting of deg. 0 divisor classes supported on the vertex set $V(G)$
$\mathcal{T}(G)$	the set of spanning trees of a graph G
$\mathcal{C}(G)$	the set of cycles of a graph G
$M^\perp(G)$	cographic matroid of G

ABSTRACT

We investigate two constructions on metric graphs, using the framework of tropical geometry. On a metric circle, i.e. a genus 1 tropical curve, each of these constructions produces a set of n points which are evenly spaced around the circle.

In the first part, we study Weierstrass points for a divisor on a metric graph (i.e. tropical curve). On a smooth algebraic curve, these are points which have "special" tangency behavior with respect to a given projective embedding. The Weierstrass locus on a metric graph may fail to be a finite set; we define a stable Weierstrass locus which is always finite. The stable locus agrees with the "naive" Weierstrass locus for a generic divisor class. We then investigate the distribution of Weierstrass points for a high-degree divisor. We show that in high degree, the distribution of Weierstrass points converges to Zhang's canonical measure. This measure can be described by probabilities of weighted spanning trees, or alternatively by current flows in an electrical resistor network. This distribution result is a tropical analogue of a theorem of Neeman concerning Weierstrass points on a complex algebraic curve.

In the second part, we consider how a metric graph under the Abel–Jacobi embedding intersects torsion points of its Jacobian. The Manin–Mumford conjecture states that this intersection is finite for a smooth algebraic curve of genus $g \geq 2$; this conjecture was proved by Raynaud. For a metric graph, this conjecture fails when the edge lengths are all rational numbers. However, we show that the Manin–Mumford conjecture does hold for metric graphs (of genus $g \geq 2$) which are biconnected and have edge lengths which are "sufficiently irrational" in a precise sense. Under these assumptions we prove a bound on the size of the intersection which depends only on the genus, namely $\#(AJ(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}) \leq 3g - 3$. Next we consider higher-degree analogues of the Manin–Mumford conjecture, concerning the maps sending d-tuples of points to the Jacobian. This motivates the definition of the "independent girth" of a graph, which gives a strict upper bound for d such that the higher-degree Manin–Mumford property holds. For a metric graph with large genus g, the independent girth is bounded above by $O(\log g)$.

CHAPTER 1

Introduction

In this thesis, we study two generalizations of a simple construction: dividing a circle in N equal parts. These division points are called *torsion points* of the circle.

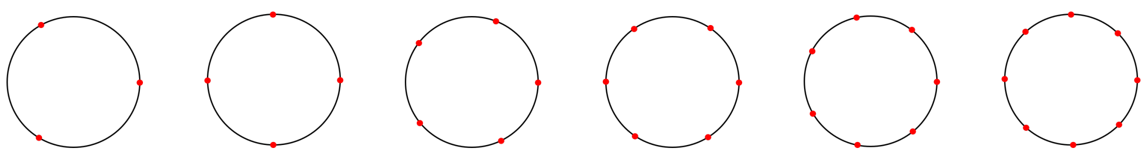

Figure 1.1: Torsion points on a circle.

The word "torsion" comes from algebraic terminology—if we equip the circle with the additive structure $\mathbb{R}/\mathbb{Z}$, i.e. how we usually think of adding angles together, then the N-torsion points are points x which satisfy $N \cdot x = x + \cdots + x = 0$. (There are N such points.)

A circle is a simple example of a metric graph. A metric graph captures the structure of a network, meaning something made up of nodes and edges, where additionally each edge is assigned a positive real length. If we take just one node and one edge, with the edge joined to the node at both ends, then we get a circle. If we use more nodes and edges, we can get a more complicated metric graph.

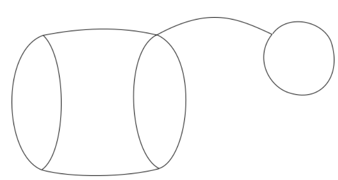

Figure 1.2: A metric graph, with 5 nodes and 8 edges.

For an arbitrary metric graph, we can ask: How does one divide this object into n "equal parts"? There is probably no single good answer to such a question, but we consider two constructions which generalize N-torsion points of a circle to arbitrary metric graphs. Both constructions are taken from the study of complex algebraic curves, via the framework of tropical geometry.

In algebraic geometry, the analogue of a circle is an elliptic curve. An elliptic curve over the complex numbers is topologically equivalent to a parallelogram with opposite sides glued together. The elliptic curve also has an additive structure of $\mathbb{R}^2/\Lambda$, coming from addition of vectors in $\mathbb{R}^2$ modulo integer combinations of vectors forming the sides of the parallelogram. The N-torsion points are the points which satisfy the equation $Nx = x + \cdots + x = 0$ with respect to this addition law. In this case there are N^2 such points.

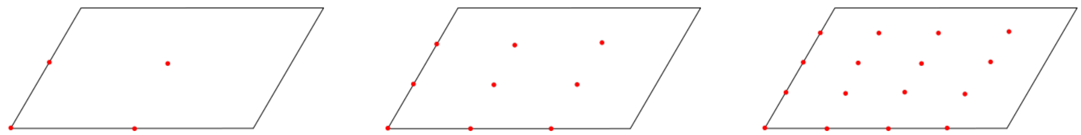

Figure 1.3: Torsion points on a complex elliptic curve.

The torsion points may also be constructed without reference to an addition law, as follows. Given a curve X in projective space $\mathbb{P}^r$, a *flex point* is a point p on X such that some hyperplane intersects X at p with multiplicity at least $r + 1$. If we embed an elliptic curve into projective space $\mathbb{P}^r$ using a *complete linear system* of degree N divisors, then the set of flex points is in fact a set of N-torsion points (for some choice of 0 on the elliptic curve).

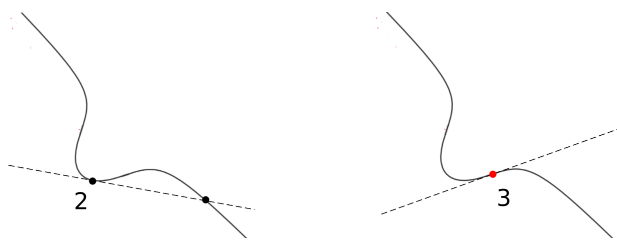

Figure 1.4: Flex point (right) on an embedded curve in $\mathbb{P}^2$.

There are two ways to take this concept of torsion points on an elliptic curve (genus $g = 1$) and generalize it to a smooth algebraic curve of higher genus ($g \geq 2$). In the first perspective, torsion points come from some additive group law, and this leads to the study of torsion points of the *Jacobian*—a g-dimensional variety with group structure, which naturally contains the higher-genus curve as a 1-dimensional subvariety (up to a choice of translation). The second perspective of torsion points, as the flex points of some projective embedding, leads to the study of (generalized) *Weierstrass points* in higher genus.

1.1 Tropical geometry

Tropical geometry is a relatively new area of mathematics which allows one to translate statements about algebraic curves to graph theory, and vice versa. For a thorough introduction to tropical geometry and tropical curves, we refer the reader to [18, 38, 41].

Algebraic geometry is the study of solutions to polynomial equations such as $x^4 + y^4 = 1$. Over the complex numbers, the set of solutions is known as a Riemann surface. Tropical geometry allows us to turn a Riemann surface into a graph. This

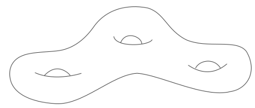

Figure 1.5: A genus three Riemann surface.

may be achieved from either an "embedded" or "non-embedded" perspective.

In the embedded perspective, given a complex algebraic variety in $\mathbb{C}^n$ we may consider the image in $\mathbb{R}^n$ under the logarithm map $(z_1, \dots, z_n) \mapsto (\log|z_1|, \dots, \log|z_n|)$. This image is called the *amoeba* of the algebraic variety. Bergman [12] observed that the limit of this amoeba, when "zooming out to infinity," forms a polyhedral complex. This polyhedral complex is known as the *logarithmic limit set* of the variety.

Example 1.1 (Logarithmic limit set)**.** Consider the solutions to $x^3+y^3+4xy+1 = 0$ where $x, y \in \mathbb{C}^2$. The amoeba of this complex curve is shown on the left side of Figure 1.6, shaded in gray, and the logarithmic limit set is on the right.

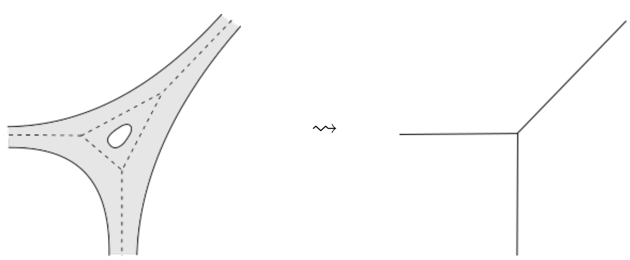

Figure 1.6: Amoeba and logarithmic limit set of $x^3 + y^3 + 4xy + 1 = 0$.

Under the process of sending $z_i \mapsto \log|z_i|$ and "zooming out to infinity," the effect is that

$$\lim_{|z|\to\infty} \frac{\log|a_n z^n + \text{ (lower-order terms)}|}{\log|z|} = \lim_{|z|\to\infty} \frac{\log|a_n| + n\log|z|}{\log|z|} = n.$$

In words, a polynomial with leading term $a_n z^n$ is replaced with $n \log|z|$. One way to algebraically formalize, or mimic, this process is to work with the non-Archimedean valuation

$$\mathrm{val}(a_n t^n + \text{ (lower-order terms in } t)) = n$$

applied to the field of Laurent series $K = \mathbb{C}((t^{-1}))$ (or Puiseux series $\cup_{n\geq 1}\mathbb{C}((t^{-1/n}))$), and consider varieties over the ground field K rather than $\mathbb{C}$. Then given a variety $X \subset K^n$ cut out by polynomials in $K[x_1, \ldots, x_n]$, its *tropicalization* (or *non-Archimedean amoeba*) is the image of X under $(z_1(t), \ldots, z_n(t)) \mapsto (\operatorname{val} z_1(t), \ldots, \operatorname{val} z_n(t))$. It turns out that this image is a polyhedral complex in $\mathbb{R}^n$, no zooming out needed.

A fundamental theorem of tropical geometry is that the tropicalization of X, as defined above, may be computed via the following process on the polynomials cutting out X. (For simplicity, we describe the case of polynomials in two variables.) Given a polynomial $f = \sum_{i,j\geq 0} a_{i,j}(t)x^i y^j \in K[x, y]$, its tropicalization is defined as

$$\operatorname{trop}(f) = \max_{i,j\geq 0}\{\operatorname{val}(a_{i,j}(t)) + ix + jy\}.$$

This expression $\operatorname{trop}(f)$ defines a piecewise-linear function $(x, y) \mapsto \operatorname{trop}(f)$ on $\mathbb{R}^2$. The *break locus* of $\operatorname{trop}(f)$ is the subset of $\mathbb{R}^2$ where the function is not linear.

Theorem 1.2 (Fundamental theorem of tropical geometry)**.** *Given some polynomial $f \in K[x_1, \ldots, x_n]$, suppose $(z_1(t), \ldots, z_n(t)) \in K^n$ lies on the variety cut out by $f = 0$. Then the point* $(\operatorname{val} z_1(t), \ldots, \operatorname{val} z_n(t)) \in \mathbb{R}^n$ *lies in the break locus of* $\operatorname{trop}(f)$.

The converse of Theorem 1.2 is not true, but there is some sense in which the converse holds for a "sufficiently general" $f \in K[x_1, \ldots, x_n]$.

Example 1.3. The tropicalization of the polynomial $f = x^3 + y^3 + t^N xy + 1$ is the piecewise-linear function

$$\operatorname{trop}(f) = \max\{3x,\, 3y,\, N + x + y,\, 0\}.$$

If $N > 0$, this tropicalized function has four domains of linearity. We illustrate its break locus in Figure 1.7, with the domains of linearity labelled by the corresponding linear function. The break locus consists of three bounded segments and three unbounded segments. The bounded segments have endpoints (N, N), $(-N, 0)$, and $(0, -N)$.

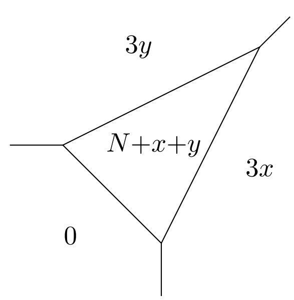

Figure 1.7: Break locus of $\operatorname{trop}(f) = \max\{3x,\, 3y,\, N + x + y,\, 0\}$.

In the abstract (non-embedded) perspective, tropicalization is achieved via *degenerating* a smooth algebraic curve to a curve with nodal singularities, along a one-parameter family, then taking the dual graph of the nodal curve. This degeneration process turns meromorphic (i.e. rational) functions on the Riemann surface (i.e. complex algebraic curve) to piecewise linear functions on the dual graph. These tools were developed by Baker–Norine [9] and others [32, 20].

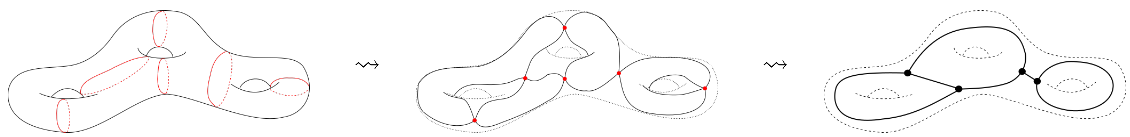

Figure 1.8: Tropicalizing a Riemann surface (left) to a graph (right).

The non-embedded perspective of tropicalization has been used to prove powerful results that relate moduli spaces of smooth algebraic curves, and their compactifications by stable curves, to corresponding moduli spaces of tropical curves; see e.g. Caporaso [14] and Abramovich et. al. [1].

1.2 Summary of results

Recall the above discussion of torsion points on an elliptic curve, as (1) flex points of a projective embedding, or (2) algebraically torsion with respect to the additive structure of the Jacobian. This thesis is concerned with studying the analogous constructions in the tropical setting. We state the main results and discuss related work in the sections below.

For background on complex algebraic curves, see [33]. In the following we assume all algebraic curves are proper and smooth, unless stated otherwise explicitly. We restrict our attention to tropical curves Γ have no "hidden genus" at vertices and no infinite legs, i.e. to those Γ arising as the skeleton of X^{an} with totally degenerate reduction and no punctures.

1.2.1 Weierstrass points

Suppose X is a smooth, proper complex algebraic curve. The *Weierstrass points* of a divisor D on X are the flex points of the projective embedding $X \to \mathbb{P}^r$ corresponding to the complete linear system of D. This defines a finite subset of X.

Historically, mathematicians were first interested in studying the Weierstrass points of the canonical divisor on a curve of genus $g \geq 2$. Hurwitz [25] showed that an algebraic curve of genus $g \geq 2$ has finite automorphism group by using the Weierstrass points of the canonical divisor. A generic curve has $g^3 - g$ such points. In the literature, the Weierstrass points of a divisor D, which is not the canonical divisor, are sometimes referred to as "higher Weierstrass points." (Sometimes, "higher Weierstrass points" refers to Weierstrass points of nK, where K is the canonical divisor and $n \geq 2$ is an integer.) See [16] for a well-written historical survey of the study of Weierstrass points.

In [34], Mumford notes that the Weierstrass points associated to a divisor of degree n should be viewed as a higher-genus analogue of the n-torsion points on an elliptic curve. The fact that n-torsion points on a complex elliptic curve become "evenly distributed" as n grows large leads one to ask whether the same phenomenon holds for Weierstrass points on other algebraic curves.

An answer was given by Neeman [35], who showed that for a complex curve (i.e. Riemann surface) of genus $g \geq 2$, when $n \to \infty$ the Weierstrass points of degree n divisors become distributed according to the Bergman measure.

Theorem 1.4 (Neeman [35])**.** *Let X be a compact Riemann surface of genus $g \geq 2$, and let $\{D_n : n \geq 1\}$ be a sequence of divisors on X with $\deg D_n = n$. Let W_n denote the Weierstrass locus of the divisor D_n, and let $\delta_n = \frac{1}{gn^2} \sum_{x \in W_n} \delta_x$ denote the normalized discrete measure on X associated to W_n (where δ_x is the Dirac measure at x). Then as $n \to \infty$, the measures δ_n converge weakly to the Bergman measure on X.*

Before Neeman's result, Olsen [36] showed that given a positive-degree divisor D on a complex algebraic curve X, the union of the Weierstrass points of the multiples nD, over all $n \geq 1$, is dense in X in the complex topology.

If one replaces the ground field $\mathbb{C}$ with a non-Archimedean field, one may consider the same question of how Weierstrass points are distributed inside the Berkovich analytification X^{an} of an algebraic curve, say after retracting to a compact skeleton Γ. This was addressed by Amini in [3]. Here the Weierstrass points are distributed according to the *Zhang canonical admissible measure*, constructed by Zhang in [42].

Theorem 1.5 (Amini [3])**.** *Let X be a smooth proper curve of genus $g \geq 1$ over a complete, algebraically closed, non-Archimedean field K with non-trivial valuation and residue characteristic 0. Let Γ be a skeleton of the Berkovich analytification X^{an} with retraction map $\rho : X^{an} \to \Gamma$. Let D be a positive-degree divisor on $X(K)$. Let W_n denote the Weierstrass locus of the divisor nD, and let $\delta_n = \frac{1}{\# W_n} \sum_{x \in W_n} \delta_{\rho(x)}$ denote the normalized discrete measure on Γ associated to W_n (where δ_x is the Dirac measure at x). Then as $n \to \infty$, the measures δ_n converge weakly to the Zhang canonical measure on Γ.*

Zhang's canonical measure does not have support on bridge edges, so it is independent of the choice of skeleton. Zhang's construction was motivated by Arakelov's pairing for divisors on a Riemann surface [5], for the purpose of answering questions in arithmetic geometry. Here we follow an approach of Chinburg–Rumely [17] and Baker–Faber [7] along more elementary lines, describing μ in terms of electric potential and current flow in a network of resistors.

In [6], Baker studies ordinary Weierstrass points on graphs and on metric graphs, and mentions several applications of number theoretic significance. These results are stated for Weierstrass points associated to the canonical divisor, without discussion of generalized Weierstrass points for other divisors. In [3], Amini raises the question of whether the distribution of (generalized) Weierstrass points is possibly intrinsic to the metric graph Γ, without needing to identify Γ with the skeleton of some Berkovich curve X^{an}. One major obstacle to this idea is that on a metric graph, the Weierstrass locus for a divisor may fail to be a finite set of points.

We give two approaches to get around this obstacle. One approach is to sidestep the issue by showing that on a tropical curve, the Weierstrass locus is finite for a generic divisor class. We also define a *stable* Weierstrass locus which is finite for an arbitrary divisor class. The stable Weierstrass locus is nicely compatible with the non-stable locus when D is a *non-special* divisor in the sense of the Riemann–Roch theorem; the relation is complicated for divisors which are Riemann–Roch special.

We compute the cardinality of the stable Weierstrass locus of a generic divisor class, by showing that if we introduce a notion of multiplicity, the stable locus has constant cardinality along a family of divisor classes. This cardinality depends only on the degree of the divisor and the genus of the underlying curve. The agreement of the stable and non-stable Weierstrass locus for a non-special divisor class allows us to extend the same result, generically, to the number of (non-stable) Weierstrass points.

With the assumption of genericity, we also show that there is a limiting distribution of Weierstrass points of high degree that is intrinsic to the tropical curve Γ. This gives a tropical result analogous Theorems 1.4 and 1.5, and answers Amini's question in [3]. The tropical Weierstrass points become distributed according to the same measure μ that appears in Amini's theorem.

We now state our results in more detail. Given a connected metric graph Γ and a divisor D of degree n and rank $r = r(D)$, we define the Weierstrass locus $W(D)$ as

$$W(D) = \{x \in \Gamma : D \sim (r+1)x + E \text{ for some } E \geq 0\},$$

where $\sim$ denotes linear equivalence. We define the stable Weierstrass locus of D as

$$W^{\mathrm{st}}(D) = \{x \in \Gamma : \mathrm{br}[D - (n-g)x] = x + E \text{ for some } E \geq 0\}$$

if the degree $n \geq g$ and $W^{\mathrm{st}}(D) = \varnothing$ otherwise, where $\mathrm{br}[D]$ denotes the break divisor representative of a degree g divisor D. See Chapter 2 for definitions of linear equivalence, rank, and break divisor.

The locus $W(D)$ may fail to be finite; in some cases it contains all of Γ. The stable Weierstrass locus is finite for any divisor. If D has rank $r(D) = n - g$, i.e. D is non-special, then we have the containment $W^{\mathrm{st}}(D) \subset W(D)$. In particular, this containment holds when the degree $n \geq 2g - 1$.

Our first result addresses the question of counting the number of Weierstrass points. Here "generic" means on a dense open subset of the space of divisor classes.

Theorem 4.20. *Let Γ be a connected metric graph of genus g.*

(a) For a generic divisor class of degree $n \geq g$, the Weierstrass locus $W(D)$ is finite with cardinality

$$\#W(D) = g(n - g + 1).$$

For a generic divisor class of degree $n < g$, $W(D)$ is empty.

(b) For an arbitrary divisor class of degree $n \geq g$, the stable Weierstrass locus $W^{\mathrm{st}}(D)$ is a finite set with cardinality

$$\#W^{\mathrm{st}}(D) \leq g(n - g + 1),$$

and equality holds for a generic divisor class.

Parts (a) and (b) of Theorem 4.20 are connected by showing that $W(D) = W^{\mathrm{st}}(D)$ for a generic divisor class.

The next main theorem describes the distribution of tropical Weierstrass points. Here, note that the condition "$W(D)$ is a finite set" is satisfied for generic $[D] \in \mathrm{Pic}^n(\Gamma)$ by Theorem 4.20.

Theorem 4.24. *Let Γ be a metric graph of genus g, and let $\{D_n : n \geq 1\}$ be a sequence of divisors on Γ with $\deg D_n = n$. Let W_n be the Weierstrass locus of D_n. Suppose each W_n is a finite set, and let*

$$\delta_n = \frac{1}{n} \sum_{x \in W_n} \delta_x$$

denote the normalized discrete measure on Γ associated to W_n (where δ_x is the Dirac measure at x). Then as $n \to \infty$, the measures δ_n converge weakly to the Zhang canonical measure μ on Γ.

The Zhang canonical measure is defined in Section 3.4. We use a different normalization for μ than previous authors; namely we have total measure $\mu(\Gamma) = g$ rather than $\mu(\Gamma) = 1$. We also obtain a quantitative version of this distribution result which specifies a bound on the rate of convergence; see Theorem 4.26.

1.2.2 Torsion points of the Jacobian

As discussed above, a smooth algebraic curve X of genus one with a chosen basepoint $x_0 \in X$ is equipped with a natural additive structure on points of X. Given an algebraic curve with fixed basepoint x_0, we say that x is a *torsion point* if the divisor $n(x - x_0)$ is linearly equivalent to 0 for some positive n. Equivalently, x is a torsion point if the Abel–Jacobi embedding $AJ : X \to \mathrm{Jac}(X)$ (with respect to x_0) sends x to the torsion subgroup $\mathrm{Jac}(X)_{\mathrm{tors}}$ of the Jacobian. The Jacobian of a genus g algebraic curve (over $\mathbb{C}$) is a compact abelian group, isomorphic to $\mathbb{C}^g/\mathbb{Z}^{2g} \cong H^1(X, \mathbb{C})/H_1(X, \mathbb{Z})^\vee$; its torsion subgroup is isomorphic to $\mathbb{Q}^{2g}/\mathbb{Z}^{2g}$.

Faltings's theorem, previously known as Mordell's conjecture, states that a smooth algebraic curve of genus $g \geq 2$ has finitely many rational points, i.e. points whose coordinates are in $\mathbb{Q}$. Motivated by analogy with Mordell's conjecture, Manin and Mumford conjectured that an algebraic curve of genus $g \geq 2$ has finitely many torsion points. This conjecture was proved by Raynaud [39].

Theorem 1.6 (Raynaud; formerly the Manin–Mumford conjecture)**.** *For any smooth algebraic curve X of genus $g \geq 2$, the intersection $AJ(X) \cap \mathrm{Jac}(X)_{\mathrm{tors}}$ is finite.*

The following stronger result remains open, though it is suspected to be true.

Problem 1.7 (Uniform Manin–Mumford bound)**.** Is there a function $N(g)$ such that any smooth algebraic curve X of genus $g \geq 2$ has $\#(AJ(X) \cap \mathrm{Jac}(X)_{\mathrm{tors}}) \leq N(g)$?

See Baker and Poonen [10, p. 111] for discussion of related problems and results; they use the equivalent language of *torsion packets* on curves.

Katz, Rabinoff, and Zureick-Brown made important progress towards resolving Problem 1.7 in [28], where they consider curves defined over a number field K and their K-rational torsion points.

Theorem 1.8 (Katz–Rabinoff–Zureick-Brown [28, Theorem 1.2])**.** *Suppose X is a smooth algebraic curve of genus $g \geq 3$ over a number field K of degree $d = [K : \mathbb{Q}]$. There is an explicit function $N(g, d)$ such that $\#(AJ(\,X(K)\,) \cap \mathrm{Jac}(X)_{\mathrm{tors}}) \leq N(g, d)$.*

When $K = \mathbb{Q}$, for example, they prove that $N(g, 1) = 84g^2 - 98g + 28$ satisfies the above bound. They also prove a more complicated explicit bound $N_\dagger(g, d)$ on the number of torsion points[1] $\#(AJ(X) \cap \mathrm{Jac}(X)_{\mathrm{tors}})$, but only conditional on added assumptions concerning the reduction of X modulo a prime. In the same paper [28], Katz et. al. make progress towards a uniform version of Faltings's' theorem, for curves which satisfy a condition on their Mordell–Weil rank. Tropical geometry (as outlined above and in Chapter 2) plays a major role in their proofs, as well as the

[1] where X is defined over K but $AJ(X)$ is not restricted to K-rational points

p-adic integration theory of Chabauty and Coleman. For more details and related work, we refer the reader to [28, 29] and the references therein.

While Problem 1.7 remains open algebraic curves, we show that—conditionally—there is a nice uniform bound for the number of torsion points on a tropical curve.

Theorem 5.24 (Uniform tropical Manin–Mumford bound)**.** *Let Γ be a connected metric graph of genus $g \geq 2$. If the intersection $AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}$ is finite, then we have the uniform bound*

$$\#(AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}) \leq 3g - 3.$$

However, not all higher-genus metric graphs satisfy the hypotheses of Theorem 5.24. In particular, the finiteness condition is not satisfied if all edge lengths of Γ are rational. This observation is a consequence of the fact that on a graph with unit edge lengths, the degree-0 divisor classes supported on vertices form a finite abelian group, known as the *critical group* of the graph. This means that vertex-supported divisor classes are always torsion; this reasoning can then be repeated on the vertex sets obtained from taking uniform edge-subdivisions of the original graph.

Say a metric graph Γ satisfies the *Manin–Mumford condition* (or is *Manin–Mumford finite*) if the intersection $AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}$ is finite, for every $q \in \Gamma$. We prove the following tropical version of the Manin–Mumford conjecture.

Theorem 5.30. *Let G be a biconnected graph of genus $g \geq 2$. For a very general choice of edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ satisfies the Manin–Mumford condition.*

Recall that a graph G is *biconnected* (or *two-connected*) if G is connected after deleting any vertex. We say that a property holds for a *very general* point of some real parameter space if it holds outside of a countable collection of proper Zariski-closed subsets. In this theorem, it suffices that the edge lengths of Γ do not satisfy any integer-coefficient polynomial relation of degree at most $g - 1$.

We can ask the same question about torsion points in the image of the *higher-degree Abel–Jacobi map* $AJ_D^{(d)} : \Gamma^d \to \mathrm{Jac}(\Gamma)$, defined by

$$(x_1, \ldots, x_d) \mapsto [\sum_{i=1}^{d} x_i - D].$$

We say a metric graph Γ satisfies the *degree d Manin–Mumford condition* (or, is *Manin–Mumford finite in degree d*) if $AJ_D^{(d)}(\Gamma^d)$ intersects only finitely many torsion points of $\mathrm{Jac}(\Gamma)$, for every $D \in \mathrm{Sym}^d(\Gamma)$. When $d = 1$, this is the usual Manin–Mumford condition. If the degree d Manin–Mumford condition holds, then it also holds in degree d' for any $1 \leq d' \leq d$.

Given that a metric graph Γ of genus $g \geq 1$ satisfies the degree d Manin–Mumford condition, it is straightforward to show that $d < g$. We show that this naive bound can be improved to $d < C \log g$ for an explicit[2] constant C (Corollary 5.39). The argument is to find a combinatorial invariant $\gamma^{\text{ind}} = \gamma^{\text{ind}}(\Gamma)$ such that the Manin–Mumford degree must satisfy $d < \gamma^{\text{ind}}$, and to show that $\gamma^{\text{ind}} < C \log g$.

We call γ^{ind} the *independent girth* of a graph, which we define as

$$\gamma^{\text{ind}}(G) = \min_C \{ \operatorname{rk}^{\perp}(E(C)) \}$$

where the minimum is taken over all cycles C in G, and $\operatorname{rk}^{\perp}$ denotes the rank function of the cographic matroid $M^{\perp}(G)$. Recall that the *girth* of a graph is the minimal length of a cycle, $\gamma(G) = \min_C \{\# E(C)\}$; it follows that $\gamma^{\text{ind}} \leq \gamma$. In contrast to girth, the independent girth is invariant under subdivision of edges, so $\gamma^{\text{ind}}(\Gamma)$ is well-defined for a metric graph Γ.

