

Degeneration of Linear Series from the Tropical Point of View and Applications

Matthew Baker and David Jensen

Abstract We discuss linear series on tropical curves and their relation to classical algebraic geometry, describe the main techniques of the subject, and survey some of the recent major developments in the field, with an emphasis on applications to problems in Brill–Noether theory and arithmetic geometry.

Keywords Chip firing • Tropical curves • Brill–Noether theory

1 Introduction

Algebraic curves play a central role in the field of algebraic geometry. Over the past century, curves have been the focus of a significant amount of research, and despite being some of the most well-understood algebraic varieties, there are still many important open questions. The goal of classical Brill–Noether theory is to study the geometry of a curve C by examining all of its maps to projective space, or equivalently the existence and behavior of all line bundles on C . Thus, we have classical results such as Max Noether’s Theorem [12, p. 117] and the Enriques–Babbage Theorem [13] that relate the presence of linear series on a curve to its geometric properties. A major change in perspective occurred during the twentieth century, as the field shifted from studying *fixed* to *general* curves—that is, general points in the moduli space of curves M_g . Many of the major results in the field, such as the Brill–Noether [66] and Gieseker–Petri [65] theorems, remained open for nearly a century as they awaited this new point of view.

A major milestone in the geometry of general curves was the development of limit linear series by Eisenbud and Harris [58]. This theory allows one to study linear series on general curves by studying one-parameter degenerations where the central fiber is a special kind of singular curve, known as a curve of compact type.

M. Baker (✉)

School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160, USA
e-mail: mbaker@math.gatech.edu

D. Jensen

Department of Mathematics, University of Kentucky, Lexington, KY 40506-0027, USA
e-mail: dave.jensen@uky.edu

One property of curves of compact type is that if they have positive genus then they must have components of positive genus. Shortly after the development of limit linear series, many researchers became interested in a different type of nodal curve, which has only rational components, and where the interesting geometric data is encoded by the combinatorics of how the components meet each other. Early examples using the so-called graph curves to establish properties of general curves include Bayer and Eisenbud's work on Green's conjecture for the general curve [28], and the Ciliberto–Harris–Miranda result on the surjectivity of the Wahl map [47].

Much like the theory of limit linear series does for curves of compact type, the recent development of tropical Brill–Noether theory provides a systematic approach to this kind of degeneration argument [5, 17, 19]. A major goal of this survey is to introduce the basic definitions and techniques of this theory, as well as describing some recent applications to the geometry of general curves and the behavior of Weierstrass points in degenerating families. Degeneration arguments also play a major role in arithmetic geometry, and we also survey how linear series on tropical curves can be used to study rational points on curves.

Here are just a few of the interesting theorems which have been proved in recent years with the aid of the theory of linear series on tropical curves.

1. *The Maximal Rank Conjecture for quadrics.* In [79], Jensen and Payne prove that for fixed g, r , and d , if C is a general curve of genus g and $V \subset \mathcal{L}(D)$ is a general linear series on C of rank r and degree d , then the multiplication map $\mu_2 : \text{Sym}^2 V \rightarrow \mathcal{L}(2D)$ is either injective or surjective.
2. *Uniform boundedness for rational points of curves of small Mordell–Weil rank.* In [81], Katz, Rabinoff, and Zureick–Brown prove that if C/\mathbb{Q} is a curve of genus g with Mordell–Weil rank at most $g - 3$, then $\#C(\mathbb{Q}) \leq 76g^2 - 82g + 22$. This is the first such bound depending only on the genus of C .
3. *Non-Archimedean equidistribution of Weierstrass points.* In [4], Amini proves that if C is an algebraic curve over $\mathbb{C}((t))$ and L is an ample line bundle on C , then the Weierstrass points of $L^{\otimes n}$ become equidistributed with respect to the Zhang measure on the Berkovich analytic space C^{an} as n goes to infinity. This gives precise asymptotic information on the limiting behavior of Weierstrass points in degenerating one-parameter families.
4. *Mnev universality for the lifting problem for divisors on graphs.* In [40], Cartwright shows that if X is a scheme of finite type over $\text{Spec } \mathbb{Z}$, there exist a graph G and a rank 2 divisor D_0 on G such that, for any infinite field k , there are a curve C over $k((t))$ and a rank 2 divisor D on C tropicalizing to G and D_0 , respectively, if and only if X has a k -point.

We will discuss the proofs of these and other results after going through the foundations of the basic theory. To accommodate readers with various interests, this survey is divided into three parts. The first part covers the basics of tropical Brill–Noether theory, with an emphasis on combinatorial aspects and the relation to classical algebraic geometry. The second part covers more advanced topics, including the non-Archimedean Berkovich space perspective, tropical moduli spaces, and the theory of metrized complexes. Each of these topics is an important part of the

theory, but is not strictly necessary for many of the applications discussed in Part 3, and the casual reader may wish to skip Part 2 on the first pass. The final part covers applications of tropical Brill–Noether theory to problems in algebraic and arithmetic geometry. For the most part, the sections in Part 3 are largely independent of each other and can be read in any order. Aside from a few technicalities, the reader can expect to follow the applications in Sects. 9, 10, and 11, as well as most of Sect. 12, without reading Part 2.

Part 1: Introductory Topics

2 Jacobians of Finite Graphs

2.1 Degeneration of Line Bundles in One-Parameter Families

A recurring theme in the theory of linear series on curves is that it is very important to understand the behavior of line bundles on generically smooth one-parameter families of curves. One of the key facts about such families is the semistable reduction theorem, which asserts that after a finite base change, one can guarantee that the singularities of the family are “as nice as possible,” i.e., the total space is regular and the fibers are reduced and have only nodal singularities (see [72, Chap. 3C] or [94, Sect. 10.4]). The questions we want to answer about such families are all local on the base, so it is convenient to consider the following setup. Let R be a discrete valuation ring with field of fractions K and algebraically closed residue field κ , let C be a smooth proper and geometrically integral curve over K , and let \mathcal{C} be a *regular strongly semistable model* for C , that is, a proper flat R -scheme with general fiber C satisfying:

1. The total space \mathcal{C} is regular.
2. The central fiber C_0 of \mathcal{C} is *strongly semistable*, i.e., the irreducible components of C_0 are all smooth and C_0 has only nodes as singularities.¹

By the semistable reduction theorem, a regular strongly semistable model for C always exists after passing to a finite extension of K .

Let L be a line bundle on the general fiber C . Because the total space \mathcal{C} is regular, there exists an extension \mathcal{L} of L to the family \mathcal{C} . One can easily see, however, that this extension is not unique—one can obtain other extensions by twisting the components of the central fiber. More concretely, if \mathcal{L} is an extension of L and $Y \subset C_0$ is an irreducible component of the central fiber, then $\mathcal{L}(Y) = \mathcal{L} \otimes \mathcal{O}_{\mathcal{C}}(Y)$ is also an extension of L .

¹We say C_0 is *semistable* if it is reduced and has only nodes as singularities. The term *strongly semistable* is not completely standard.

To understand the effect of the twisting operation, we consider the *dual graph* of the central fiber C_0 . One constructs this graph by first assigning a vertex v_Z to each irreducible component Z of C_0 , and then drawing an edge between two vertices for every node at which the corresponding components intersect.

Example 2.1. If C_0 is a union of m general lines in \mathbb{P}^2 , its dual graph is the complete graph on m vertices. Indeed, every pair of lines meets in one point, so between every pair of vertices in the dual graph, there must be an edge.

Example 2.2. If $C_0 \subset \mathbb{P}^1 \times \mathbb{P}^1$ is a union of a lines in one ruling and b lines in the other ruling, then the dual graph is the complete bipartite graph $K_{a,b}$. This is because a pair of lines in the same ruling do not intersect, whereas a pair of lines in opposite rulings intersect in one point.

Example 2.3. Let C_0 be the union of the -1 curves on a del Pezzo surface of degree 5. In this case, the dual graph of C_0 is the well-known *Petersen graph*.

