EQUIDISTRIBUTION OF WEIERSTRASS POINTS ON TROPICAL CURVES

DAVID HARRY RICHMAN

Abstract. We show that for a sequence of generic divisors on a metric graph whose degrees grow to infinity, the associated Weierstrass points become equidistributed according to the Arakelov–Zhang canonical measure. This is a tropical analogue of a result of Neeman, for equidistribution of Weierstrass points on a compact Riemann surface. This work is closely connected to and inspired by work of Amini, who proved a non-Archimedean analogue for equidistribution of Weierstrass points on a Berkovich curve. However, the results in this paper are proved using combinatorial arguments rather than algebraic or analytic geometry.

Contents

1. Introduction

In this paper we study the set of Weierstrass points on an abstract tropical curve, which is assosicated to an (arbitrary) given divisor class on the curve. In particular, we ask

(A) When is the set of Weierstrass points finite?

and

(B) How are these points distributed as the degree approaches infinity?

We show that for any abstract tropical curve Γ, the Weierstrass locus is finite for a generic divisor class, and we prove that for any degree-increasing sequence of such generic divisors the Weierstrass points become equidistributed according to the Arakelov–Zhang canonical measure on Γ. This measure can be described via interpreting Γ as an electrical network of resistors.

Date: September 20, 2018.
Work of the author was partially supported by NSF grant DMS-1600223.

1.1. **Statement of results and outline.** Given a compact, connected metric graph Γ and a divisor D of rank $r = r(D)$, we define the Weierstrass locus $W_D = W_{\Gamma,D}$ as

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

where $\sim$ denotes linear equivalence and $r(D)$ is the Baker-Norine rank. The set W_D may fail to be finite; in some cases it contains all of Γ (see Example 4.6). Our first result addresses when the Weierstrass locus is finite. Here "generic" means on a dense open subset of the space of divisor classes.

Theorem A. *Let Γ be a compact, connected metric graph. For a generic divisor class $[D]$ on Γ, the associated Weierstrass locus W_D is finite.*

The next theorem is the main result of our paper, which describes the equidistribution of tropical Weierstrass points.

Theorem B. *Let Γ be a metric graph of genus g, and let $\{[D_N] : N \geq 1\}$ be a sequence of generic divisor classes on Γ with $\deg D_N = N$. Let W_N denote the Weierstrass locus of the divisor D_N, 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 Arakelov-Zhang canonical measure μ on Γ.

(Warning: we use a different normalization for μ than previous authors, namely we have $\mu(\Gamma) = g$ rather than $\mu(\Gamma) = 1$.) We also obtain a quantitative version of this result which specifies a bound on the rate of convergence.

Theorem C. *Let Γ be a metric graph of genus g, let $[D_N]$ be a generic degree N divisor class, and let W_N denote the Weierstrass locus of D_N. Let μ denote the Arakelov-Zhang canonical measure on Γ.*

(a) For any segment e in Γ,

$$N\mu(e) - 3g - 1 \leq \#(W_N \cap e) \leq N\mu(e) + g + 2.$$

(b) As $N \to \infty$,

$$\frac{\#(W_N \cap e)}{N} = \mu(e) + O\left(\frac{1}{N}\right).$$

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

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

(The big-O constant may depend on f.)

(c) If e is a segment of Γ with $\mu(e) > \frac{3g+1}{N}$, then e contains at least one D_N-Weierstrass point.

It is likely that the bounds in part (a) are not sharp.

In Section 2 we review background material on metric graphs and their divisor theory. In Section 3 we review the interpretation of a metric graph as an electrical resistor network, and define the canonical measure. In Section 4 we define the

Weierstrass locus for a metric graph, give examples, and prove that W_D is generically finite (Theorem A). In Section 5, we prove our new results regarding the distribution of Weierstrass points on metric graphs (Theorems B and C).

1.2. **Previous work.** The set of (canonical) Weierstrass points on a complex algebraic curve of genus $g \geq 2$ has been a classical object of study (see e.g. [14]). This is a set of $g^3 - g$ points (counting with multiplicity) on X which are instrinsic to X as an abstract curve, without reference to any (non-canonical) embedding of X into an ambient space. They form a useful tool, e.g. for proving that the automorphism group of such a curve is finite. This notion naturally extends to (higher) Weierstrass points, which is a finite set of points on X associated to a choice of divisor class on X. The number of Weierstrass points (counting with multiplicity) grows quadratically as a function of the degree of $[D]$.

The following useful intution is given by Mumford [15]: the Weierstrass points associated to a divisor of degree N form a higher-genus analogue of the set of N-torsion points on an elliptic curve. (Just as choosing a different origin for the group law on a genus 1 curve leads to a different set of torsion points, choosing different degree N divisors will give you different sets of Weierstrass points.) 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.

This was answered in the affirmative in 1984 by Neeman [16] (a student of Mumford), who showed that for a complex algebraic curve of genus $g \geq 2$, when $N \to \infty$ the Weierstrass points of a degree N divisor become equidistributed according to the Bergman measure.

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}, say after retracting to a (compact) skeleton Γ. This was addressed by Amini in the preprint [2]. Here the answer is that the Weierstrass points are equidistributed according to the "Arakelov–Zhang canonical admissable measure" μ_Γ, constructed by Zhang in [17]. (This measure does not have support on bridge edges, so it is indepedent of the choice of skeleton.) Zhang's construction was motivated by Arakelov's pairing for divisors on a Riemann surface [3], for the purpose of answering questions in arithmetic geometry. Here we give a definition of μ along more elementary lines following Baker and Faber [4], using the notions of current flow and electric potential in a (1-dimensional) network of resistors.

In his preprint, Amini raises the question of whether the distribution of Weierstrass points is possibly intrisic 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. Our approach is to sidestep this issue entirely by showing that finiteness does hold for a generic choice of divisor class. With this assumption of genericity, we are able to show that equidistribution of Weierstrass points is intrinsic to Γ.

(Technical note: our tropical curves Γ have no "hidden genus" and no infinite legs, i.e. we restrict our attention to X^{an} with totally degenerate reduction and no punctures.)