If $\Gamma = (G, \ell)$ is Manin–Mumford finite in degree d, we first observe that $d < \gamma(G)$ and then improve this bound to $d < \gamma^{\text{ind}}(G)$. The bound $d < \gamma^{\text{ind}}$ is sharp in the following sense.

Theorem 5.38. *Let G be a finite connected graph of genus $g \geq 1$ with independent girth γ^{ind}. For a very general choice of edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ is Manin–Mumford finite in degree d if and only if $1 \leq d < \gamma^{\text{ind}}$.*

We also prove a conditional uniform Manin–Mumford bound in the higher-degree case, see Theorem 5.25.

1.3 Outline

In Chapter 2, we review background material on metric graphs and their divisor theory. In Chapter 3, we review the interpretation of a metric graph as an electrical resistor network, define Zhang's canonical measure, and give Kirchhoff's formulas for the voltage function in terms of weighted sums over spanning trees. In Chapter 4, we define the Weierstrass locus and stable Weierstrass locus for a divisor on a metric graph, give examples, and we prove that $W(D)$ is generically finite and compute its cardinality. We then prove results on the distribution of Weierstrass points on a metric graph. The results in this chapter appeared earlier in the preprint [40]. In Chapter 5, we prove results on the Jacobian torsion points of a metric graph. We give a tropical analogue of Raynaud's theorem, and give a uniform bounds on the number of torsion points assuming very general edge lengths.

[2] $C = 4/\log 2 \approx 5.771$

CHAPTER 2

Tropical Curves

In this section we define metric graphs and linear equivalence of divisors on metric graphs. We use the terms "metric graph" and "abstract tropical curve" interchangeably. We recall the Baker–Norine rank of a divisor, and state the Riemann–Roch theorem which is satisfied by this rank function.

2.1 Metric graphs and divisors

A *metric graph* is a compact, connected metric space which comes from assigning positive real edge lengths to a finite connected combinatorial graph. Namely, we construct a metric graph Γ by taking a finite set of edges $E = \{e_i\}$, each isometric to a real interval $e_i = [0, L_i]$ of length $L_i > 0$, gluing their endpoints to a finite set of vertices V, and imposing the path metric. The underlying combinatorial graph $G = (E, V)$ is called a *combinatorial model* for Γ. We allow loops and parallel edges in a combinatorial graph G. We say e is a *segment* of Γ if it is an edge in some combinatorial model.

The *valence* $\mathrm{val}(x)$ of a point x on a metric graph Γ is defined to be the number on connected components of a sufficiently small punctured neighborhood of x. Points in the interior of a segment of Γ always have valence 2. All points x with $\mathrm{val}(x) \neq 2$ are contained in the vertex set of any combinatorial model.

The *genus* of a metric graph Γ is its first Betti number as a topological space,

$$g(\Gamma) = b_1(\Gamma) = \dim_{\mathbb{R}} H_1(\Gamma, \mathbb{R}).$$

If G is a combinatorial model for Γ, the genus is equal to $g(\Gamma) = \#E(G) - \#V(G) + 1$.

Example 2.1. The metric graph on the left of Figure 2.1 has genus 0. A minimal combinatorial model has 8 vertices and 7 edges.

Example 2.2. The metric graph on the right of Figure 2.1 has genus 2. A minimal combinatorial model has 2 vertices and 3 edges.

Figure 2.1: Metric graphs of genus 0 (left) and genus 2 (right).

A *divisor* on a metric graph Γ is a finite formal sum of points of Γ with integer coefficients. The *degree* of a divisor is the sum of its coefficients; i.e. for the divisor $D = \sum_{x \in \Gamma} a_x x$, we have $\deg(D) = \sum_{x \in \Gamma} a_x$. We let $\mathrm{Div}(\Gamma)$ denote the set of all divisors on Γ, and let $\mathrm{Div}^d(\Gamma)$ denote the divisors of degree d. We say a divisor is *effective* if all of its coefficients are non-negative; we write $D \geq 0$ to indicate that D is effective. More generally, we write $D \geq E$ to indicate that $D - E$ is an effective divisor. We let $\mathrm{Sym}^d(\Gamma)$ denote the set of effective divisors of degree d on Γ. $\mathrm{Sym}^d(\Gamma)$ inherits from Γ the structure of a polyhedral cell complex of dimension d.

We let $\mathrm{Div}_{\mathbb{R}}(\Gamma)$ denote the set of divisors on Γ with coefficients in $\mathbb{R}$. In other words, $\mathrm{Div}_{\mathbb{R}}(\Gamma) = \mathrm{Div}(\Gamma) \otimes_{\mathbb{Z}} \mathbb{R}$.

2.2 Principal divisors and linear equivalence

We define linear equivalence for divisors on metric graphs, following Gathmann–Kerber [20] and Mikhalkin–Zharkov [32]. This notion is analogous to linear equivalence of divisors on an algebraic curve, where rational functions are replaced with piecewise $\mathbb{Z}$-linear functions.

A *piecewise linear function* on Γ is a continuous function $f : \Gamma \to \mathbb{R}$ such that there is some combinatorial model for Γ such that f restricted to each edge is a linear function, i.e. a function of the form

$$f(x) = ax + b, \qquad a, b \in \mathbb{R},$$

where x is a length-preserving parameter on the edge. We let $\mathrm{PL}_{\mathbb{R}}(\Gamma)$ denote the set of all piecewise linear functions on Γ.

A *piecewise $\mathbb{Z}$-linear function* on Γ is a piecewise linear function such that all its slopes are integers, i.e. f restricted to each edge has the form

$$f(x) = ax + b, \qquad a \in \mathbb{Z},\ b \in \mathbb{R}$$

(for some combinatorial model). We let $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ denote the set of all piecewise $\mathbb{Z}$-linear functions on Γ. The functions $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ are closed under the operations of addition, multiplication by $\mathbb{Z}$, and taking pairwise max and min.

We let $UT_x\Gamma$ denote the unit tangent fan of Γ at x, which is the set of "directions going away from x" on Γ. For $v \in UT_x\Gamma$, the symbol ϵv for sufficiently small $\epsilon \geq 0$ means the point in Γ that is distance ϵ away from x in the direction v. For $v \in UT_x\Gamma$ and a function $f : \Gamma \to \mathbb{R}$ we let

$$D_v f(x) = \lim_{\epsilon \to 0^+} \frac{f(x + \epsilon v) - f(x)}{\epsilon}$$

denote the slope of f while travelling away from x in the direction v (if it exists).

Given $f \in \mathrm{PL}_{\mathbb{Z}}(\Gamma)$, we define the *principal divisor* $\Delta(f) \in \mathrm{Div}^0(\Gamma)$ by

$$\Delta(f) = \sum_{x \in \Gamma} a_x x \qquad \text{where} \qquad a_x = \sum_{v \in UT_x\Gamma} D_v f(x). \tag{2.1}$$

In words, the coefficient in $\Delta(f)$ of a point x is equal to the sum of the outgoing slopes of f at x. On a given segment, this divisor is supported on the finite set of points at which f is not linear, sometimes called the "break locus" of f. If $\Delta(f) = D - E$ where D, E are effective divisors with disjoint support, then we call $D = \Delta^+(f)$ the *divisor of zeros of* f and $E = \Delta^-(f)$ the *divisor of poles* of f.

We say two divisors D, E are *linearly equivalent*, denoted $D \sim E$, if there exists a piecewise $\mathbb{Z}$-linear function f such that

$$\Delta(f) = D - E.$$

Note that linearly equivalent divisors must have the same degree. We let $[D]$ denote the linear equivalence class of divisor D, i.e.

$$[D] = \{E \in \mathrm{Div}(\Gamma) : E \sim D\} = \{D + \Delta(f) : f \in \mathrm{PL}_{\mathbb{Z}}(\Gamma)\}.$$

We say a divisor class $[D]$ is *effective*, or write $[D] \geq 0$, if there is an effective representative $E \sim D$, $E \geq 0$ in the equivalence class.

We let $|D|$ denote the *(complete) linear system* of D, which is the set of effective divisors linearly equivalent to D. We have

$$\begin{aligned} |D| &= \{E \in \mathrm{Div}(\Gamma) : E \sim D,\, E \geq 0\} \\ &= \{D + \Delta(f) : f \in \mathrm{PL}_{\mathbb{Z}}(\Gamma),\, \Delta(f) \geq -D\}. \end{aligned}$$

Unlike $[D]$, the linear system $|D|$ is naturally a compact polyhedral complex, with topology induced by the inclusion $|D| \subset \mathrm{Sym}^d(\Gamma)$.

Remark 2.3. The map $\Delta : \mathrm{PL}_{\mathbb{Z}}(\Gamma) \to \mathrm{Div}(\Gamma)$ is also known as the *metric graph Laplacian* on Γ. This comes from identifying $\mathrm{Div}(\Gamma)$ with the space of integer-valued discrete measures on Γ, via

$$D = \sum_{i=1}^n a_i x_i \quad \longleftrightarrow \quad \delta = \sum_{i=1}^n a_i \delta_{x_i}$$

so that $\Delta(f)$ coincides with the (distributional) second derivative $-\frac{d^2}{dx^2}f(x)$. (The second derivative here must be extended to accommodate points of valence $\neq 2$.) The definition of the Laplacian on piecewise $\mathbb{Z}$-linear functions, equation (2.1), naturally extends to arbitrary piecewise linear functions on Γ with real slopes, if we also allow real-valued coefficients in the divisor $\Delta(f)$. This yields a map

$$\mathrm{PL}_{\mathbb{R}}(\Gamma) \xrightarrow{\Delta} \mathrm{Div}_{\mathbb{R}}(\Gamma)$$

from the space of all piecewise-linear functions of Γ, to $\mathbb{R}$-valued discrete measures on Γ. The kernel of this map is the set of constant functions. The cokernel of this map is simply the degree function $\mathrm{Div}_{\mathbb{R}}(\Gamma) \xrightarrow{\deg} \mathbb{R}$. We will see why this is the cokernel in Section 3.1 on voltage functions. This fits in the short exact sequence

$$0 \to \mathbb{R} \xrightarrow{\text{const}} \mathrm{PL}_{\mathbb{R}}(\Gamma) \xrightarrow{\Delta} \mathrm{Div}_{\mathbb{R}}(\Gamma) \xrightarrow{\deg} \mathbb{R} \to 0.$$

(Compare to the integral case, where the short exact sequence is

$$0 \to \mathbb{R} \xrightarrow{\text{const}} \mathrm{PL}_{\mathbb{Z}}(\Gamma) \xrightarrow{\Delta} \mathrm{Div}(\Gamma) \to \mathrm{Pic}(\Gamma) \to 0,$$

with $\mathrm{Pic}(\Gamma) \cong \mathbb{Z} \times (S^1)^g$; see Section 2.3 below.)

Remark 2.4 (Linear equivalence as chip firing)**.** We sometimes speak of a degree n effective divisor on Γ as a collection of n "chips" placed on Γ. Changing the divisor D to a linearly equivalent divisor D' can be achieved through a sequence of "chip firing moves" where we choose and simple cut[1] of Γ consisting of m segments of length ϵ, and on each edge move a chip from one end to the other. The piecewise-linear

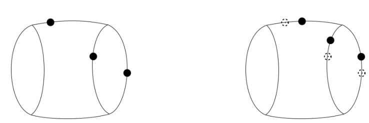

Figure 2.2: Chip firing across an elementary cut.

function associated to such a chip firing move has slope 0 outside the cut segments, and slope 1 on the cut segments. For more discussion of chip-firing see [4, Remark 2.2], [9, Section 1.5] and the references therein.

Remark 2.5 (Linear interpolation along f)**.** Given a function $f \in \mathrm{PL}_{\mathbb{Z}}(\Gamma)$, we may associate to f a 1-parameter family of effective divisors which "linearly interpolate" between the zeros $\Delta^+(f)$ and poles $\Delta^-(f)$. We can think of this construction as specifying a unique "geodesic path" between any two points in the complete linear system $|D|$. This notion previously appeared in [31] under the name *t-path*.

[1] A *simple cut* is a collection of segments of Γ such that removing the interiors of these segments disconnects Γ into exactly two components.

Namely, for $\lambda \in \mathbb{R}$ we let $\lambda \in \mathrm{PL}_{\mathbb{Z}}(\Gamma)$ also denote the constant function on Γ by abuse of notation, and we define the effective divisor $f_{\Delta}^{-1}(\lambda)$ by

$$f_{\Delta}^{-1}(\lambda) = \Delta^{-}(f) + \Delta(\max\{f, \lambda\}).$$

See Figure 2.3 for an illustration. Note that according to this definition, $f_{\Delta}^{-1}(\lambda) = \Delta^{-}(f)$ for λ sufficiently large and $f_{\Delta}^{-1}(\lambda) = \Delta^{+}(f)$ for λ sufficiently small. It is clear from definition that for any λ, $f_{\Delta}^{-1}(\lambda)$ is linearly equivalent to $\Delta^{+}(f)$ and to $\Delta^{-}(f)$.

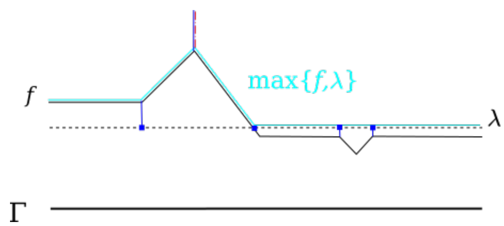

Figure 2.3: Linear interpolation showing the divisor $f_{\Delta}^{-1}(\lambda)$.

2.3 Picard group and Jacobian

Let $\mathrm{Pic}(\Gamma)$ denote the *Picard group* of Γ, which is the abelian group of all linear equivalence classes of divisors on Γ. The addition operation on $\mathrm{Pic}(\Gamma)$ is induced from addition of divisors in $\mathrm{Div}(\Gamma)$. In other words, $\mathrm{Pic}(\Gamma)$ is the cokernel of the map Δ sending a piecewise $\mathbb{Z}$-linear function to its associated principal divisor:

$$\mathrm{PL}_{\mathbb{Z}}(\Gamma) \xrightarrow{\Delta} \mathrm{Div}(\Gamma) \to \mathrm{Pic}(\Gamma) \to 0.$$

The kernel of Δ is the set of constant functions on Γ.

Since the degree of a divisor class is well-defined, we have a disjoint union decomposition

$$\mathrm{Pic}(\Gamma) = \bigsqcup_{d \in \mathbb{Z}} \mathrm{Pic}^{d}(\Gamma),$$

where $\mathrm{Pic}^{d}(\Gamma)$ consists of divisor classes of degree d. The degree-0 component $\mathrm{Pic}^{0}(\Gamma)$ is known as the *Jacobian* of Γ, denoted $\mathrm{Jac}(\Gamma) := \mathrm{Pic}^{0}(\Gamma)$. The Jacobian $\mathrm{Jac}(\Gamma)$ is a compact abelian group.

Theorem 2.6 (Abel–Jacobi theorem for metric graphs)**.** *Let Γ be a metric graph of genus g. Then there is an isomorphism of compact abelian topological groups.*

$$\mathrm{Jac}(\Gamma) \cong (S^1)^{\times g} = \overbrace{S^1 \times \cdots \times S^1}^{g}.$$

Proof. See Mikhalkin–Zharkov [32]. The proof follows the same idea as the classical Abel–Jacobi theorem, to show that $\mathrm{Pic}^{0}(\Gamma) = H^1(\Gamma, \mathbb{R})/H_1(\Gamma, \mathbb{Z})^{\vee} \cong \mathbb{R}^g/\mathbb{Z}^g$. $\square$

Addition of divisor classes induces an action of $\mathrm{Jac}(\Gamma)$ on $\mathrm{Pic}^d(\Gamma)$, for any fixed degree $d \in \mathbb{Z}$. Since $\mathrm{Pic}^d(\Gamma)$ is a torsor (or principal homogeneous space) for $\mathrm{Jac}(\Gamma)$ under this action, the Abel–Jacobi theorem also implies there are homeomorphisms $\mathrm{Pic}^d(\Gamma) \cong (S^1)^{\times g}$.

We let $\mathrm{Eff}^d(\Gamma)$ denote the set of divisor classes on Γ of degree d which have an effective representative. In other words, $\mathrm{Eff}^d(\Gamma)$ is the image of $\mathrm{Sym}^d(\Gamma)$ under the (degree-d restriction of the) cokernel map $\mathrm{Div}(\Gamma) \to \mathrm{Pic}(\Gamma)$:

$$\begin{tikzcd}
\mathrm{Sym}^d(\Gamma) \arrow[r] \arrow[d] & \mathrm{Div}^d(\Gamma) \arrow[d, "\mathrm{coker}\,\Delta"] \\
\mathrm{Eff}^d(\Gamma) \arrow[r] & \mathrm{Pic}^d(\Gamma).
\end{tikzcd}$$

The space $\mathrm{Eff}^d(\Gamma)$ is naturally a polyhedral complex of pure dimension d when $0 \leq d \leq g$ (see Gross et. al. [22]). In degree $d \geq g$, we have $\mathrm{Eff}^d(\Gamma) = \mathrm{Pic}^d(\Gamma)$, i.e. every divisor class has an effective representative. This fact follows from the theory of break divisors; see Section 2.5 below.

As a particularly important case, the *theta divisor* $\Theta = \Theta(\Gamma)$ is $\Theta = \mathrm{Eff}^{g-1}(\Gamma)$, which lives inside $\mathrm{Pic}^{g-1}(\Gamma)$ as a codimension 1 polyhedral complex. Another important case is in degree 1; $\mathrm{Eff}^1(\Gamma)$ is the image of the map $\Gamma \to \mathrm{Pic}^1(\Gamma)$ which sends a point x to the divisor class $[x]$. If Γ has no bridge edge, then the map $\Gamma \to \mathrm{Eff}^1(\Gamma)$ is a homeomorphism. This allows us to think of the metric graph Γ as a subset of $\mathrm{Pic}^1(\Gamma)$ in a canonical way.

There is a standard way to map a metric graph to its Jacobian, which depends on a choice of basepoint. Given a choice of basepoint $q \in \Gamma$, the *Abel–Jacobi map* is defined by

$$AJ_q : \Gamma \to \mathrm{Jac}(\Gamma) \qquad x \mapsto [x - q]. \tag{2.2}$$

2.4 Reduced divisors

A divisor class $[D]$ is typically very large, so it is convenient to have a method of choosing a (somewhat-)canonical representative divisor inside $[D]$. When D has arbitrary degree, we can do so after fixing a basepoint q on our metric graph Γ, using the q-reduced divisor construction.

Given a point $q \in \Gamma$, the *q-reduced divisor* $\mathrm{red}_q[D]$ is the unique divisor in $[D]$ which is effective away from q, and which minimizes a certain energy function among such representatives. Intuitively, $\mathrm{red}_q[D]$ is the divisor in $[D]$ whose chips are "as close as possible" to the basepoint q. We defer giving the full definition until Section

3.2, following [11, Appendix A]. For now, we state these important properties of the reduced divisor:

(RD1) $[D] \geq 0$ if and only if $\mathrm{red}_q[D] \geq 0$

(RD2) for any integer m, $\mathrm{red}_q[mq + D] = mq + \mathrm{red}_q[D]$

(RD3) the degree of $\mathrm{red}_q[D]$ away from q is at most g, the genus of Γ (follows from Riemann's inequality, Corollary 2.15)

(RD4) for a fixed effective divisor D, the map $\Gamma \to |D|$ sending $q \mapsto \mathrm{red}_q[D]$ is continuous (due to Amini [2, Theorem 3]).

2.5 Break divisors and ABKS decomposition

When a divisor D has degree g, there is a canonical representative of $[D]$ without any choice of basepoint, using the concept of break divisor. This notion was introduced by Mikhalkin–Zharkov [32] and studied extensively by An–Baker–Kuperberg–Shokrieh [4]. We review some of their results in this section.

A *break divisor* is an effective divisor of degree g (the genus) which can be constructed in the following manner: choose a combinatorial model $G = (V, E)$ for Γ and choose a spanning tree T of G, then place one chip on each edge in the complement $E \backslash E(T)$. (Note that $E \backslash E(T)$ contains exactly g edges.) Placing a chip on the endpoint of an edge is allowed.

The set of break divisors does not depend on the choice of combinatorial model. We use $\mathrm{Br}^g(\Gamma)$ to denote the set of all break divisors on Γ. We may view $\mathrm{Br}^g(\Gamma)$ as a topological space, using the topology induced from the inclusion in $\mathrm{Sym}^g(\Gamma)$.

Example 2.7. In Figure 2.4 we show three examples of break divisors, on the left, and three examples of non-break divisors, on the right, on a genus 3 metric graph.

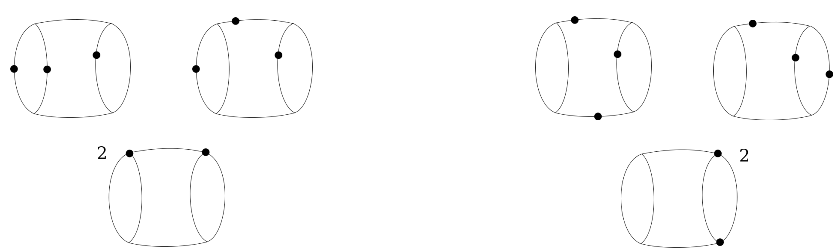

Figure 2.4: Break divisors and non-break divisors.

For a divisor class $[D]$ whose degree is g, the genus of the underlying curve, there is a unique representative of $[D]$ which is a break divisor.

Theorem 2.8 (see [4, Theorem 1.1], [32, Corollary 6.6])**.** *Let Γ be a metric graph of genus g.*

(a) Every divisor class $[D] \in \mathrm{Pic}^g(\Gamma)$ *contains a unique break divisor, which we denote* $\mathrm{br}[D]$.

(b) The map $\mathrm{br} : \mathrm{Pic}^g(\Gamma) \to \mathrm{Sym}^g(\Gamma)$ *sending a divisor class to its break divisor representative is continuous and injective. Its image is the space of all break divisors* $\mathrm{Br}^g(\Gamma)$.

(c) The map $\mathrm{br} : \mathrm{Pic}^g(\Gamma) \to \mathrm{Sym}^g(\Gamma)$ *is the unique continuous section of the map* $[-] : \mathrm{Sym}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma)$ *taking an effective divisor to its linear equivalence class. Namely,* br *is the unique continuous map such that the composition*

$$\mathrm{Pic}^g(\Gamma) \xrightarrow{\mathrm{br}} \mathrm{Sym}^g(\Gamma) \xrightarrow{[-]} \mathrm{Pic}^g(\Gamma)$$

is the identity homeomorphism.

If we choose a combinatorial model (G, ℓ) for the metric graph Γ, An–Baker–Kuperberg–Shokrieh [4] showed that the theory of break divisors implies a nice combinatorial decomposition of $\mathrm{Pic}^g(\Gamma)$. ($\mathrm{Pic}^g(\Gamma)$ is defined in Section 2.3.)

Theorem 2.9 (ABKS decomposition, see [4, Section 3.2])**.** *Suppose* $\Gamma = (G, \ell)$ *is a metric graph with a combinatorial model. Let* $\mathcal{T}(G)$ *denote the set of spanning trees of* G*. Then*

$$\mathrm{Pic}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} C_T$$

where

$$C_T = \{[x_1 + \cdots + x_g] : E(G) \backslash E(T) = \{e_1, \ldots, e_g\},\ x_i \in e_i\}$$

denotes the set of divisor classes represented by summing a point from each edge of G *not in* T*. The cells* C_T *have disjoint interiors, as* $T \in \mathcal{T}(G)$ *varies.*

For fixed T*, if we parametrize each edge* $e_i \notin E(T)$ *as the closed real interval* $[0, \ell(e_i)]$*, there is a natural surjective map* $\prod_{i=1}^g [0, \ell(e_i)] \to C_T$*. This map always restricts to a homeomorphism on the respective interiors* $\prod_{i=1}^g (0, \ell(e_i)) \to C_T^\circ$*, but may be non-injective on the boundary.*

The proof is to combine Theorem 2.8 with the definition of break divisor, using the auxiliary data of the spanning tree. Since $\mathrm{Pic}^g(\Gamma)$ is canonically homeomorphic to $\mathrm{Br}^g(\Gamma)$, we may view Theorem 2.9 as a decomposition of $\mathrm{Br}^g(\Gamma)$.

Remark 2.10. If we take the combinatorial model for Γ to be sufficiently subdivided, then for each $T = G \backslash \{e_1, \ldots, e_g\}$, the surjection $\prod_{i=1}^g [0, \ell(e_i)] \to C_T$ is a (global) homeomorphism. In particular, for this to hold it suffices that G has girth $> g$ (i.e. every cycle contains more than g edges). A necessary condition is that G has no loops or parallel edges (if $g \geq 2$).

Example 2.11. Consider the metric graph shown on the left side of Figure 2.5. Its minimal combinatorial model $\Gamma = (G, \ell)$ contains two vertices and three edges. The associated ABKS decomposition of $\mathrm{Pic}^2(\Gamma)$ is shown on the right side of Figure 2.5; segments on the boundary are glued to the parallel boundary segment. There are three cells, corresponding to the three spanning trees in G.

Here $\mathrm{Pic}^2(\Gamma)$ is homeomorphic to a torus (cf. Theorem 2.6). Each cell C_T is homeomorphic to a rectangle with a pair of opposite vertices glued together.

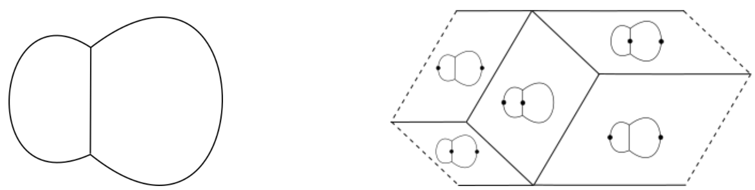

Figure 2.5: Metric graph of genus 2 (left), with ABKS decomposition of $\mathrm{Pic}^2(\Gamma)$.

Proposition 2.12. *Let $q \in \Gamma$ be an arbitrary basepoint on a genus g metric graph.*

(a) For a generic divisor class $[D]$ of degree g, the reduced divisor $\mathrm{red}_q[D]$ is equal to the break divisor $\mathrm{br}[D]$.

(b) For a generic divisor class $[D]$ of degree n, the reduced divisor $\mathrm{red}_q[D]$ is equal to

$$\mathrm{red}_q[D] = (n - g)q + E$$

where E is a break divisor.

Proof. (a) This follows from [4, Lemma 3.5], which states that inside an open cell of the ABKS decomposition, a divisor class $[D]$ has only a single effective representative.

(b) This follows from (a) and the property (2.4) that taking the q-reduced representative is equivariant with respect to adding a multiple of q. $\square$

A *semibreak divisor* is an effective divisor which is a "partial sum" of a break divisor, in the sense that $E \in \mathrm{Sym}^d(\Gamma)$ is a semibreak divisor if

$$E + E' \text{ is a break divisor, for some } E' \in \mathrm{Sym}^{g-d}(\Gamma).$$

In contrast to the case $d = g$, when $0 \le d < g$ an effective divisor class $[E] \in \mathrm{Eff}^d(\Gamma)$ may have more than one semibreak representative. However, every divisor class in this range has at least one semibreak representative.

Theorem 2.13. *On a metric graph Γ of genus g, suppose $0 \le d \le g$. Any effective divisor of degree d is linearly equivalent to a semibreak divisor of degree d.*

Proof. This is a result of Gross–Shokrieh–Tóthmérész; see [22, Theorem A]. $\square$

2.6 Rank and Riemann–Roch

We recall the definition of the rank of a divisor on a metric graph, originally due to Baker and Norine [9] for divisors on a combinatorial graph, and extended to metric graphs by Gathmann–Kerber [20] and Mikhalkin–Zharkov [32]. The rank function is a natural way to extend the important distinction between effective and non-effective divisor classes on a metric graph. Divisor classes with larger rank are in a sense "further away" from the set of non-effective divisor classes, where distance between divisors is given by adding or subtracting single points.

The *rank* $r(D)$ of a divisor D on Γ is defined as

$$r(D) = \max\{r \geq 0 : [D - E] \geq 0 \text{ for all } E \in \mathrm{Sym}^r(\Gamma)\}$$

if $[D]$ is effective, and $r(D) = -1$ otherwise. Equivalently,

$$r(D) = \begin{cases} -1 & \text{if } [D] \text{ is not effective,} \\ 1 + \min_{x \in \Gamma}\{r(D - x)\} & \text{if } [D] \text{ is effective.} \end{cases}$$

This second definition inductively gives the rank of a divisor in terms of divisors of smaller degree; the base case is the set of non-effective divisor classes.[2] Note that the rank of a divisor D depends only on its linear equivalence class.

The *canonical divisor* on a metric graph Γ is defined as

$$K = \sum_{x \in \Gamma} (\mathrm{val}(x) - 2) \cdot x.$$

The degree of the canonical divisor is $\deg K = 2g - 2$, which agrees with the canonical divisor on an algebraic curve.

Theorem 2.14 (Riemann-Roch for metric graphs)**.** *Let Γ be a metric graph of genus g, and let K be the canonical divisor on Γ. For any divisor D on Γ,*

$$r(D) - r(K - D) = \deg(D) + 1 - g.$$

Proof. See Gathmann–Kerber [20, Proposition 3.1] and Mikhalkin–Zharkov [32, Theorem 7.3], which both adapt the arguments of Baker–Norine [9] for the case of combinatorial graphs. □

Corollary 2.15 (Riemann's inequality for metric graphs)**.** *For a divisor D on a metric graph of genus g,*

$$r(D) \geq \deg(D) - g.$$

[2] By Riemann's inequality, Corollary 2.15, a non-effective divisor class has degree at most $g - 1$.