2.2 Divisors and Linear Equivalence on Graphs

In this paper, by a *graph* we will always mean a finite connected graph which is allowed to have multiple edges between pairs of vertices but is not allowed to have any loop edges. Given a graph G , we write $\text{Div}(G)$ for the free abelian group on the vertices of G . An element D of $\text{Div}(G)$ is called a *divisor* on G , and is written as a formal sum

$$D = \sum_{v \in V(G)} D(v)v,$$

where the coefficients $D(v)$ are integers. The *degree* of a divisor D is defined to be the sum

$$\text{deg}(D) = \sum_{v \in V(G)} D(v).$$

Returning now to our family of curves, let us fix for a moment an extension \mathcal{L} of our line bundle L . We define a corresponding divisor $\text{mdeg}(\mathcal{L})$ on G , called the *multidegree* of \mathcal{L} , by the formula

$$\text{mdeg}(\mathcal{L}) = \sum_Z (\text{deg } \mathcal{L}|_Z) v_Z,$$

where the sum is over all irreducible components Z of C_0 . The quantity $\text{deg } \mathcal{L}|_Z$ can also be interpreted as the intersection multiplicity of \mathcal{L} with Z considered as a (vertical) divisor on the surface \mathcal{C} . Note that the degree of $\text{mdeg}(\mathcal{L})$ is equal to $\text{deg}(L)$.

We now ask ourselves how $\text{mdeg}(\mathcal{L})$ changes if we replace \mathcal{L} by a different extension. Since any two such extensions differ by a sequence of twists by components of the central fiber, it suffices to study the effect of twisting by one such component Y . Given a vertex v of a graph, let $\text{val}(v)$ denote its valence. As the central fiber C_0 is a principal divisor on \mathcal{C} , we have

$$\text{mdeg}(\mathcal{L}(Y)) = \text{mdeg}(\mathcal{L}) + \sum_Z (Y \cdot Z) v_Z$$

with

$$Y \cdot Z = \begin{cases} -\text{val}(v_Y) & \text{if } Z = Y \\ |Z \cap Y| & \text{if } Z \neq Y \end{cases}.$$

The corresponding operation on the dual graph is known as a *chip-firing move*. This is because we may think of a divisor on the graph as a configuration of chips (and anti-chips) on the vertices. In this language, the effect of a chip-firing move is that a vertex v “fires” one chip along each of the edges emanating from v . This decreases the number of chips at v by the valence of v , and increases the number of chips at each of the neighbors w of v by the number of edges between v and w .

Motivated by these observations, we say that two divisors D and D' on a graph G are *equivalent*, and we write $D \sim D'$, if one can be obtained from the other by a sequence of chip-firing moves. We define the *Picard group* $\text{Pic}(G)$ of G to be the group of divisors on G modulo equivalence. Note that the degree of a divisor is invariant under chip-firing moves, so there is a well-defined homomorphism

$$\text{deg} : \text{Pic}(G) \rightarrow \mathbb{Z}.$$

We will refer to the kernel $\text{Pic}^0(G)$ of this map as the *Jacobian* $\text{Jac}(G)$ of the graph G . This finite abelian group goes by many different names in the mathematical literature—in combinatorics, it is commonly referred to as the *sandpile group* or the *critical group* of G (see, for example, [30, 56, 92, 109]).

Remark 2.4. There is a tremendous amount of combinatorial literature concerning Jacobians of graphs. As our focus is on applications in algebraic geometry, however, we will not go into many details here—we refer the reader to [19, 76] and the references therein. We cannot resist mentioning one remarkable fact, however: the cardinality of $\text{Jac}(G)$ is the number of *spanning trees* in G . (This is actually a disguised form of Kirchhoff’s celebrated *Matrix-Tree Theorem*.)

Our discussion shows that for any two extensions of the line bundle L , the corresponding multidegrees are equivalent divisors on the dual graph G . There is therefore a well-defined degree-preserving map

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(G),$$

which we refer to as the *specialization* or *tropicalization* map.

For later reference, we note that there is a natural homomorphism $\rho : \text{Div}(C) \rightarrow \text{Div}(G)$ defined by setting $\rho(D) = \text{mdeg}(D)$, where \mathcal{D} is the Zariski closure of D in \mathcal{C} . The map ρ takes principal (resp., effective) divisors to principal (resp., effective) divisors, and the map $\text{Jac}(C) \rightarrow \text{Jac}(G)$ induced by ρ coincides with Trop (cf. [17, Sect. 2.1]).

It is useful to reformulate the definition of equivalence of divisors on G as follows. For any function $f : V(G) \rightarrow \mathbb{Z}$, we define

$$\text{ord}_v(f) := \sum_{w \text{ adjacent to } v} f(v) - f(w)$$

and

$$\text{div}(f) := \sum_{v \in V(G)} \text{ord}_v(f)v.$$

Divisors of the form $\text{div}(f)$ are known as *principal* divisors, and two divisors are equivalent if and only if their difference is principal. The reason is that the divisor $\text{div}(f)$ can be obtained, starting with the 0 divisor, by firing each vertex v exactly $f(v)$ times.

2.3 Limit Linear Series and Néron Models

The Eisenbud–Harris theory of limit linear series focuses on the case where the dual graph of C_0 is a tree. In this case, the curve C_0 is said to be of *compact type*.² Although they would not have stated it this way, a key insight of the Eisenbud–Harris theory is that the Jacobian of a tree is trivial. Given a line bundle \mathcal{L} of degree d on \mathcal{C} , if the dual graph of C_0 is a tree, then one can repeatedly twist to obtain a line bundle with any degree distribution summing to d on the components of the central fiber. In particular, given a component $Y \subset C_0$, there exists a twist \mathcal{L}_Y such that

$$\text{deg } \mathcal{L}_Y|_Z = \begin{cases} d & \text{if } Z = Y \\ 0 & \text{if } Z \neq Y \end{cases}.$$

This observation is the jumping-off point for the basic theory of limit linear series.

At the other end of the spectrum is the *maximally degenerate* case where all of the components of C_0 have genus 0. In this case, the first Betti number of the dual graph G (which we refer to from now on as the *genus* of G) is equal to the

²The reason for the name *compact type* is that the Jacobian of a nodal curve C_0 is an extension of the Jacobian of the normalization of C_0 by a torus of dimension equal to the first Betti number of the dual graph. It follows that the Jacobian of C_0 is compact if and only if its dual graph is a tree.

geometric genus of the curve and essentially all of the interesting information about degenerations of line bundles is combinatorial. At its core, tropical Brill–Noether theory studies the behavior of line bundles on the curve C using the combinatorics of their specializations to the graph G .

This discussion can also be understood in the context of Néron models. An important theorem of Raynaud [111] asserts (in the language of this paper) that the group Φ of connected components of the special fiber $\tilde{\mathcal{J}}$ of the Néron model of $J = \text{Jac}(C)$ is canonically isomorphic to $\text{Jac}(G)$, where G is the dual graph of the special fiber of \mathcal{C} . This result can be summarized by the commutativity of the following diagram:

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \text{Prin}(C) & \longrightarrow & \text{Div}^0(C) & \longrightarrow & J(K) \longrightarrow 0 \\
 & & \downarrow \rho & & \downarrow \rho & & \downarrow \\
 0 & \longrightarrow & \text{Prin}(G) & \longrightarrow & \text{Div}^0(G) & \longrightarrow & \text{Jac}(G) = \Phi \longrightarrow 0
 \end{array} \tag{1}$$

where the right vertical arrow is the canonical quotient map

$$J(K) \rightarrow J(K)/\mathcal{J}^0(R) = \Phi.$$

The tropical approach to Brill–Noether theory and the approach via the theory of limit linear series are in some sense orthogonal. The former utilizes the component group Φ , or its analytic counterpart, the tropical Jacobian, whereas classical limit linear series are defined only when Φ is trivial. On the other hand, limit linear series involve computations in the compact part of $\tilde{\mathcal{J}}$, which is trivial in the maximally degenerate case.

Recent developments have led to a sort of hybrid of tropical and limit linear series that can be used to study degenerations of line bundles to arbitrary nodal curves. We will discuss these ideas in Sect. 8.

3 Jacobians of Metric Graphs

3.1 Behavior of Dual Graphs Under Base Change

The field of fractions of a DVR is never algebraically closed. For many applications, we will be interested in $\text{Pic}(C_{\bar{K}})$ rather than $\text{Pic}(C_K)$, and we must therefore study the behavior of the specialization map under base change.