1.3. **Notation.** Here we collect some notation which will be used throughout the paper.

Γ	a compact, connected metric graph
$\mathrm{PL}_{\mathbb{R}}(\Gamma)$	continuous, piecewise linear functions on Γ
$\mathrm{PL}_{\mathbb{Z}}(\Gamma)$	continous, piecewise $\mathbb{Z}$-linear functions on Γ
$\mathrm{S}(\Gamma)$	"well-behaved" piecewise smooth (or polynomial?) functions on Γ
$\Delta(f)$	the principal divisor associated to a piecewise ($\mathbb{Z}$-)linear function f
D	a divisor on a metric graph (or algebraic curve)
D_N	a divisor of degree N
$K = K_\Gamma$	the canonical divisor on Γ
$r(D)$	the Baker-Norine rank of D
$\mathrm{Div}(\Gamma)$	divisors on Γ (with $\mathbb{Z}$-coefficients)
$\mathrm{Div}_{\mathbb{R}}(\Gamma)$	divisors with $\mathbb{R}$-coefficients on Γ ($= \mathrm{Div}(\Gamma) \otimes_{\mathbb{Z}} \mathbb{R}$)
$\mathrm{Div}^d(\Gamma)$	divisors of degree d on Γ
$[D]$	the set of divisors linearly equivalent to D
$\lvert D\rvert$	the set of effective divisors linearly equivalent to D
$\mathrm{red}_x[D]$	the x-reduced divisor equivalent to D, where $x \in \Gamma$
$\mathrm{Eff}^d(\Gamma)$	effective divisors of degree d on Γ
$\mathrm{Pic}^d(\Gamma)$	divisor classes of degree d on Γ
$W_0^d(\Gamma)$	effective divisor classes of degree d on Γ
$\mu = \mu_\Gamma$	the Arakelov-Zhang canonical measure on Γ

2. Abstract 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" interchageably. 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 in Figure 1 has genus 0. A minimal combinarial model has 8 vertices and 7 edges.

Example 2.2. The metric graph in Figure 2 has genus 2. A minimal combinatorial model has 2 vertices and 3 edges.

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

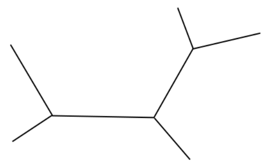

FIGURE 1. Metric graph of genus 0.

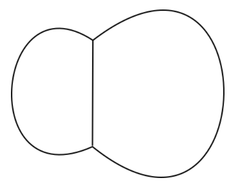

FIGURE 2. Metric graph of genus 2.

$D = \sum_{x\in\Gamma} a_x x$, 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 its coefficients are non-negative; we write $D \geq 0$ to indicate that D is effective. We let $\mathrm{Eff}^d(\Gamma)$ denote the set of effective divisors of degree d on Γ. $\mathrm{Eff}^d(\Gamma)$ 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 [12] and Mikhalkin–Zharkov [13]. 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 (affine) 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

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

In words, the coefficient in $\Delta(f)$ of a point x is equal to the sum of the outgoing slopes of f at x. 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 $|D| \subset \mathrm{Eff}^d(\Gamma)$.

Remark 2.3 (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 contruction as specifying a unique "geodesic path" between any two points in the complete linear system $|D|$.)

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 3 for an illustration. Note that according to this definition, $f_\Delta^{-1}(\lambda) = \Delta^-(f)$

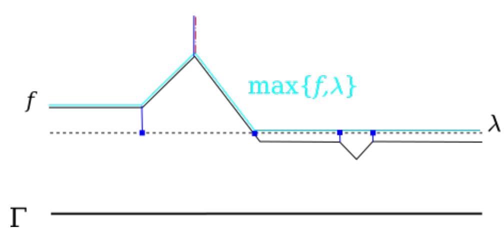

FIGURE 3. Linear interpolation showing the divisor $f_\Delta^{-1}(\lambda)$

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 $\Delta^-(f)$.

A divisor class $[D]$ is typically very large, so it is convenient to have a method of choosing a somewhat-canonical representative divisor inside $[D]$. We can do so after fixing a basepoint q on our metric graph Γ. The *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 effective part is "as close as possible" to the basepoint q. We defer giving the full definition until Section 3.2, following [Baker–Shokrieh, 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) the degree of $\mathrm{red}_q[D]$ away from q is at most g
(RD3) for a fixed effective divisor D, the map $\Gamma \to |D|$ sending $q \mapsto \mathrm{red}_q[D]$ is continuous (due to Amini [1, Theorem 3])

2.3. **Picard group and Abel-Jacobi.** We 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.)

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

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

The degree-0 component $\mathrm{Pic}^0(\Gamma)$ is an abelian group, and each $\mathrm{Pic}^d(\Gamma)$ is a torsor for $\mathrm{Pic}^0(\Gamma)$.

Theorem 2.4 (Abel-Jacobi for metric graphs)**.** *Let Γ be a metric graph of genus g. Then for any degree d, there is a homeomorphism of topological spaces*

$$\mathrm{Pic}^d \Gamma \cong (S^1)^{\times g}.$$

When $d = 0$, this is a isomorphism of (compact abelian) topological groups.

Proof. See Mikhalkin–Zharkov [13]. The proof follows the same idea as the classical Abel-Jacobi theorem, to show that $\mathrm{Pic}^0(\Gamma) = H^1(\Gamma, \mathbb{R})/H_1(\Gamma, \mathbb{Z})^\vee \cong \mathbb{R}^g/\mathbb{Z}^g$. □

We let $W_0^d(\Gamma)$ denote the set of divisor classes on Γ of degree d which have an effective representative. In other words, $W_0^d(\Gamma)$ is the image of $\mathrm{Eff}^d(\Gamma)$ under the (degree-d restriction of the) cokernel map $\mathrm{Div}(\Gamma) \to \mathrm{Pic}(\Gamma)$:

$$\begin{array}{ccc} \mathrm{Eff}^d(\Gamma) & \longrightarrow & \mathrm{Div}^d(\Gamma) \\ \downarrow & & \downarrow{\scriptstyle \mathrm{coker}\,\Delta} \\ W_0^d(\Gamma) & \longrightarrow & \mathrm{Pic}^d(\Gamma). \end{array}$$

The space W_0^d is naturally a polyhedral complex of pure dimension d, when $0 \leq d \leq g$. As a particularly important case, the *theta divisor* $\Theta = \Theta(\Gamma)$ is $\Theta = W_0^{g-1}(\Gamma)$, which lives inside $\mathrm{Pic}^{g-1}(\Gamma) \cong (S^1)^g$ as a codimension 1 polyhedral complex.