Proof. This follows from Riemann–Roch since $r(K - D) \geq -1$. □

By Riemann's inequality, combined with the bound $r(D) \geq -1$ immediate from the definition of rank, any divisor D satisfies $r(D) \geq \max\{\deg(D) - g, -1\}$. We say D is *nonspecial* if $r(D) = \max\{\deg(D) - g, -1\}$, and *special* otherwise.

2.7 Matroids

In this section we review the definition of a matroid. In particular, we recall the graphic matroid and cographic matroid associated to a connected graph. Cographic matroids will be useful for understanding the structure of the Jacobian of a metric graph. For a complete reference on matroids, see [37] or [27].

A *matroid* $M = (E, \mathcal{B})$ is a finite set E equipped with a collection $\mathcal{B} \subset 2^E$ of subsets of E, called the *bases* of the matroid, satisfying the basis exchange axiom: for distinct subsets $B_1, B_2 \in \mathcal{B}$, there exists some $x \in B_1 \backslash B_2$ and $y \in B_2 \backslash B_1$ such that $(B_1 \backslash x) \cup y \in \mathcal{B}$. In other words, we can produce a new basis by exchanging an element of B_1 with an element of B_2.

An *independent set* of a matroid $M = (E, \mathcal{B})$ is a subset of E which is a subset of some basis. A *cycle* of M is a subset of E which is minimal among non-independent sets, under the inclusion relation. The *rank* of a subset $A \subset E$ is the cardinality of a maximal independent set contained in A; we denote this by $\mathrm{rk}(A)$ or $\mathrm{rk}_M(A)$.

Given a graph $G = (V, E)$, the *graphic matroid* $M(G)$ is the matroid on the ground set $E = E(G)$ with bases $\mathcal{B} = \{E(T) : T \text{ is a spanning tree of } G\}$. An independent set in $M(G)$ is a subset of edges which span an acyclic subgraph. (i.e. $h^1(G|A) = 0$.) A cycle in $M(G)$ is a cycle in the graph-theoretic sense, i.e. a subset of edges which span a subgraph homeomorphic to a circle. The graphic matroid $M(G)$ is also known as the *cycle matroid* of G.

Example 2.16. Suppose G is the *Wheatstone graph* shown in Figure 2.6. The bases of $M(G)$ are $\{abd, abe, acd, ace, ade, bcd, bce, bde\}$. The cycles are $\{abc, abde, cde\}$. (Here abc is shorthand for the set $\{a, b, c\}$.)

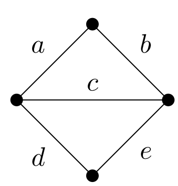

Figure 2.6: Wheatstone graph.

Given a graph $G = (V, E)$, the *cographic matroid* $M^{\perp}(G)$ is the matroid on the ground set $E = E(G)$ whose bases are complements of spanning trees of G. An independent set in $M^\perp(G)$ is a set of edges whose removal does not disconnect G (i.e. a set $A \subset E$ such that $G\backslash A$ is connected, equivalently $h^0(G\backslash A) = 1$).[3] A cycle in $M^\perp(G)$ is a minimal set of edges A such that $h^0(G\backslash A) = 2$; this is called a *simple cut* or a *bond* of G. The cographic matroid is also known as the *cocycle matroid* or *bond matroid* of G. For more on cographic matroids, see [37, Chapter 2.3].

Note: when discussing the graphic or cographic matroid of a graph G, we always use "cycle of G" to refer to a cycle in the *graphic* matroid sense. We use the terms "cycle" and "simple cycle" for a graph interchangeably.

Example 2.17. Suppose G is the Wheatstone graph, shown in Figure 2.6. The bases of the cographic matroid $M^\perp(G)$ are $\{ac, ad, ae, bc, bd, be, cd, ce\}$. The cycles of $M^\perp(G)$ are $\{ab, acd, ace, bcd, bce, de\}$.

A consequence of Mikhalkin and Zharkov's proof [32] of the tropical Abel–Jacobi theorem (Theorem 2.6) is that the Abel–Jacobi map $\Gamma \to \mathrm{Jac}(\Gamma)$ is linear on each edge of Γ. The universal cover of $\mathrm{Jac}(\Gamma)$ is naturally identified with $H^1(\Gamma, \mathbb{R})$. The Abel–Jacobi map, restricted to a single edge $e \subset \Gamma$, lifts locally to $e \to H^1(\Gamma, \mathbb{R})$. The structure of the edge-vectors in the image $\Gamma \to \mathrm{Jac}(\Gamma)$ is exactly recorded by the cographic matroid $M^\perp(G)$, for any combinatorial model $\Gamma = (G, \ell)$.

Definition 2.18. Let $\Gamma = (G, \ell)$ be a metric graph. Given edges $e_1, \dots, e_k \in E(G)$, let $\mathrm{Div}(e_1, \dots, e_k) \subset \mathrm{Div}^k(\Gamma)$ denote the set of effective divisors formed by adding together one point from each edge e_i. Let $\mathrm{Eff}(e_1, \dots, e_k)$ denote the corresponding set of effective divisor classes,

$$\mathrm{Eff}(e_1, \dots, e_k) = \{[x_1 + \cdots + x_k] : x_i \in e_i\} \subset \mathrm{Pic}^k(\Gamma).$$

Theorem 2.19. *Let $\Gamma = (G, \ell)$ be a metric graph. The dimension of $\mathrm{Eff}(e_1, \dots, e_k)$ is equal to the rank of $\{e_1, \dots, e_k\}$ in the cographic matroid $M^\perp(G)$.*

Proof. For each edge $e_i \in E(G)$, let $v_i \in H^1(\Gamma, \mathbb{R})$ denote a vector parallel to the Abel–Jacobi image of e_i in $\mathrm{Jac}(\Gamma)$. Then according to Definition 5.1.3 of [15, p. 156], the set of vectors $\{v_i : e_i \in E(G)\}$ form a realization of the cographic matroid $M^\perp(G)$. This means that the cographic rank of $\{e_1, \dots, e_k\}$ agrees with the dimension of the linear span of $\{v_1, \dots, v_k\}$.

The subset $\mathrm{Eff}(e_1, \dots, e_k) \subset \mathrm{Pic}^k(\Gamma)$ is naturally identified with the Minkowski sum of the corresponding vectors $v_1, \dots, v_k \in H^1(\Gamma, \mathbb{R})$, so the claim follows. $\square$

Corollary 2.20. *Let $\Gamma = (G, \ell)$ be a metric graph of genus g. For any integer d in the range $0 \le d \le g$, the space $\mathrm{Eff}^d(\Gamma)$ of degree d effective divisor classes has the structure of a cellular complex whose top-dimensional cells are indexed by independent sets of size d in the cographic matroid $M^\perp(G)$.*

[3] $A \subset E(G)$ is called a *cut* of G if $G\backslash A$ is disconnected.

CHAPTER 3

Resistor Networks

In this section we view a metric graph as a resistor network, where each edge is a resistor whose resistance is equal to the length of the edge. This allows us to derive useful properties of the local and global structure of the metric graph.

We define the Zhang canonical measure on a metric graph (due to Zhang [42]) via the perspective of resistor networks following Baker–Faber [7].

3.1 Voltage function

We view a metric graph Γ as a resistor network by interpreting an edge of length L as a resistor of resistance L. Note that this is well-defined on a metric graph due to the series rule for combining resistances, so we have compatibility with subdividing an edge into edges of shorter length. This interpretation is not only mathematically convenient, but physically honest—the electrical resistance of a wire is directly proportional to its length, a fact known as Pouillet's law.

On a resistor network we may send current from one point to another. On a given segment, the voltage drop across the segment is equal to the resistance (i.e. length) of the segment multiplied by the amount of current passing through the segment—this is Ohm's law.

Under an externally-applied current, the flow of current within the network is determined by Kirchhoff's circuit laws: the current law says that the sum of directed currents out of any point is equal to zero (accounting for external currents), and the voltage law says that the sum of directed voltage differences around any closed loop is equal to zero. Our convention is that current flows from higher voltage to lower voltage.

It is a well-known empirical fact that Kirchhoff's circuit laws can be solved uniquely for any externally-applied current flow which satisfies conservation of current (i.e. internal current flows are unique). To some, it is also a well-known mathematical result. This is expressed in the following two definitions.

Definition 3.1 (Physics version)**.** Given points $y, z \in \Gamma$, the *voltage function* (or electric potential function) $j_z^y : \Gamma \to \mathbb{R}$ is defined by

$$j_z^y(x) = \text{voltage at } x \text{ when sending one unit of current from } y \text{ to } z,$$

such that $j_z^y(z) = 0$, i.e. the network is "grounded" at z.

Recall that $\Delta : \mathrm{PL}_{\mathbb{R}}(\Gamma) \to \mathrm{Div}_{\mathbb{R}}(\Gamma)$ is defined by $\Delta(f) = \sum_x a_x x$ where a_x is the "sum of outgoing slopes" at x, i.e. $a_x = \sum_{v \in UT_x\Gamma} D_v f(x)$.

Definition 3.2 (Math version; definition–theorem)**.** Given points $y, z \in \Gamma$, the *voltage function* j_z^y is the unique function in $\mathrm{PL}_{\mathbb{R}}(\Gamma)$ satisfying the conditions

$$\Delta(j_z^y) = z - y \quad \text{and} \quad j_z^y(z) = 0.$$

Proof. For the existence and uniqueness of j_z^y, see Theorem 6 and Corollary 3 of Baker–Faber [7]. Note that they use the notation $j_z(y, -)$ for $j_z^y(-)$. □

Note that j_z^y satisfies the following properties:

(V1) for any $x \in \Gamma$, $0 = j_z^y(z) \le j_z^y(x) \le j_z^y(y)$,

(V2) $j_z^y(x)$ is piecewise linear in x,

(V3) $j_z^y(x)$ is continuous in x, y, and z.

Example 3.3 (Voltage function on a graph)**.** Consider the metric graph shown in Figure 3.1, where one unit of current is sent from y (top left) to z (bottom left). The left side of the figure indicates the values of j_z^y at trivalent points of Γ; at all other points, j_z^y linearly interpolates between the values at the endpoints.

The right side of the figure indicates the magnitude of the slope of j_z^y along each edge. Arrows point in the direction of negative slope.

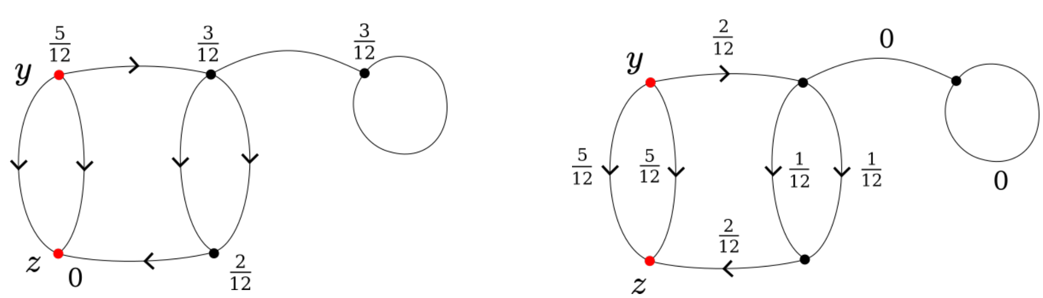

Figure 3.1: Voltage function and currents on a metric graph.

Proposition 3.4. *The voltage function j_z^y obeys the following symmetries.*

(a) For any three points $x, y, z \in \Gamma$,

$$j_z^y(x) = j_z^x(y)$$

(b) For any four points $x, y, z, w \in \Gamma$,

$$j_z^y(x) - j_z^y(w) = j_w^x(y) - j_w^x(z).$$

Proof. See Baker–Faber [7, Theorem 8]; they refer to (b) as the "Magical Identity". Note that (a) follows from (b) by setting $z = w$. $\square$

Remark 3.5. The existence of $j_z^y \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ for any $y, z \in \Gamma$ implies that the principal divisor map $\Delta : \mathrm{PL}_{\mathbb{R}}(\Gamma) \to \mathrm{Div}^0_{\mathbb{R}}(\Gamma)$ is surjective. This verifies the claim made in Remark 2.3 concerning the exactness of the sequence

$$0 \to \mathbb{R} \xrightarrow{\text{const}} \mathrm{PL}_{\mathbb{R}}(\Gamma) \xrightarrow{\Delta} \mathrm{Div}_{\mathbb{R}}(\Gamma) \xrightarrow{\deg} \mathbb{R} \to 0.$$

We many interpret any function $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ as a voltage function on Γ, which results from the externally applied current $\Delta(f) \in \mathrm{Div}_{\mathbb{R}}(\Gamma)$. In other words, the voltage f results from sending current from $\Delta^-(f)$ to $\Delta^+(f)$ in Γ.

Proposition 3.6 (Slope-current principle)**.** *Suppose $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ has zeros $\Delta^+(f)$ and poles $\Delta^-(f)$ of degree $d \in \mathbb{R}$. Then the slope of f is bounded by d, i.e.*

$$|f'(x)| \le d \qquad \textit{for any } x \textit{ where } f \textit{ is linear.}$$

(This bound is sharp; it is attained only on bridge edges, and only when all zeros are on one side of the bridge and all poles are on the other side.)

Proof. Let $\lambda = f(x)$. Then the "tropical preimage"

$$f_\Delta^{-1}(\lambda) := \Delta^-(f) + \Delta(\max\{f, \lambda\})$$

has multiplicity $|f'(x)|$ at x, since the outgoing slopes of $\max\{f, \lambda\}$ at x are $|f'(x)|$ and 0. (Note x cannot be in $\Delta^-(f)$ since f is linear at x.) Since the divisor $f_\Delta^{-1}(\lambda)$ is effective of degree d, this implies $|f'(x)| \le d$ as desired. $\square$

Remark 3.7. The above proposition is obvious from its "physical interpretation": f gives the voltage in the resistor network Γ when subjected to an external current described by $\Delta^-(f)$ units flowing into the network and $\Delta^+(f)$ units flowing out. The slope $|f'(x)|$ is equal to the current flowing through the wire containing x, which must be no more than the total in-flowing (or out-flowing) current.

Next we address how the voltage function $j_z^y \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ may be approximated by a sequence of functions in $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ (up to rescaling), which depend on reduced divisors. We only use property (RD3) of reduced divisors.

Proposition 3.8 (Discrete approximation of voltage function)**.** *Let $\{D_n : n \geq 1\}$ be a sequence of divisors on Γ with $\deg D_n = n$. Fix two points $y, z \in \Gamma$. Let $\mathrm{red}_y[D_n]$ and $\mathrm{red}_z[D_n]$ denote the y– and z–reduced representatives in the divisor class $[D_n]$, and let f_n be the unique function in $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ satisfying*

$$\Delta(f_n) = \mathrm{red}_z[D_n] - \mathrm{red}_y[D_n]$$

and $f_n(z) = 0$. Then the functions $\frac{1}{n}f_n$ converge uniformly to j_z^y as $n \to \infty$.

Proof. If the sequence $\frac{1}{n}h_n$ converges to a limit, then the sequence $\frac{1}{n+c}h_n$ must also converge to the same limit as $n \to \infty$, for any constant c. Thus it suffices to show that the functions $\frac{1}{n}f_{n+g}$ converge uniformly to j_z^y.

Let $\phi_n = \frac{1}{n}f_{(n+g)} - j_z^y$. We claim that the sequence of functions $\{\phi_n \in \mathrm{PL}_{\mathbb{R}}(\Gamma) : n \geq 1\}$ converges uniformly to 0. Note that each ϕ_n is a continuous, piecewise-differentiable function with $\phi_n(z) = 0$, so for an arbitrary $x \in \Gamma$ we may calculate the value of $\phi_n(x)$ by integrating the derivative of ϕ_n along some path in Γ from z to x. The length of such a path is bounded uniformly in x, since Γ is compact, so to show that $\phi_n \to 0$ uniformly it suffices to show that the magnitude of the derivative $|\phi_n'|$ approaches 0 uniformly.

Claim: For any $x \in \Gamma$, $|\phi_n'(x)| \leq \frac{g}{n}$.

This follows from the slope-current principle (Proposition 3.6). By Riemann's inequality, the y-reduced representative in $[D_{(n+g)}]$ may be expressed as

$$\mathrm{red}_y[D_{n+g}] = ny + E_n$$

for some effective divisor E_n of degree g. Similarly, $\mathrm{red}_z[D_{n+g}] = nz + F_n$ for some effective F_n of degree g. Thus the principal divisor associated to $\frac{1}{n}f_{n+g}$ is

$$\Delta(\frac{1}{n}f_{n+g}) = z + \frac{1}{n}F_n - y - \frac{1}{n}E_n.$$

Recall that $\Delta(j_z^y) = z - y$; it follows that the principal $\mathbb{R}$-divisor associated to ϕ_n is

$$\Delta(\phi_n) = \Delta\left(\frac{1}{n}f_{n+g} - j_z^y\right) = \frac{1}{n}F_n - \frac{1}{n}E_n.$$

In particular, $\Delta(\phi_n)$ is a difference of effective $\mathbb{R}$-divisors of degree $\frac{g}{n}$, so the zeros $\Delta^+(\phi_n)$ and poles $\Delta^-(\phi_n)$ each have degree at most $\frac{g}{n}$. By Proposition 3.6, this implies $|\phi_n'(x)| \leq \frac{g}{n}$ as claimed. $\square$

We separate the central claim in the above proof to a named proposition, for future reference.

Proposition 3.9 (Quantitative version of voltage approximation)**.** *Let* Γ *be a metric graph of genus* g*, and let* D_n *be a degree* n *divisor on* Γ*. Fix two points* y *and* z *on* Γ*, and let* f_n *be the unique function in* $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ *satisfying*

$$\Delta(f_n) = \mathrm{red}_z[D_n] - \mathrm{red}_y[D_n]$$

and $f_n(z) = 0$*. Then for* $n > g$ *and any* $x \in \Gamma$*,* $|(\frac{1}{n-g} f_n - j_z^y)'(x)| \leq \frac{g}{n-g}$*.*

Remark 3.10. We can interpret Proposition 3.8 as follows: the existence of the voltage function $j_z^y : \Gamma \to \mathbb{R}$ follows from Riemann's inequality for divisors on Γ.

3.2 Energy and reduced divisors

Here we give a definition of q-reduced divisors on a metric graph. We will only need to use q-reduced divisors for effective divisor classes, so we restrict our discussion here to the effective case.

Definition 3.11. Given a basepoint q on Γ, we define the *q-energy* $\mathcal{E}_q : \Gamma \to \mathbb{R}$ by

$$\mathcal{E}_q(y) = j_q^y(y) = r(y, q).$$

Given an effective divisor $D = \sum_i y_i$, we define the *q-energy* $\mathcal{E}_q(D)$ by

$$\mathcal{E}_q(D) = \sum_i \sum_j j_q^{y_i}(y_j).$$

Note that

- $\mathcal{E}_q(D) \geq 0$,
- $\mathcal{E}_q(D)$ is strictly positive if D has support outside of q,
- $\mathcal{E}_q(D) \geq \sum_i \mathcal{E}_q(y_i)$, and in general this inequality is strict.

Theorem 3.12 (Baker–Shokrieh [11, Theorem A.7])**.** *Fix a basepoint* $q \in \Gamma$*, and let* D *be an effective divisor on* Γ*. There is a unique divisor* $D_0 \in |D|$ *which minimizes the* q*-energy, i.e. such that*

$$\mathcal{E}_q(D_0) < \mathcal{E}_q(E) \qquad \text{for all} \qquad E \in |D|,\ E \neq D_0.$$

Using this result, we define the *q-reduced divisor* $\mathrm{red}_q[D]$ as the unique divisor in $|D|$ which minimizes the q-energy $\mathcal{E}_q$.

Note that this definition is non-standard; the standard definition for reduced divisor is a combinatorial condition which can be phrased in the language of chip-firing, see [2, p. 4854], [4, Definition 2.3].

Example 3.13. In Figure 3.2 we show a degree 4 divisor, on the left, and its reduced representative with respect to basepoint q, on the right.

Figure 3.2: A divisor and its reduced divisor representative.

3.3 Resistance function

In this section we recall the definition of the (Arakelov–Zhang–Baker–Faber) canonical measure μ on a metric graph.

Definition 3.14. Let $r : \Gamma \times \Gamma \to \mathbb{R}$ denote the *effective resistance* function on the metric graph Γ. Namely, viewing Γ as a resistor network

$$\begin{aligned} r(x, y) &= \text{effective resistance between } x \text{ and } y \\ &= \text{total voltage drop when sending 1 unit of current from } x \text{ to } y \end{aligned}$$

If we wish to emphasize the underlying graph, we write $r(x, y; \Gamma)$. In terms of the voltage function from Section 3.1, $r(x, y) = j^x_y(x)$.

It is straightforward to verify that the resistance function satisfies the following properties:

1. $r(x, x) = 0$,
2. $r(x, y) > 0$ if $x \neq y$,
3. $r(x, y)$ is continuous with respect to x and y
4. $r(x, y) = r(y, x)$

In contrast with the voltage function j^y_z, the function $x \mapsto r(x, y)$ is not piecewise linear; we will see that it is instead piecewise quadratic.

There is a special case of effective resistance which will be particularly useful in the following sections.

Definition 3.15. Given a segment e in a metric graph Γ, the *deleted effective resistance* $\ell_{\text{eff}}(\Gamma \backslash e)$ is the effective resistance between endpoints of e in the e-deleted subgraph; that is, if s, t are the endpoints of e

$$\ell_{\text{eff}}(\Gamma \backslash e) = r(s, t; \Gamma \backslash e).$$

Note that $\ell_{\text{eff}}(\Gamma \backslash e) = 0$ when e is a loop, and $\ell_{\text{eff}}(\Gamma \backslash e) = +\infty$ when e is a bridge. The rule for combining resistances in parallel implies that for a segment e with endpoints s and t,

$$r(s, t; \Gamma) = \left(\frac{1}{\ell(e)} + \frac{1}{\ell_{\text{eff}}(\Gamma \backslash e)} \right)^{-1} = \frac{\ell(e)\ell_{\text{eff}}(\Gamma \backslash e)}{\ell(e) + \ell_{\text{eff}}(\Gamma \backslash e)}.$$

Example 3.16. Let Γ be a circle of circumference L. By choosing a basepoint which we denote as 0, we may parametrize Γ with the interval $[0, L]$. Identifying points in this way, we have

$$\begin{aligned} r(x, 0) &= \text{parallel combination of resistances } x \text{ and } L - x \\ &= \frac{x(L-x)}{x + (L-x)} = x - \frac{1}{L}x^2. \end{aligned}$$

The effective resistance is maximized when $x = \frac{1}{2}L$, with maximum value $\frac{1}{4}L$. The effective resistance is minimized when $x = 0$ or $x = L$, with effective resistance 0.

3.4 Canonical measure

Definition 3.17. The *canonical measure* $\mu = \mu_\Gamma$ on a metric graph Γ is the continuous measure defined by

$$\mu = \mu(dx) = -\frac{1}{2}\frac{d^2}{dx^2} r(x, y_0)\, dx,$$

where x is a length-preserving parameter on Γ, dx is the Lebesgue measure, and y_0 is a fixed point in Γ. This defines μ on the open dense subset of Γ where the second derivative exists; at the finite set of points where $r(-, y_0)$ is not differentiable, or where the valence of x differs from 2, we let $\mu_\Gamma = 0$.

Remark 3.18. The first derivative of a smooth function on Γ is only well-defined up to a choice of sign, since there are two directions in which we could parametrize any segment. The second derivative, however, is well-defined on each segment (without choosing an orientation) because $(\pm 1)^2 = 1$ so either choice of direction yields the same second derivative.

Remark 3.19. The definition of canonical measure is independent of the choice of basepoint y_0 because of the "Magical Identity" in Proposition 3.4 (b). Namely, for two basepoints y_0, z_0 we have $j^x_{y_0}(x) - j^x_{y_0}(z_0) = j^x_{z_0}(x) - j^x_{z_0}(y_0)$ which implies

$$\begin{aligned} r(x, y_0) - r(x, z_0) &= j^x_{y_0}(x) - j^x_{z_0}(x) \\ &= j^x_{y_0}(z_0) - j^x_{z_0}(y_0) = j^{z_0}_{y_0}(x) - j^{y_0}_{z_0}(x). \end{aligned}$$

Since the voltage functions $j^{z_0}_{y_0}, j^{y_0}_{z_0}$ are piecewise linear, we have

$$\frac{d^2}{dx^2}(r(x, y_0) - r(x, z_0)) = \frac{d^2}{dx^2}(j^{z_0}_{y_0}(x) - j^{y_0}_{z_0}(x)) = 0.$$

Remark 3.20. The definition of canonical measure given here differs from that used by Baker–Faber [7], in that our μ does not have a discrete part supported at the points of Γ with valence different from 2.

Remark 3.21. The definition of canonical measure given here is equal to Zhang's canonical measure [42, Section 3, Theorem 3.2 c.f. Lemma 3.7] associated to the canonical divisor $D = K$, up to a multiplicative factor. Our canonical measure is normalized to satisfy $\mu(\Gamma) = g$ rather than $\mu(\Gamma) = 1$.

The canonical measure of Baker–Faber is equal to Zhang's canonical measure associated to $D = 0$.

Example 3.22 (Canonical measure in genus one and two)**.**

(a) If Γ is a circle of circumference L, by Example 3.16 we have $r(x, 0) = x - \frac{1}{L}x^2$ so the canonical measure is $\mu = \frac{1}{L}dx$. The total measure on the metric graph is $\mu(\Gamma) = 1$.

(b) Consider the metric graph Γ of genus 2 shown in Figure 3.3, with edge lengths a, b, c.

Figure 3.3: Genus 2 metric graph with edge lengths a, b, c.

On the edge of length a, we have $\ell(e) = a$ and $\ell_{\text{eff}}(\Gamma\backslash e) = \frac{bc}{b+c}$. When measuring effective resistance between points in the interior of e, we can think of Γ as a circle of total length $\ell(e) + \ell_{\text{eff}}(\Gamma\backslash e) = \frac{ab+ac+bc}{b+c}$. Thus the canonical measure on this edge is $\mu = \frac{b+c}{ab+ac+bc}dx$, by the computation for a circle in Example 3.16. The total measure on this edge is $\mu(e) = \frac{ab+ac}{ab+ac+bc}$. By symmetry, the total measure on the metric graph is $\mu(\Gamma) = 2$.

Proposition 3.23. *The canonical measure μ on a metric graph Γ is a piecewise-constant multiple of the Lebesgue measure which vanishes on all bridge segments.*

On a non-bridge segment e in Γ,

$$\mu|_e = \frac{1}{\ell(e) + \ell_{\mathit{eff}}(\Gamma\backslash e)}dx \tag{3.1}$$

where $\ell(e)$ denotes the length of e and $\ell_{\mathit{eff}}(\Gamma\backslash e)$ denotes the effective resistance between the endpoints of e on the graph after removing the interior of e.

For a bridge segment, $\mu|_e = 0$.

Proof. See Baker–Faber [7, Theorem 12]; note that our μ is defined to be the continuous part of Baker–Faber's μ_{can}.

The proof idea is that when x, y lie on the segment e, the resistance function $r(x, y)$ behaves as if Γ were a circle of length $\ell(e) + \ell_{\text{eff}}(\Gamma\backslash e)$; (see Example 3.22(b).) $\square$

Note that the local expression (3.1) for the canonical measure is preserved under subdividing an edge. If a segment e is subdivided into $e_1 \sqcup e_2$, the formula for $\mu|_e$ agrees with $\mu|_{e_1}$ because

$$\ell(e_1) = \ell(e) - \ell(e_2) \qquad \text{and} \qquad \ell_{\text{eff}}(\Gamma \backslash e_1) = \ell_{\text{eff}}(\Gamma \backslash e) + \ell(e_2).$$

Corollary 3.24. *Let Γ be a metric graph with canonical measure μ, and let e be a segment in Γ (i.e. e is subspace isometric to a closed interval, whose interior points all have valence 2 in Γ). Then*

(a) $0 \le \mu(e) \le 1$*;*

(b) $\mu(e) = 0 \quad \Leftrightarrow \quad$ *e is a bridge edge;*

(c) $\mu(e) = 1 \quad \Leftrightarrow \quad$ *e is a loop edge.*

Proof. By Proposition 3.23, we have $\mu(e) = 0$ for bridges and $\mu(e) = \frac{\ell(e)}{\ell(e)+\ell_{\text{eff}}(\Gamma \backslash e)}$ otherwise. $\square$

Proposition 3.25 (Foster's theorem)**.** *Let Γ be a metric graph of genus g, and let μ be the canonical measure on Γ. Then the total measure on Γ is*

$$\mu(\Gamma) = g.$$

Proof. See Baker–Faber [7, Corollary 5 and Corollary 6]. An equivalent statement, using different terminology, appeared in Foster [19]. $\square$

3.5 Kirchhoff formulas

In this section we review Kirchhoff's formulas for the currents and voltage drops in a resistor network. These formulas were published (in some equivalent form) by Kirchhoff in [30]. The argument is combinatorial, but can be expressed as linear algebra and is essentially equivalent to what is known as the *matrix-tree theorem.* Expositions of this material are found in Bollobás [13, §II.1] and Grimmet [21, §1.2].

The material in this section will be used in Chapter 5.