Let K' be a finite extension of K , let R' be the valuation ring of K' , and let $C_{K'} = C \times_K K'$. An important issue is that the new total space $\mathcal{C}_{K'} = \mathcal{C} \times_K K'$ may not be regular; it can pick up singularities at the nodes of the central fiber. More specifically, if a point z on \mathcal{C} corresponding to a node of the central fiber has a local analytic equation of the form $xy = \pi$, where π is a uniformizer for

R , then a local analytic equation for z over R' will be $xy = (\pi')^e$, where π' is a uniformizer for R' and e is the ramification index of the extension K'/K . A standard computation shows that we can resolve such a singularity by a chain of $e - 1$ blowups. Repeating this procedure for each singular point of the special fiber, we obtain a regular strongly semistable model \mathcal{C}' for $C_{K'}$. The dual graph G' of the central fiber of \mathcal{C}' is obtained by subdividing each edge of the original dual graph G $e - 1$ times. In other words, if we assign a length of 1 to each edge of G , and a length of $\frac{1}{e}$ to each edge of G' , then G and G' are isomorphic as *metric graphs*.

A *metric graph* Γ is, roughly speaking, a finite graph G in which each edge e has been identified with a real interval I_e of some specified length $\ell_e > 0$. The points of Γ are the vertices of G together with all points in the relative interiors of the intervals I_e . More precisely, a metric graph is an equivalence class of finite edge-weighted graphs, where two weighted graphs G and G' give rise to the same metric graphs G and G' if they have a common length-preserving refinement. A finite weighted graph G representing the equivalence class of Γ is called a *model* for Γ .

Example 3.1. Let $K = \mathbb{C}((t))$, and consider the family $\mathcal{C} : xy = tz^2$ of smooth conics degenerating to a singular conic in \mathbb{P}^2 . The dual graph G of the central fiber consists of two vertices v, v' connected by a single edge. Let $D = P + Q$ be the divisor on the general fiber cut out by the line $y = x$. In homogeneous coordinates, we have $P = (\sqrt{t} : \sqrt{t} : 1)$ and $Q = (-\sqrt{t} : -\sqrt{t} : 1)$. Although the divisor D itself is K -rational, P and Q are not, and both points specialize to the node of the special fiber. It is not hard to check that $\rho(D) = v + v' \in \text{Div}(G)$. If $K' = \mathbb{C}((\sqrt{t}))$, then the total space of the family $\mathcal{C} \times_K K'$ has a singularity at the node of the central fiber. This singularity can be resolved by blowing up the node, and the dual graph G' of the new central fiber is a chain of 3 vertices connected by two edges, with the new vertex v'' corresponding to the exceptional divisor of the blowup (see Fig. 1). The base change D' of D to K' specializes to a sum of two smooth points on the exceptional divisor, and in particular $\rho_{\mathcal{C}'}(D') = 2v''$ is not the image of $\rho_{\mathcal{C}}(D)$ with respect to the natural inclusion map $\text{Div}(G) \hookrightarrow \text{Div}(G')$. Note, however, that $\rho_{\mathcal{C}'}(D')$ and $\rho_{\mathcal{C}}(D)$ are linearly equivalent on G' (this turns out to be a general phenomenon).

Remark 3.2. A similar example, with a semistable family of curves of genus 3, is given in [17, Sect. 4.4].




Fig. 1 The dual graph of the central fiber in Example 3.1 initially (*on the left*), and after base change followed by resolution of the singularity (*on the right*). If we give the segment *on the left* a length of 1 and each of the segments *on the right* a length of 1/2, then both weighted graphs are models for the same underlying metric graph Γ , which is a closed segment of length 1

There are two possible ways to address the lack of functoriality with respect to base change illustrated in Example 3.1. One is to only consider the induced maps on Picard groups, rather than divisors. The other is to replace $\rho : \text{Div}(C) \rightarrow \text{Div}(G)$ with a map $\text{Div}(C_{\bar{K}}) \rightarrow \text{Div}(\Gamma)$, where Γ is the metric graph underlying G , defined by first base-changing to an extension K' over which all points in the support of D are rational and then applying ρ . The most natural way to handle these simultaneous base changes and prove theorems about the resulting map is to work on the Berkovich analytification of C ; see Sect. 6 for details.

3.2 Divisors and Linear Equivalence on Metric Graphs

Let Γ be a metric graph. A *divisor* D on Γ is a formal linear combination

$$D = \sum_{v \in \Gamma} D(v)v$$

with $D(v) \in \mathbb{Z}$ for all $v \in \Gamma$ and $D(v) = 0$ for all but finitely many $v \in \Gamma$. Let $\text{PL}(\Gamma)$ denote the set of continuous, piecewise linear functions $f : \Gamma \rightarrow \mathbb{R}$ with integer slopes. The *order* $\text{ord}_v(f)$ of a function f at a point $v \in \Gamma$ is the sum of its incoming³ slopes along the edges containing v . As in the case of finite graphs, we write

$$\text{div}(f) := \sum_{v \in \Gamma} \text{ord}_v(f)v.$$

A divisor is said to be *principal* if it is of the form $\text{div}(f)$ for some $f \in \text{PL}(\Gamma)$, and two divisors D, D' are *equivalent* if $D - D'$ is principal. We let $\text{Prin}(\Gamma)$ denote the subgroup of $\text{Div}^0(\Gamma)$ (the group of degree-zero divisors on Γ) consisting of principal divisors. By analogy with the case of finite graphs, the group

$$\text{Jac}(\Gamma) := \text{Div}^0(\Gamma)/\text{Prin}(\Gamma)$$

is called the (tropical) *Jacobian* of Γ .

Example 3.3. The Jacobian of a metric tree Γ is trivial. To see this, note that given any two points $P, Q \in \Gamma$, there is a unique path from P to Q . One can construct a continuous, piecewise linear function f that has slope 1 along this path, and has slope 0 everywhere else. We then see that $\text{div}(f) = Q - P$, so any two points on the tree are equivalent (Fig. 2).

Example 3.4. A circle Γ is a torsor for its own Jacobian. To see this, fix a point $O \in \Gamma$. Given two points $P, Q \in \Gamma$, there exists a continuous, piecewise linear function f that has slope 1 on the interval from O and P , slope -1 on the interval from some 4th

³As a caution, some authors use the opposite sign convention.

Fig. 2 Slopes of a function f on a tree with $\text{div}(f) = Q - P$





Fig. 3 A function f on the circle with $\text{div}(f) = O + Q - (P + R)$



point R to Q , and slope 0 everywhere else. We then have $\text{div}(f) = O + Q - (P + R)$, so $P - Q \sim R - O$. It follows that every divisor of degree zero is equivalent to a divisor of the form $R - O$ for some point R . It is not difficult to see that the point R is in fact unique (Fig. 3).

Note that if two divisors are equivalent in the finite graph G , then they are also equivalent in the corresponding metric graph Γ , called the *regular realization* of G , in which every edge of G is assigned a length of 1. It follows that there is a natural inclusion $\iota : \text{Pic}(G) \hookrightarrow \text{Pic}(\Gamma)$. As we saw above, the multidegree of a line bundle L on $C_{\bar{K}}$ can be identified with a divisor on some subdivision of G , which can in turn be identified with a divisor on Γ . One can show that this yields a well-defined map

$$\text{Trop} : \text{Pic}(C_{\bar{K}}) \rightarrow \text{Pic}(\Gamma)$$

whose restriction to $\text{Pic}(C)$ coincides with the previously defined map

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(G)$$

via the inclusion ι .

One can see from this construction that principal divisors on C specialize to principal divisors on Γ . More precisely, there is a natural way to define a map

$$\text{trop} : \bar{K}(C)^* \rightarrow \text{PL}(\Gamma)$$

on rational functions such that

$$\text{Trop}(\text{div}(f)) = \text{div}(\text{trop}(f))$$

for every $f \in \bar{K}(C)^*$. This is known as the *Slope Formula*, cf. Theorem 6.4 below. We refer the reader to Sect. 6.3 for a formal definition of the map trop .

As already discussed in the remarks following Example 3.1, there is also a natural way to define a map $\text{Div}(C_{\overline{K}}) \rightarrow \text{Div}(\Gamma)$ which induces the map $\text{Trop} : \text{Pic}(C_{\overline{K}}) \rightarrow \text{Pic}(\Gamma)$. As with the map trop on rational functions, the map Trop on divisors is most conveniently described using Berkovich’s theory of non-Archimedean analytic spaces, and we defer a detailed discussion to Sect. 6.