Remark 2.5. 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 identitying $\mathrm{Div}(\Gamma)$ with the space of integer-valued discrete measures on Γ, via

$$D = \sum_{i=1}^{N} a_i x_i \quad \leftrightarrow \quad \delta = \sum_{i=1}^{N} a_i \delta_{x_i}$$

so that $\Delta(f)$ coincides with the (distributional) second derivative $-\frac{d^2}{dx^2} f(x)$, at least for x in the interior of an edge. The definition of metric graph Laplacian naturally extends to piecewise linear functions on Γ with arbitrary real slopes, if we also allow real-valued coefficients in the divisor $\Delta(f)$. This yields a map

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

The cokernel of this map is less interesting (e.g. it does not tell us the genus of Γ); it 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

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

where $\mathrm{Pic}(\Gamma) \cong \mathbb{Z} \times (S^1)^g$.)

2.4. **Rank and Riemann-Roch.** We recall the definition of the rank of a divisor on a metric graph, which is due to Baker and Norine [6] (originally for divisors on a combinatorial graph) and extended to metric graphs by Gathmann-Kerber [12] and Mikhalkin–Zharkov [13]. 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

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

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

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

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, which have degree at most $g - 1$. 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 2.6 (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 [12, Proposition 3.1] and Mikhalkin–Zharkov [13, Theorem 7.3], which both adapt the arguments of Baker–Norine [6] for the case of combinatorial graphs. □

Corollary 2.7 (Riemann's inequality for metric graphs)**.** *For a divisor D on a metric graph of genus g,*

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

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

3. Canonical measure and resistor networks

In this section we define the Arakelov–Zhang canonical measure on a metric graph (due to Zhang [17]) via the perspective of resistor networks following Baker–Faber [4]. We may view this construction as a one-dimensional analogue of Gaussian curvature on a closed two-dimensional surface.

3.1. **Voltage/potential 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 Kirchoff'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. It is a well-known empirical fact that Kirchoff'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.

Our convention is that current flows from higher voltage to lower voltage.

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.

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

$$\Delta(j_z^y) = z - y \in \mathrm{Div}_{\mathbb{R}}^0(\Gamma) \quad \text{and} \quad j_z^y(z) = 0.$$

Proof. For the existence and uniqueness of j_z^y, see Theorem 6 and Corollary 3 of Baker–Faber [4]. Note that they use the notation $j_z(y, -)$ for $j_z^y(-)$. □

Remark 3.3. 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.

The existence of $j_z^y \in \mathrm{PL}_{\mathbb{R}}(\Gamma)$ for any $y, z \in \Gamma$ implies that the principal divisor map $\Delta : \mathrm{PL}_{\mathbb{R}}(\Gamma) \to \mathrm{Div}_{\mathbb{R}}^0(\Gamma)$ is surjective. This verifies the claim made in Remark 2.5 concerning the exactness of the sequence

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

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 continuous in x, y, and z.

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 [4, Theorem 8]; they refer to (b) as the “Magical Identity”. Note that (a) follows from (a) by setting $z = w$. □

Proposition 3.5 (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 for any $x \in \Gamma$ where f is linear, the slope of f is bounded by d, i.e.*

$$|f'(x)| \le d.$$

(This bound is sharp; it is attained only on bridge edges.)

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. □

Remark 3.6. 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 (RD2) of reduced divisors.)

Proposition 3.7 (Discrete approximation of voltage function). *Let $\{D_N : N \ge 1\}$ 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 linear equivalence 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. Let $\phi_N = \frac{1}{N} f_N - j_z^y$. 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 $x \in \Gamma$, $|\phi'_N(x)| \leq \frac{2g}{N}$.

This follows from the slope-current principle (Proposition 3.5). By Riemann's inequality, the y-reduced representative in $[D_N]$ may be expressed as

$$\mathrm{red}_y[D_N] = (N-g)y + E_N$$

for some effective divisor E_N of degree g. Similary, $\mathrm{red}_z[D_N] = (N-g)z + F_N$ for some effective F_N of degree g. Thus the principal divisor associated to $\frac{1}{N} f_N$ is

$$\Delta(\frac{1}{N} f_N) = (1 - \frac{g}{N})z + \frac{1}{N} F_N - (1 - \frac{g}{N})y - \frac{1}{N} E_N.$$

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

$$\Delta(\phi_N) = \Delta\left(\frac{1}{N} f_N - j_z^y\right) = \frac{g}{N} y + \frac{1}{N} F_N - \left(\frac{g}{N} z + \frac{1}{N} E_N\right). \tag{1}$$

In particular, this is a difference of effective divisors of degree $\frac{2g}{N}$, so the zeros $\Delta^+(\phi_N)$ and poles $\Delta^-(\phi_N)$ have degree at most $\frac{2g}{N}$. By Proposition 3.5, this implies $|\phi'_N| \leq \frac{2g}{N}$ which proves the claim. □

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

Proposition 3.8 (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 $|(\frac{1}{N} f_N - j_z^y)'(x)| \leq \frac{2g}{N}$ for any $x \in \Gamma$.

In particular, if $|(j_z^y)'(x)| > \frac{2g}{N}$ then $f'_N(x) \neq 0$ (for any $D_N \in \mathrm{Div}^N(\Gamma)$).