Theorem 3.26 (Kirchhoff)**.** *Suppose $\Gamma = (G, \ell)$ is a resistor network (metric graph) with resistance function $\ell : E(G) \to \mathbb{R}_{>0}$. For vertices $y, z \in V(G)$, let $j_z^y : \Gamma \to \mathbb{R}$ denote the voltage function which sends one unit of current from y to z.*

(a) The current across a directed edge $\vec{e} = (e_+, e_-)$ is

$$\frac{j_z^y(e_+) - j_z^y(e_-)}{\ell(e)} = \frac{\sum_{T \in \mathcal{T}(G)} sgn(T, y, z, \vec{e}) w(T)}{\sum_{T \in \mathcal{T}(G)} w(T)} \tag{3.2}$$

where $\mathcal{T}(G)$ denotes the spanning trees of G, the weight $w(T)$ *of a spanning tree is defined as*

$$w(T) = \prod_{e_i \notin E(T)} \ell(e_i),$$

and

$$sgn(T, y, z, \vec{e}) = \begin{cases} +1 & \textit{if the path in } T \textit{ from } y \textit{ to } z \textit{ passes through } \vec{e} \\ -1 & \textit{if the path in } T \textit{ from } y \textit{ to } z \textit{ passes through } -\vec{e} \\ 0 & \textit{otherwise.} \end{cases}$$

(b) The total voltage drop between y and z is

$$j_z^y(y) - j_z^y(z) = \frac{\sum_{T \in \mathcal{T}(G_0)} w(T)}{\sum_{T \in \mathcal{T}(G)} w(T)} \tag{3.3}$$

in the same notation as above, and where the graph G_0 (in the numerator) is the graph obtained from G by identifying vertices y and z.

Proof. For part (a), see Bollobás [13, Theorem 2, §II.1]. Part (b) follows from consideration of the graph G_+ obtained by adding an auxiliary edge to G between y and z, and then applying part (a) to G_+ with respect to the auxiliary edge. □

The expressions (3.2), (3.3) for the current, resp. voltage drop, are both a ratio of homogeneous polynomials[1] in the variables $\{\ell(e_i) : e_i \in E(G)\}$. In (3.2), the numerator and denominator are homogeneous of degree g; in (3.3), the denominator has degree g while the numerator has degree $g + 1$. As a result, the current (3.2) is invariant under simultaneous rescaling of edge lengths, while the voltage drop (3.3) scales linearly with respect to simultaneously rescaling all edge lengths. This should agree with physical intuition.

Example 3.27. Consider the theta graph shown in Figure 3.4, where $a = \ell(e_1)$, $b = \ell(e_2)$, $c = \ell(e_3)$ are edge lengths (resistances). The spanning trees are $\{e_3, e_2, e_1\}$ which have respective weights $\{ab, ac, bc\}$. The current along edge e_1 is

$$\frac{j_z^y(y) - j_z^y(z)}{a} = \frac{bc}{ab + ac + bc},$$

according to (3.2). We have

$$j_y^x(x) - j_y^x(y) = a \left(\frac{bc}{ab + ac + bc} \right) = \frac{abc}{ab + ac + bc}$$

in agreement with (3.3); G_0 consists of three loop edges. Note the symmetry in a, b, c.

[1] moreover, polynomials whose nonzero coefficients are all ± 1

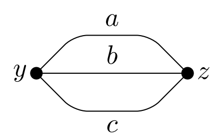

Figure 3.4: Theta graph with variable edge lengths.

Example 3.28. Let G be the Wheatstone graph in Figure 3.5 (left), with edge lengths $a = \ell(e_1), \ldots, f = \ell(e_5)$. The spanning trees are

$$\mathcal{T} = \{345, 245, 234, 145, 135, 125, 124, 123\},$$

where 123 shorthand for spanning tree $\{e_1, e_2, e_3\}$, and the corresponding weights are $\{ab, ac, af, bc, bd, cd, cf, df\}$. The current along edge e_3 is

$$\frac{j_z^y(e_{3,+}) - j_z^y(e_{3,-})}{c} = \frac{ab + af}{ab + ac + af + bc + bd + cd + cf + df},$$

while the current along e_1 is

$$\frac{j_z^y(y) - j_z^y(z)}{a} = \frac{bc + bd + cd + cf + df}{ab + ac + af + bc + bd + cd + cf + df}.$$

The total voltage drop from y to z is

$$j_z^y(y) - j_z^y(z) = \frac{abc + abd + acd + acf + adf}{ab + ac + af + bc + bd + cd + cf + df},$$

in agreement with (3.3); the quotient graph G_0 is shown to the right in Figure 3.5.

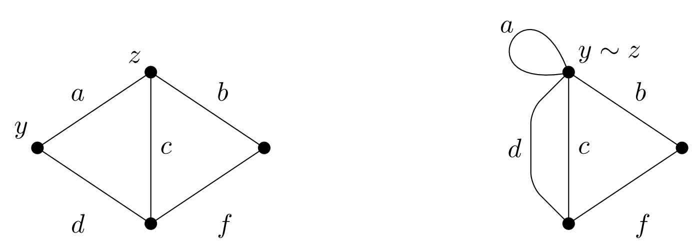

Figure 3.5: Wheatstone graph with variable edge lengths, and a quotient graph.

If we let $d = f = 0$, then we recover the formulas of Example 3.27.

CHAPTER 4

Weierstrass Points

In this chapter we define the Weierstrass locus and the stable Weierstrass locus of an arbitrary divisor D on a metric graph Γ. We first review the notion of Weierstrass point on an algebraic curve. We then prove theorems regarding the distribution of Weierstrass points.

The results in this chapter first appeared in the preprint [40].

4.1 Classical Weierstrass points

Recall that for an algebraic curve X of genus g, the *ordinary Weierstrass points* are defined as follows. The canonical divisor K on X determines a canonical map to projective space $\varphi_K : X \to \mathbb{P}^{g-1}$. Generically, a point on $\varphi_K(X)$ will have an osculating hyperplane in $\mathbb{P}^{g-1}$ which intersects $\varphi_K(X)$ with multiplicity $g-1$. For finitely many "exceptional" points on $\varphi_K(X)$, the osculating hyperplane will intersect the curve with higher multiplicity; the preimages of these exceptional points are the ordinary Weierstrass points of X. These are also known as the *flex points* of the embedded curve $\varphi_K(X) \subset \mathbb{P}^{g-1}$.

This notion may be generalized by replacing K with an arbitrary (basepoint-free) divisor. Given a divisor D on X, there is an associated map to projective space $\varphi_D : X \to \mathbb{P}^r$, known as the complete linear embedding defined by D. The set of flex points of the embedded curve $\varphi_D(X)$, where the osculating hyperplane intersects the curve with multiplicity greater than r, are the *(generalized) Weierstrass points* associated to the divisor D. If D has degree $n \geq 2g-1$, the number of Weierstrass points of D counted with multiplicity is $g(n-g+1)^2$.

The existence of an osculating hyperplane of multiplicity greater than r, at the point $\varphi_D(x) \in \varphi_D(X)$, is equivalent to the existence of a non-zero global section of the line bundle $\mathcal{L}(X, D-(r+1)x)$, i.e. to having $h^0(X, D-(r+1)x) \geq 1$.

4.2 Tropical Weierstrass points

Given a divisor D on a metric graph, we define the set of Weierstrass points of D using the Baker–Norine rank function $r(D)$, which is the analogue of $h^0(D) - 1$.

Definition 4.1. Let D be a divisor on a metric graph Γ, with rank $r = r(D)$. A point $x \in \Gamma$ is a *Weierstrass point for D* if

$$[D - (r+1)x] \geq 0.$$

The *Weierstrass locus* $W(D) \subset \Gamma$ of D is the set of its Weierstrass points. An *ordinary Weierstrass point* is a Weierstrass point for the canonical divisor K.

Note that the Weierstrass locus of D depends only on the divisor class $[D]$.

Remark 4.2. If the divisor class $[D]$ is not effective, i.e. $r(D) = -1$, then the set of Weierstrass points of D is empty. Thus we may restrict our attention to studying Weierstrass points for effective divisor classes.

Example 4.3. Suppose Γ is a genus 1 graph and D is a divisor of degree 6, indicated by the black dots in Figure 4.1 with multiplicities. This divisor has rank $r = 5$ since it is in the nonspecial range of Riemann–Roch. The Weierstrass locus of D consists of 6 points evenly spaced around Γ, indicated in red.

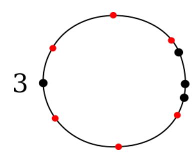

Figure 4.1: Weierstrass points, in red, on a genus 1 metric graph.

Example 4.4. Suppose Γ is a complete graph on 4 vertices, with distinct edge lengths. This graph has genus 3. Consider the canonical divisor K on Γ, which is supported on the four trivalent vertices. The Weierstrass locus of K consists of 8 distinct points on Γ, shown in red in Figure 4.2.

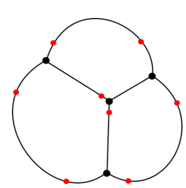

Figure 4.2: Weierstrass locus on a genus 3 metric graph.

In the following examples, we use "chip firing" language to describe linear equivalence of divisors; see Remark 2.4.

Example 4.5 (Wedge of circles)**.** Suppose Γ is a wedge of g circles, and let x_0 denote the point of Γ lying on all g circles. For a generic divisor class $[D_n]$ of degree n (meaning generic inside of $\mathrm{Pic}^n(\Gamma)$), the x_0-reduced representative of $[D_n]$ consists of $n-g$ chips at x_0 and one chip in the interior of each circle. The Weierstrass locus $W(D_n)$ contains $n-g+1$ evenly-spaced points on each circle of Γ, for a total of $g(n-g+1)$ points.

Example 4.6 (Failure of $W(D)$ to be finite)**.** Consider the genus 3 graph shown in Figure 4.3. Suppose $D = K$ is the canonical divisor. By Riemann–Roch, K has rank $r = 2$. It is possible to move all 4 chips to lie on the middle loop, so any point in the middle loop has $\mathrm{red}_x[D] \geq 3x$. The Weierstrass locus $W(K)$ contains the middle loop, but not the two outer loops.

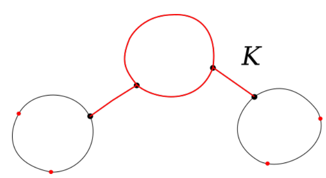

Figure 4.3: Weierstrass locus, in red, which is not finite.

Example 4.7 (Failure of $W(D)$ to be finite)**.** Consider the genus 3 graph shown in Figure 4.4. Suppose D is a degree 4 divisor supported on one of the bridge edges as shown. (Note that $D \sim K$.) This divisor has rank $r \leq 2$, since we cannot move the chips in D to lie on three distinct loops freely. However, for any point x, the reduced divisor $\mathrm{red}_x[D]$ has at least 3 chips at x.

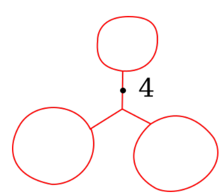

Figure 4.4: Weierstrass locus which contains Γ.

Remark 4.8. For any metric graph with a bridge edge, it can be shown that the entire bridge edge is contained in the Weierstrass locus of the canonical divisor so in particular $W(K)$ is not finite. We omit the details.

4.2.1 Stable tropical Weierstrass points

In this section we define the stable Weierstrass locus $W^{\mathrm{st}}(D)$ of a divisor D on a metric graph. This definition is meant to fix undesirable behavior of the naive Weierstrass locus $W(D)$. In particular, $W^{\mathrm{st}}(D)$ is always a finite set.

For the definition of break divisor, see Section 2.5.

Definition 4.9. Let D be a divisor of degree n on a metric graph Γ. If $n \geq g$, the *stable Weierstrass locus* $W^{\mathrm{st}}(D) \subset \Gamma$ is the set of all points $x \in \Gamma$ such that

$$\mathrm{br}[D - (n-g)x] \geq x$$

where $\mathrm{br}[E]$ is the break divisor representative of the divisor class $[E]$. In other words, x is a stable Weierstrass point of D if

$$\text{there exists a break divisor } E \geq x \text{ such that} \qquad E + (n-g)x \in [D].$$

Note that if D has degree $n = g$, then $W^{\mathrm{st}}(D)$ is exactly the support of $\mathrm{br}[D]$.

If D has degree $n < g$, we define $W^{\mathrm{st}}(D)$ to be empty.

In the above definition, if $n \geq g$ then $n - g$ is the rank of a generic divisor class in $\mathrm{Pic}^n(\Gamma)$. If a divisor class $[D]$ in $\mathrm{Pic}^n(\Gamma)$ has rank $r(D) = n - g$, then $W^{\mathrm{st}}(D) \subset W(D)$; otherwise, this containment may fail to hold. In particular, we have $W^{\mathrm{st}}(D) \subset W(D)$ for all divisors of degree $n \geq 2g - 1$.

Example 4.10 (Divisor with $W^{\mathrm{st}}(D) \not\subset W(D)$)**.** Consider the genus 3 metric graph shown in Figure 4.5. The canonical divisor K is indicated in black. This divisor has degree $n = 4$ and rank $r(K) = 2$. The divisor is special, because $r(K) > n - g = 1$. On the left side, the Weierstrass locus is shown in red; the right side shows the stable Weierstrass locus. The stable Weierstrass locus consists of the midpoint of each edge. The sets $W(K)$ and $W^{\mathrm{st}}(K)$ are disjoint.

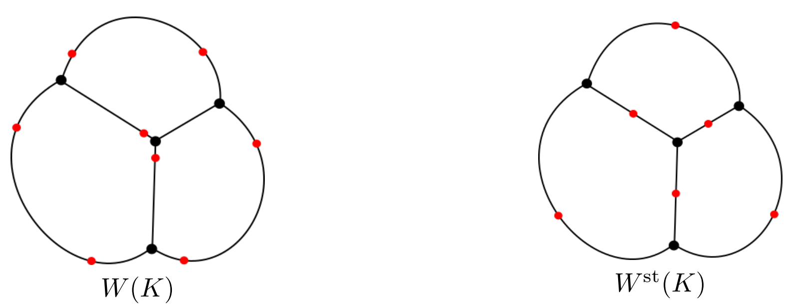

Figure 4.5: Divisor with Weierstrass locus and stable Weierstrass locus.

4.3 Finiteness of Weierstrass points

In this section we show that the Weierstrass locus of a generic divisor class $[D]$ on a metric graph is a finite set whose cardinality is $\#W(D) = g(n - g + 1)$. We do so by studying the stable Weierstrass locus $W^{\mathrm{st}}(D)$, defined in Section 4.2.1.

4.3.1 Setup

Our main technical tool is to consider the ABKS decomposition of $\mathrm{Pic}^g(\Gamma)$ (see Section 2.5) and the topology of certain branched covering spaces.

As the divisor class $[D]$ varies over $\mathrm{Pic}^n(\Gamma)$, we realize the stable Weierstrass loci $W^{\mathrm{st}}(D)$ as the fibers of a surjective map $X \to \mathrm{Pic}^n(\Gamma)$. We are able to study the cardinality of $W^{\mathrm{st}}(D)$ by imposing a nice topology on X and analyzing topological properties of the map $X \to \mathrm{Pic}^n(\Gamma)$.

Recall that $\mathrm{Br}^g(\Gamma)$ denotes the space of break divisors on Γ, viewed as a subspace of $\mathrm{Sym}^g(\Gamma)$.

Definition 4.11. Let $\widetilde{\mathrm{Br}}^g(\Gamma)$ denote the space

$$\widetilde{\mathrm{Br}}^g(\Gamma) = \{(x,E) \in \Gamma \times \mathrm{Sym}^{g-1}(\Gamma) : x + E \text{ is a break divisor}\}.$$

This defines a closed subset of the compact Hausdorff space $\Gamma \times \mathrm{Sym}^{g-1}(\Gamma)$, so $\widetilde{\mathrm{Br}}^g(\Gamma)$ is compact and Hausdorff.

Remark 4.12. We may think of $\widetilde{\mathrm{Br}}^g(\Gamma)$ as the space of "pointed break divisors" on Γ, i.e. $\widetilde{\mathrm{Br}}^g(\Gamma)$ is homeomorphic to $\{(x,D) \in \Gamma \times \mathrm{Br}^g(\Gamma)$ such that $x \leq D\}$.

Let $\sigma : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$ denote the "summation" map $(x,E) \mapsto x + E$, and let $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ denote the "summation with multiplicity" map defined by

$$\sigma_m : (x,E) \mapsto [mx + E].$$

Let $\pi_1 : \widetilde{\mathrm{Br}}^g(\Gamma) \to \Gamma$ denote projection to the first factor, i.e. $\pi_1(x,E) = x$.

Lemma 4.13. *Suppose $[D] \in \mathrm{Pic}^{m+g-1}(\Gamma)$, and let σ_m and π_1 be defined as above.*

(a) The stable Weierstrass locus $W^{\mathrm{st}}(D)$ is equal to $\pi_1(\sigma_m^{-1}[D])$.

(b) We have $\#W^{\mathrm{st}}(D) = \#\sigma_m^{-1}[D]$.

Proof. (a) This follows from the definition of the stable Weierstrass locus.

(b) The claim is that π_1 is injective on the preimage $\sigma_m^{-1}[D]$. To see this, consider two points (x,E) and $(x',E') \in \widetilde{\mathrm{Br}}^g(\Gamma)$ in the same fiber $\sigma_m^{-1}[D]$. This means that $[mx + E] = [mx' + E'] = [D]$. Suppose $\pi_1(x,E) = \pi_1(x',E')$, i.e. that $x = x'$. Then

$$[D - (m-1)x] = [x + E] = [x + E'] \in \mathrm{Pic}^g(\Gamma).$$

Since both $(x + E)$ and $(x + E')$ are break divisors, the uniqueness of break divisor representatives (Theorem 2.8) implies that $E = E'$. This shows that the restriction of π_1 to $\sigma_m^{-1}[D]$ is injective, as desired. □

Let (G, ℓ) be a combinatorial model for Γ, which induces a decomposition of break divisors $\mathrm{Br}^g(\Gamma)$ into a union of cells

$$\mathrm{Br}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} C_T \tag{4.1}$$

indexed by spanning trees of G, where the interior of each cell C_T is homeomorphic to an open hypercube. (See Section 2.5 or [4].) Note that $\mathrm{Br}^g(\Gamma)$ is homeomorphic to $\mathrm{Pic}^g(\Gamma)$. The ABKS decomposition (4.1) of $\mathrm{Br}^g(\Gamma)$ induces a decomposition

$$\widetilde{\mathrm{Br}}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} \left(\bigcup_{e \notin E(T)} \widetilde{C}_{T,e} \right) \tag{4.2}$$

where the second union is over edges e of G not contained in the spanning tree T. There are g such edges for any T. Namely,

$$\widetilde{C}_{T,e} = \{(x, E) \in \widetilde{\mathrm{Br}}^g(\Gamma) : x + E \in C_T,\ x \in e\}$$

The map $\widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$ sends the cell $\widetilde{C}_{T,e}$ surjectively to C_T. On the interior C_T° of each cell, each fiber of $\widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$ contains exactly g points.

If $\kappa(G) = \#\mathcal{T}(G)$ denotes the number of spanning trees of G, the ABKS decomposition (4.2) decomposes $\widetilde{\mathrm{Br}}^g(\Gamma)$ into a union of $g \cdot \kappa(G)$ cells.

Example 4.14. In Figure 4.6, we show the decomposition of $\widetilde{\mathrm{Br}}^2(\Gamma)$ into six cells $\widetilde{C}_{T,e}$, where Γ is a theta graph. This graph has genus $g = 2$ and $\kappa(G) = 3$ spanning trees. In this case $\mathrm{Br}^2(\Gamma) \cong \mathrm{Pic}^2(\Gamma) \cong \mathbb{R}^2/\mathbb{Z}^2$ is a genus 1 surface (cf. Example 2.11, Theorem 2.6), and $\widetilde{\mathrm{Br}}^2(\Gamma)$ is a surface of genus 2. The map $\widetilde{\mathrm{Br}}^2(\Gamma) \to \mathrm{Br}^2(\Gamma)$ is a branched double cover ramified at two points, corresponding to the two break divisors which consist of two chips at a trivalent vertex of Γ.

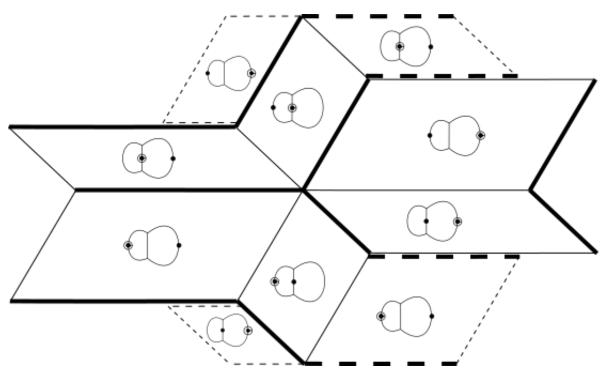

Figure 4.6: ABKS decomposition of $\widetilde{\mathrm{Br}}^2(\Gamma)$.

In Figure 4.6, each cell $\widetilde{C}_{T,e}$ shows a representative break divisor $x + E$ where the point $x \in e$ is marked with an extra outline. Edges of $\widetilde{C}_{T,e}$ which have x on an endpoint of e are marked in bold. Edges on the boundary are glued to the parallel boundary edge which has the same weighting (bold or unbold).

4.3.2 Point-set topology

Definition 4.15. Let M and N be compact Hausdorff spaces, and let N be path-connected. We say $p : M \to N$ is a *branched covering map* if

(i) p is continuous and surjective

(ii) p is an open map (the image of an open set is open)

(iii) $p^{-1}(y)$ is finite for each $y \in N$

and there exists a closed subset $R \subset N$ such that

(iv) $N\backslash R$ is path-connected

(v) R has empty interior in N

(vi) the restriction of p to $M\backslash p^{-1}(R) \to N\backslash R$ is a topological covering map.

The subspace R is a *ramification locus* of p, and the preimage $p^{-1}(R)$ is a *branch locus*. (Note that properties (ii) and (v) imply $p^{-1}(R)$ has empty interior in M.)

It is straightforward to verify that the map $\widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$ from Section 4.3.1 is a branched covering. We show below, in Proposition 4.19, that in fact each $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$, for $m \geq 1$, is a branched covering.

Recall that a map is *proper* if the preimage of a compact set is compact. Recall that a map $f : X \to Y$ is a *local homeomorphism* if, for any $x \in X$ there is an open neighborhood U containing x such that $f(U)$ is open in Y and the restriction $U \to f(U)$ is a homeomorphism. A covering map is always a local homemorphism, but the converse is not true.

The following lemma will be used to check the last condition (vi) in Definition 4.15, that the restriction $M\backslash p^{-1}(R) \to N\backslash R$ is a covering map.

Lemma 4.16. *Suppose $p : X \to Y$ is a local homeomorphism between locally compact, Hausdorff spaces. If p is proper and surjective, then p is a covering map.*

This is a standard exercise in point-set topology; see e.g. Ho [24, Lemma 2].

Lemma 4.17. *Suppose $p : M \to N$ is a branched covering with ramification locus $R \subset N$ such that the restriction $p : M\backslash p^{-1}(R) \to N\backslash R$ is a covering map of degree d. Then for any $y \in N$, the preimage $p^{-1}(y)$ has cardinality at most d.*

Note: the restriction of p to $M\backslash p^{-1}(R) \to N\backslash R$ has constant degree d because in the definition of branched cover, $N\backslash R$ is assumed to be path connected.

Proof of Lemma 4.17. Let $y \in R$ be a point in the ramification locus, and let $x_1, \dots, x_k$ be the points in the preimage $p^{-1}(y)$. Since M is Hausdorff, we may choose open neighborhoods $U_1, \dots, U_k$ with $x_i \in U_i$ which are disjoint, $U_i \cap U_j = \varnothing$.

Let $C = M \backslash (U_1 \cup \cdots \cup U_k)$ be the complement of these neighborhoods, which is closed in M. Since M is compact and N is Hausdorff, the image $p(C)$ is closed in N. Thus $V = N \backslash p(C)$ is open and nonempty since $y \in V$. Note that by construction $p^{-1}(V) = M \backslash p^{-1}(p(C)) \subset M \backslash C = U_1 \cup \cdots \cup U_k$.

Let U_i' be the intersection of $p^{-1}(V)$ with U_i, which is open and nonempty because $x_i \in U_i'$. Since the U_i were chosen to be disjoint, $p^{-1}(V) = U_1' \sqcup \cdots \sqcup U_k'$.

Note that p is an open map (by definition of branched cover), so the intersection $p(U_1') \cap \cdots \cap p(U_k')$ is an open neighborhood of y in N. Since R has empty interior in N, we can choose some point

$$z \in (p(U_1') \cap \cdots \cap p(U_k')) \backslash R \subset V \backslash R.$$

By the assumption that $M \backslash p^{-1}(R) \to N \backslash R$ is a degree d covering map, the preimage $p^{-1}(z)$ contains d points $w_1, \ldots, w_d$. Since $z \in V$ by construction, each $w_i \in p^{-1}(V) = U_1' \sqcup \cdots \sqcup U_k'$ so w_i lies within U_j' for some unique $j \in \{1, \ldots, k\}$. This relation defines a map $\pi : \{1, \ldots d\} \to \{1, \ldots, k\}$. Moreover, the map π is surjective because $z \in p(U_j')$ for each $j \in \{1, \ldots, k\}$. This proves that $k \leq d$, so the preimage $p^{-1}(y)$ has cardinality at most d as desired. □

4.3.3 Proofs

Proposition 4.18. *For any divisor D, the stable Weierstrass locus $W^{\mathrm{st}}(D)$ is a finite subset of Γ.*

Proof. If D has degree $n < g$, the stable Weierstrass locus is defined to be empty. Thus we assume below that D has degree $n \geq g$.

Recall that $\widetilde{\mathrm{Br}}^g(\Gamma) = \{(x, E) \in \Gamma \times \mathrm{Sym}^{g-1}(\Gamma) : \ x + E \text{ is a break divisor}\}$ and that $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ is defined by

$$\sigma_m : (x, E) \mapsto [mx + E].$$

Recall that π_1 denotes the projection $\pi_1(x, E) = x$. (See Section 4.3.1.) By Lemma 4.13, for a divisor D of degree $m + g - 1$ we have $W^{\mathrm{st}}(D) = \pi_1(\sigma_m^{-1}[D])$. Hence it suffices to show that the preimage $\sigma_m^{-1}[D]$ is a finite set.

Let (G, ℓ) be a combinatorial model for Γ, which induces the ABKS decomposition $\mathrm{Br}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} C_T$, where the cells C_T are indexed by spanning trees of G. The ABKS decomposition of $\mathrm{Br}^g(\Gamma)$ induces a decomposition

$$\widetilde{\mathrm{Br}}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} \left(\bigcup_{e \notin E(T)} \widetilde{C}_{T,e} \right).$$

Let $\sigma_m^{(T,e)} : \widetilde{C}_{T,e} \to \mathrm{Pic}^{m+g-1}(\Gamma)$ denote the restriction of σ_m to $\widetilde{C}_{T,e}$.

Claim: The preimage of $[D]$ under $\sigma_m^{(T,e)} : \widetilde{C}_{T,e} \to \operatorname{Pic}^{m+g-1}(\Gamma)$ is finite.
This Claim implies that the preimage $\sigma_m^{-1}[D]$ is a finite set, since $\widetilde{\operatorname{Br}}^g(\Gamma)$ is covered by finitely many $\widetilde{C}_{T,e}$.

Proof of Claim: The map $\sigma_m^{T,e} : \widetilde{C}_{T,e} \to \operatorname{Pic}^{m+g-1}(\Gamma)$ is locally defined by a linear map, which we show is full rank. For a spanning tree $T = G \backslash \{e, e_2, \dots, e_g\}$, there is a natural surjective parametrization $\prod_{i=1}^g [0, \ell(e_i)] \to \widetilde{C}_{T,e}$.

Let $f_m^{T,e}$ denote the lift of $\prod_{i=1}^g [0, \ell(e_i)] \to \widetilde{C}_{T,e} \to \operatorname{Pic}^{m+g-1}(\Gamma)$ to the universal cover $\mathbb{R}^g \to \operatorname{Pic}^{m+g-1}(\Gamma)$.

$$\begin{tikzcd}
\prod_{i=1}^g [0, \ell(e_i)] \arrow[r, dashed, "f_m^{T,e}"] \arrow[d] & \mathbb{R}^g \arrow[d, "\pi"] \\
\widetilde{C}_{T,e} \arrow[r, "\sigma_m^{T,e}"] & \operatorname{Pic}^{m+g-1}(\Gamma)
\end{tikzcd}$$

When $m = 1$, coordinates may be chosen on $\mathbb{R}^g$ such that $f_1^{T,e}$ is represented by the identity matrix. Using these same coordinates on $\mathbb{R}^g$ (up to a translation from Pic^g to $\operatorname{Pic}^{m+g-1}$), for $m \geq 1$ the definition $\sigma_m(x, E) = [mx + E]$ implies that $f_{T,e}^m$ is represented by the diagonal matrix

$$\begin{pmatrix} m & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}.$$

This shows that $f_m^{T,e}$ is locally injective, which implies $\sigma_m^{T,e}$ is locally injective as well. Thus for any $[D] \in \operatorname{Pic}^{m+g-1}(\Gamma)$, the preimage under $\sigma_m^{T,e}$ is a discrete subset of $\widetilde{C}_{T,e}$. Since $\widetilde{C}_{T,e}$ is compact, the preimage of $[D]$ is finite as claimed. $\square$

In the following proposition, "generic" means the statement holds for $[D] \in \operatorname{Pic}^n(\Gamma)$ outside of a nowhere dense exceptional set.