3.3 The Tropical Abel–Jacobi Map

A 1-form on a graph G is an element of the real vector space generated by the formal symbols de , as e ranges over the oriented edges of G , subject to the relations that if e, e' represent the same edge with opposite orientations then $de' = -de$. After fixing an orientation on G , a 1-form $\omega = \sum \omega_e de$ is called *harmonic* if, for all vertices v , the sum $\sum \omega_e$ over the outgoing edges at v is equal to 0. Denote by $\Omega(G)$ the space of harmonic 1-forms on G . It is well-known that $\Omega(G)$ is a real vector space of dimension equal to the *genus* g of G , which can also be defined combinatorially as the number of edges of G minus the number of vertices of G plus one, or topologically as the dimension of $H_1(G, \mathbb{R})$.

If G and G' are models for the same metric graph Γ , then $\Omega(G')$ is canonically isomorphic to $\Omega(G)$. We may therefore define the space $\Omega(\Gamma)$ of harmonic 1-forms on Γ as $\Omega(G)$ for any weighted graph model G . (We also define the *genus* of a metric graph Γ to be the genus of any model for Γ .) Given an isometric path $\gamma : [a, b] \rightarrow \Gamma$, any harmonic 1-form ω on Γ pulls back to a classical 1-form on the interval, and we can thus define the integral $\int_{\gamma} \omega$. Note that the definition of a harmonic 1-form does not depend on the metric, but the integral $\int_{\gamma} \omega$ does.

Fix a base point $v_0 \in \Gamma$. For any point $v \in \Gamma$, the integral

$$\int_{v_0}^v \omega$$

is well-defined up to a choice of path from v_0 to v . This gives a map


$$AJ_{v_0} : \Gamma \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

known as the *tropical Abel–Jacobi map*. Extending linearly to $\text{Div}(\Gamma)$ and then restricting to $\text{Div}^0(\Gamma)$, we obtain a map

$$AJ : \text{Div}^0(\Gamma) \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

which does not depend on the choice of a base point. As in the classical case of Riemann surfaces, we have (cf. [101]):

Fig. 4 Two metric graphs Γ and Γ' of genus 2



Tropical Abel–Jacobi Theorem. *The map AJ is surjective and its kernel is precisely $\text{Prin}(\Gamma)$. Thus there is a canonical isomorphism*

$$AJ : \text{Div}^0(\Gamma)/\text{Prin}(\Gamma) \cong \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

between the Jacobian of Γ and a g -dimensional real torus.

Example 3.5. We consider the two metric graphs of genus 2 pictured in Fig. 4. In the first case, we can choose a basis ω_1, ω_2 of harmonic 1-forms by assigning the integer 1 to one of the (oriented) loops, and the integer 0 to the other. We let η_1, η_2 be the elements of the dual basis. We then see from the tropical Abel–Jacobi theorem that

$$\text{Jac}(\Gamma) \cong \mathbb{R}^2/(\mathbb{Z}\eta_1 + \mathbb{Z}\eta_2).$$

A similar argument in the second case yields

$$\text{Jac}(\Gamma') \cong \mathbb{R}^2/ \left(\mathbb{Z}[\eta_1 + \frac{1}{2}\eta_2] + \mathbb{Z}[\frac{1}{2}\eta_1 + \eta_2] \right).$$

Although the Jacobians of the metric graphs Γ and Γ' from Example 3.5 are isomorphic as abstract real tori, they are non-isomorphic as *principally polarized* real tori in the sense of Mikhalkin and Zharkov [101]. In fact, there is an analogue of the Torelli theorem in this context saying that up to certain “Whitney flips,” the Jacobian as a principally polarized real torus determines the metric graph Γ ; see [36]. The map $AJ_{v_0} : \Gamma \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$ is *harmonic* (or *balanced*) in a certain natural sense; see, e.g., [18, Theorem 4.1]. This fact is used in the paper [81], about which we will say more in Sect. 11.2. Basic properties of the tropical Abel–Jacobi map are also used in important ways in [41, 52].


4 Ranks of Divisors

4.1 Linear Systems

By analogy with algebraic curves, a divisor $D = \sum D(v)v$ on a (metric) graph is called *effective* if $D(v) \geq 0$ for all v , and we write $D \geq 0$. The complete linear series of a divisor D is defined to be

$$|D| = \{E \geq 0 \mid E \sim D\}.$$

Fig. 5 The linear system $|2v + w|$ is not equidimensional



Similarly, we write

$$R(D) = \{f \in \text{PL}(\Gamma) \mid \text{div}(f) + D \geq 0\}$$

for the set of tropical rational functions with poles along the divisor D .

As explained in [63], the complete linear series $|D|$ has the structure of a compact polyhedral complex. However, this polyhedral complex often fails to be equidimensional, as the following example shows.

Example 4.1. Consider the metric graph pictured in Fig. 5, consisting of two loops attached at a point v . Let $D = 2v + w$, where w is a point on the interior of the first loop. The complete linear system is the union of two tori. The first, which consists of divisors equivalent to D that are supported on the first loop, has dimension two, while the other, which consists of divisors equivalent to D that have positive degree on the second loop, has dimension one.


We now wish to define the rank of a divisor on a graph. As the previous example shows, the appropriate definition should not be the dimension of the linear system $|D|$, considered as a polyhedral complex.⁴ Instead, we note that a line bundle L on an algebraic curve C has rank at least r if and only if, for every collection of r points of C , there is a nonzero section of L that vanishes at those points. This motivates the following definition.

Definition 4.2. Let D be a divisor on a (metric) graph. If D is not equivalent to an effective divisor, we define its rank to be -1 . Otherwise, we define $r(D)$ to be the largest nonnegative integer r such that $|D - E| \neq \emptyset$ for all effective divisors E of degree r .

Example 4.3. Even when a linear series is equidimensional, its dimension may not be equal to the rank of the corresponding divisor. For example, consider the metric graph pictured in Fig. 6, consisting of a loop meeting a line segment in a vertex v . The linear system $|v|$ is 1-dimensional, because v is equivalent to any point on the line segment. The rank of v , however, is 0, because if w lies in the interior of the loop, then v is not equivalent to w .

⁴Another natural idea would be to try to define $r(D)$ as one less than the “dimension” of $R(D)$ considered as a semimodule over the tropical semiring \mathbf{T} consisting of $\mathbb{R} \cup \{\infty\}$ together with the operations of min and plus. However, this approach also faces significant difficulties. See [73] for a detailed discussion of the tropical semimodule structure on $R(D)$.

Fig. 6 The vertex v moves in a one-dimensional family, but has rank zero



Remark 4.4. There are several other notions of rank in the literature related to our setup of a line bundle on a degenerating family of curves. We mention, for example, the generalized rank functions of Katz and Zureick-Brown [80] and the algebraic rank of Caporaso [35]. The rank as defined here is sometimes referred to as the *combinatorial rank* to distinguish it from these other invariants.

As in the case of curves, we write

$$W_d^r(\Gamma) := \{D \in \text{Pic}^d(\Gamma) \mid r(D) \geq r\}.$$

We similarly define the *gonality* of a graph to be the smallest degree of a divisor of rank at least one. The *Clifford index* of the graph is

$$\text{Cliff}(\Gamma) := \min\{\deg(D) - 2r(D) \mid r(D) > \max\{0, \deg(D) - g + 1\}\}.$$

Remark 4.5. The definition of gonality above is sometimes called the *divisorial gonality*, to distinguish it from the *stable gonality*, which is the smallest degree of a harmonic morphism from a modification of the given metric graph to a tree. The divisorial gonality is always less than or equal to the stable gonality, see, e.g., [9].

4.2 Specialization

One of the key properties of the combinatorial rank is its behavior under specialization. Note that the specialization map takes effective line bundles to effective divisors. Combining this with the fact that it takes principal divisors to principal divisors, we see that, for any divisor D on C , we have

$$\text{Trop } |D| \subseteq |\text{Trop}(D)| \text{ and}$$

$$\text{trop } \mathcal{L}(D) \subseteq R(\text{Trop}(D)).$$

Combining this with the definition of rank yields the following semicontinuity result.

Specialization Theorem. [17] *Let D be a divisor on C . Then*

$$r(D) \leq r(\text{Trop}(D)).$$

Another way of stating this is that $\text{Trop}(W_d^r(C)) \subseteq W_d^r(\Gamma)$. The power of the Specialization Theorem lies in the fact that the rank of the divisor D is an algebro-geometric invariant, whereas the rank of $\text{Trop}(D)$ is a combinatorial invariant. We can therefore use techniques from each field to inform the other. For example, an immediate consequence of specialization is the following fact.