Remark 3.9. We can interpret Proposition 3.7 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. **Total potential 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 3.10. 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

$$\mathcal{E}_q(D) = \sum_i j_q^D(y_i) \qquad \text{where} \qquad j_q^D = \sum_i j_q^{y_i} : \Gamma \to \mathbb{R}.$$

Note that $\mathcal{E}_q(D) \geq 0$, and $\mathcal{E}_q(D)$ is strictly positive if D has support outside of q.

Theorem 3.11 (Baker–Shokrieh)**.** *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*

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

Proof. See Baker–Shokrieh [8, Theorem A.7]. □

Definition 3.12. The *q-reduced divisor* $\mathrm{red}_q[D]$ is the unique divisor in $|D|$ which minimizes the q-energy $\mathcal{E}_q$.

3.3. **Resistance function.** In this section we define the (Arakelov–Zhang–Baker–Faber) canonical measure μ on a metric graph.

Definition 3.13. 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

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

In terms of the voltage function from Section 3.1, $r(x,y) = j_y^x(x)$.

It is straighforward to verify that the resistance function satisfies the following properties

(1) $r(x,x) = 0$,
(2) $r(x,y) > 0$ when $x \neq y$,
(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_z^y, the function $x \mapsto r(x,y)$ is not piecewise linear. We will see that it is instead piecewise quadratic.

Example 3.14. 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.

Definition 3.15. 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 a segment, 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 3.16. This definition 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

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

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

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

Remark 3.17. 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 3.18. This definition of canonical measure differs from that used by Baker–Faber [4], in that our μ does not have a discrete part supported at the points of Γ with valence different from 2.

Remark 3.19. This definition of canonical measure is equal to Zhang's canonical measure [17, Section 3, Theorem 3.2 c.f. Lemma 3.7] associated to the canonical divisor $D = K$, except 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 3.20. If Γ is a circle of circumference L, by Example 3.14 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 $\mu(\Gamma) = 1$.

Proposition 3.21. *Let Γ be a metric graph with canonical measure μ. For an edge e in Γ,*

$$\mu|_e = \frac{1}{R_e + R_{\Gamma\backslash e}} dx$$

where R_e denotes the length of e and $R_{\Gamma\backslash e}$ denotes the effective resistance between the endpoints of e on the graph after removing the interior of e. Equivalently, for any combinatorial model

$$\mu = \sum_{e\in E} \frac{1}{R_e + R_{\Gamma\backslash e}} dx|_e$$

Proof. See Baker–Faber [4, Theorem 12]; note that our μ is defined to be the continuous part of Baker–Faber's μ_{can}.

(The proof idea should be clear from the following example.) □

Example 3.22. Consider the metric graph Γ of genus 2 shown in Figure 4, with edge lengths a, b, c. On the edge of length a, we have $R_e = a$ and $R_{\Gamma\backslash e} = \frac{bc}{b+c}$. When measuring effective resistance $r(x,y)$ between points on e, we can effectively think of Γ as a circle of length $R_e + R_{\Gamma\backslash e} = \frac{ab+ac+bc}{b+c}$. Thus the canonical measure on this edge is $\mu = \frac{b+c}{ab+ac+bc}dx$, by the computation for a circle in Example 3.14.

The total measure on this edge is $\mu(e) = \frac{ab+ac}{ab+ac+bc}$, and by symmetry the total measure on the metric graph is $\mu(\Gamma) = 2$.

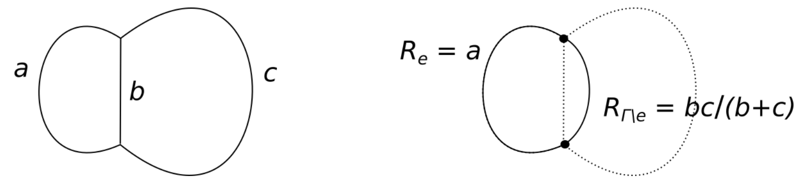

FIGURE 4. Genus 2 metric graph

Corollary 3.23. *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) $0 \leq \mu(e) \leq 1$*;*
(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 3.21, $\mu(e) = \int_e \frac{dx}{R_e + R_{\Gamma \setminus e}} = \frac{R_e}{R_e + R_{\Gamma \setminus e}}$. □

Proposition 3.24 (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 [4, Corollary 5 and Corollary 6] and Foster [11]. □

4. Weierstrass points

In this section we define the Weierstrass locus of an arbitrary divisor D on a metric graph Γ. We show that for a generic divisor class $[D]$ on Γ, the associated Weierstrass locus is a finite set of points, whose cardinality we can specify in terms of the genus of Γ and degree of D.

4.1. **Classical Weierstrass points.** Recall that for an algebraic curve X of genus g, the (canonical) 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 *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 to Weierstrass points when replacing K with an arbitrary divisor. Given a (basepoint-free) divisor D on X, there is an associated map to projective space $\varphi_D : X \to \mathbb{P}^r$. (If the degree of D is at least $2g-1$ then $r = \deg D - g$ by Riemann-Roch.) 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 *(higher) Weierstrass points* associated to the divisor D.

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 associated D-Weierstrass points 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.

A *canonical 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 D-Weierstrass points 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 the figure below with multiplicities. This divisor has rank $r = 5$ since it is in the "non-special range"(?) of Riemann-Roch. The D-Weierstrass locus consists of 6 points evenly spaced around Γ, indicated in red.

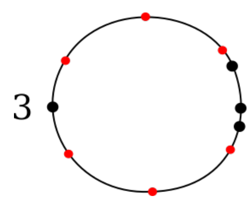

FIGURE 5. D_6-Weierstrass points on a genus 1 metric graph

Example 4.4. Suppose Γ is the genus 3 metric graph with edge lengths as below. Consider the canonical divisor K on Γ, which is supported on the four trivalent vertices. For a generic choice of edge lengths, the Weierstrass locus consists of 8 distinct points on Γ.

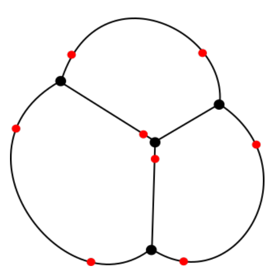

FIGURE 6. Metric graph with W_K finite.