Proposition 4.19. *For any divisor class $[D]$ of degree $n \geq g$, we have*

$$\# W^{\mathrm{st}}(D) \leq g(n - g + 1).$$

For a generic divisor class $[D]$ of degree $n \geq g$, the stable Weierstrass locus $W^{\mathrm{st}}(D)$ has cardinality $\# W^{\mathrm{st}}(D) = g(n - g + 1)$.

Proof. Let $\widetilde{\operatorname{Br}}^g(\Gamma)$, $\sigma_m : \widetilde{\operatorname{Br}}^g(\Gamma) \to \operatorname{Pic}^{m+g-1}(\Gamma)$, and $\pi_1 : \widetilde{\operatorname{Br}}^g(\Gamma) \to \Gamma$ be defined as in Section 4.3.1. Recall that for a divisor D of degree $m+g-1$, we have $\# W^{\mathrm{st}}(D) = \#(\sigma_m^{-1}[D])$ by Lemma 4.13. Thus it suffices to show that $\sigma_m : \widetilde{\operatorname{Br}}^g(\Gamma) \to \operatorname{Pic}^{m+g-1}(\Gamma)$ is a branched covering map of degree gm, for any $m \geq 1$. From this, Lemma 4.17

implies the inequality $\#W^{\mathrm{st}}(D) \le gm$ and Definition 4.15 implies that equality holds for $[D]$ outside of the ramification locus.

(If D has degree $n = m + g - 1$, then $gm = g(n - g + 1)$.)

Claim 1: The map $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ is open, for any $m \ge 1$.

Proof of Claim 1: As above, let (G, ℓ) be a combinatorial model for Γ, and

$$\widetilde{\mathrm{Br}}^g(\Gamma) = \bigcup_{T \in \mathcal{T}(G)} \bigcup_{e \notin E(T)} \widetilde{C}_{T,e}$$

the induced ABKS decomposition. (See Section 4.3.1.) The map σ_m is naturally a piecewise affine map with domains of linearity $\widetilde{C}_{T,e}$.

To show that σ_m is open, it suffices to check that for any $(x_0, E_0) \in \widetilde{\mathrm{Br}}^g(\Gamma)$, the image of a neighborhood contains points in all tangent directions around $\sigma_m(x_0, E_0) \in \mathrm{Pic}^{m+g-1}(\Gamma)$. To check this, we observe how σ_m restricts to each domain of linearity $\widetilde{C}_{T,e}$ containing (x_0, E_0). We will show that the behavior of σ_m on tangent directions does not depend on the integer m.

For a point (x_0, E_0) in $\widetilde{C}_{T,e}$, let $\mathrm{cone}(\sigma_m^{T,e}(x_0, E_0))$ denote the positive cone in $\mathbb{R}^g$ spanned by

$$\sigma_m(x, E) - \sigma_m(x_0, E_0) \quad \text{for } (x, E) \text{ in a neighborhood of } (x_0, E_0) \text{ in } \widetilde{C}_{T,e}.$$

(Here we identify $\mathbb{R}^g$ with the tangent space of $\mathrm{Pic}^0(\Gamma)$ at the identity.) Since σ_m is affine on $\widetilde{C}_{T,e}$, this cone does not depend on the neighborhood chosen. Since $m \ge 1$, the positive span of

$$\sigma_m(x, E) - \sigma_m(x_0, E_0) = m[x - x_0] + [E - E_0] \quad \text{for } (x, E) \text{ in } \widetilde{C}_{T,e}$$

is equal to the positive span of

$$\sigma_1(x + E) - \sigma_1(x_0 + E_0) = [x - x_0] + [E - E_0] \quad \text{for } (x, E) \text{ in } \widetilde{C}_{T,e},$$

so $\mathrm{cone}(\sigma_m^{T,e}(x_0, E_0)) = \mathrm{cone}(\sigma_1^{T,e}(x_0, E_0))$. This holds for all cells $\widetilde{C}_{(T,e)}$ containing (x_0, E_0).

Hence to show that σ_m is open, it suffices to show that $\sigma_1 : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma)$ is open. This is clear from the construction of $\widetilde{\mathrm{Br}}^g(\Gamma)$ as a branched cover $\widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$, and from Theorem 2.8 which states that $\mathrm{Br}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma)$ is a homeomorphism.

Claim 2: The map $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ is a branched cover, for any $m \ge 1$.

Proof of Claim 2: In the definition of branched cover, Definition 4.15, condition (ii) was verified by Claim 1 and condition (iii) was verified by Proposition 4.18. Condition (i) is clear.[1]

[1] The map σ_m is surjective because it is an open map from a compact space to a connected, Hausdorff space.

We first identify a ramification locus R for σ_m, and then apply Lemma 4.16 to show that the restriction of σ_m away from R is a covering map.

Let $\mathrm{Br}^g(\Gamma) = \bigcup_{T\in\mathcal{T}(G)} C_T$ be the ABKS decomposition induced by a combinatorial model $\Gamma = (G,\ell)$ (see Section 2.5). Let $Z^{(2)} \subset \mathrm{Br}^g(\Gamma)$ denote the union of faces of C_T of codimension at least 2, and let $U^{(2)} = \mathrm{Br}^g(\Gamma)\backslash Z^{(2)}$. In other words,

$$U^{(2)} = \bigcup_{T\in\mathcal{T}(G)} \{\text{interior } C_T^\circ \text{ of } C_T\} \cup \{\text{interiors of facets of } \partial C_T\}.$$

More concretely in terms of break divisors, given a set of edges $e_1,\dots,e_g$ in G whose complement is a spanning tree, $U^{(2)}$ contains break divisors which are a sum of g points taken from the interior of each $e_1, e_2, \dots, e_g$, and divisors which are a sum of one endpoint of e_1 and a point in the interior of each $e_2,\dots,e_g$. We assume our combinatorial model (G,ℓ) is chosen to have no loops, so that each cell C_T in the ABKS decomposition has $2g$ distinct boundary facets.

Note that for a break divisor E,

$$\text{if } E \in U^{(2)}, \text{ the support of } E \text{ consists of } g \text{ distinct points.} \tag{4.3}$$

We let $\widetilde{Z^{(2)}}$ and $\widetilde{U^{(2)}}$ denote the preimages of $Z^{(2)}$ and $U^{(2)}$ under $\sigma : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Br}^g(\Gamma)$. Note that with respect to the ABKS decomposition

$$\widetilde{\mathrm{Br}}^g(\Gamma) = \bigcup_{T\in\mathcal{T}(G)} \bigcup_{e\notin E(T)} \widetilde{C}_{T,e},$$

$\widetilde{Z^{(2)}}$ is the union of codimension 2 faces of $\widetilde{C}_{T,e}$, and $\widetilde{U^{(2)}} = \widetilde{\mathrm{Br}}^g(\Gamma)\backslash\widetilde{Z^{(2)}}$. Thus $\widetilde{Z^{(2)}}$ is a closed subset of codimension 2 and $\widetilde{U^{(2)}}$ is a dense open subset of $\widetilde{\mathrm{Br}}^g(\Gamma)$.

Next, let $R = R_m = \sigma_m(\widetilde{Z^{(2)}})$. We will show that R is a valid ramification locus for the branched cover σ_m. The conditions (iv) and (v) hold because R is a codimesion 2 submanifold of the connected manifold $\mathrm{Pic}^{m+g-1}(\Gamma)$. It remains to check condition (vi), that the restriction

$$\sigma_m|_{\widetilde{\mathrm{Br}}^g(\Gamma)\backslash\sigma_m^{-1}(R)} : \widetilde{\mathrm{Br}}^g(\Gamma)\backslash\sigma_m^{-1}(R) \to \mathrm{Pic}^{m+g-1}(\Gamma)\backslash R \tag{4.4}$$

away from ramification is a covering map. To check this condition, we apply Lemma 4.16. It is clear that the domain and codomain of (4.4) are locally compact Hausdorff spaces.[2] The map in (4.4) is surjective by construction; it is proper because σ_m is a map from a compact space to a Hausdorff space, hence proper. It remains to check

[2] The domain is locally compact and Hausdorff because it is an open subspace of $\widetilde{\mathrm{Br}}^g(\Gamma)$ which is a finite CW complex, hence compact and Hausdorff. The same holds for the codomain, as an open subspace of $\mathrm{Pic}^{m+g-1}(\Gamma) \cong \mathbb{R}^g/\mathbb{Z}^g$.

that (4.4) is a local homeomorphism, which we leave for the next claim. Note that the domain of (4.4) is contained in $\widetilde{U^{(2)}}$:

$$\widetilde{\mathrm{Br}}^g(\Gamma)\backslash\sigma_m^{-1}(R) = \widetilde{\mathrm{Br}}^g(\Gamma)\backslash\sigma_m^{-1}(\sigma_m(\widetilde{Z^{(2)}})) \subset \widetilde{\mathrm{Br}}^g(\Gamma)\backslash Z^{(2)} = \widetilde{U^{(2)}}.$$

Assuming Claim 3, Lemma 4.16 implies that σ_m is a covering map away from the ramification locus R, which completes the proof of Claim 2.

Claim 3: The restriction of σ_m to $\widetilde{U^{(2)}} \to \mathrm{Pic}^{m+g-1}(\Gamma)$ is a local homeomorphism, for any $m \geq 1$.

Proof of Claim 3: First consider $m = 1$. Observation (4.3) implies that

$$\text{(4.5)} \quad \text{the restriction } \sigma_1|_{\widetilde{U^{(2)}}} : \widetilde{U^{(2)}} \to U^{(2)} \text{ is a (unbranched) covering of degree } g.$$

Since $U^{(2)} \subset \mathrm{Pic}^g(\Gamma)$ is open, it follows that $\sigma_1 : \widetilde{U^{(2)}} \to \mathrm{Pic}^g(\Gamma)$ is a local homeomorphism.

Recall that $\widetilde{U^{(2)}}$ is the union of the interior of $\widetilde{C}_{T,e}$ and the interiors of facets of $\partial\widetilde{C}_{T,e}$, over all (T, e). In the interior of $\widetilde{C}_{T,e}$, σ_m can be expressed as a full-rank linear map so it is a local homeomorphism. Now consider how σ_m acts near the interior of a facet of $\partial\widetilde{C}_{T,e}$. We claim that each facet is shared by exactly two cells.

Suppose $T = G\backslash\{e = e_1, e_2, \ldots, e_g\}$. There are $2g$ facets of the boundary $\partial\widetilde{C}_{T,e}$, indexed by choosing an edge e_j and choosing one of its two endpoints. For a fixed index j in $\{1, \ldots, g\}$ and $v(e_j)$ a fixed endpoint of e_j, the corresponding facet of $\partial\widetilde{C}_{T,e}$ consists of pairs $(x, E) \in \widetilde{\mathrm{Br}}^g(\Gamma)$ of the form

$$\text{(4.6)} \qquad \widetilde{F}^{(j,v)}_{(T,e)} = \{(x = x_1, E = x_2 + \cdots + x_g) : \begin{array}{l} x_j = v(e_j), \\ x_i \in e_i^\circ \text{ for } i = 1, \ldots g,\, i \neq j\} \end{array}$$

Let $G_j = T \cup e_j$. Since $e_j \notin T$, the graph G_j contains a unique cycle, which must contain $v(e_j) \in e_j$. Let e_j' be the unique edge $\neq e_j$ in this cycle which also borders $v(e_j)$, and let $T' = G_j\backslash e_j' = (T \cup e_j)\backslash e_j'$. Then $\widetilde{C}_{T',e'}$ is the only other cell containing the facet (4.6), where $e' = e_1'$ if $j = 1$, and $e' = e$ otherwise. The facet (4.6) is then the relative interior of $\widetilde{C}_{T,e} \cap \widetilde{C}_{T',e'}$

As before, let $f_m^{T,e}$ denote the lift of $\widetilde{C}_{T,e} \to \mathrm{Pic}^{m+g-1}(\Gamma)$ in the diagram

$$\begin{tikzcd}
\prod_{i=1}^g [0, \ell(e_i)] \arrow[r, dashed, "f_m^{T,e}"] \arrow[d] & \mathbb{R}^g \arrow[d, "\pi"] & \prod_{i=1}^g [0, \ell(e_i')] \arrow[l, dashed, "f_m^{T',e'}"'] \arrow[d] \\
\widetilde{C}_{T,e} \arrow[r] & \mathrm{Pic}^{m+g-1}(\Gamma) & \widetilde{C}_{T',e'} \arrow[l]
\end{tikzcd}$$

and define $f_m^{T',e'}$ analogously.

We may choose coordinates (depending on T) on $\mathbb{R}^g$ such that

$$\text{the matrix representing } f_m^{T,e} \text{ is } \begin{pmatrix} m & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}.$$

In these same coordinates, the matrix representing $f_m^{T',e'}$ is

$$\begin{pmatrix} -m & & & \\ * & 1 & & \\ * & & \ddots & \\ * & & & 1 \end{pmatrix} \text{ if } j = 1, \quad \text{or} \quad \begin{pmatrix} m & & * & \\ & \ddots & * & \\ & & -1 & \\ & & * & \ddots \end{pmatrix} \text{ if } j \in \{2, \ldots, g\}.$$

(Recall that j is the index specifying which edge $e_j \in G \backslash T$ has a break divisor chip on one of its endpoints; e_j is the unique edge in $T' \backslash T$.) This shows that σ_m is a local homeomorphism in a neighborhood of the chosen facet of $\partial \widetilde{C}_{T,e}$.

<u>Claim 4</u>: The branched cover $\sigma_m : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ has degree gm.

Proof of Claim 4: When $m = 1$, it is clear that $\sigma_1 : \widetilde{\mathrm{Br}}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma) \cong \mathrm{Br}^g(\Gamma)$ is a degree g branched cover. When $m > 1$, we note that σ_m differs from σ_1 by a scaling factor of m, i.e. on a sufficiently small neighborhood $U \subset \widetilde{\mathrm{Br}}(\Gamma)$, the Haar measure of $\sigma_m(U)$ is m-times as large as the Haar measure of $\sigma_1(U)$. (The space $\mathrm{Pic}^{m+g-1}(\Gamma)$ carries a Haar measure since it is a torsor for the compact topological group $\mathrm{Pic}^0(\Gamma)$.) This implies that the degree of σ_m as a branched cover must be m times the degree of σ_1, so σ_m must have degree gm as desired. $\square$

Theorem 4.20. *Let Γ be a compact, connected metric graph of genus g.*

(a) For a generic divisor class of degree $n \geq g$, the Weierstrass locus $W(D)$ is finite with cardinality $\#W(D) = g(n - g + 1)$. For a generic divisor class of degree $n < g$, $W(D)$ is empty.

(b) For an arbitrary divisor class of degree $n \geq g$, the stable Weierstrass locus $W^{\mathrm{st}}(D)$ is finite with cardinality

$$\#W^{\mathrm{st}}(D) \leq g(n - g + 1),$$

and equality holds for a generic divisor class.

Proof. Part (b) is a restatement of Proposition 4.19.

For part (a), first suppose $n < g$. The space $\mathrm{Pic}^n(\Gamma)$ has dimension g, while the subspace of effective divisor classes has dimension at most n. Thus a generic divisor class in $\mathrm{Pic}^n(\Gamma)$ is not effective, assuming $n < g$. By Remark 4.2, the Weierstrass locus is empty for a non-effective divisor class.

Now suppose $n \geq g$. To prove (a), it suffices to show that $W(D) = W^{\mathrm{st}}(D)$ for a generic divisor class, since then part (b) applies. To compare $W(D)$ with $W^{\mathrm{st}}(D)$, we construct a map $X \to \mathrm{Pic}^n(\Gamma)$ whose fiber over $[D]$ is the Weierstrass locus $W(D)$; this parallels our construction in Section 4.3.1 for $W^{\mathrm{st}}(D)$.

For $m \geq 1$, let $s_m : \Gamma \times \mathrm{Sym}^{g-1}(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ denote the map

$$s_m(x, E) = [mx + E].$$

Let $\pi_1 : \Gamma \times \mathrm{Sym}^{g-1}(\Gamma) \to \Gamma$ denote projection to the first factor.

The Riemann–Roch formula, Theorem 2.14, implies that a generic divisor class $[D] \in \mathrm{Pic}^{m+g-1}(\Gamma)$ has rank $r(D) = (m + g - 1) - g = m - 1$. For such a divisor,

$$W(D) = \{x \in \Gamma : [D - mx] \geq 0\} = \pi_1(s_m^{-1}[D]).$$

Recall that $W^{\mathrm{st}}(D) = \pi_1(\sigma_m^{-1}[D])$, where σ_m is defined to be the restriction of s_m to the subset $\widetilde{\mathrm{Br}}^g(\Gamma) \subset \Gamma \times \mathrm{Sym}^{g-1}(\Gamma)$; note that

$$\sigma_m^{-1}[D] = s_m^{-1}[D] \cap \widetilde{\mathrm{Br}}^g(\Gamma) \subset s_m^{-1}[D]. \tag{4.7}$$

Under the genericity assumption on $[D]$, we have

$$W^{\mathrm{st}}(D) = \pi_1(\sigma_m^{-1}[D]) \subset \pi_1(s_m^{-1}[D]) = W(D).$$

Using part (b), this observation implies that a generic Weierstrass locus $W(D)$ contains at least $g(n - g + 1)$ points.

We consider when $W(D)$ can be strictly larger than $W^{\mathrm{st}}(D)$. By (4.7), this happens only if $s_m^{-1}[D]$ is not contained in $\widetilde{\mathrm{Br}}(\Gamma)$; equivalently, only if $[D]$ lies in the image of $(\Gamma \times \mathrm{Sym}^{g-1}(\Gamma)) \backslash \widetilde{\mathrm{Br}}(\Gamma)$ under s_m.

Claim: The image $s_m(\,(\Gamma \times \mathrm{Sym}^{g-1}(\Gamma)) \backslash \widetilde{\mathrm{Br}}(\Gamma)\,)$ has dimension $g-1$ in $\mathrm{Pic}^{m+g-1}(\Gamma)$.

It is clear that s_m is piecewise affine on $\Gamma \times \mathrm{Sym}^{g-1}(\Gamma)$, with domains of linearity indexed by g-tuples of edges $(e_1; e_2, \ldots, e_g)$, up to reordering the edges $e_2, \ldots, e_g$. (Here we choose an arbitrary combinatorial model (G, ℓ) for Γ.) The edges e_i are not necessarily distinct.

If the edges $(e_1; e_2, \ldots, e_g)$ form the complement of a spanning tree T in G, then the corresponding domain is in $\widetilde{\mathrm{Br}}^g(\Gamma)$; namely, it is the cell $\widetilde{C}_{T,e_1}$ in the notation of Section 4.3.1. Conversely, if the edges $(e_1; e_2, \ldots, e_g)$ are not the complement of a spanning tree in G, then either some edge is repeated or the edges contain a cut set of G. In either case, the fibers of $s_m : \Gamma \times \mathrm{Sym}^{g-1}(\Gamma) \to \mathrm{Pic}^{m+g-1}(\Gamma)$ have dimension at least 1 over the interior of the corresponding domain (see [23, Proposition 13]), so the image of this domain under s_m has dimension at most $g-1$. This proves the claim.

The claim implies that for a generic divisor class $[D]$, the preimage $s_m^{-1}[D]$ is contained in $\widetilde{\mathrm{Br}}^g(\Gamma)$. By (4.7) this implies $W(D) = W^{\mathrm{st}}(D)$, as desired. $\square$

4.4 Distribution of Weierstrass points

In this section we prove Theorem 4.24, which states that for a degree-increasing sequence of generic divisors on a metric graph, the Weierstrass points become distributed with respect to the Zhang canonical measure (defined in Section 3.3). We also give a quantitative version of this distribution result, Theorem 4.26.

Our proofs of Theorems 4.24 and 4.26 work unchanged when $W(D)$ is replaced by the stable Weierstrass locus $W^{\mathrm{st}}(D)$.

4.4.1 Examples

First we consider some low genus examples of Weierstrass points converging to a limiting distribution.

Example 4.21 (Genus 0 metric graph)**.** Let Γ be a genus 0 metric graph. For any divisor D_n, the associated Weierstrass locus $W(D_n)$ is empty so $\delta_n = 0$. All edges are bridges, so the canonical measure is $\mu = 0$.

Example 4.22 (Genus 1 metric graph)**.** Let Γ be a genus 1 metric graph which consists of a loop of length L. For a divisor D_n of degree n, the Weierstrass locus $W_n = W(D_n)$ consists of n evenly-spaced points ("torsion points") around the loop. The distance between adjacent points is L/n, so on a segment e of length $\ell(e)$ the number of Weierstrass points is bounded by

$$\frac{\ell(e)}{L/n} - 1 \leq \#(W_n \cap e) \leq \frac{\ell(e)}{L/n} + 1.$$

This means the associated discrete measure $\delta_n = \frac{1}{n}\sum_{x \in W_n} \delta_x$ satisfies

$$\delta_n(e) = \frac{\#(W_n \cap e)}{n} \qquad \Rightarrow \qquad \frac{\ell(e)}{L} - \frac{1}{n} \leq \delta_n(e) \leq \frac{\ell(e)}{L} + \frac{1}{n}.$$

Hence $\delta_n(e) \to \frac{\ell(e)}{L} = \mu(e)$ as $n \to \infty$.

4.4.2 Proofs

We now address the limiting distribution of Weierstrass points $W(D_n)$ as $n \to \infty$ in the case of an arbitrary metric graph Γ.

Lemma 4.23. *Suppose the Weierstrass locus $W(D)$ is finite. Let $r = r(D)$.*

(a) If x is in the interior of a segment, $\mathrm{red}_x[D]$ *contains at most* $r+1$ *chips at* x.

(b) If x is in the interior of a segment $e \subset \Gamma$, $\mathrm{red}_x[D]$ *contains at most* $r+1$ *chips on e (including its endpoints).*

Proof. (a) Suppose $\mathrm{red}_x[D]$ contains $r+2$ chips at x. Then for sufficiently small ϵ we can move $r+1$ of these chips together for a distance ϵ in one direction, while moving 1 chip a distance $(r+1)\epsilon$ in the other. This gives a positive-length interval in $W(D)$, a contradiction.

(b) Suppose $\mathrm{red}_x[D]$ contains $r+2$ chips on the closed segment e. Note that at least r of these chips must be at x, in the interior of e. By chip-firing, we may move all $r+2$ chips to a single point x' in the interior of e. Then part (a) applies. □

Theorem 4.24. *Let $\{D_n : n \geq 1\}$ be a sequence of divisors on Γ with* $\deg D_n = n$. *Let W_n be the Weierstrass locus of D_n. Suppose each W_n is a finite set, and let*

$$\delta_n = \frac{1}{n} \sum_{x \in W_n} \delta_x$$

denote the normalized discrete measure on Γ associated to W_n. Then as $n \to \infty$, the measures δ_n converge weakly to the Zhang canonical measure μ on Γ.

Recall that by definition of weak convergence, Theorem 4.24 says that for any continuous function $f : \Gamma \to \mathbb{R}$, as $n \to \infty$ we have convergence

$$\frac{1}{n} \sum_{x \in W_n} f(x) =: \int_\Gamma f(x) \delta_n(dx) \quad \to \quad \int_\Gamma f(x) \mu(dx).$$

Proof of Theorem 4.24. To show weak convergence of measures on Γ it suffices to show convergence when integrated against step functions. Hence it suffices to integrate the measures against the indicator function of an arbitrary segment of Γ.

Let e be a segment in the metric graph Γ of length $\ell(e)$, with endpoints s and t. Let $W_n \cap e$ denote the set of Weierstrass points of D_n lying on the segment e. It suffices to show that

$$\lim_{n \to \infty} \frac{\#(W_n \cap e)}{n} = \mu(e). \tag{4.8}$$

Recall that by Proposition 3.23,

$$\mu(e) = \frac{\ell(e)}{\ell(e) + \ell_{\text{eff}}(\Gamma \backslash e)}$$

where $\ell_{\text{eff}}(\Gamma \backslash e)$ denotes the effective resistance between the endpoints of e when the interior of e is removed from Γ. (If $\Gamma \backslash e$ is disconnected, $\ell_{\text{eff}}(\Gamma \backslash e) = +\infty$ and

$\mu(e) = 0$.) We prove (4.8) by relating each side to the slope of a piecewise linear function on Γ.

For the right-hand side of (4.8), consider the voltage function $j_t^s : \Gamma \to \mathbb{R}$ (see Section 3.1). The voltage drop in Γ between endpoints of e is the effective resistance

$$j_t^s(s) - j_t^s(t) = r(s,t) = \frac{\ell(e)\ell_{\text{eff}}(\Gamma\backslash e)}{\ell(e) + \ell_{\text{eff}}(\Gamma\backslash e)},$$

by the parallel rule for effective resistance. Thus we have

$$\frac{j_t^s(s) - j_t^s(t)}{\ell(e)} = \frac{\ell_{\text{eff}}(\Gamma\backslash e)}{\ell(e) + \ell_{\text{eff}}(\Gamma\backslash e)} = 1 - \frac{\ell(e)}{\ell(e) + \ell_{\text{eff}}(\Gamma\backslash e)} = 1 - \mu(e). \tag{4.9}$$

(Recall that this slope can be interpreted as the current flowing along the segment e from s to t, since current $= \frac{\text{voltage drop}}{\text{resistance}}$.)

To connect j_t^s to the left-hand side of (4.8), we consider a sequence of piecewise-linear functions which are "discrete approximations" of j_t^s, and show that certain slopes in these functions are related to the number of Weierstrass points.

Let f_n be the piecewise $\mathbb{Z}$-linear function on Γ satisfying

$$\Delta(f_n) = \mathrm{red}_t[D_n] - \mathrm{red}_s[D_n] \qquad \text{and} \qquad f_n(t) = 0.$$

(Recall that $\mathrm{red}_x[D]$ denotes the x-reduced divisor linearly equivalent to D.) By Proposition 3.8, as $n \to \infty$ we have uniform convergence

$$\frac{1}{n} f_n \to j_t^s. \tag{4.10}$$

Thus to show (4.8) using (4.9) and (4.10), it suffices to show that

$$\lim_{n\to\infty} \frac{1}{n} \left(\frac{f_n(s) - f_n(t)}{\ell(e)} \right) = 1 - \lim_{n\to\infty} \frac{\#(W_n \cap e)}{n}. \tag{4.11}$$

We first give an intuitive explanation for (4.11): the slope of the function f_n on a directed segment is equal to the net flow of chips across the segment, as we move from $\mathrm{red}_s[D_n]$ to $\mathrm{red}_t[D_n]$ along any path in the linear system $|D_n|$. If we follow $\mathrm{red}_x[D_n]$ as x varies from s to t, we have $n - g$ chips moving in the "forward" direction of e (following x) and some number of chips moving in the reverse direction one-by-one. The number of "reverse-moving" chips is equal to $\#(W_n \cap e)$, since x is in W_n exactly when $\mathrm{red}_x[D_n]$ has an "extra" chip at x, i.e. when the $n - g$ chips on x collide with a reverse-moving chip. Thus the net number of chips moving across the segment e is equal to $(n - g) - \#(W_n \cap e)$, up to some bounded error due to boundary behavior. This yields (4.11) after dividing by n and taking $n \to \infty$.

Now we give a rigorous argument. Let $w_1, w_2, \ldots, w_m$ denote the Weierstrass points on e, ordered from s to t, so that $m = \#(W_n \cap e)$. Here we use the hypothesis that W_n is finite. (Note that $m = m_n$ depends on n.)

We partition the segment $e = [s, t]$ into subintervals $[s, w_1], [w_1, w_2], \ldots, [w_m, t]$. (It is possible that the intervals $[s, w_1]$ and $[w_m, t]$ are degenerate.) Let $\ell([w_i, w_{i+1}])$ denote the length of the segment $[w_i, w_{i+1}] \subset e$. We have

$$\ell(e) = \ell([s, w_1]) + \ell([w_1, w_2]) + \cdots + \ell([w_{m-1}, w_m]) + \ell([w_m, t]).$$

For each $i = 1, 2, \ldots, m-1$, let $g_n^{(i)}$ denote the function in $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ satisfying

$$\Delta(g_n^{(i)}) = \mathrm{red}_{w_{i+1}}[D_n] - \mathrm{red}_{w_i}[D_n],$$

and let $g_n^{(0)}$ and $g_n^{(m)}$ denote functions satisfying

$$\Delta(g_n^{(0)}) = \mathrm{red}_{w_1}[D_n] - \mathrm{red}_s[D_n], \qquad \text{and} \qquad \Delta(g_n^{(m)}) = \mathrm{red}_t[D_n] - \mathrm{red}_{w_m}[D_n].$$

By adding an appropriate constant, we may assume that $g_n^{(i)}(t) = 0$ for each $i = 0, 1, \ldots, m$. By telescoping of poles and zeros, we have

$$\Delta(f_n) = \Delta(g_n^{(0)}) + \Delta(g_n^{(1)}) + \cdots + \Delta(g_n^{(m)}).$$

With the additional constraint that $f_n(t) = \sum_i g_n^{(i)}(t) = 0$, this implies that

$$f_n = g_n^{(0)} + g_n^{(1)} + \cdots + g_n^{(m)}. \tag{4.12}$$

Thus we can compute $f_n(s) - f_n(t)$ by summing $\sum_{i=0}^{m} \left(g^{(i)}(s) - g^{(i)}(t) \right)$.

To analyze the slopes of $g^{(i)}$ on segment e, we make use of Lemma 4.23. This information is sufficient to deduce all slopes over e. We may assume without loss of generality that $r(D_n) = n - g$, since this holds for $n \geq 2g - 1$.