Theorem 4.6. *Let Γ be a metric graph of genus g , and let d, r be positive integers such that $g - (r + 1)(g - d + r) \geq 0$. Then $W_d^r(\Gamma) \neq \emptyset$.*

Proof. As we will see in Corollary 7.2, there exist a curve C over a discretely valued field K , and a semistable R -model \mathcal{C} of C such that the dual graph of the central fiber is isometric to Γ . A well-known theorem of Kempf and Kleiman–Laksov [84, 86] asserts that $W_d^r(C) \neq \emptyset$. It follows from Theorem 4.2 that $W_d^r(\Gamma) \neq \emptyset$ as well.

Corollary 4.7. *A metric graph of genus g has gonality at most $\lceil \frac{g+2}{2} \rceil$.*

Remark 4.8. We are unaware of a purely combinatorial proof of Theorem 4.6. There are many reasons that such a proof would be of independent interest. For example, a combinatorial proof could shed some light on whether the analogous statement is true for finite graphs, as conjectured in [17, Conjectures 3.10, 3.14].⁵

4.3 Riemann–Roch

Another key property of the combinatorial rank is that it satisfies a tropical analogue of the Riemann–Roch theorem:

Tropical Riemann–Roch Theorem ([63, 101]). *Let Γ be a metric graph of genus g , and let $K_\Gamma := \sum_{v \in \Gamma} (\text{val}(v) - 2)v$ be the canonical divisor on Γ . Then for every divisor D on Γ ,*

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

The first result of this kind was the discrete analogue of Theorem 4.3 for finite (non-metric) graphs proved in [19]. In the Baker–Norine Riemann–Roch theorem, one defines $r(D)$ for a divisor D on a graph G exactly as in Definition 4.2, the subtle difference being that the effective divisor E is restricted to the vertices of G . Gathmann and Kerber [63] showed that one can deduce Theorem 4.3 from the Baker–Norine theorem using a clever approximation argument, whereas Mikhalkin and Zharkov [101] generalized the method of proof from [19] to the metric graph

⁵Sam Payne has pointed out that there is a gap in the proof of Conjectures 3.10 and 3.14 of Baker [17] given in [34]. The claim on page 82 that $W_{d,\phi}^r$ has nonempty fiber over b_0 does not follow from the discussion that precedes it. A priori, the fiber of $\overline{W_{d,\phi}^r}$ over b_0 might be contained in the boundary $\overline{P_\phi^d} \setminus \text{Pic}_\phi^d$ of the compactified relative Picard scheme.

setting. Later on, Hladky–Kral–Norine [75] and Luo [96] proved theorems which imply that one can also deduce Riemann–Roch for graphs from tropical Riemann–Roch. (We will discuss Luo’s theory of rank-determining sets in Sect. 5.2.) So in retrospect, one can say that in some sense the Baker–Norine theorem and the theorem stated above are equivalent.

Baker and Norine’s strategy of proof for Theorem 4.3, as modified by Mikhalkin and Zharkov, is to first show that if \mathcal{O} is an orientation of the graph (i.e., a choice of a head vertex and tail vertex for each edge of G), then

$$D_{\mathcal{O}} := \sum_{v \in V(G)} (\text{indeg}_{\mathcal{O}}(v) - 1)v$$

is a divisor of degree $g - 1$, and this divisor has rank -1 if and only if the orientation \mathcal{O} is acyclic. This fact helps to establish the Riemann–Roch theorem in the case of divisors of degree $g - 1$, which serves as the base case for the more general argument. It is interesting to note that the tropical Riemann–Roch theorem has thus far resisted attempts to prove it via classical algebraic geometry. At present, neither the tropical Riemann–Roch theorem nor the classical Riemann–Roch theorem for algebraic curves is known to imply the other.

If \mathcal{C} is a strongly semistable R -model for a curve C over a discretely valued field K with the property that all irreducible components of the special fiber C_0 have genus 0, then the multidegree of the relative dualizing sheaf $\Omega_{C/R}^1$ is equal to the canonical divisor of the graph G . This is a simple consequence of the adjunction formula, which shows more generally that $\text{mdeg}(\Omega_{C/R}^1) = K_{(G,\omega)}^\#$ in the terminology of Sect. 4.4 below. If K is not discretely valued, this is still true with the right definition of the sheaf $\Omega_{C/R}^1$ (see [81]). This “explains” in some sense why there is a canonical *divisor* on a metric graph while on an algebraic curve there is merely a canonical *divisor class*.

It is clear from the definition of rank that if D and E are divisors on a metric graph Γ having nonnegative rank, then $r(D + E) \geq r(D) + r(E)$. Combining this with tropical Riemann–Roch, one obtains a tropical version of Clifford’s inequality:

Tropical Clifford’s Theorem. *Let D be a special divisor on a metric graph Γ , that is, a divisor such that both D and $K_\Gamma - D$ have nonnegative rank. Then*

$$r(D) \leq \frac{1}{2} \text{deg}(D).$$

Remark 4.9. The classical version of Clifford’s theorem is typically stated in two parts. The first part is the inequality above, while the second part states that, when equality holds, either $D \sim 0$, $D \sim K_C$ or the curve C is hyperelliptic and the linear equivalence class of D is a multiple of the unique g_2^1 . The same conditions for equality hold in the tropical case as well, by a recent theorem of Coppens, but the proof is quite subtle as the classical methods do not work in the tropical context. See [52] for details.




Fig. 7 Two metric graphs of genus 3, the first of which is hyperelliptic, and the second of which is not. (All edges have length 1)

Note that, as in the case of curves, the Riemann–Roch theorem significantly limits the possible ranks that a divisor of fixed degree on a metric graph may have. For example, a divisor of negative degree necessarily has rank -1 , so a divisor of degree $d > 2g - 2$ must have rank $d - g$. It is only in the intermediate range $0 \leq d \leq 2g - 2$ where there are multiple possibilities for the rank.

Example 4.10. The canonical divisor on a circle is trivial, and it is the only divisor of degree 0 with nonnegative rank. If D is a divisor of degree $d > 0$, then by Riemann–Roch D has rank $d - 1$. This can also be seen using the fact that the circle is a torsor for its Jacobian, as in Example 3.4: if E is an effective divisor of degree $d - 1$, then there is a unique point P such that $D - E \sim P$, and hence $r(D) \geq d - 1$.

Example 4.11. The smallest genus for which the rank of an effective divisor is not completely determined by the degree is genus 3. Pictured in Fig. 7 are two examples of genus 3 metric graphs, the first of which is *hyperelliptic*, meaning that it admits a divisor of degree 2 and rank 1, and the second of which is not. For the first graph, one can check by hand that the sum of the two vertices on the left has rank at least 1, and it cannot have rank higher than 1 by Clifford’s theorem. We will show that the second graph is not hyperelliptic in Example 5.4, as the argument will require some techniques for computing ranks of divisors that we will discuss in the next section.

4.4 Divisors on Vertex-Weighted Graphs

In [6], Amini and Caporaso formulate a refinement of the Specialization Theorem which takes into account the genera of the components of the special fiber. In this section we describe their result following the presentation in [5].

A *vertex-weighted metric graph* is a pair (Γ, ω) consisting of a metric graph Γ and a weight function $\omega : \Gamma \rightarrow \mathbb{Z}_{\geq 0}$ such that $\omega(x) = 0$ for all but finitely many $x \in \Gamma$. Following [6], we define a new metric graph $\Gamma^\#$ by attaching $\omega(x)$ loops of arbitrary positive length at each point $x \in \Gamma$. There is a natural inclusion of Γ into $\Gamma^\#$. The *canonical divisor* of (Γ, ω) is defined to be

$$K^\# = K_\Gamma + \sum_{x \in \Gamma} 2\omega(x),$$

which can naturally be identified with the canonical divisor of $K_{\Gamma^\#}$ restricted to Γ . Its degree is $2g^\# - 2$, where $g^\# = g(\Gamma) + \sum_{x \in \Gamma} \omega(x)$ is the genus of $\Gamma^\#$.

Following [6], the *weighted rank* $r^\#$ of a divisor D on Γ is defined to be $r^\#(D) := r_{\Gamma^\#}(D)$. By [5, Corollary 4.12], we have the more intrinsic description

$$r^\#(D) = \min_{0 \leq E \leq \mathcal{W}} (\deg(E) + r_\Gamma(D - 2E)),$$

where $\mathcal{W} = \sum_{x \in \Gamma} \omega(x)(x)$.