Example 4.5. 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 (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. Thus the Weierstrass locus of D_N will contain $N-g+1$ 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 7. 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.

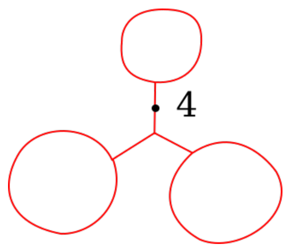

FIGURE 7. Divisor with $W_D = \Gamma$.

Example 4.7 (Failure of W_D to be finite, v2)**.** Consider the genus 3 graph shown in Figure 8. Suppose $D = K$ is the canonical divisor on K. By Riemann-Roch K has rank $r = 2$. Since it is possible to move all 4 chips to lie on the middle loop, any point in the middle loop is a Weierstrass point for K.

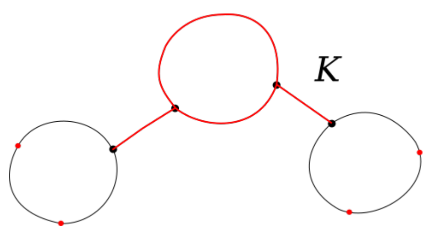

FIGURE 8. Metric graph with W_K not finite.

Example 4.8 (Failure of W_D to be finite, v3)**.** Suppose Γ is a metric graph constructed by adding a bridge edge connecting two disjoint metric graphs Γ_1, Γ_2 of positive genus g_1, g_2 respectively. Let $K = K_\Gamma$ denote the canonical divisor on Γ; it is clear that

$$K = K_{\Gamma_1} + K_{\Gamma_2} + y_1 + y_2,$$

where y_1 and y_2 are the endpoints of the bridge edge.

By considering PL functions which vary on Γ_1, we have $K_{\Gamma_1} \sim_\Gamma (g_1 - 1)y_1 + E_1$ for some effective divisor E_1, and similary $K_{\Gamma_2} \sim_\Gamma (g_2 - 1)y_2 + E_2$. Thus on Γ there is a linear equivalence

$$K \sim g_1 y_1 + g_2 y_2 + E_1 + E_2.$$

Since chips may move freely on a bridge edge, $g_1 y_1 + g_2 y_2 \sim (g_1 + g_2)y$ for any $y \in e$. This shows that e is contained in the Weierstrass locus since

$$r(K) = g - 1 = g_1 + g_2 - 1.$$

4.3. **Intersection with Θ.** In this section we describe the Weierstrass locus as an intersection of two polyhedral subcomplexes of complementary dimension in $\mathrm{Pic}^{g-1}(\Gamma)$. This allows us to prove Theorem A concerning when W_D is finite.

Given a divisor D, let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ denote the map

$$\Phi_D : x \mapsto [D - (r+1)x] \quad \text{where } r = r(D).$$

By definition of the theta divisor $\Theta \subset \mathrm{Pic}^{g-1}(\Gamma)$ as

$$\Theta = \{[D] \in \mathrm{Pic}^{g-1}(\Gamma) : [D] \geq 0\},$$

the Weierstrass locus of D is equal to the intersection $\Phi_D(\Gamma) \cap \Theta$, pulled back to Γ from $\mathrm{Pic}^{g-1}(\Gamma)$.

Proposition 4.9. *If $r(D) \geq 0$, the map $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ is locally injective (i.e. an immersion), except on bridge edges.*

Proof. The map Φ_D may be expressed as a composition of three maps

$$\Phi_D : \Gamma \xrightarrow{\alpha} \mathrm{Pic}^1(\Gamma) \xrightarrow{\beta} \mathrm{Pic}^{r+1}(\Gamma) \xrightarrow{\gamma} \mathrm{Pic}^{g-1}(\Gamma),$$

where α sends $x \mapsto [x]$, β sends $[E] \mapsto [-(r+1)E]$, and γ sends $[E] \mapsto [D+E]$. The map γ is simply a translation, so it is an isomorphism of topological spaces. (isometry of polyhedral complexes?). The map β is a $(r+1)^g$-fold covering map, so it is a local isomorphism if $r = r(D) \geq 0$. Thus it suffices to show that the first map α is a local isomorphism on non-bridge edges.

This follows from the Abel-Jacobi theorem for metric graphs, see e.g. Baker–Faber [5, Theorem 4.1 (3)(4)]. Note that $\mathrm{Pic}^1(\Gamma)$ is (non-canonically) isomorphic to the Jacobian $\mathrm{Jac}(\Gamma) \cong \mathrm{Pic}^0(\Gamma)$ by choosing a basepoint x_0 to subtract. □

Theorem A. *For a dense subset of divisor classes $[D]$ 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 $\mathrm{Eff}^N(\Gamma) \to \mathrm{Pic}^N(\Gamma)$ is a polyhedral complex with positive codimension. In this case, for a generic 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 $r(D) = N - g$ (since generically $r(K - D) = -1$ by the above paragraph). Thus, it suffices to show that W_D is finite for divisors satisfying $r(D) = N - g$.

Let $\Phi_D : \Gamma \to \mathrm{Pic}^{g-1}(\Gamma)$ be the map defined above, sending

$$x \mapsto [D - (r+1)x] = [D - (N - g + 1)x].$$

Note that as $[D]$ varies, the image $\Phi_D(\Gamma)$ varies by translation inside $\mathrm{Pic}^{g-1}(\Gamma)$. Recall that the Weierstrass locus W_D is equal to

$$W_D = \Phi_D^{-1}(\Phi_D(\Gamma) \cap \Theta) \subset \Gamma$$

where $\Theta = \{[D] \in \mathrm{Pic}^{g-1}(\Gamma) : [D] \geq 0\}$ is the theta divisor. For W_D to be finite, it suffices that

(i) $\Phi_D(\Gamma) \cap \Theta$ is finite, and

(ii) $\Phi_D^{-1}(x)$ is finite for each $x \in \Phi_D(\Gamma) \cap \Theta$.