For $i = 1, 2, \ldots, m-1$, the function $g_n^{(i)}$ has slope $-(n-g)$ on the interval $[w_i, w_{i+1}]$, and slope 1 on e outside of this interval. See Figure 4.7.

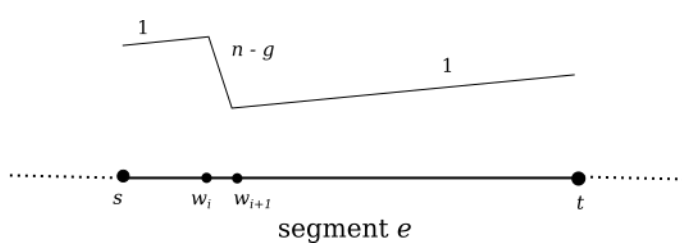

Figure 4.7: Function $g_n^{(i)}$ having zeros $\mathrm{red}_{w_{i+1}}[D_n]$ and poles $\mathrm{red}_{w_i}[D_n]$.

Thus we have

$$\begin{aligned} g_n^{(i)}(s) - g_n^{(i)}(t) &= (n-g)\ell([w_i, w_{i+1}]) - \ell([s, w_i]) - \ell([w_{i+1}, t]) \\ &= (n-g+1)\ell([w_i, w_{i+1}]) - \ell(e). \end{aligned} \tag{4.13}$$

For $i = 0$ and $i = m$, to write an expression for $g_n^{(i)}(x) - g_n^{(i)}(t)$ we need to set additional notation. If $\mathrm{red}_s[D_n]$ has a chip in the interior of e, let y be the position of this chip (which is unique by Lemma 4.23); otherwise, let $y = t$. Similarly, let z be the position of the unique chip of $\mathrm{red}_t[D_n]$ in the interior of e if it exists; otherwise let $z = s$. We have

$$\begin{aligned} g_n^{(0)}(s) - g_n^{(0)}(t) &= (n-g)\ell([s,w_1]) - \ell([w_1,y]) \\ &= (n-g+1)\ell([s,w_1]) - \ell([s,y]) \end{aligned} \tag{4.14}$$

and

$$\begin{aligned} g_n^{(m)}(s) - g_n^{(m)}(t) &= (n-g)\ell([w_m,t]) - \ell([z,w_m]) \\ &= (n-g+1)\ell([w_m,t]) - \ell([z,t]) \end{aligned} \tag{4.14'}$$

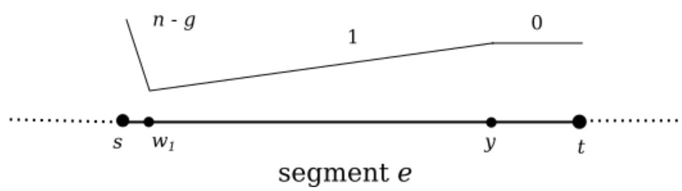

Figure 4.8: Function $g_n^{(0)}$ having zeros $\mathrm{red}_{w_1}[D_n]$ and poles $\mathrm{red}_s[D_n]$.

Thus adding the expressions (4.13) and (4.14) together, by (4.12) we have

$$\begin{aligned} f_n(s) - f_n(t) &= (n-g+1)\big(\ell([s,w_1]) + \ell([w_1,w_2]) + \cdots + \ell([w_{m-1},w_m]) + \ell([w_m,t])\big) \\ &\qquad - \ell([s,y]) - (m-1)\ell(e) - \ell([z,t]) \\ &= (n-g+1)\ell(e) - (m-1)\ell(e) - \ell([s,y]) - \ell([z,t]) \\ &= (n-g-m+2)\ell(e) - \ell([s,y]) - \ell([z,t]) \\ &= (n-g-m)\ell(e) + (\ell(e) - \ell([s,y])) + (\ell(e) - \ell([z,t])) \\ &= (n-g-m)\ell(e) + \ell([y,t]) + \ell([s,z]). \end{aligned}$$

Since $0 \le \ell([y,t]) + \ell([s,z]) \le 2\ell(e)$ and $m = \#(W_n \cap e)$, this shows that

$$n - g - \#(W_n \cap e) \le \frac{f_n(s) - f_n(t)}{\ell(e)} \le n - g + 2 - \#(W_n \cap e).$$

Dividing by n and taking the limit $n \to \infty$ yields (4.11) as desired. $\square$

Theorem 4.25. *Consider the setup of Theorem 4.24.*

(a) Suppose each $[D_n]$ is generic in $\mathrm{Pic}^n(\Gamma)$. Then each W_n is finite and we have weak convergence $\delta_n \to \mu$.

(b) Let $W_n^{\mathrm{st}} = W^{\mathrm{st}}(D_n)$ be the stable Weierstrass locus, and define δ_n^{st} analogously to δ_n. For any divisors $\{D_n : n \geq 1\}$ we have weak convergence $\delta_n^{st} \to \mu$.

Proof. (a) This is part of Theorem 4.20.

(b) We may follow the same argument used in Theorem 4.24, except in place of $\mathrm{red}_x[D_n]$ we consider the "stable reduced divisor"

$$\mathrm{red}_x^{\mathrm{st}}[D_n] := (n-g)x + \mathrm{br}[D_n - (n-g)x].$$

With this change in the definitions of f_n and $g_n^{(i)}$, equations (4.13) and (4.14) still hold, as does the convergence (4.10). □

Theorem 4.26 (Quantitative distribution of $W(D)$). *Let Γ be a metric graph of genus g, let D_n be a divisor class of degree $n > g$ and let W_n denote the Weierstrass locus of D_n. Suppose W_n is finite. Let μ denote the Zhang canonical measure on Γ.*

(a) For any segment e in Γ,

$$n\mu(e) - 2g \leq \#(W_n \cap e) \leq n\mu(e) + g + 2.$$

(b) If e is a segment of Γ with canonical measure $\mu(e) > \frac{2g}{n}$, then e contains at least one Weierstrass point of D_n.

(c) For a fixed continuous function $f : \Gamma \to \mathbb{R}$,

$$\frac{1}{n}\sum_{x \in W_n} f(x) = \int_\Gamma f(x)\mu(dx) + O\left(\frac{1}{n}\right).$$

Proof. It is clear that part (b) follows from part (a), since $\#(W_n \cap e)$ must be an integer. Part (c) is a straightforward extension of (a).

We now prove part (a). Let f_n be the piecewise linear function satisfying $\Delta(f_n) = \mathrm{red}_t[D_n] - \mathrm{red}_s[D_n]$ and $f_n(t) = 0$, where s and t are the endpoints of e. By Proposition 3.9, we have

$$|(f_n - (n-g)j_t^s)'(x)| \leq g$$

so

$$|f_n'(x)| \leq (n-g)|j'(x)| + g.$$

Recall that for x on the segment e, $|j'(x)| = 1 - \mu(e)$. Thus we have the bound

$$|f_n'(x)| \leq n - n\mu(e) + \mu(e)g.$$

Moreover the proof of Theorem 4.24 shows that

$$n - g - \#(W_n \cap e) \leq |f_n'(x)|.$$

Combining these inequalities gives

$$n\mu(e) - (1 + \mu(e))g \leq \#(W_n \cap e).$$

Finally, the inequality $\mu(e) \leq 1$ from Corollary 3.24 yields the lower bound in (a).

We similarly obtain the upper bound

$$\#(W_n \cap e) \leq n\mu(e) + g + 2$$

by combining the inequalities

$$n - n\mu(e) - (2 - \mu(e))g \leq |f_n'(x)| \quad \text{and} \quad |f_n'(x)| \leq n - g - \#(W_n \cap e) + 2$$

and $\mu(e) \geq 0$ from Corollary 3.24. □

4.5 Tropicalizing Weierstrass points

In this section, we describe how the Weierstrass locus for a tropical curve can be related to the Weierstrass locus for an algebraic curve. The key result is Baker's Specialization Lemma [6, Lemma 2.8]; here we use a more general version given by Jensen–Payne [26] in the language of Berkovich analytic spaces.

Throughout this section, let K denote an algebraically closed field equipped with a nontrivial non-Archimedean valuation $v : K^\times \to \mathbb{R}$; we assume K is complete with respect to v.

Theorem 4.27 (Specialization Lemma [26, Lemma 2.4])**.** *Suppose X is a smooth projective algebraic curve over K. Let Γ be a skeleton on the Berkovich analytification X^{an}, let $\rho : X^{an} \to \Gamma$ be the retraction to the skeleton and let $\rho_* : \mathrm{Div}(X) \to \mathrm{Div}(\Gamma)$ denote the induced map on divisors. Then for any divisor $D \in \mathrm{Div}(X)$,*

$$r_X(D) \leq r_\Gamma(\rho_*(D)).$$

Here r_X denotes the dimension of a complete linear system $|D|$ on X, and r_Γ denotes the Baker–Norine rank on Γ (see Section 2.6).

Theorem 4.28. *Consider the setup of Theorem 4.27. For any divisor $D \in \mathrm{Div}(X)$ such that $\rho_*(D) \in \mathrm{Div}(\Gamma)$ is Riemann–Roch nonspecial, we have*

$$\rho_*(W_X(D)) \subseteq W_\Gamma(\rho_*(D)).$$

Proof. The map ρ_* respects degree; let $n = \deg(D) = \deg(\rho_*(D))$. Recall that $\rho_*(D)$ is *nonspecial* means that

$$r_\Gamma(\rho_*(D)) = \max\{n - g,\, -1\}.$$

In this case, Theorem 4.27 implies $r_X(D) \le \max\{n-g, -1\}$ while Riemann–Roch implies $r_X(D) \ge \max\{n-g, -1\}$ for any divisor. Thus $r_X(D) = r_\Gamma(\rho_*(D))$.

Let r denote the rank in either sense. If $x \in W_X(D)$, we have

$$r_X(D - (r+1)x) \ge 0.$$

By Theorem 4.27 and linearity of ρ_*, this implies

$$r_\Gamma(\rho_*(D - (r+1)x)) = r_\Gamma(\rho_*(D) - (r+1)\rho_*(x)) \ge 0.$$

This means $\rho_*(x) \in W_\Gamma(\rho_*(D))$ as claimed. $\square$

The conclusion of Theorem 4.28 also holds for $D = K_X$ the canonical divisor, and $\rho_*(K_X) \sim K_\Gamma$. This was observed by Baker in [6, Corollary 4.9].

CHAPTER 5

Torsion Points of the Jacobian

In this chapter we study torsion points in the Jacobian of a tropical curve. Given a metric graph Γ of genus g, we are specifically interested in torsion points which lie in the image of the Abel–Jacobi map $AJ : \Gamma \to \mathrm{Jac}(\Gamma)$, which embeds a metric graph (of genus $g \geq 1$) in its Jacobian. In other words, we are interested in studying the intersection

$$AJ(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}.$$

The Jacobian $\mathrm{Jac}(\Gamma)$ is a compact abelian group isomorphic to $\mathbb{R}^g/\mathbb{Z}^g$. The subgroup of torsion points $\mathrm{Jac}(\Gamma)_{\mathrm{tors}}$ is isomorphic to $\mathbb{Q}^g/\mathbb{Z}^g$.

5.1 The Manin–Mumford conjecture for tropical curves

The Manin–Mumford conjecture states that for a smooth algebraic curve X of genus $g \geq 2$, the analogous intersection $AJ(X) \cap \mathrm{Jac}(X)_{\mathrm{tors}}$ is a finite set. This statement was proved by Raynaud in [40]. This gives us motivation to ask whether the analogous finiteness statement holds for a metric graph; we consider this a "tropical" Manin–Mumford conjecture.

It turns out that the tropical Manin–Mumford conjecture fails for a fairly large class of metric graphs—namely, those graphs which have rational edge lengths (Proposition 5.21). For a tropical analogue to work, additional constraints are needed on the metric graphs.

Our first main result of this chapter is that the tropical Manin–Mumford conjecture does hold for a metric graph whose edge lengths are "sufficiently irrational." We then prove a higher-degree generalization of this theorem: assuming sufficiently general edge lengths, we determine the values of d such that the map $AJ^{(d)} : \Gamma^d \to \mathrm{Jac}(\Gamma)$ has finitely many torsion points in its image.

Definition 5.1. We say a metric graph Γ satisfies the *Manin–Mumford condition* if $\#(AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}})$ is finite for every choice of basepoint $q \in \Gamma$.

(The Abel–Jacobi map $AJ_q : \Gamma \to \mathrm{Jac}(\Gamma)$ is defined in Section 2.3). Our first main theorem verifies a tropical version of the Manin–Mumford conjecture.

Theorem 5.30. *Suppose G is a biconnected graph of genus $g \geq 2$. For a very general choice of edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ satisfies the Manin–Mumford condition.*

Recall that a graph G is *biconnected* (or *two-connected*) if G is connected after deleting any vertex. We say that a property holds for a *very general* point of some real parameter space if it holds outside of a countable collection of proper Zariski-closed subsets. See Section 5.3 for further discussion of these conditions.

5.1.1 Higher-degree Manin–Mumford

Given an effective divisor Q of degree d, there is an associated Abel–Jacobi map

$$\begin{aligned} AJ_Q^{(d)} &: \Gamma^d \to \mathrm{Jac}(\Gamma) \\ (x_1, \dots, x_d) &\mapsto [\sum_{i=1}^{d} x_i - Q]. \end{aligned} \tag{5.1}$$

We may ask whether the image contains finitely many torsion points of $\mathrm{Jac}(\Gamma)$.

Definition 5.2. We say a metric graph Γ satisfies the *degree-d Manin–Mumford condition* if the image of the d-dimensional Abel–Jacobi map

$$AJ_Q^{(d)} : \Gamma^d \to \mathrm{Jac}(\Gamma)$$

intersects finitely many torsion points of $\mathrm{Jac}(\Gamma)$, for every choice of effective base-divisor $Q \in \mathrm{Eff}^d(\Gamma)$. We abbreviate this condition as $\mathrm{MM}(d)$.

When $d = 1$ the condition $\mathrm{MM}(1)$ is the usual Manin–Mumford condition on Γ. When $g = g(\Gamma) \geq 1$ and $d \geq g$, then $\mathrm{MM}(d)$ cannot hold, since the higher Abel–Jacobi map $AJ_Q^{(d)}$ is surjective and $\mathrm{Jac}(\Gamma)_{\mathrm{tors}}$ is infinite. If a metric graph Γ satisfies $\mathrm{MM}(d)$, then it also satisfies $\mathrm{MM}(d')$ for every $1 \leq d' \leq d$.

Theorem 5.38. *Let G be a connected graph of genus $g \geq 1$ and independent girth γ^{ind}. For a very general choice of edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ satisfies* $\mathrm{MM}(d)$ *if and only if* $1 \leq d < \gamma^{\mathrm{ind}}$.

The *independent girth* of G is a combinatorial invariant which is defined in Section 5.3. This invariant satisfies $\gamma^{\mathrm{ind}} \leq \gamma$, where γ denotes the usual girth, i.e. the minimal length of a cycle. We show that in relation to the genus, $\gamma^{\mathrm{ind}} < C \log g$.

5.2 The classical Manin–Mumford conjecture

Given an algebraic curve X and choice of basepoint x_0, we say that $x \in X$ is a *torsion point* if the divisor $n(x - x_0)$ is linearly equivalent to 0 for some positive integer n. Equivalently, x is a torsion point if the Abel–Jacobi embedding (with respect to x_0) sends x to the torsion subgroup of the Jacobian. The Jacobian of a genus g smooth algebraic curve over $\mathbb{C}$ is a compact abelian group, isomorphic to $\mathbb{C}^g/\mathbb{Z}^{2g} \cong H^1(X, \mathbb{C})/H_1(X, \mathbb{Z})^\vee$.

Faltings's theorem (previously known as Mordell's conjecture) states that a smooth curve of genus $g \geq 2$ has finitely many rational points, i.e. points whose coordinates are all rational numbers.

By analogy with Mordell's conjecture, Manin and Mumford conjectured that an algebraic curve of genus 2 or more has finitely many torsion points. The Manin–Mumford Conjecture was proved by Raynaud [39], which inspired several generalizations concerning torsion points in abelian varieties.

5.3 Definitions and setup

Given an abelian group A, the *torsion subgroup* A_{tors} is the set of elements $a \in A$ such that $na = a + \cdots + a = 0$ for some positive integer n. It may be checked that this defines a subgroup of A. For example, the torsion subgroup of $\mathbb{R}/\mathbb{Z}$ is $\mathbb{Q}/\mathbb{Z}$ and the torsion subgroup of $\mathbb{R}$ is $\{0\}$. Recall that the *Jacobian* $\mathrm{Jac}(\Gamma)$ of a metric graph is the abelian group on the set of degree 0 divisor classes; we have

$$\mathrm{Jac}(\Gamma)_{\text{tors}} = \{[D] : D \in \mathrm{Div}^0(\Gamma), \quad n[D] = 0 \text{ for some } n \in \mathbb{Z}_{>0}\}.$$

We say points $x, y \in \Gamma$ are *torsion equivalent* if there exists a positive integer n such that $n[x - y] = 0$ in $\mathrm{Jac}(\Gamma)$. If two points x, y represent the same divisor class $[x] = [y]$, then x and y are torsion equivalent; hence this relation descends to a relation on $\mathrm{Eff}^1(\Gamma) = \{[x] : x \in \Gamma\}$. It will be convenient for us to consider this relation on $\mathrm{Eff}^1(\Gamma)$ rather than on Γ.

Lemma 5.3. *Torsion equivalence defines an equivalence relation on* $\mathrm{Eff}^1(\Gamma)$.

Proof. It is clear that torsion equivalence is reflexive and symmetric. Suppose n, m are positive integers such that $n[x - y] = 0$ and $m[y - z] = 0$ in $\mathrm{Jac}(\Gamma)$. Then $mn[x - z] = mn([x - y] + [y - z]) = 0$. This shows that torsion equivalence is transitive. $\square$

It is natural to extend this relation to divisor classes of higher degree: we say effective classes $D, E \in \mathrm{Eff}^d(\Gamma)$ are *torsion equivalent* if $n[D - E] = 0$ for some positive integer n. We call an equivalence class under this relation a *torsion packet*.

Definition 5.4. The *torsion packet* of $[E] \in \mathrm{Eff}^d(\Gamma)$ is the set of divisor classes

$$\{[E]\}_{\mathrm{tors}} = \{[D] \in \mathrm{Eff}^d(\Gamma) \text{ such that } [D-E] \in \mathrm{Jac}(\Gamma)_{\mathrm{tors}}\}.$$

The terminology of torsion packets allows us to restate the Manin–Mumford condition in a basepoint-free manner.

Proposition 5.5.

(a) Given an effective divisor class $[D] \in \mathrm{Eff}^d(\Gamma)$, there is a canonical bijection

$$\{[D]\}_{\mathrm{tors}} \quad \leftrightarrow \quad AJ_D^{(d)}(\Gamma^d) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}$$

where $AJ_D^{(d)} : \Gamma^d \to \mathrm{Jac}(\Gamma)$ is the Abel–Jacobi map (5.1)*.*

(b) A metric graph Γ satisfies the degree d Manin–Mumford condition if and only if every torsion packet of degree d is finite.

Proof. For part (a), we have the diagram

$$\begin{tikzcd}
 & \{[D]\}_{\mathrm{tors}} \arrow[r] \arrow[d] & \mathrm{Jac}(\Gamma)_{\mathrm{tors}} \arrow[d] \\
\Gamma^d \arrow[r] \arrow[rr, bend right, "AJ^{(d)}"'] & \mathrm{Eff}^d(\Gamma) \arrow[r] & \mathrm{Jac}(\Gamma)
\end{tikzcd}$$

where the torsion packet $\{[D]\}_{\mathrm{tors}}$ is the pullback of the two inclusions $\mathrm{Eff}^d(\Gamma) \to \mathrm{Jac}(\Gamma)$ and $\mathrm{Jac}(\Gamma)_{\mathrm{tors}} \to \mathrm{Jac}(\Gamma)$, and $\Gamma^d \to \mathrm{Eff}^d(\Gamma)$ is surjective.

Part (b) follows directly from (a) and the definitions above. $\square$

Recall that the voltage function j_y^x is the piecewise $\mathbb{R}$-linear function satisfying

$$\Delta(j_y^x) = y - x \qquad \text{and} \qquad j_y^x(y) = 0.$$

Lemma 5.6. *Suppose x, y are two points on a metric graph Γ. Then $[x-y]$ is torsion in the Jacobian of Γ if and only if all slopes of the voltage function j_y^x are rational.*

The above lemma is the special case $d = 1$ of the following statement.

Lemma 5.7. *Suppose $D = x_1 + \cdots + x_d$ and $E = y_1 + \cdots + y_d$ are effective divisors of degree d on a metric graph Γ. Let $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ be a function satisfying $\Delta(f) = D - E$. (Up to an additive constant, $f = \sum_{i=1}^d j_{x_i}^{y_i}$.)*

(a) The divisor class $[D-E] = 0$ if and only if all slopes of f are integers.

(b) The divisor class $[D-E]$ is torsion if and only if all slopes of f are rational.

Proof. Part (a) is a restatement of the definition of linear equivalence (Section 2.2). Part (b) follows from part (a) by linearity of the Laplacian Δ: $[D-E]$ is torsion of order n iff $[n(D-E)] = [n\,\Delta(f)] = [\Delta(n \cdot f)] = 0$ iff all slopes of $n \cdot f$ lie in $\mathbb{Z}$ iff all slopes of f lie in $\frac{1}{n}\mathbb{Z}$. $\square$

5.3.1 Very general subsets

A *very general* subset of $\mathbb{R}^n$ is one whose complement is contained in a countable union of distinguished Zariski-closed sets. A distinguished Zariski-closed set is the set of zeros of a polynomial function which is not identically zero[1]. Given a polynomial $f \in \mathbb{R}[x_1, \ldots, x_n]$, we denote

$$Z(f) = \{(a_1, \ldots, a_n) \in \mathbb{R}^n : f(a) = 0\} \quad \text{and} \quad U(f) = \{(a_1, \ldots, a_n) \in \mathbb{R}^n : f(a) \neq 0\}.$$

In this notation, a very general subset $S \subset \mathbb{R}^n$ is one which can be expressed as

$$S \supset \mathbb{R}^n \setminus \left(\bigcup_{i \in I} Z(f_i) \right) = \bigcap_{i \in I} U(f_i)$$

where I is a countable index set and each f_i is nonzero. Note that the zero locus $Z(f)$ has Lebesgue measure zero if f is nonzero. Thus the complement of a (measurable) very general subset of $\mathbb{R}^n$ has Lebesgue measure zero. However, it is still possible that the complement of a very general subset is dense in $\mathbb{R}^n$.

If $D \subset \mathbb{R}^n$ is some parameter space with nonempty interior (with respect to the Euclidean topology), we say that a subset of D is *very general* if it has the form $D \cap S$ for a very general subset $S \subset \mathbb{R}^n$. In our applications, the relevant parameter space will be the positive orthant $D = (\mathbb{R}_{>0})^n$. We say that a property holds for a *very general* point of some real parameter space if it holds on a very dense subset.

Example 5.8.

(a) For a fixed nonconstant polynomial $f \in \mathbb{Z}[x_1, \ldots, x_n]$, the set

$$U(f - \mathbb{Q}) = \{(a_1, \ldots, a_n) \in \mathbb{R}^n : f(a_1, \ldots, a_n) \notin \mathbb{Q}\} \tag{5.2}$$

is very general, since $\{f - \lambda : \lambda \in \mathbb{Q}\}$ is a countable collection of nonzero polynomials.

(b) For polynomials $f, g \in \mathbb{Z}[x_1, \ldots, x_n]$ with $g \neq 0$ and f/g nonconstant, the set

$$U(\frac{f}{g} - \mathbb{Q}) = \{(a_1, \ldots, a_n) \in \mathbb{R}^n : \frac{f(a_1, \ldots, a_n)}{g(a_1, \ldots, a_n)} \notin \mathbb{Q}\} \tag{5.3}$$

is very general, since $\{f - \lambda g : \lambda \in \mathbb{Q}\}$ is a countable collection of nonzero polynomials.

(c) The set

$$U^n_{\text{tr.}} = \{(a_1, \ldots, a_n) \in \mathbb{R}^n : f(a_1, \ldots, a_n) \neq 0 \quad \text{for every } f \in \mathbb{Z}[x_1, \ldots, x_n] \setminus \{0\}\} \tag{5.4}$$

[1] More generally, a Zariski-closed set is the set of common zeros of a finite collection of polynomials

is very general, since $\mathbb{Z}[x_1, \ldots, x_n]$ is countable. We call $U^n_{\text{tr.}}$ the set of *transcendental* points of $\mathbb{R}^n$. In particular, U^1_{tr} is the set of transcendental real numbers.

Note that in the above examples, the subsets (5.2) and (5.3) contain the transcendental points $U^n_{\text{tr.}}$. Conversely, $U^n_{\text{tr.}}$ is the intersection of (5.2) over all choices of f (resp. (5.3) over all choices of f and g).

In the later theorem statements (5.30 and 5.38) which concern very general edge lengths, the stated property holds when the edge lengths are transcendental (in the sense of (5.4), $n = \#E(G)$). More precisely, these conditions will hold on a finite intersection of sets of the form (5.3). The polynomials f, g will come from Kirchhoff's formulas (see Theorem 3.26 in Section 3.5).

5.3.2 Critical group

The *critical group* $\text{Jac}(G)$ of a combinatorial graph G is a finite abelian group related to the Jacobian construction as follows. A combinatorial graph G can be viewed as a metric graph $\Gamma_{\mathbf{1}} = (G, \mathbf{1})$ with unit edge lengths. In the metric graph Jacobian, $\text{Jac}(G)$ is the subgroup of divisor classes supported on vertices of G,

$$\text{Jac}(G) := \{[D] : D \in \text{Div}^0(V(G))\} \subset \text{Jac}(\Gamma_{\mathbf{1}}).$$

The size of the critical group is equal to the number of spanning trees of G.

For more on the critical group, see Baker–Norine [9] and the references therein.

Example 5.9. Let G be the theta graph, shown below (left), which has two vertices x and y connected by three edges. The critical group $\text{Jac}(G)$ has order three and is generated by the divisor class $[x - y]$. The multiples $n[x - y]$ inside the metric graph Jacobian are illustrated in Figure 5.1, to the right.

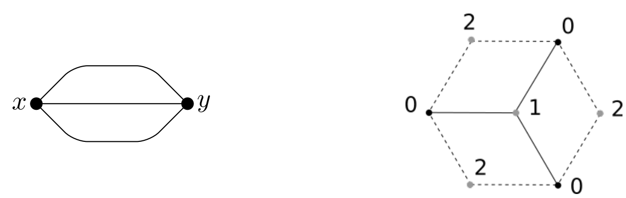

Figure 5.1: Graph with critical group of order 3.

Example 5.10. Let G be the graph shown on the left of Figure 5.2. The critical group has order 11 and is generated by the divisor class $[x-y]$. The multiples $n[x-y]$ for $n = 0, 1, \ldots, 10$ are shown in Figure 5.2 on the right.

In contrast to the examples above, the critical group is not always cyclic. The graph G obtained from the theta graph by subdividing each edge into m edges has

$$\text{Jac}(G) \cong \mathbb{Z}/(m) \times \mathbb{Z}/(3m).$$

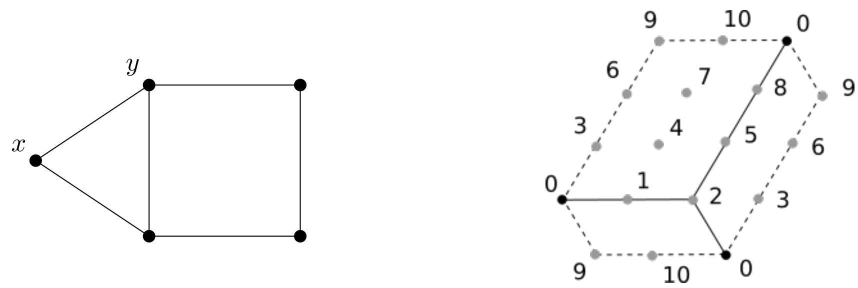

Figure 5.2: Graph with critical group of order 11.

In general, $\mathrm{Jac}(G)$ decomposes as a direct sum of k cyclic groups (e.g. in the invariant factor decomposition), where k is bounded above by the genus of G.

5.3.3 Stabilization of metric graphs

A connected combinatorial graph G is *stable* if every vertex $v \in V(G)$ has valence at least 3, and *semistable* if every vertex has $\mathrm{val}(v) \geq 2$. A metric graph Γ is *semistable* if every point $x \in \Gamma$ has valence at least 2. Note that a nontrivial metric graph cannot be stable, since points in the interior of an edge will have valence 2.

This notion is useful for our purposes because questions about Abel–Jacobi maps $AJ : \Gamma \to \mathrm{Jac}(\Gamma)$ maybe be reduced to $AJ : \Gamma' \to \mathrm{Jac}(\Gamma')$ where Γ' is a semistable metric graph. This allows us to find explicit bounds on the number of points

$$\#(AJ(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}) \qquad \text{and} \qquad \#(AJ^{(d)}(\Gamma^d) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}),$$

when these numbers are finite, see Theorems 5.24 and 5.25.

Proposition 5.11 (Metric graph stabilization)**.** *Suppose Γ has genus $g \geq 1$.*

(a) There is a canonical semistable subgraph $\Gamma' \subset \Gamma$ and a retract map $r : \Gamma \to \Gamma'$ such that r is a homotopy inverse to the inclusion $\Gamma' \to \Gamma$.