The Riemann–Roch theorem for $\Gamma^\#$ implies the following “vertex-weighted” Riemann–Roch theorem for Γ :

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

If C is a curve over K , together with a semistable model \mathcal{C} over R , we define the associated vertex-weighted metric graph (Γ, ω) by taking Γ to be the skeleton of \mathcal{C} and defining the weight function ω by $\omega(v) = g_v$. With this definition, the genus of the weighted metric graph (Γ, ω) is equal to the genus of C and we have $\text{trop}(K_C) \sim K^\#$ on Γ [5, Sect. 4.7.1]. The following weighted version of the Specialization Theorem, inspired by the results of [6], is proved in [5, Theorem 4.13]:

Weighted Specialization Theorem. *For every divisor $D \in \text{Div}(C)$, we have $r_C(D) \leq r^\#(\text{trop}(D))$.*

5 Combinatorial Techniques

The tropical approach to degeneration of line bundles in algebraic geometry derives its power from the combinatorial tools which one has available, many of which have no classical analogues. We describe some of these tools in this section.

5.1 Reduced Divisors and Dhar’s Burning Algorithm

Definition 5.1. Let G be a finite graph. Given a divisor $D \in \text{Div}(G)$ and a vertex v of G , we say that D is *v-reduced* if

- (RD1) $D(w) \geq 0$ for every $w \neq v$ and
- (RD2) for every nonempty set $A \subseteq V(G) \setminus \{v\}$ there is a vertex $w \in A$ such that $\text{outdeg}_A(w) > D(w)$.

Here $\text{outdeg}_A(w)$ denotes the outdegree of w with respect to A , i.e., the number of edges connecting $w \in A$ to a vertex not in A . The following important result (cf. [19, Proposition 3.1]) shows that v -reduced divisors form a distinguished set of representatives for linear equivalence classes of divisors on G :

Lemma 5.2. *Every divisor on G is equivalent to a unique v -reduced divisor.*

If D has nonnegative rank, the v -reduced divisor equivalent to D is the divisor in $|D|$ that is lexicographically “closest” to v . It is a discrete analogue of the unique divisor in a classical linear series $|D|$ with the highest possible order of vanishing at a given point $p \in C$.

There is a simple algorithm for determining whether a given divisor satisfying (RD1) above is v -reduced, known as *Dhar’s burning algorithm*. For $w \neq v$, imagine that there are $D(w)$ buckets of water at w . Now, light a fire at v . The fire starts spreading through the graph, burning through an edge as soon as one of its endpoints is burnt, and burning a vertex w if the number of burnt edges adjacent to w is greater than $D(w)$ (that is, there is not enough water to fight the fire). The divisor D is v -reduced if and only if the fire consumes the whole graph. For a detailed account of this algorithm, we refer to [55] and the more recent [22, Sect. 5.1].

There is a completely analogous set of definitions and results for metric graphs. Let Γ be a metric graph, let D be a divisor on Γ , and choose a point $v \in \Gamma$ (which need not be a vertex). We say that D is *v -reduced* if the two conditions from Definition 5.1 hold, with the second condition replaced by

(RD2’) for every closed, connected, nonempty set $A \subseteq \Gamma \setminus \{v\}$ there is a point $w \in A$ such that $\text{outdeg}_A(w) > D(w)$.

It is not hard to see that condition (RD2’) is equivalent to requiring that for every non-constant tropical rational function $f \in R(\Gamma)$ with a global maximum at v , the divisor $D' := D + \text{div}(f)$ does not satisfy (RD1), i.e., there exists $w \neq v$ in Γ such that $D'(w) < 0$.

The analogue of Lemma 5.2 remains true in the metric graph context: every divisor on Γ is equivalent to a unique v -reduced divisor. Moreover, Dhar’s burning algorithm as formulated above holds almost *verbatim* for metric graphs: the fire starts spreading through Γ , getting blocked at a point $w \in \Gamma$ iff the number of burnt tangent directions at w is less than or equal to $D(w)$; the divisor D is v -reduced if and only if the fire consumes all of Γ .

We note the following important fact, which is useful for computing ranks of divisors.

Lemma 5.3. *Let D be a divisor on a finite or metric graph, and suppose that D is v -reduced for some v . If D has nonnegative rank, then $D(v) \geq r(D)$.*

Example 5.4. Let Γ be the complete graph on 4 vertices endowed with arbitrary edge lengths. We can use the theory of reduced divisors to show that Γ is not hyperelliptic, justifying one of the claims in Example 4.11. Suppose that there exists a divisor D on Γ of degree 2 and rank 1 and choose a vertex v . Since D has rank 1, $D \sim D' := v + v'$ for some $v' \in \Gamma$. Now, let $w \neq v, v'$ be a vertex. Note that there




Fig. 8 Using Dhar's burning algorithm to compute the v_5 -reduced divisor equivalent to $v_1 + v_2$. The burnt vertices after each iteration are colored *white*

are at least two paths from w to v that do not pass through v' , and if $v = v'$, there are three. It follows by Dhar's burning algorithm that D' is w -reduced. But $D'(w) = 0$, contradicting the fact that D (and hence D') has rank 1.

If a given divisor is not v -reduced, Dhar's burning algorithm provides a method for finding the unique equivalent v -reduced divisor. In the case of finite graphs, after performing Dhar's burning algorithm, if we fire the vertices that are left unburnt we obtain a divisor that is lexicographically closer to the v -reduced divisor, and after iterating the procedure a finite number of times, it terminates with the v -reduced divisor (cf. [22]). For metric graphs, there is a similar procedure but with additional subtleties—we refer the interested reader to [96, Algorithm 2.5] and [14].

Example 5.5. Consider the finite graph depicted in Fig. 8, consisting of two triangles meeting at a vertex v_3 . We let $D = v_1 + v_2$, and compute the v_5 -reduced divisor equivalent to D . After performing Dhar's burning algorithm once, we see that vertices v_1 and v_2 are left unburnt. Firing these, we see that D is equivalent to $2v_3$. Performing Dhar's burning algorithm a second time, all three of the vertices v_1, v_2, v_3 are not burnt. Firing these, we obtain the divisor $v_4 + v_5$. A third run of Dhar's burning algorithm shows that $v_4 + v_5$ is v_5 -reduced.

5.2 Rank-Determining Sets

The definition of the rank of a divisor on a finite graph G implies easily that there is an algorithm for computing it.⁶ Indeed, since there are only a finite number of effective divisors E of a given degree on G , we are reduced to the problem of determining whether a given divisor is equivalent to an effective divisor or not. This problem can be solved in polynomial time by using the iterated version of Dhar's algorithm described above to compute the v -reduced divisor D' equivalent to D for some vertex v . If $D'(v) < 0$, then $|D| = \emptyset$, and otherwise $|D| \neq \emptyset$.

⁶Although there is no known *efficient* algorithm; indeed, it is proved in [85] that this problem is NP-hard.

If one attempts to generalize this algorithm to the case of metric graphs, there is an immediate problem, since there are now an infinite number of effective divisors E to test. The idea behind rank-determining sets is that it suffices, in the definition of $r(D)$, to restrict to a finite set of effective divisors E .

Definition 5.6. Let Γ be a metric graph. A subset $A \subseteq \Gamma$ is a *rank-determining set* if, for any divisor D on Γ , D has rank at least r if and only if $|D - E| \neq \emptyset$ for every effective divisor E of degree r supported on A .

In [96], Luo provides a criterion for a subset of a metric graph to be rank-determining. A different proof of Luo’s criterion has been given recently by Backman [15]. Luo defines a *special open set* to be a connected open set $U \subseteq \Gamma$ such that every connected component $X \subseteq \Gamma \setminus U$ contains a boundary point v such that $\text{outdeg}_X(v) \geq 2$.

Theorem 5.7 ([96]). A subset $A \subseteq \Gamma$ is rank-determining if and only if all nonempty special open subsets of Γ intersect A .

Corollary 5.8. Let G be a model for a metric graph Γ . If G has no loops, then the vertices of G are a rank-determining set.

There are lots of other interesting rank-determining sets besides vertices of models.