We verify that each of the above conditions is satisfied by a dense subset of $\mathrm{Pic}^N(\Gamma)$, which implies the desired result.

First, consider when $\Phi_D(\Gamma) \cap \Theta$ is not finite. Recall that $\Phi_D(\Gamma)$ is a 1-dimensional polyhedral complex with finitely many segments, and Θ is a $(g-1)$-dimensional polyhedral complex with finitely many facets. There can only be finitely many transverse intersections between $\Phi_D(\Gamma)$ and Θ since the ambient space $\mathrm{Pic}^{g-1}(\Gamma)$ is compact. Hence the intersection $\Phi_D(\Gamma) \cap \Theta$ is not finite only if there is some non-transverse intersection between a segment of $\Phi_D(\Gamma)$ and a facet of Θ. If we fix some choice of segment and facet, then there is at most a $(g-1)$-dimensional space of translations which allow these to intersect non-transversely. Taking a union over all (finitely many) choices of segment of $\Phi_D(\Gamma)$ and facet of Θ, we see that the set

$$\{[D] \in \mathrm{Pic}^N(\Gamma) \text{ such that } \Phi_D(\Gamma) \cap \Theta \text{ is not finite } \}$$

has dimension at most $g-1$; hence its complement is dense in $\mathrm{Pic}^N(\Gamma)$.

Next we consider when (ii) holds. Let Γ^{br} denote the closed subset of Γ consisting of all bridge edges, and let

$$\Phi_D(\Gamma)^\circ = \Phi_D(\Gamma) \setminus \Phi_D(\Gamma^{\mathrm{br}}).$$

Note that $\Phi_D(\Gamma^{\mathrm{br}})$ is a finite set of points, and $\Phi_D(\Gamma)^\circ$ is dense in $\Phi_D(\Gamma)$.

For $x \in \Phi_D(\Gamma)^\circ$ the preimage $\Phi_D^{-1}(x)$ is finite, of size at most $(r+1)^g = (N - g+1)^g$. Thus for (ii) to hold it suffices that $\Phi_D(\Gamma^{\mathrm{br}}) \cap \Theta$ is empty. By an argument like in the paragraph above, the space

$$\{[D] \text{ such that } \Phi_D(\Gamma^{\mathrm{br}}) \cap \Theta \text{ is non-empty}\}$$

has dimension at most $g-1$, so its complement is dense in $\mathrm{Pic}^N(\Gamma)$ as desired. □

5. Equidistribution

In this section we prove our main result. We show that for a degree-increasing sequence of divisors with discrete Weierstrass locus (i.e. Weierstrass-generic divisors), the Weierstrass points become equidistributed in Γ with respect to the canonical measure μ, defined in Section 3.3.

First we consider some low genus examples of this phenomenon.

Example 5.1 (Genus 0 metric graph)**.** For any divisor D_N, the associated D_N-Weierstrass locus W_N is empty so $\delta_N = 0$. All edges are bridges, so the canonical measure is $\mu = 0$.

Example 5.2 (Genus 1 metric graph)**.** Let Γ be a genus 1 metric graph whose unique closed (simple) cycle has length L. For a divisor D_N of degree N, the Weierstrass locus consists of N evenly-spaced points ("torsion points") around the loop. The distance between adjacent points is $\frac{L}{N}$, so on a segment e of length R the number of Weierstrass points is bounded by

$$\frac{R}{L/N} - 1 \le \#(W_N \cap e) \le \frac{R}{L/N} + 1.$$

Normalizing by $\frac{1}{N}$, the associated discrete measure δ_N satisfies

$$\frac{R}{L} - \frac{1}{N} \le \delta_N(e) \le \frac{R}{L} + \frac{1}{N}.$$

Hence $\delta_N(e) \to \frac{R}{L} = \mu(e)$ as $N \to \infty$.

We now address the general case.

Theorem B. *Let $\{D_N : N \geq 1\}$ be a sequence of Weierstrass-generic divisors on Γ with $\deg D_N = N$. Let W_N be the Weierstrass locus of the divisor D_N, 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 canonical measure μ on Γ.

In other words, for any continuous function $f : \Gamma \to \mathbb{R}$, as $N \to \infty$ we have convergence

$$\int_\Gamma f(x)\delta_N(dx) = \frac{1}{N} \sum_{x \in W_N} f(x) \to \sum_{e \in E(\Gamma)} \frac{1}{R_e + R_{\Gamma\backslash e}} \int_e f(x)dx = \int_\Gamma f(x)\mu(dx).$$

Proof. 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 R_e, with endpoints s and t. Let $W_N \cap e$ denote the set of D_N-Weierstrass points lying on the segment e. It suffices to show that

$$\lim_{N \to \infty} \frac{\#(W_N \cap e)}{N} = \mu(e). \tag{2}$$

Recall that by Proposition 3.21

$$\mu(e) = \frac{R_e}{R_e + R_{\Gamma\backslash e}}$$

where $R_{\Gamma\backslash e}$ denotes the effective resistance between the endpoints of e when the interior of e is removed from Γ. (If $\Gamma\backslash e$ is disconnected, then $R_{\Gamma\backslash e} = +\infty$ and $\mu(e) = 0$.) We prove the convergence (2) by relating each side to slopes of piecewise linear functions on Γ.