(b) The retract $r : \Gamma \to \Gamma'$ induces an isomorphism $\mathrm{Jac}(\Gamma) \to \mathrm{Jac}(\Gamma')$ on Jacobians.

For a proof and further motivation, see Caporaso [14].

Example 5.12. Figure 5.3 shows the stabilization Γ' of a metric graph Γ of genus two. The retract map $\Gamma \to \Gamma'$ sends a point of Γ to the closest point of Γ' in the path metric.

Figure 5.3: A metric graph (left) and its stabilization (right).

Proposition 5.13. *A semistable metric graph Γ with genus $g \geq 2$ has a unique stable model (G, ℓ), (i.e. a model such that G is stable).*

Proof. The unique stable model has vertex set $V(G) = \{x \in \Gamma : \mathrm{val}(x) \geq 3\}$. The edges $E(G)$ correspond to connected components of $\Gamma \setminus V(G)$, which is isometric to a disjoint union of open intervals of finite length. $\square$

Proposition 5.14. *Suppose G is a stable graph of genus g. Then the number of edges in G is at most $3g - 3$.*

Proof. Since every vertex has valence at least 3, we have

$$\#V(G) \leq \frac{1}{3} \sum_{v \in V(G)} \mathrm{val}(v) = \frac{2}{3} \cdot \#E(G).$$

By the genus formula $g = \#E(G) - \#V(G) + 1$, this implies

$$\#E(G) = g - 1 + \#V(G) \leq g - 1 + \frac{2}{3} \cdot \#E(G)$$

which is equivalent to the desired inequality $\#E(G) \leq 3g - 3$. $\square$

It follows from the previous proposition that a stable graph has genus $g \geq 2$.

5.3.4 Girth and independent girth

Recall that the *girth* $\gamma = \gamma(G)$ of a graph is the minimal length of a cycle; a *(simple) cycle* is a subgraph homeomorphic to a circle.[2] In other words,

$$\gamma(G) = \min_{C \in \mathcal{C}(G)} \{\#E(C)\} \tag{5.5}$$

where $\mathcal{C}(G)$ denotes the set of cycles of G.

Definition 5.15. The *independent girth* γ^{ind} of a graph is defined as

$$\gamma^{\mathrm{ind}}(G) = \min_{C \in \mathcal{C}(G)} \{ \mathrm{rk}^{\perp}(E(C)) \} \tag{5.6}$$

where $\mathrm{rk}^{\perp}$ is the rank function of the cographic matroid $M^{\perp}(G)$. (See Section 2.7 for discussion of cographic matroids). If G has genus zero, we let $\gamma^{\mathrm{ind}}(G) = \gamma(G) = +\infty$.

Equivalently,

$$\gamma^{\mathrm{ind}}(G) = \min_{C \in \mathcal{C}(G)} \{ \#E(C) + 1 - h_0(G \backslash E(C)) \}$$

where $G \backslash E(C)$ denotes deleting the interior of each edge in C, and h_0 denotes the number of connected components of a topological space.

[2] We use the terms "cycle" and "simple cycle" of a graph interchangeably.

Proposition 5.16. *(a) For any graph G, $\gamma^{\text{ind}}(G) \le \gamma(G)$.*

(b) If (G, ℓ) and (G', ℓ') are combinatorial models for the same metric graph Γ, then $\gamma^{\text{ind}}(G) = \gamma^{\text{ind}}(G')$.

Proof. (a) The rank function of any matroid satisfies $\text{rk}(A) \le \#A$. The claim follows from comparing definitions (5.5) and (5.6).

(b) The independent girth does not change under subdivision of edges, and any two combinatorial models of Γ have a common refinement by edge subdivisions. $\square$

Proposition 5.16(b) implies that γ^{ind} is a well-defined invariant for a metric graph; given a metric graph Γ we have

$$\gamma^{\text{ind}}(\Gamma) := \gamma^{\text{ind}}(G) \quad \text{for any choice of model } \Gamma = (G, \ell). \tag{5.7}$$

Note that γ^{ind} is also invariant under stabilization.

Example 5.17. Consider Figure 5.4. The graph on the left has seven simple cycles; their lengths are $\{4, 4, 4, 6, 6, 6, 6\}$, and their ranks in the cographic matroid are all 3. For this graph, $\gamma = 4$ and $\gamma^{\text{ind}} = 3$. After deleting a central edge, the resulting graph on the right has three simple cycles with lengths $\{4, 6, 6\}$ and cographic rank 2; hence $\gamma = 4$ and $\gamma^{\text{ind}} = 2$.

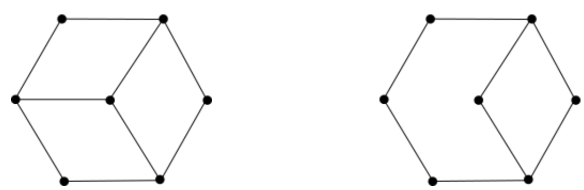

Figure 5.4: Graphs with independent girth 3, resp. independent girth 2.

Example 5.18. Consider Figure 5.5. This graph has $\gamma = 4$ and $\gamma^{\text{ind}} = 3$, with the minimum achieved on the 4-cycle in the middle. After deleting one of the horizontal edges in the middle cycle, the resulting graph has $\gamma = 4$ and $\gamma^{\text{ind}} = 4$.

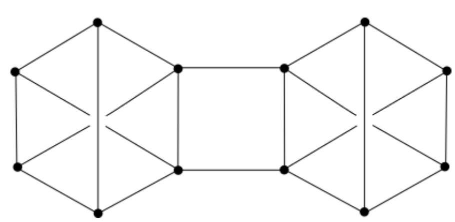

Figure 5.5: Graph with girth 4 and independent girth 3.

In general, under edge deletion we have $\gamma(G \backslash e) \ge \gamma(G)$ since $\mathcal{C}(G \backslash e) \subset \mathcal{C}(G)$. The examples above demonstrate that $\gamma^{\text{ind}}(G \backslash e)$ can increase or can decrease, relative to $\gamma^{\text{ind}}(G)$.

Theorem 5.19. *Let G be a stable graph of genus $g \geq 2$ and girth γ. Then*

$$\gamma < C \log g$$

for some constant C.

Proof. Recall that the girth γ of a graph G is the minimal length of a (simple) cycle in G. Let v be a vertex in $V(G)$. Let $N_r(v)$ denote the neighborhood of radius r around v, in the graph G. For any radius $r < \frac{1}{2}\gamma$, the neighborhood $N_r(v)$ is a tree (i.e. $N_r(v)$ is connected and acyclic).

Recall that G is stable if every vertex has valence ≥ 3. Since G is stable, we may calculate a simple lower bound for the number of edges in $N_r(v)$. Namely,

$$\#E(N_r(v)) \geq 3 + 6 + \cdots + 3 \cdot 2^{r-1} = 3(2^r - 1).$$

This quantity is clearly a lower bound for the total number of edges $\#E(G)$. Moreover, by Proposition 5.14 we have $\#E(G) \leq 3g - 3$. Thus

$$3(2^r - 1) \leq \#E(G) \leq 3g - 3 \qquad \Rightarrow \qquad 2^r \leq g$$

for any integer $r < \frac{1}{2}\gamma$. Hence

$$2^{\gamma/2-1} < g \qquad \Leftrightarrow \qquad \gamma < 2\log_2 g + 2.$$

By the assumption $g \geq 2$, this bound implies $\gamma < 4 \log_2 g$, as desired. □

Corollary 5.20. *Let Γ be a metric graph of genus g and independent girth γ^{ind}. Then $\gamma^{\text{ind}} < C \log g$ for some constant C.*

Proof. Combine Theorem 5.19 with Proposition 5.16(a) and (5.7). □

5.4 Failure of Manin–Mumford condition

In this section, we consider cases when a metric graph fails to satisfy the Manin–Mumford condition, in degree one and in higher degree.

Proposition 5.21. *If $\Gamma = (G, \ell)$ is a metric graph whose edge lengths are all rational, then the Manin–Mumford condition fails to hold.*

Proof. Rescaling all edge lengths of Γ by the same factor does not change the validity of the Manin–Mumford condition, so we may assume that all edge lengths are integers. This means Γ has a combinatorial model $(G, \mathbf{1})$ with unit edge lengths. On a graph with unit edge lengths, the degree-0 divisor classes supported on vertices form a finite abelian group, known as the *critical group* of the graph,. This implies that all vertices of G lie in the same torsion packet.

Now consider taking the k-th subdivision graph $G^{(k)}$ of G, meaning every edge if G is subdivided into k edges of equal length; the number of vertices is

$$\#V(G^{(k)}) = \#V(G) + (k-1)\#E(G).$$

The same reasoning implies that these new vertices are also in the same torsion packet of Γ. Taking $k \to \infty$ shows that Γ has an infinite torsion packet. $\square$

Proposition 5.21 can also be proved using part (a) of the following lemma. Recall that given edges $e_i \in E(G)$, $\mathrm{Eff}(e_1, \ldots, e_k)$ denotes the set of effective divisor classes $[x_1 + \cdots + x_k]$ which sum a point $x_i \in e_i$ from each edge (x_i is allowed to be an endpoint of e_i).

Lemma 5.22. *Let $\Gamma = (G, \ell)$ be a metric graph.*

(a) If an edge $e \in E(G)$ contains two points x, y such that $[x], [y]$ are distinct but in the same torsion packet, then the torsion packet $\{[x]\}_{\mathrm{tors}}$ is infinite.

(b) If $\mathrm{Eff}(e_1, \ldots, e_d)$ contains distinct divisor classes $[D], [E]$ in the same degree d torsion packet, then the torsion packet $\{[D]\}_{\mathrm{tors}}$ is infinite.

Proof. (a) Suppose that an edge e contains distinct points x, y such that $[x - y]$ is torsion. Let z denote the midpoint of x and y; we claim $[x - z]$ is also torsion. The midpoint satisfies $[2z] = [x + y]$, hence $2[x - z] = [x - z] + [z - y] = [x - y]$. If n is a positive integer such that $n[x - y] = 0$, then $2n[x - z] = n[x - y] = 0$. This proves the claim that $[x - z]$ is torsion. By repeating this argument on the midpoint of x and z, we obtain infinitely many points on e in the same torsion packet $\{[x]\}_{\mathrm{tors}}$.

(b) Since the cell $\mathrm{Eff}(e_1, \ldots, e_d)$ is convex, it contains a line segment connecting $[D]$ and $[E]$; this segment is nontrivial by the assumption $[D] \neq [E]$. Moreover, for

$$[F] = (\text{any rational affine combination of } [D] \text{ and } [E] \text{ along this line}),$$

the class $[D - F]$ is torsion. This guarantees infinitely many divisor classes $[F]$ in the torsion packet $\{[D]\}_{\mathrm{tors}}$, as claimed. $\square$

Proposition 5.23. *Suppose G has a simple cycle with d edges. Then for any edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ fails to satisfy the degree d Manin–Mumford condition.*

Proof. Let C be a simple cycle in G with edges $e_1, e_2, \ldots, e_d$ and vertices $v_1, v_2, \ldots, v_d$ in cyclic order, where edge e_i has endpoints v_i and v_{i+1} (indices taken modulo d). Consider the effective divisors $D = v_1 + \cdots + v_d$ and $E = x_1 + \cdots + x_d$ where x_i is the midpoint on edge e_i.

To show that $[D-E]$ is torsion, we construct a piecewise linear function f with $\Delta(f) = D - E$. Let $f : \Gamma \to \mathbb{R}$ be zero-valued outside of the cycle C, and $f(v_i) = 0$ for each vertex (potentially required by continuity of f). On each edge e_i, let f have slope $\frac{1}{2}$ in the directions away from v_i, so that at the midpoint $f(x_i) = \frac{1}{2}\ell(e_i)$. It is straightforward to verify that $\Delta(f) = D - E$ as desired.

By Lemma 5.7, the slopes $\pm\frac{1}{2}$ of f imply that $[D-E]$ is a nonzero, torsion divisor class. Moreover, both $[D]$ and $[E]$ lie in the same cell $\mathrm{Eff}(e_1, \ldots, e_d)$. Then Lemma 5.22(b) implies that the torsion packet $\{[D]\}_{\mathrm{tors}}$ is infinite, which violates the degree d Manin–Mumford condition. $\square$

5.5 Uniform Manin–Mumford bounds

In this section, we show that a metric graph which is Manin–Mumford finite satisfies a bound on $\#(AJ(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}})$ which depends only on the genus of Γ.

Theorem 5.24. *Suppose Γ is a metric graph of genus $g \geq 2$. If $AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}$ is finite, then*

$$\#(AJ_q(\Gamma) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}) \leq 3g - 3.$$

Proof. The retract map $r : \Gamma \to \Gamma'$ from a metric graph to its stabilization induces an isomorphism on Jacobians $\mathrm{Jac}(\Gamma) \xrightarrow{\sim} \mathrm{Jac}(\Gamma')$ and on $AJ_q(\Gamma) \xrightarrow{\sim} AJ_{r(q)}(\Gamma')$, so we may assume that Γ is semistable and that (G, ℓ) is a stable combinatorial model for Γ. Proposition 5.14 states that $\#E(G) \leq 3g - 3$ since G is stable. Lemma 5.22(a) implies that a finite torsion packet has at most one point on a given edge of G. This proves that the size of a finite, degree 1 torsion packet is at most $3g - 3$. By Proposition 5.5, we are done. $\square$

We next generalize the above argument to the higher-degree case.

Theorem 5.25. *Let $\Gamma = (G, \ell)$ be a connected metric graph of genus $g \geq 2$. If Γ satisfies the Manin–Mumford condition in degree d, then*

$$\#(AJ_D^{(d)}(\Gamma^d) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}}) \leq \binom{3g-3}{d}.$$

Proof. The number $\#(AJ_D^{(d)}(\Gamma^d) \cap \mathrm{Jac}(\Gamma)_{\mathrm{tors}})$ does not change under replacing Γ with its stabilization, so we may assume Γ is semistable and (G, ℓ) is a stable model. This means that the number of edges $\#E(G)$ is bounded above by $3g - 3$.

The image of $AJ_D^{(d)}(\Gamma^d)$ is homeomorphic to $\mathrm{Eff}^d(\Gamma)$. (They differ by a translation sending $\mathrm{Pic}^d(\Gamma)$ to $\mathrm{Pic}^0(\Gamma)$.) The maximal cells in the ABKS decomposition of $\mathrm{Eff}^d(\Gamma)$ are indexed by independent sets of size d in the cographic matroid $M^\perp(G)$, c.f. Corollary 2.20. The number of maximal cells is clearly bounded above by $\binom{\#E(G)}{d}$, the number of all size-d subsets of edges. Since we assumed G is stable, we have $\binom{\#E(G)}{d} \leq \binom{3g-3}{d}$.

From Lemma 5.22(b), we know that a finite degree d torsion packet contains at most one element from a given maximal cell of $AJ_D^{(d)}(\Gamma^d)$, which finishes the proof. $\square$

5.6 Manin–Mumford for generic edge lengths, degree one

In this section we prove our first main theorem, which gives conditions on when a metric graph satisfies the Manin–Mumford condition in degree 1. In this section, "torsion packet" will always mean a degree 1 torsion packet (c.f. Definition 5.4). Before addressing the general case, we demonstrate an example in small genus.

Example 5.26. Let G be the theta graph (see Figure 3.4) with vertices x, y and edges e_1, e_2, e_3, and consider the metric graph $\Gamma = (G, \ell)$ with edge lengths $a = \ell(e_1), b = \ell(e_2), c = \ell(e_3)$.

If a torsion packet contains two points on e_1, then Proposition 5.27 implies that $[x-y]$ is torsion on the deleted subgraph $\Gamma_1 = \Gamma \backslash e_1$. By Lemma 5.6, this would imply the voltage function which sends current from x to y on the subgraph Γ_1 has rational slopes. We can compute these slopes directly: Γ_1 is a parallel combination of wires with resistances b and c, so the slope along e_2 is $\frac{c}{b+c}$. (This calculation also follows from Theorem 3.26.) To summarize:

$$\text{(some torsion packet contains} \geq \text{two points of } e_1) \qquad \Rightarrow \qquad \frac{c}{b+c} \in \mathbb{Q}.$$

The contrapositive statement is that

$$\frac{c}{b+c} \notin \mathbb{Q}. \qquad \Rightarrow \qquad \text{(every torsion packet contains at most one point of } e_1).$$

To satisfy the Manin–Mumford condition, it suffices that every torsion packet $\{[x]\}_{\text{tors}} \subset \text{Eff}^1(\Gamma)$ contains at one point of each edge e_1, e_2, e_3. Thus the Manin–Mumford condition holds for Γ if the edge lengths are in set

$$\{(a,b,c) \in \mathbb{R}^3_{>0} : \frac{b}{a+b} \notin \mathbb{Q} \text{ and } \frac{c}{a+c} \notin \mathbb{Q} \text{ and } \frac{c}{b+c} \notin \mathbb{Q}\}.$$

This is very general subset of $\mathbb{R}^3_{>0}$, c.f. Example 5.8(b).

Proposition 5.27. *Suppose Γ is a metric graph and points $x, y \in \Gamma$ lie on the same edge. Let Γ_0 denote the metric graph with the open segment between x and y removed. If $[x-y]$ is torsion on Γ and $[x-y] \neq 0$, then $[x-y]$ is torsion on Γ_0.*

Proof. Suppose $[x-y]$ is torsion on Γ. Let j_x^y denote the voltage function on Γ when one unit of current is sent from y to x. By Lemma 5.6, all slopes of j_x^y are rational.

In particular, the slope of j_x^y on the segment between x and y is rational; let s denote this slope. Since $[x-y] \neq 0$, we have $s < 1$.

Let Γ_0 denote the metric graph obtained from Γ by deleting the interior of edge e. It is clear that the restriction of j_x^y to Γ_0 has Laplacian $\Delta(j_x^y\big|_{\Gamma_0}) = (1-s)x-(1-s)y$.

Let $j_{x,0}^y$ denote the voltage function on Γ_0 when one unit of current is sent from y to x. Since $j_{x,0}^y = (1-s)^{-1} j_x^y$, all slopes of $j_{x,0}^y$ are rational. By Lemma 5.6, this implies $[x-y]$ is torsion on Γ_0 as desired. □

Proposition 5.28. *Suppose x, y are two vertices on a graph G. Let j_x^y be the voltage function on $\Gamma = (G, \ell)$, depending on variable edge lengths $\ell : E(G) \to \mathbb{R}$. Either:*

(1) all slopes of j_x^y are 1 or 0, independent of edge lengths; or

(2) there exists some edge e such that the slope of j_x^y along e is a non-constant rational function of the edge lengths.

Proof. Suppose there is a unique simple path in G from x to y. Then the slope of j_x^y is 1 along this path, and 0 away from this path, since all current flowing from y to x must follow this path. Thus we are in case (1).

On the other hand, suppose there are two distinct simple paths π_1, π_2 in G from x to y. Let e be an edge of G which lies on π_1 but not π_2. If we fix the lengths of edges in π_1 and send all other edge lengths to infinity, then the slope of j_x^y along e approaches 1. If we send the length $\ell(e)$ to infinity while keeping all other edge lengths fixed, then the slope of j_x^y along e approaches zero. Thus the slope of j_x^y along e is a non-constant function of the edge lengths. By Kirchhoff's formulas, Theorem 3.26, the slope (i.e. current) is a rational polynomial function of the edge lengths. This is case (2). □

Proposition 5.29. *Suppose x, y are two vertices on a graph G. Then for the metric graph $\Gamma = (G, \ell)$, either*

(1) $[x-y] = 0$ in $\mathrm{Jac}(\Gamma)$ for any edge lengths ℓ, or

(2) $[x-y]$ is non-torsion in $\mathrm{Jac}(\Gamma)$ for very general edge lengths ℓ.

Proof. If none of the slopes of j_x^y vary as a function of edge lengths, then by Proposition 5.28 all slopes of j_x^y are zero or one. This implies that $[x-y] = 0$.

On the other hand, suppose for some edge e the slope of j_x^y along e is a non-constant rational function $\frac{p(\ell_1,\ldots,\ell_m)}{q(\ell_1,\ldots,\ell_m)}$. Then the subset

$$U = \left\{ (\ell_1, \ldots, \ell_m) \in \mathbb{R}^m_{>0} : \frac{p(\ell_1, \ldots, \ell_m)}{q(\ell_1, \ldots, \ell_m)} \notin \mathbb{Q} \right\}$$

parametrizing edge-lengths where the slope at e take irrational values is very general, c.f. Example 5.8(b). By Lemma 5.6, $[x-y]$ is nontorsion on U, as desired. □

Theorem 5.30. *Suppose G is a biconnected metric graph of genus $g \geq 2$. For a very general choice of edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ satisfies the Manin–Mumford condition.*

Proof. Let $m = \#E(G)$ and choose an ordering $E(G) = \{e_1, e_2, \ldots, e_m\}$, which induces a homeomorphism from the space of edge-lengths $\{\ell : E(G) \to \mathbb{R}_{>0}\}$ to the positive orthant $\mathbb{R}^m_{>0}$. We claim that for each edge e_i, there is a corresponding very general subset $U_i \subset \mathbb{R}^m_{>0}$ such that

$$\text{when edge lengths are chosen in } U_i \text{, every torsion packet of } \Gamma = (G, \ell) \text{ contains at most one point of } e_i. \tag{5.8}$$

Let e_i^+, e_i^- denote the endpoints of e_i, and let $G_i = G \backslash e_i$ denote the graph with edge e_i deleted. If the endpoints e_i^+, e_i^- are not connected by any path in G_i, this contradicts our assumption that G is biconnected. If the endpoints are connected by only one path π in G_i, then the union $\pi \cup \{e_i\}$ is a genus 1 biconnected component of G, which contradicts our assumption that G is biconnected and has genus $g \geq 2$. Thus e_i^+, e_i^- are connected by at least two distinct paths in G_i.

Therefore, the divisor class $[e_i^+ - e_i^-] \neq 0$ in $\mathrm{Jac}(\Gamma_i)$ where $\Gamma_i = (G_i, \ell_i)$. By Proposition 5.29, $[e_i^+ - e_i^-]$ is nontorsion in $\mathrm{Jac}(\Gamma_i)$ on a very general subset $V_i \subset \mathbb{R}^{m-1}_{>0}$ of edge-length space. (Note that G_i has $m-1$ edges.) Finally, we let U_i be the preimage of V_i under the coordinate projection $\mathbb{R}^m_{>0} \to \mathbb{R}^{m-1}_{>0}$ forgetting coordinate i. The subset U_i is very general, and satisfies the claimed condition (5.8).

For any edge lengths in the intersection $U = \bigcap_{i=1}^m U_i$ a torsion packet of the corresponding $\Gamma = (G, \ell)$ can have at most one point on each edge e_i, giving the bound $\#\{[x]\}_{\mathrm{tors}} \leq m$. The subset U is very general, since it is a finite intersection of very general subsets. This completes the proof. $\square$

5.7 Manin–Mumford for generic edge lengths, higher degree

In this section we address when a metric graph with very general edge lengths satisfies the Manin–Mumford condition in higher degree.

The next proposition is a strengthening of Proposition 5.23. Recall that $M^\perp(G)$ denotes the cographic matroid of G

Proposition 5.31. *Suppose G contains a cycle C which has rank $d = \mathrm{rk}^\perp(E(C))$ in the cographic matroid $M^\perp(G)$. Then for any edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$, the metric graph $\Gamma = (G, \ell)$ fails the degree d Manin–Mumford condition.*

Proof. Suppose the given cycle of G consists of the edges $\{e_1, \ldots, e_k\}$ and vertices $\{v_1, \ldots, v_k\}$ in cyclic order; note that $k \geq d$. Let $D = v_1 + \cdots + v_k$ be the sum of the cycle's vertices. In the proof of Proposition 5.23, we showed that the degree-k torsion packet $\{[D]\}_{\mathrm{tors}}$ has infinite intersection with the cell $\mathrm{Eff}(e_1, \ldots, e_k)$, for any choice of edge lengths ℓ.

Recall that $\mathrm{Eff}(e_1, \ldots, e_k)$ is the image of $\mathrm{Div}(e_1, \ldots, e_k)$ under the linear equivalence map $\mathrm{Div}^k(\Gamma) \to \mathrm{Pic}^k(\Gamma)$. The map $\mathrm{Div}(e_1, \ldots, e_k) \to \mathrm{Pic}^k(\Gamma)$ lifts to a linear map ϕ in the diagram

$$\begin{tikzcd} \prod_{i=1}^k [0, \ell(e_i)] \arrow[r, "\phi"] \arrow[d] & \mathbb{R}^g \arrow[d] \\ \mathrm{Div}(e_1, \ldots, e_k) \arrow[r] & \mathrm{Pic}^k(\Gamma), \end{tikzcd}$$

where $\prod_{i=1}^k [0, \ell(e_i)] \to \mathrm{Div}(e_1, \ldots, e_k)$ is the product of isometries $[0, \ell(e_i)] \to e_i$ and $\mathbb{R}^g \to \mathrm{Pic}^k(\Gamma)$ is an isometric universal cover. By Theorem 2.19, $\mathrm{Eff}(e_1, \ldots, e_k)$ has dimension $d = \mathrm{rk}^\perp(\{e_1, \ldots, e_k\})$ (where $d \leq k$). This implies that ϕ has rank d, so the image of ϕ is covered by the restrictions of ϕ to the d-faces of $\prod_{i=1}^k [0, \ell(e_i)]$.

Thus $\mathrm{Eff}(e_1, \ldots, e_k)$ is covered by the corresponding images of the d-faces of $\mathrm{Div}(e_1, \ldots, e_k)$, which have the form

$$\mathrm{Eff}(e_i : i \in I) + [\sum_{i \notin I} v_i^\pm] \quad \subset \quad \mathrm{Eff}(e_1, \ldots, e_k), \tag{5.9}$$

where I is a size-d subset of $\{1, \ldots, k\}$ and $v_i^\pm \in \{v_i, v_{i+1}\}$ is an endpoint of e_i. (There are $\binom{k}{d} 2^{k-d}$ such choices.)

Since $\mathrm{Eff}(e_1, \ldots, e_k)$ has infinite intersection with the torsion packet $\{[D]\}_{\mathrm{tors}}$, there is some choice of $I, v_i^\pm$ such that the subset (5.9) of $\mathrm{Eff}(e_1, \ldots, e_k)$ has infinite intersection with $\{[D]\}_{\mathrm{tors}}$. This implies that the degree-d torsion packet $\{[D - \sum_{i \notin I} v_i^\pm]\}_{\mathrm{tors}}$ has infinite intersection with $\mathrm{Eff}(e_i : i \in I)$, thus violating the degree d Manin–Mumford condition. $\square$

Next, we consider the converse situation of Proposition 5.31, i.e. when an edge set is acyclic after taking the closure in $M^\perp(G)$. Recall from Section 2.7 the notation $\mathrm{Div}(e_1, \ldots, e_k)$ and $\mathrm{Eff}(e_1, \ldots, e_k)$. Here we introduce a slight variation: let $\mathrm{Div}(e_1, \ldots, e_k)^\circ$ denote the set of effective divisors of the form $D = x_1 + \cdots + x_k$ where x_i is in the *interior* e_i° of edge e_i; respectively let $\mathrm{Eff}(e_1, \ldots, e_k)^\circ$ denote the divisor classes of the form $[x_1 + \cdots + x_k]$, where $x_i \in e_i^\circ$.

Proposition 5.32. *Suppose $e_1, \ldots, e_k$ are edges in G such that $\{e_1, \ldots, e_k\}$ is independent in $M^\perp(G)$ and the closure of $\{e_1, \ldots, e_k\}$ in $M^\perp(G)$ spans an acyclic subgraph of G. Then for very general edge lengths on $\Gamma = (G, \ell)$, distinct divisor classes in $\mathrm{Eff}(e_1, \ldots, e_k)^\circ \subset \mathrm{Pic}^k(\Gamma)$ are in distinct torsion packets.*

Before proving this statement, we introduce some lemmas and definitions.

Definition 5.33. Given a piecewise linear function f on Γ, say an edge of G is *current-active* with respect to f if the slope f' is nonzero in a neighborhood of the endpoints[3]; let $E^{\text{c.a.}}(G, f)$ denote the current-active edges,

$$E^{\text{c.a.}}(G, f) = \{e \in E(G) : f' \neq 0 \text{ in a neighborhood of } e^+, e^- \text{ in } e\}.$$

Say an edge is *voltage-active* with respect to f if the net change in f across e is nonzero; let $E^{\text{v.a.}}(G, f)$ denote the voltage-active edges,

$$E^{\text{v.a.}}(G, f) = \{e \in E(G) : f(e^+) - f(e^-) \neq 0 \quad \text{where } e = (e^+, e^-)\}.$$

Recall that a *cut* of G is a set of edges $\{e_1, \ldots, e_k\}$ such that the deletion $G \setminus \{e_1, \ldots, e_k\}$ is disconnected.