Example 5.9. Let G be a model for a metric graph Γ . Choose a spanning tree of G and let e_1, \dots, e_g be the edges in the complement of the spanning tree. For each such edge e_i , choose a point v_i in its interior, and let w be any other point of Γ . Then the set $A = \{v_1, \dots, v_g\} \cup \{w\}$ is rank-determining. This construction is used, for example, in [52].

Example 5.10. Let G be a bipartite finite graph, and let Γ be a metric graph having G as a model. If we fix a 2-coloring of the vertices of G , then the vertices of one color are a rank-determining set. This is the key observation in [77], in which the second author shows that the Heawood graph admits a divisor of degree 7 and rank 2, regardless of the choice of edge lengths. The interest in this example arises because it shows that there is a nonempty open subset of the (highest-dimensional component of the) moduli space M_8^{trop} containing no Brill–Noether general metric graph. (See Sect. 7.1 for a description of M_g^{trop} .)

5.3 Tropical Independence

Many interesting questions about algebraic curves concern the ranks of linear maps between the vector spaces $\mathcal{L}(D)$. For example, both the Gieseker–Petri theorem and the Maximal Rank Conjecture are statements about the rank of the multiplication maps

$$\mu : \mathcal{L}(D) \otimes \mathcal{L}(D') \rightarrow \mathcal{L}(D + D')$$

for certain pairs of divisors D, D' on a general curve.

One simple strategy for showing that a map, such as μ , has rank at least k is to carefully choose k elements of the image, and then check that they are linearly independent. To this end, we formulate a notion of *tropical independence*, which gives a sufficient condition for linear independence of rational functions on a curve C in terms of the associated piecewise linear functions on the metric graph Γ .

Definition 5.11 ([78]). A set of piecewise linear functions $\{f_1, \dots, f_k\}$ on a metric graph Γ is *tropically dependent* if there are real numbers b_1, \dots, b_k such that the minimum

$$\min\{f_1(v) + b_1, \dots, f_k(v) + b_k\}$$

occurs at least twice at every point v in Γ .

If there are no such real numbers b_1, \dots, b_k , then we say $\{f_1, \dots, f_k\}$ is *tropically independent*. We note that linearly dependent functions on C specialize to tropically dependent functions on Γ . Although the definition of tropical independence is merely a translation of linear dependence into tropical language, one can often check tropical independence using combinatorial methods. The following lemma illustrates this idea.

Lemma 5.12 ([78]). Let D be a divisor on a metric graph Γ , with f_1, \dots, f_k piecewise linear functions in $R(D)$, and let

$$\theta = \min\{f_1, \dots, f_k\}.$$

Let $\Gamma_j \subset \Gamma$ be the closed set where $\theta = f_j$, and let $v \in \Gamma_j$. Then the support of $\text{div}(\theta) + D$ contains v if and only if v belongs to either

1. the support of $\text{div}(f_j) + D$ or
2. the boundary of Γ_j .

5.4 Break Divisors

Another useful combinatorial tool for studying divisor classes on graphs and metric graphs is provided by the theory of *break divisors*, which was initiated by Mikhalkin–Zharkov in [101] and studied further by An–Baker–Kuperberg–Shokrieh in [10]. Given a metric graph Γ of genus g , fix a model G for Γ . For each spanning tree T of G , let Σ_T be the image of the canonical map

$$\prod_{e \notin T} e \rightarrow \text{Div}_+^g(\Gamma)$$

sending (p_1, \dots, p_g) to $p_1 + \dots + p_g$. (Here $\text{Div}_+^g(\Gamma)$ denotes the set of effective divisors of degree g on Γ and e denotes a *closed* edge of G , so the points p_i are allowed to be vertices of G .) We call $B(\Gamma) := \bigcup_T \Sigma_T$ the set of *break divisors* on Γ . The set of break divisors does not depend on the choice of the model G . The following result shows that the natural map $B(\Gamma) \subset \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ is bijective:

Theorem 5.13 ([10, 101]). *Every divisor of degree g on Γ is linearly equivalent to a unique break divisor.*

Since $B(\Gamma)$ is a compact subset of $\text{Div}_+^g(\Gamma)$ and $\text{Pic}^g(\Gamma)$ is also compact, it follows from general topology that there is a canonical continuous section σ to the natural map $\pi : \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ whose image is precisely the set of break divisors. In particular, every degree g divisor class on a metric graph Γ has a *canonical* effective representative. The analogue of this statement in algebraic geometry is *false*: when $g = 2$, for example, the natural map $\text{Sym}^2(C) = \text{Div}_+^2(C) \rightarrow \text{Pic}^2(C)$ is a birational isomorphism which blows down the \mathbb{P}^1 corresponding to the fiber over the unique \mathfrak{g}_2^1 , and this map has no section. This highlights an interesting difference between the algebraic and tropical settings.


The proof of Theorem 5.13 in [101] utilizes the theory of tropical theta functions and the tropical analogue of Riemann’s theta constant. A purely combinatorial proof based on the theory of *q-connected orientations* is given in [10], and the combinatorial proof yields an interesting analogue of Theorem 5.13 for finite graphs. If G is a finite graph and Γ is its *regular realization*, in which all edges are assigned a length of 1, define the set of *integral break divisors* on G to be $B(G) = B(\Gamma) \cap \text{Div}(G)$. In other words, $B(G)$ consists of all break divisors for the underlying metric graph Γ which are supported on the vertices of G .

Theorem 5.14 ([10]). *Every divisor of degree g on G is linearly equivalent to a unique integral break divisor.*

Since $\text{Pic}^g(G)$ and $\text{Pic}^0(G) = \text{Jac}(G)$ have the same cardinality (the former is naturally a torsor for the latter), it follows from Remark 2.4 that the number of integral break divisors on G is equal to the number of spanning trees of G , though there is in general no canonical bijection between the two. A family of interesting combinatorial bijections is discussed in [23].

If we define $C_T = \pi(\Sigma_T)$, then $\text{Pic}^g(\Gamma) = \bigcup_T C_T$ by Theorem 5.13. It turns out that the relative interior of each cell C_T is (the interior of) a parallelotope, and if $T \neq T'$, then the relative interiors of C_T and $C_{T'}$ are disjoint. Thus $\text{Pic}^g(\Gamma)$ has a polyhedral decomposition depending only on the choice of a model for Γ . The maximal cells in the decomposition correspond naturally to spanning trees, and the minimal cells (i.e., vertices) correspond naturally to integral break divisors, as illustrated in Fig. 9.

Fig. 9 A polyhedral decomposition of $\text{Pic}^2(\Gamma)$ for the metric realization of the graph G obtained by deleting an edge from the complete graph K_4



Remark 5.15. Mumford’s non-Archimedean analytic uniformization theory for degenerating abelian varieties [102], as recently refined by Gubler and translated into the language of tropical geometry [67, 68], shows that if G is the dual graph of the special fiber of a regular R -model \mathcal{C} for a curve C , then the canonical polyhedral decomposition $\{C_T\}$ of $\text{Pic}^g(\Gamma)$ gives rise to a canonical proper R -model for $\text{Pic}^g(C)$. Sam Payne has asked (Payne, 2013, Personal communication) whether (up to the identification of Pic^g with Pic^0) this model coincides with the compactification of the Néron model of $\text{Jac}(C)$ introduced by Caporaso in [32].

Break divisors corresponding to the relative interior of some cell C_T are called *simple break divisors*. They can be characterized as the set of degree g effective divisors D on Γ such that $\Gamma \setminus \text{supp}(D)$ is connected and simply connected. Dhar’s algorithm shows that such divisors are *universally reduced*, i.e., they are q -reduced for all $q \in \Gamma$. A consequence of this observation and the Riemann–Roch theorem for metric graphs is the following result, which is useful in tropical Brill–Noether theory (cf. Sect. 9.2).

Proposition 5.16. *Let Γ be a metric graph and let D be a simple break divisor (or more generally any universally reduced divisor) on Γ . Then D has rank 0 and $K_\Gamma - D$ has rank -1 . Therefore:*

- (1) *The set of divisor classes in $\text{Pic}^g(\Gamma)$ having rank at least 1 is contained in the codimension one skeleton of the polyhedral decomposition $\bigcup_T C_T$.*
- (2) *If T is a spanning tree for some model G of Γ and D, E are effective divisors with $D + E$ linearly equivalent to K_Γ , then there must be an open edge e° in the complement of T such that D has no chips on e° .*

Remark 5.17. The set $B(\Gamma)$ of all break divisors on Γ can be characterized as the topological closure in $\text{Div}_+^g(\Gamma)$ of the set of universally reduced divisors.