First, consider the voltage function $j_t^s : \Gamma \to \mathbb{R}$ (see Section 3.1 for details). Note that the voltage drop in Γ between endpoints of e is

$$j_t^s(s) - j_t^s(t) = r(s,t) = \frac{R_e R_{\Gamma\backslash e}}{R_e + R_{\Gamma\backslash e}},$$

by the parallel rule for effective resistance, so the slope of j_t^s along the segment e is equal to

$$\begin{aligned} \text{slope of } j_t^s \text{ on } e = \frac{j_t^s(s) - j_t^s(t)}{R_e} &= \frac{R_{\Gamma\backslash e}}{R_e + R_{\Gamma\backslash e}} \\ &= 1 - \frac{R_e}{R_e + R_{\Gamma\backslash e}} = 1 - \mu(e). \end{aligned} \tag{3}$$

(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 $\mu(e)$ to the left-hand side of (2), we consider functions in $\mathrm{PL}_\mathbb{Z}(\Gamma)$ which are "discrete approximations" of j_t^s, and show that the current flow in these functions is 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 3.7, as $N \to \infty$ we have uniform convergence

$$\frac{1}{N} f_N \to j_t^s. \tag{4}$$

Thus to show (2) using (3) and (4), it suffices to show that

$$\lim_{N\to\infty} \left(\text{slope of } \frac{1}{N} f_N \text{ on } e \right) = 1 - \lim_{N\to\infty} \frac{\#(W_N \cap e)}{N}. \tag{5}$$

First we explain why "slope of $\frac{1}{N} f_N$ on e" converges in the limit, even though the function may have distinct slopes on e. Note that $\mathrm{red}_s[D_N]$ has at most one chip on the interior of the segment e; otherwise it could be further s-reduced by Dhar's burning algorithm. For the same reason $\mathrm{red}_t[D_N]$ has at most one chip on the interior of e. Thus f_N has at most three distinct slopes on e and adjacent slopes differ by 1.

We give an intuitive reason why (5) is true: The slope of the function f_N on a directed segment is equal to the net number of chips moving 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 reverse one-by-one. The number of "reverse-moving" chips is equal to $\#(W_N \cap e)$, since x is Weierstrass exactly when $\mathrm{red}_x[D_N]$ has an "extra" chip, 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 since the first and last reverse chips may move only partially along e). When we normalize by dividing by N, we have $\frac{g}{N} \to 0$ as $N \to \infty$ so we obtain (5).

Now we give a more detailed rigorous argument. We parametrize the segment e by the real interval $[0, L]$, sending $0 \mapsto s$ and $L \mapsto t$. Consider the map

$$[0, L] \to \mathrm{Eff}^N\,\Gamma$$

defined by sending a point x to the x-reduced representative of $[D_N]$, i.e.

$$x \mapsto \mathrm{red}_x[D_N].$$

The divisor $\mathrm{red}_x[D_N]$ has at least $N - g$ chips at x by Riemann-Roch, and it has at least $N - g + 1$ chips x if and only if x is a D_N-Weierstrass point. (Since D_N is Weierstrass-generic, we cannot have more than $N - g + 1$ chips at x in $\mathrm{red}_x[D_N]$.)

We partition the interval $[0, L]$ into two types of sub-intervals: we say x is "Weierstrass-far" if the divisor $\mathrm{red}_x[D_N]$ contains only $N - g$ chips on the interior of the segment e, and say x "Weierstrass-close" if there are more than $N - g$ chips on e.

Note that each Weierstrass-close interval contains exactly one Weierstrass point in its interior, and each Weierstrass-far interval contains none. The boundaries between Weierstrass–close and –far intervals happen exactly when $\mathrm{red}_x[D_N]$ contains a chip on one endpoint of e.

There are four cases to consider: $\mathrm{red}_s[D_N]$ could be Weierstrass-close or Weierstrass-far, and similarly for $\mathrm{red}_t[D_N]$. We will address the case where $\mathrm{red}_s[D_N]$ and $\mathrm{red}_t[D_N]$ are Weierstrass-far; the other cases can be handled easily with small modifications.

Let $s = y_0 < y_1 < y_2 < \ldots < y_m$ be the set of points in $[0, L]$ such that $\mathrm{red}_{y_i}[D_N]$ has a chip on the endpoint s, and let $z_0 < z_1 < z_2 < \ldots < z_m = t$ be the set of points such that $\mathrm{red}_{z_i}[D_N]$ has a chip on the endpoint t. Each Weierstrass-far interval is of the form $[y_i, z_i]$ for $i = 0, 1, \ldots m$ and each Weierstrass-close interval is of the form $[z_{i-1}, y_i]$ for $i = 1, 2, \ldots, m$. For each Weierstrass-far interval $[y_i, z_i]$ let $g_N^{(i)}$ denote a function in $\mathrm{PL}_{\mathbb{Z}}(\Gamma)$ such that

$$\Delta(g_N^{(i)}) = \mathrm{red}_{z_i}[D_N] - \mathrm{red}_{y_i}[D_N],$$

and let $h_N^{(i)}$ be a function such that

$$\Delta(h_N^{(i)}) = \mathrm{red}_{y_i}[D_N] - \mathrm{red}_{z_{i-1}}[D_N]$$

for each Weierstrass-close interval $[z_{i-1}, y_i]$. In Figures 9 and 10 we illustrate the graphs of these functions over the segment e. (All slopes in the figures refer to the "uphill" direction.)

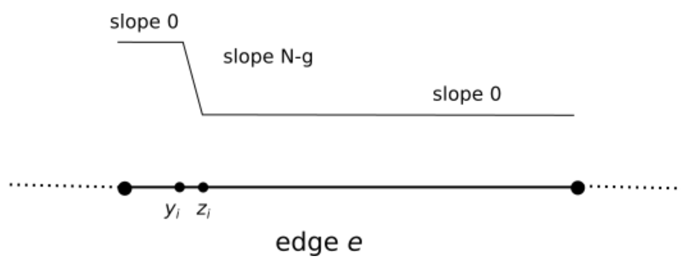

FIGURE 9. Function $g_N^{(i)}$ corresponding to a Weierstrass-far interval

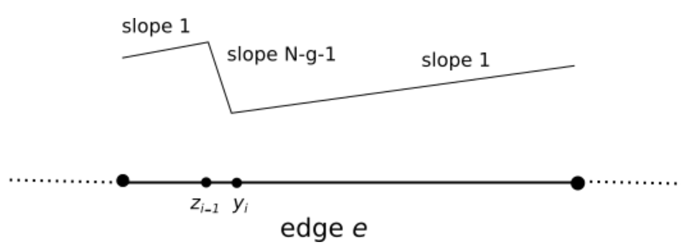

FIGURE 10. Function $h_N^{(i)}$ corresponding to a Weierstrass-close interval

By telescoping of the principal divisors in the right-hand sum, we have the identity