Lemma 5.34. *Consider a metric graph $\Gamma = (G, \ell)$ and $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$. If $E^{\text{v.a.}}(G, f)$ is nonempty, it contains a cut of G.*

Proof. Suppose $e = (e^+, e^-)$ is voltage-active with respect to f, so that $f(e^+) > f(e^-)$ for some ordering of endpoints. Then we may partition $V(G)$ into two nonempty sets $V^+ \cup V^-$, where

$$V^+ = \{v \in V(G) : f(v) \geq f(e^+)\} \qquad \text{and} \qquad V^- = \{v \in V(G) : f(v) < f(e^+)\}.$$

It is clear that $E^{\text{v.a.}}(G, f)$ contains all edges between V^+ and V^-; such edges form a cut of G. □

Lemma 5.35. *On $\Gamma = (G, \ell)$, consider $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ such that $\Delta(f) = E - D$ for $D, E \in \mathrm{Div}(e_1, \ldots, e_k)^\circ$. If $E^{\text{c.a.}}(G, f)$ is nonempty, then it contains a cycle of G.*

Proof. Suppose $D = x_1 + \cdots + x_k$ and $E = y_1 + \cdots + y_k$ where $x_i, y_i \in e_i$. Since the divisor $\Delta(f)$ restricted to e_i has the form $y_i - x_i$, the slopes of f along e_i are as shown in Figure 5.6, where slopes are indicated in the rightward direction.

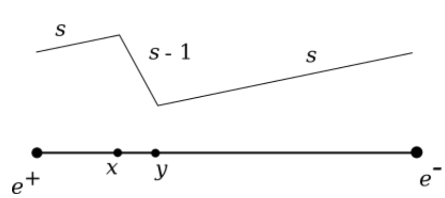

Figure 5.6: Slopes on edge e where $\Delta(f) = y - x$.

Edge e_i is current-active iff the corresponding slope s ($= s_i$) is nonzero. In particular, if $e_i \in E^{\text{c.a}}(G, f)$ it is current-active *at both endpoints*.

[3]if $e \cong [0, 1]$, here a "neighborhood of the endpoints" means $[0, \epsilon) \cup (1 - \epsilon, 1]$ for some $\epsilon > 0$

On the other hand, consider an edge $e \in E(G) \setminus \{e_1, \dots, e_k\}$. Then $\Delta(f)$ is not supported on e, so f does not change slope on e. Again in this case, if $e \in E^{\text{c.a.}}(G, f)$ then it is current-active at both endpoints.

By assumption that divisors $D, E \in \mathrm{Div}(e_1, \dots, e_k)^\circ$, $\Delta(f)$ is supported away from the vertex set $V(G)$. This means that around a vertex v, the outward slopes of f sum to zero. The number of nonzero terms in the sum must be 0 or ≥ 2, and each nonzero term corresponds to a current-active edge incident to v. Thus

$$\begin{gathered} E^{\text{c.a.}}(G, f) \text{ spans a subgraph of } G \text{ where} \\ \text{every vertex has } \operatorname{val}(v) = 0 \text{ or } \operatorname{val}(v) \geq 2. \end{gathered}$$

The claim follows. □

Lemma 5.36. *Consider $D, E \in \mathrm{Div}(e_1, \dots, e_k)^\circ$ and $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ such that $\Delta(f) = E - D$. If $D \neq E$, then $E^{\text{v.a.}}(G, f)$ or $E^{\text{c.a.}}(G, f)$ is nonempty (or both are).*

Proof. If $D = x_1 + \cdots + x_k$ is not equal to $E = y_1 + \cdots + y_k$, then there is some index i such that $x_i \neq y_i$. Consider the illustration of f in Figure 5.6, applied to the edge with $x_i \neq y_i$. We have

$$f(e_i^-) - f(e_i^+) = s \cdot \ell(e_i) - \ell([x_i, y_i]), \tag{5.10}$$

where $\ell([x_i, y_i])$ is the distance between x_i and y_i on e_i. If $s = 0$, then e_i is not current-active but is voltage-active. If $s = \ell([x_i, y_i])/\ell(e_i)$, then e_i is not voltage-active but is current-active. □

Lemma 5.37. *Consider a fixed vertex-supported $\mathbb{R}$-divisor $D = \lambda_1 v_1 + \cdots + \lambda_r v_r$ of degree zero on G, so $v_i \in V(G)$, $\lambda_i \in \mathbb{R}$ and $\sum \lambda_i = 0$. On $\Gamma = (G, \ell)$, suppose $f \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ satisfies $\Delta(f) = D$ and f has nonzero slope on $e \in E(G)$. If e is not a bridge, then the slope on e is a nonconstant rational function of edge lengths of Γ.*

Proof. Suppose we let $\ell(e) \to \infty$ and fix the lengths of all edges $e' \neq e$; we claim that the slope of f across e approaches zero.

The slope-current principle, Proposition 3.6, states that the slope of f is bounded above in magnitude by Λ, where $\Lambda = \frac{1}{2} \sum_i |\lambda_i|$ does not depend on the edge lengths.[4] Since $e = (e^+, e^-)$ is not a bridge edge, there is a simple path π from e^+ to e^- which does not contain e. By integration along π, $|f(e^-) - f(e^+)|$ is bounded above by $\Lambda \cdot \ell(\pi)$, which implies the bound

$$|f'(e)| = \left| \frac{f(e^-) - f(e^+)}{\ell(e)} \right| \leq \frac{\Lambda \cdot \ell(\pi)}{\ell(e)}.$$

[4] Since $\sum \lambda_i = 0$, we have $\Lambda = \sum\{\lambda_i : \lambda_i > 0\} = -\sum\{\lambda_i : \lambda_i < 0\}$.

If we let $\ell(e) \to \infty$ and keep $\ell(e')$ constant for each $e' \in E(G) \setminus \{e\}$, this upper bound approaches zero as claimed.

Thus the slope of f along e is a non-constant function of the edge lengths. It is a rational function by Kirchhoff's formulas, Theorem 3.26. $\square$

Proof of Proposition 5.32. Suppose $D = x_1 + \cdots + x_k$ and $E = y_1 + \cdots + y_k$ are divisors in $\mathrm{Div}(e_1, \ldots, e_k)^\circ$. Let f be a piecewise linear function such that $\Delta(f) = E - D$. By Lemma 5.7, $[D]$ and $[E]$ lie in the same torsion packet if and only if all slopes of f are rational.

Let Γ_0 (resp. G_0) denote the metric graph (resp. combinatorial graph) obtained from deleting the interiors of edges $e_1, \ldots, e_k$ from Γ (resp. G). Let $f_0 = f\big|_{\Gamma_0}$ denote the restriction of f to Γ_0. We have

$$\Delta(f_0) = \lambda_1 w_1 + \cdots + \lambda_r w_r, \tag{5.11}$$

where $\{w_1, \ldots, w_r\} \subset V(G)$ is the set of endpoints of edges $e_1, \ldots, e_k$ and $\lambda_i \in \mathbb{R}$.

First, suppose the tuple $(\lambda_1, \ldots, \lambda_r) = (0, \ldots, 0)$. Then f_0 is constant, so every edge of G_0 is neither current-active nor voltage-active with respect to f. Since the edges $\{e_1, \ldots, e_k\}$ are assumed independent in $M^\perp(G)$, they do not contain a cut of G so the inclusion $E^{\text{v.a.}}(G, f) \subset \{e_1, \ldots, e_k\}$ implies that $E^{\text{v.a.}}(G, f) = \varnothing$ by Lemma 5.34. Since the edges $\{e_1, \ldots, e_k\}$ do not contain a cycle of G, the inclusion $E^{\text{c.a.}}(G, f) \subset \{e_1, \ldots, e_k\}$ implies that $E^{\text{c.a.}}(G, f) = \varnothing$ by Lemma 5.35. Then Lemma 5.36 implies that $D = E$.

Next, suppose the tuple $(\lambda_1, \ldots, \lambda_r) \in \mathbb{R}^r$ from (5.11) is nonzero. This means that some edge of G_0 must be current-active, so $E^{\text{c.a.}}(G, f)$ is nonempty. By Lemma 5.35, $E^{\text{c.a.}}(G, f)$ contains a cycle of G. The closure of $\{e_1, \ldots, e_k\}$ with respect to the cographic matroid $M^\perp(G)$ is equal to

$$\{e_1, \ldots, e_k\} \cup \{b_1, \ldots, b_j\} \quad \text{where } \{b_1, \ldots, b_j\} \text{ are the bridge edges of } G_0.$$

By assumption that $\{e_1, \ldots, e_k\} \cup \{b_1, \ldots, b_j\}$ is acyclic, $E^{\text{c.a.}}(G, f)$ must contain an edge $e_* \notin \{e_1, \ldots, e_k\}$ which is not a bridge in G_0.[5]

Now consider applying Lemma 5.37 to the graph G_0, the divisor (5.11), and the edge $e_* \in E(G_0)$. The lemma concludes that as a function of the edge-lengths of Γ_0, the slope of f_0 (equivalently f) on e_* is a nonconstant ratio of polynomials. In particular,

$$V(\lambda_1, \ldots, \lambda_r) = \{\text{edge lengths of } \Gamma_0 \text{ such that } f_0' \text{ is irrational on } e_*\} \tag{5.12}$$

is a very general subset of $\mathbb{R}^{m-k}_{>0} \cong \{\ell_0 : E(G_0) \to \mathbb{R}_{>0}\}$, and on this subset we have $[D]$ and $[E]$ are in distinct torsion packets.

[5] the edge e_* depends on the tuple $(\lambda_1, \ldots, \lambda_r)$

Finally, let $U(\lambda_1, \ldots, \lambda_r)$ be the preimage of $V(\lambda_1, \ldots, \lambda_r)$ under the projection $\mathbb{R}^m_{>0} \to \mathbb{R}^{m-k}_{>0}$, which is very general, and let

$$U = \bigcap_{\substack{(\lambda_1,\ldots,\lambda_r)\\ \in \mathbb{Q}^r \setminus (0,\ldots,0)}} U(\lambda_1, \ldots, \lambda_r) \quad \subset \quad \mathbb{R}^m_{>0}.$$

The subset U is very general, as a countable intersection of very general subsets.

If edge lengths of $\Gamma = (G, \ell)$ are chosen such that there are distinct divisors $D, E \in \mathrm{Eff}(e_1, \ldots, e_k)^\circ$ where $[D]$ and $[E]$ are in the same torsion packet, then the tuple $(\lambda_1, \ldots, \lambda_r)$ as in (5.11) must be rational and nonzero. Then the chosen edge lengths on $G_0 \subset G$ are excluded from the subset (5.12), hence the edge lengths are excluded also from U, as desired. □

Theorem 5.38. *Let G be a connected graph of genus $g \geq 1$ and independent girth γ^{ind}. The metric graph $\Gamma = (G, \ell)$ satisfies the degree d Manin–Mumford condition for very general edge lengths $\ell : E(G) \to \mathbb{R}_{>0}$ if and only if $1 \leq d < \gamma^{\mathrm{ind}}$.*

Proof. If $d \geq \gamma^{\mathrm{ind}}$, then $d \geq \mathrm{rk}^\perp(E(C))$ for some cycle C of G. Proposition 5.31 states that Γ fails the Manin–Mumford condition in degree $d' = \mathrm{rk}^\perp(E(C))$, so the condition also fails in degree $d \geq d'$.

Conversely if $d < \gamma^{\mathrm{ind}}$, then for each d-subset of edges $\{e_1, \ldots, e_d\}$, its closure in $M^\perp(G)$ does not contain a cycle of G. In particular, the edges for each maximal cell $\mathrm{Eff}(e_1, \ldots, e_d)$ of $\mathrm{Eff}^d(\Gamma)$ satisfy the hypotheses of Proposition 5.32, so there is a very general subset of edge lengths of Γ for which every degree d torsion packet has at most one element in the chosen cell $\mathrm{Eff}(e_1, \ldots, e_d)$. Since there are finitely many maximal cells (cf. Corollary 2.20), this implies that for very general edge lengths there are finitely many elements in each degree d torsion packet. □

Corollary 5.39. *Let Γ be a metric graph of genus $g \geq 1$, and suppose Γ satisfies the Manin–Mumford condition in degree d. Then*

$$d < C \log g$$

for some constant C.

Proof. This follows from Proposition 5.31, which implies that $d < \gamma^{\mathrm{ind}}$, and the bound $\gamma^{\mathrm{ind}} < C \log g$ from Corollary 5.20. □

APPENDIX

APPENDIX A

Theta Intersection

In this appendix we give an alternate description of the Weierstrass locus $W(D)$ as the intersection of two polyhedral subcomplexes of complementary dimension in $\mathrm{Pic}^{g-1}(\Gamma)$. This allows us to give an alternate proof that $W(D)$ is finite for a generic divisor class $[D]$. In this perspective, the stable Weierstrass locus $W^{\mathrm{st}}(D)$ naturally appears as the stable tropical intersection of these two subsets.

Throughout this section (including the above paragraph), we assume that the divisor class $[D]$ is *(Riemann–Roch) nonspecial*, meaning that its rank satisfies

$$r(D) = \begin{cases} \deg(D) - g & \text{if } \deg(D) \geq g, \\ -1 & \text{otherwise.} \end{cases}$$

A generic divisor class in $\mathrm{Pic}^n(\Gamma)$ is nonspecial. If $n \geq 2g-1$, all divisors in $\mathrm{Pic}^n(\Gamma)$ are nonspecial.

A.1 Intersection with Θ

Recall that the *theta divisor* $\Theta \subset \mathrm{Pic}^{g-1}(\Gamma)$ is the space of degree $g-1$ divisor classes which have an effective representative;

$$\Theta = \{[D] \in \mathrm{Pic}^{g-1}(\Gamma) : [D] \geq 0\}.$$

Given a divisor D of degree $n \geq g$, let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ denote the map

$$\Phi_D : x \mapsto [D - (n-g+1)x].$$

If D has degree $n < g$ let $\Phi_D : x \mapsto [D]$ be the constant map. Note that the map Φ_D depends only on the divisor class $[D]$. The Weierstrass locus of D may be recovered from the image of Φ_D.

Proposition A.1. *Let D be a nonspecial divisor of degree $n \geq g$, and let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ be the map $\Phi_D(x) = [D - (n-g+1)x]$. Then*

$$W(D) = \Phi_D^{-1}(\Phi_D(\Gamma) \cap \Theta).$$

Proof. This follows from the definition of Weierstrass locus, if D has rank $n-g$. $\square$

Proposition A.2. *Suppose Γ is a bridgeless metric graph. If D has degree $n \geq g$, the map $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ is locally injective (i.e. an immersion).*

Proof. The map Φ_D may be expressed as a composition of three maps

$$\Phi_D : \Gamma \xrightarrow{\alpha} \mathrm{Pic}^1(\Gamma) \xrightarrow{\beta} \mathrm{Pic}^{n-g+1}(\Gamma) \xrightarrow{\gamma} \mathrm{Pic}^{g-1}(\Gamma),$$

where α sends $x \mapsto [x]$, β sends $[E] \mapsto [(n-g+1)E]$, and γ sends $[E] \mapsto [D-E]$. The map $\gamma = \gamma_D$ is a homeomorphism. The map β is a $(n-g+1)^g$-fold covering map, so it is a local homeomorphism if $n \geq g$. Thus it suffices to verify that the first map α is locally injective.

This follows from the Abel–Jacobi theorem for metric graphs, see e.g. Baker–Faber [8, Theorem 4.1 (3)(4)]. Note that $\mathrm{Pic}^1(\Gamma)$ is (non-canonically) isomorphic to the Jacobian $\mathrm{Jac}(\Gamma) = \mathrm{Pic}^0(\Gamma)$ by subtracting a basepoint x_0. $\square$

If Γ contains bridge segments, let $\Gamma_{/(\mathrm{br})}$ denote the metric graph obtained from Γ by contracting all bridges. Let $S_{(\mathrm{br})} \subset \Gamma_{/(\mathrm{br})}$ denote the set of points which were bridges in Γ.

Lemma A.3. *Let $\pi : \Gamma \to \Gamma_{/(br)}$ denote the canonical map contracting all bridge segments of Γ, which induces $\pi_* : \mathrm{Pic}^n(\Gamma) \to \mathrm{Pic}^n(\Gamma_{/(br)})$ for all n. For any divisor D on Γ,*

$$W(D) = \pi^{-1} W(\pi_*(D)).$$

Proof. On Γ the linear equivalence map $x \mapsto [x]$ factors through $\pi : \Gamma \to \Gamma_{/(\mathrm{br})}$; i.e. we have a commuting diagram

$$\begin{array}{ccc} \Gamma & \xrightarrow{\pi} & \Gamma_{/(\mathrm{br})} \\ {\scriptstyle [x]}\downarrow & & \downarrow{\scriptstyle [x]} \\ \mathrm{Pic}^1(\Gamma) & \xrightarrow{\sim} & \mathrm{Pic}^1(\Gamma_{/(\mathrm{br})}). \end{array}$$

Using this, the result is clear from the definition of $W(D)$. $\square$

Lemma A.4. *Suppose $S \subset \Gamma$ is a finite set of points in a metric graph Γ. For a generic divisor class $[D]$, the intersection $W(D) \cap S$ is empty.*

Proof. It suffices to consider when $S = \{s\}$ contains one point. Assuming D is nonspecial, which holds for generic $[D] \in \mathrm{Pic}^n(\Gamma)$, we have $s \in W(D)$ if and only if

$$[D - (n-g+1)s] \text{ is effective} \quad \Leftrightarrow \quad [D] = [(n-g+1)s + E] \text{ for some } [E] \in \Theta.$$

Since Θ has dimension $g-1$, the space $\{[D] = [(n-g+1)s + E] : [E] \in \Theta\}$ also has dimension $g-1$. Hence a generic class $[D]$ has $s \notin W(D)$. $\square$

Theorem A.5. *For a generic divisor class $[D]$ in $\mathrm{Pic}^n(\Gamma)$, the Weierstrass locus $W(D)$ is finite.*

Proof. If $n < g$, then a generic divisor class in $\mathrm{Pic}^n(\Gamma)$ is not effective because the image of $\mathrm{Sym}^n(\Gamma) \to \mathrm{Pic}^n(\Gamma)$ has dimension at most n, while $\mathrm{Pic}^n(\Gamma)$ has dimension g. For a non-effective divisor class $[D]$, the Weierstrass locus $W(D)$ is empty.

Now suppose $n \geq g$. By Riemann–Roch, a generic divisor class in $\mathrm{Pic}^n(\Gamma)$ has rank $r(D) = n - g$. (By the above paragraph, $r(K - D) = -1$ generically.) Thus, it suffices to show that $W(D)$ is finite for a generic nonspecial divisor class.

Case 1: Γ is bridgeless. As above, let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ be the map $\Phi_D(x) = [D - (n - g + 1)x]$. Recall that the Weierstrass locus $W(D)$ is equal to

$$W(D) = \Phi_D^{-1}(\Phi_D(\Gamma) \cap \Theta) \subset \Gamma$$

where $\Theta = \{[E] \in \mathrm{Pic}^{g-1}(\Gamma) : [E] \geq 0\}$ is the theta divisor. Note that as $[D]$ varies, the image $\Phi_D(\Gamma)$ varies by translation inside $\mathrm{Pic}^{g-1}(\Gamma)$.

Recall that Θ is a $(g - 1)$-dimensional polyhedral complex with finitely many facets, and $\Phi_D(\Gamma)$ is a 1-dimensional polyhedral complex with finitely many segments. This implies that the space of translations which cause $\Phi_D(\Gamma)$ to intersect Θ non-transversally has dimension at most $g - 1$. Hence for a generic divisor class $[D]$, the intersection $\Phi_D(\Gamma) \cap \Theta$ is transverse.

Suppose all intersections in $\Phi_D(\Gamma) \cap \Theta$ are transverse, and occur in the interiors of the respective segment and facet. Recall that Φ_D is locally injective by Proposition A.2. If Φ_D sends $x \in \Gamma$ to a transverse intersection, then x must have some neighborhood $U \subset \Gamma$ such that $\Phi_D(U \backslash \{x\})$ is disjoint from Θ. This means that $W(D) = \Phi_D^{-1}(\Phi_D(\Gamma) \cap \Theta)$ is a discrete subset of Γ. Because Γ is compact, this implies $W(D)$ is finite.

Case 2: Γ has bridge segments. Let $\pi : \Gamma \to \Gamma_{/(\mathrm{br})}$ denote the map contracting all bridge segments of Γ. Let $S_{(\mathrm{br})} \subset \Gamma_{/(\mathrm{br})}$ denote the image of all bridges, which is a finite subset of $\Gamma_{/(\mathrm{br})}$. Note that π restricts to an injection away from $\pi^{-1} S_{(\mathrm{br})}$.

By Lemma A.4, a generic divisor class $[D] \in \mathrm{Pic}^n(\Gamma_{/(\mathrm{br})})$ has $W(D)$ disjoint from $S_{(\mathrm{br})}$. Since π induces a homeomorphism $\pi_* : \mathrm{Pic}^n(\Gamma) \to \mathrm{Pic}^n(\Gamma_{/(\mathrm{br})})$, this implies that a generic class $[D] \in \mathrm{Pic}^n(\Gamma)$ has $W(\pi_*[D])$ disjoint from $S_{(\mathrm{br})}$. The result then follows from Lemma A.3 and Case 1. $\square$

A.2 Stable Weierstrass locus

In this section we describe the relation of the current setup, involving the theta divisor Θ, and the stable Weierstrass locus defined in Section 4.2.1.

Proposition A.6. *Suppose Γ is a bridgeless metric graph of genus g. Let D be a divisor of degree g, and let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ send $\Phi_D(x) = [D - x]$. Then the break divisor* $\mathrm{br}[D]$ *is equal to*

$$\mathrm{br}[D] = \Phi_D^{-1}(\Phi_D(\Gamma) \cap^{\mathrm{st}} \Theta)$$

where Θ is the theta divisor and $\cap^{\mathrm{st}}$ denotes stable tropical intersection.[1]

Proof. Let us denote $\mathrm{br}^*[D] := \Phi_D^{-1}(\Phi_D(\Gamma) \cap^{\mathrm{st}} \Theta)$. For a generic divisor class $[D] \in \mathrm{Pic}^g(\Gamma)$, the intersection $\Phi_D(\Gamma) \cap \Theta$ is transverse so

$$\mathrm{br}^*[D] = \{x \in \Gamma : [D - x] \geq 0\},$$

i.e. $\mathrm{br}^*[D]$ contains the support of any effective representative of $[D]$. Generically, the class $[D]$ contains a single effective representative so $\mathrm{br}^* : \mathrm{Pic}^g(\Gamma) \to \mathrm{Sym}^g(\Gamma)$ defines a generic section of the linear equivalence map $\mathrm{Sym}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma)$.

By general properties of stable tropical intersection, the map $\mathrm{br}^* : \mathrm{Pic}^g(\Gamma) \to \mathrm{Sym}^g(\Gamma)$ is continuous. But by Theorem 2.8, the break divisor map br is the unique continuous section of $\mathrm{Sym}^g(\Gamma) \to \mathrm{Pic}^g(\Gamma)$ so we must have $\mathrm{br}^*[D] = \mathrm{br}[D]$. □

Recall that for a divisor of degree g, we have $W^{\mathrm{st}}(D) = \mathrm{br}[D]$. Proposition A.6 can be generalized to the statement that

$$W^{\mathrm{st}}(D) = \Phi_D^{-1}(\Phi_D(\Gamma) \cap^{\mathrm{st}} \Theta)$$

for a divisor of degree $n \geq g$ on a bridgeless metric graph. We omit the details here.

[1] The stable tropical intersection may have multiplicities, so here we interpret the preimage to be a multiset in Γ carrying the same multiplicities.

BIBLIOGRAPHY

[1] D. Abramovich, L. Caporaso, and S. Payne. The tropicalization of the moduli space of curves. *Ann. Sci. Éc. Norm. Supér.*, 48(4):765–809, 2015.

[2] O. Amini. Reduced divisors and embeddings of tropical curves. *Trans. Amer. Math. Soc.*, 365:4851–4880, 2013.

[3] O. Amini. Equidistribution of Weierstrass points on curves over non-Archimedean fields, preprint, 2014. arXiv:1412.0926.

[4] Y. An, M. Baker, G. Kuperberg, and F. Shokrieh. Canonical representatives for divisor classes on tropical curves and the matrix tree theorem. *Forum Math. Sigma*, 2:e25, 2014.

[5] S. J. Arakelov. Intersection theory of divisors on an arithmetic surface. *Izv. Akad. Nauk*, 8:1167–1180, 1974.

[6] M. Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008.

[7] M. Baker and X. Faber. Metrized graphs, Laplacian operators, and electrical networks. In *Quantum graphs and their applications*, pages 15–33. Amer. Math. Soc., Providence, 2006.

[8] M. Baker and X. Faber. Metric properties of the tropical Abel–Jacobi map. *J. Algebr. Comb.*, 33:349–381, 2011.

[9] M. Baker and S. Norine. Riemann–Roch and Abel–Jacobi theory on a finite graph. *Adv. in Math.*, 215:766–788, 2007.

[10] M. Baker and B. Poonen. Torsion packets on curves. *Composition Math.*, 127:109–116, 2001.

[11] M. Baker and F. Shokrieh. Chip-firing games, potential theory on graphs, and spanning trees, preprint, 2011. arXiv:1107.1313.

[12] G. Bergman. The logarithmic limit-set of an algebraic variety. *Trans. Amer. Math. Soc.*, 157:459–469, 1971.

[13] B. Bollobás. *Graph Theory: An Introductory Course*, volume 63 of *Gratuate Texts in Mathematics*. Springer-Verlag, New York, 1979.

[14] L. Caporaso. Algebraic and tropical curves: comparing their moduli spaces. In *Handbook of moduli. Vol. I*, volume 24 of *Adv. Lect. Math. (ALM)*, pages 119–160. Int. Press, Somerville, MA, 2013.

[15] L. Caporaso and F. Viviani. Torelli theorem for graphs and tropical curves. *Duke Math. J.*, 153(1):129–171, 2010.

[16] A. Del Centina. Weierstrass points and their impact in the study of algebraic curve: a historical account from the "Lückensatz" to the 1970s. *Ann. Univ. Ferrara*, 54:37–59, 2008.

[17] T. Chinburg and R. Rumely. The capacity pairing. *J. Reine Angew. Math.*, 434:1–44, 1993.

[18] M. Einsiedler, M. Kapranov, and D. Lind. Non-archimedean amoebas and tropical varieties. *J. Reine Angew. Math.*, 601:139–157, 2006.

[19] R. M. Foster. The average impedance of an electrical network. In J. W. Edwards, editor, *Reissner Anniversary Volume, Constributions to Applied Mechanics*, pages 333–340. Ann Arbor, MI, 1949.

[20] A. Gathmann and M. Kerber. A Riemann–Roch theorem in tropical geometry. *Math. Z.*, 259:217–230, 2008.

[21] G. Grimmett. *Probability on Graphs: Random Processes on Graphs and Lattices.* Cambridge University Press, second edition, 2018.

[22] A. Gross, F. Shokrieh, and L. Tóthmérész. Effective divisor classes on metric graphs, preprint, 2018. arXiv:1807.00843v1.

[23] C. Haase, G. Musiker, and J. Yu. Linear systems on tropical curves. *Math. Z.*, 270(3–4):1111–1140, 2012.

[24] C. Ho. A note on proper maps. *Proc. Amer. Math. Soc.*, 51(1):237–241, 1975.

[25] A. Hürwitz. Über algebraische Gebilde mit eindeutigen Transformationen in sich. *Math. Ann.*, 41:403–442, 1892.

[26] D. Jensen and S. Payne. Tropical independence I: Shapes of divisors and a proof of the Gieseker–Petri theorem. *Algebra Number Theory*, 8(9):2043–2066, 2014.

[27] E. Katz. Matroid theory for algebraic geometers. In M. Baker and S. Payne, editors, *Nonarchimedean and Tropical Geometry*, Simons Symposia, pages 435–517. Springer, Cham, 2016.

[28] E. Katz, J. Rabinoff, and D. Zureick-Brown. Uniform bounds for the number of rational points on curves of small Mordell–Weil rank. *Duke Math. J.*, 165:3189–3240, 2016.

[29] E. Katz, J. Rabinoff, and D. Zureick-Brown. Diophantine and tropical geometry, and uniformity of rational points on curves. In *Algebraic Geometry: Salt Lake City 2015*, pages 231–279, 2018.

[30] G. Kirchhoff. Über die Auflösung der Gleichungen, auf welche man bei der Untersuchungen der linearen Vertheilung galvanischer Ströme gefürt wird. *Ann. Phys. Chem.*, 72:497–508, 1847.

[31] Y. Luo. Tropical convexity and canonical projections, preprint, 2013. arXiv:1304.7963.

[32] G. Mikhalkin and I. Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and Abelian Varieties*, volume 465 of *Contemp. Math.*, pages 203–231. Amer. Math. Soc., Providence, RI, 2008.

[33] R. Miranda. *Algebraic Curves and Riemann Surfaces*, volume 5 of *Graduate Studies in Mathematics.* American Mathematical Society, Providence, RI, 1995.

[34] D. Mumford. *Curves and their Jacobians.* The University of Michigan Press, Ann Arbor, MI, 1977.

[35] A. Neeman. The distribution of Weierstrass points on a compact Riemann surface. *Ann. of Math.*, 120:317–328, 1984.

[36] B. A. Olsen. On higher order Weierstrass points. *Ann. of Math.*, 95(2):357–364, 1972.

[37] J. G. Oxley. *Matroid Theory.* Oxford Univ. Press, New York, 1992.

[38] N. Pflueger. Tropical curves, 2011. minor thesis.

[39] M. Raynaud. Courbes sur une variété abélienne et points de torsion. *Invent. Math.*, 71:207–233, 1983.

[40] D. H. Richman. The distribution of Weierstrass points on a tropical curve, preprint, 2019. arXiv:1809.07920.

[41] 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.

[42] S. Zhang. Admissible pairing on a curve. *Invent. Math.*, 112:171–193, 1993.