$$\Delta(f_N) = \Delta(g_N^{(0)}) + \Delta(h_N^{(1)}) + \Delta(g_N^{(1)}) + \cdots + \Delta(h_N^{(m)}) + \Delta(g_N^{(m)}),$$

which implies

$$f_N = g_N^{(0)} + h_N^{(1)} + g_N^{(1)} + \cdots + h_N^{(m)} + g_N^{(m)} + C, \tag{6}$$

where C is the additive constant that guarantees both sides evaluate to 0 at $t \in \Gamma$. By examining the slope of each $g_N^{(i)}$ and $h_N^{(i)}$, we see that the right-hand side of (6) has slope $N - g - m$ on the segment e. Since m is the number of Weierstrass-close intervals, we have $m = \#(W_N \cap e)$ (unless the endpoints could be Weierstrass, in which case $m \leq \#(W_N \cap e) \leq m + 2$). Thus (6) implies the bounds

$$N - g - \#(W_N \cap e) \leq \text{ (slope of } f_N \text{ on } e\text{) } \leq N - g - \#(W_N \cap e) + 2.$$

Normalizing by a factor of N, we obtain

$$(7)\qquad 1 - \frac{g}{N} - \frac{\#(W_N \cap e)}{N} \leq \text{ (slope of } \frac{1}{N} f_N \text{ on } e) \; \leq 1 - \frac{g-2}{N} - \frac{\#(W_N \cap e)}{N}$$

Taking the limit as $N \to \infty$ implies (5), which completes the proof. □

Theorem C (Quantitative version of Theorem B)**.** *Let Γ be a metric graph of genus g, let $[D_N]$ be a generic degree N divisor class, and let W_N denote the Weierstrass locus of D_N. Let μ denote the Arakelov–Zhang canonical measure on Γ.*

(a) For any segment e in Γ,

$$N\mu(e) - 3g - 1 \leq \#(W_N \cap e) \leq N\mu(e) + g + 2.$$

(b) For a sequence of divisors D_N, as $N \to \infty$

$$\frac{\#(W_N \cap e)}{N} = \mu(e) + O(\frac{1}{N})$$

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

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

(c) If e is a segment of Γ with $\mu(e) > \frac{3g+1}{N}$, then e contains at least one Weierstrass point of D_N.

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

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 3.8, we have

$$|(f_N - Nj_t^s)'(x)| \leq 2g$$

so

$$|f_N'(x)| \leq N|j'(x)| + 2g.$$

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

$$|f_N'(x)| \leq N - N\mu(e) + 2g.$$

Moreover the proof of Theorem B shows that

$$N - g - \#(W_N \cap e) - 1 \leq |f_N'(x)|.$$

Combining these inequalities gives

$$N\mu(e) - 3g - 1 \leq \#(W_N \cap e)$$

We similarly obtain the upper bound

$$\#(W_N \cap e) \leq N\mu(e) + g + 2$$

by combining the inequalities

$$N - N\mu(e) - 2g \leq |f_N'(x)| \quad \text{and} \quad |f_N'(x)| \leq N - g - \#(W_N \cap e) + 2.$$

□

References

[1] O. Amini, *Reduced divisors and embeddings of tropical curves*, Trans. Amer. Math. Soc. **365** (2013), 4851–4880.

[2] O. Amini, *Equidistribution of Weierstrass points on curves over non-Archimedean fields*, preprint, arXiv:1412.0926v1, 2014.

[3] S. J. Arakelov, *Intersection theory of divisors on an arithmetic surface*, Izv. Akad. Nauk **8** (1974), 1167–1180.

[4] M. Baker and X. Faber, *Metrized graphs, Laplacian operators, and electrical networks*, in Quantum graphs and their applications, volume 415 of Contemp. Math., pages 15–33. Amer. Math. Soc., Providence, RI, 2006.

[5] M. Baker and X. Faber, *Metric properties of the tropical Abel-Jacobi map*, J. Algebr. Comb. **33** (2011), 349–381.

[6] M. Baker and S. Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph*, Adv. in Math. **215** (2007), 766–788.

[7] M. Baker and R. Rumely, *Harmonic analysis on metrized graphs*, Canad. J. Math. **59** (2007), no. 2, 225–275.

[8] M. Baker and F. Shokrieh, *Chip-firing games, potential theory on graphs, and spanning trees*, J. Combin. Theory Ser. A **120** (2013), no. 1, 164–182. arXiv:1107.1313v2, 2012.

[9] T. Chinburg and R. Rumely, *The capacity pairing*, J. Reine Angew. Math. **434** (1993), 1–44.

[10] D. Eisenbud and J. Harris, *Existence, decomposition, and limits of certain Weierstrass points*, Invent. Math. **87** (1987), 495–515.

[11] R. M. Foster, *The average impedance of an electrical network*, in Reissner Anniversary Volume, Contributions to Applied Mechanics, pages 333–340, J. W. Edwards, Ann Arbor, MI, 1949.

[12] A. Gathmann and M. Kerber, *A Riemann-Roch theorem in tropical geometry*, Math. Z. **259** (2008), no. 1, 217–230.

[13] G. Mikhalkin and I. Zharkov, *Tropical curves, their Jacobians and theta functions*, preprint, arXiv:0612267v2 (2007).

[14] R. Miranda, *Algebraic Curves and Riemann Surfaces*, Graduate Studies in Mathematics vol. 5, American Mathematical Society, Providence, RI, 1995.

[15] D. B. Mumford, *Curves and their Jacobians*, The University of Michigan Press, Ann Arbor, MI, 1977.

[16] A. Neeman, *The distribution of Weierstrass points on a compact Riemann surface*, Ann. of Math. (2) **120** (1984), no. 2, 317–328,

[17] S. Zhang, *Admissible pairing on a curve*, Invent. Math. **112** (1993), no. 1, 171–193.