Remark 5.18. The real torus $\text{Pic}^g(\Gamma)$ has a natural Riemannian metric. One can compute the volume of $\text{Pic}^g(\Gamma)$ in terms of a matrix determinant associated with G , and the volume of the cell C_T is the product of the lengths of the edges not in T . Comparing the volume of the torus to the sums of the volumes of the cells C_T yields a *dual version* of a weighted form of Kirchhoff’s Matrix-Tree Theorem. See [10] for details.

Part 2: Advanced Topics

We now turn to the more advanced topics of non-Archimedean analysis, tropical moduli spaces, and metrized complexes. Each of these topics plays an important role in tropical Brill–Noether theory, and we would be remiss not to mention them here. We note, however, that most of the applications we discuss in Part 3 do not require these techniques, and the casual reader may wish to skip this part on the first pass.

6 Berkovich Analytic Theory

Rather than considering a curve over a discretely valued field and then examining its behavior under base change, we could instead start with a curve over an algebraically closed field and directly associate a metric graph to it. We do this by making use of Berkovich’s theory of analytic spaces. In addition to being a convenient bookkeeping device for changes in dual graphs and specialization maps under field extensions, Berkovich’s theory also allows for clean formulations of some essential results in the theory of tropical linear series, such as the Slope Formula (Theorem 6.4 below). The theory also furnishes a wealth of powerful tools for understanding the relationship between algebraic curves and their tropicalizations.

6.1 A Quick Introduction to Berkovich Spaces

We let K be a field which is complete with respect to a non-Archimedean valuation

$$\text{val} : K^* \rightarrow \mathbb{R}.$$

We let $R \subset K$ be the valuation ring, κ the residue field, and $|\cdot| = \exp(-\text{val})$ the corresponding norm on K .

A *multiplicative seminorm* on a nonzero ring A is a function $|\cdot| : A \rightarrow \mathbb{R}$ such that for all $x, y \in A$ we have $|0| = 0$, $|1| = 1$, $|xy| = |x| \cdot |y|$, and $|x + y| \leq |x| + |y|$. A multiplicative seminorm is *non-Archimedean* if $|x + y| \leq \max\{|x|, |y|\}$ for all $x, y \in A$, and is a *norm* if $|x| = 0$ implies $x = 0$. If L is a field, a function $|\cdot| : L \rightarrow \mathbb{R}$ is a non-Archimedean norm if and only if $-\log |\cdot| : L \rightarrow \mathbb{R} \cup \{\infty\}$ is a *valuation* in the sense of Krull.

If $X = \text{Spec}(A)$ is an affine scheme over K , we define its *Berkovich analytification* X^{an} to be the set of non-Archimedean multiplicative seminorms on the K -algebra A extending the given absolute value on K , endowed with the weakest topology such that the map $X^{\text{an}} \rightarrow \mathbb{R}$ defined by $|\cdot|_x \mapsto |f|_x$ is continuous for all $f \in A$. One can globalize this construction to give a Berkovich analytification of an arbitrary scheme of finite type over K . As we have defined it, the Berkovich analytification is merely a topological space, but it can be equipped with a structure similar to that of a locally ringed space and one can view X^{an} as an object in a larger category of (not necessarily algebraizable) *Berkovich analytic spaces*. The space X^{an} is locally compact and locally path-connected. It is Hausdorff if and only if X is separated, compact if and only if X is proper, and path-connected if and only if X is connected. We refer the reader to [29, 51] for more background information on Berkovich spaces in general, and [16, 25] for more details in the special case of curves.

There is an alternate perspective on Berkovich spaces which is often useful and highlights the close analogy with schemes. If K is a field, points of an affine K -scheme $\text{Spec}(A)$ can be identified with equivalence classes of pairs (L, ϕ) where L is a field extension of K , $\phi : A \rightarrow L$ is a K -algebra homomorphism, and two pairs (L_1, ϕ_1) and (L_2, ϕ_2) are equivalent if there are embeddings of L_1 and L_2 into a common overfield L' and a homomorphism $\phi' : A \rightarrow L'$ such that the composition $\phi_i : A \rightarrow L_i \rightarrow L$ is ϕ' . Indeed, to a pair (L, ϕ) one can associate the prime ideal $\ker(\phi)$ of A , and to a prime ideal \mathfrak{p} of A one can associate the pair $(K(\mathfrak{p}), \phi)$ where $K(\mathfrak{p})$ is the fraction field of A/\mathfrak{p} and $\phi : A \rightarrow K(\mathfrak{p})$ is the canonical map.

Similarly, if K is a complete valued field, points of $\text{Spec}(A)^{\text{an}}$ can be identified with equivalence classes of pairs (L, ϕ) , where L is a complete valued field extension of K and $\phi : A \rightarrow L$ is a K -algebra homomorphism. The equivalence relation is as before, except that L' should be complete and extend the valuation on the L_i . Indeed, to a pair (L, ϕ) one can associate the multiplicative seminorm $a \mapsto |\phi(a)|$ on A , and to a multiplicative seminorm $|\cdot|_x$ one can associate the pair $(\mathcal{H}(x), \phi)$ where $\mathcal{H}(x)$ is the completion of the fraction field of $A/\ker(|\cdot|_x)$ and $\phi : A \rightarrow \mathcal{H}(x)$ is the canonical map.

We will assume for the rest of this section that K is **algebraically closed and nontrivially valued**. This ensures, for example, that the set $X(K)$ is dense in X^{an} .

If X/K is an irreducible variety, there is a dense subset of X^{an} consisting of the set Val_X of norms⁷ on the function field $K(X)$ that extend the given norm on K . Within the set Val_X , there is a distinguished class of norms corresponding to *divisorial valuations*. By definition, a valuation v on $K(X)$ is *divisorial* if there is an R -model \mathcal{X} for X and an irreducible component Z of the special fiber of \mathcal{X} such that $v(f)$ is equal to the order of vanishing of f along Z . The set of divisorial points is known to be dense in X^{an} .

Remark 6.1. In this survey we have intentionally avoided the traditional perspective of tropical geometry, in which one considers subvarieties of the torus $(K^*)^n$, and the tropicalization is simply the image of coordinatewise valuation. We refer the reader to [98] for a detailed account of this viewpoint on tropical geometry. The Berkovich analytification can be thought of as a sort of intrinsic tropicalization—the one that does not depend on a choice of coordinates. This is reinforced by the result that the Berkovich analytification is the inverse limit of all tropicalizations [62, 108].

6.2 Berkovich Curves and Their Skeletons

If C/K is a complete nonsingular curve, the underlying set of the Berkovich analytic space C^{an} consists of the points of $C(K)$ together with the set Val_C . We write

$$\text{val}_y : K(C)^* \rightarrow \mathbb{R}$$

for the valuation corresponding to a point $y \in \text{Val}_C = C^{\text{an}} \setminus C(K)$. The points in $C(K)$ are called type-1 points, and the remaining points of C^{an} are classified into three more types. We will not define points of type 3 or 4 in this survey article; see, e.g., [24, Sect. 3.5] for a definition. Note, however, that every point of C^{an} becomes a type-1 point after base-changing to a suitably large complete non-Archimedean field extension L/K .

If the residue field of $K(C)$ with respect to val_y has transcendence degree 1 over κ , then y is called a *type-2 point*. These are exactly the points corresponding to divisorial valuations. Because it has transcendence degree 1 over κ , this residue field corresponds to a unique smooth projective curve over κ , which we denote C_y . A *tangent direction* at y is an equivalence class of continuous injections $\gamma : [0, 1] \rightarrow C^{\text{an}}$ sending 0 to y , where $\gamma \sim \gamma'$ if $\gamma([0, 1]) \cap \gamma'([0, 1]) \not\supseteq \{y\}$. Closed points of C_y are in one-to-one correspondence with tangent directions at y in C^{an} .

There is natural metric on the set Val_C which is described in detail in [25, Sect. 5.3]. This metric induces a topology that is much finer than the subspace

⁷Note that there is a one-to-one correspondence between norms on $K(X)$ and valuations on $K(X)$, hence the terminology Val_X . It is often convenient to work with (semi-)valuations rather than (semi-)norms.

topology on $\text{Val}_C \subset C^{\text{an}}$, and with respect to this metric, Val_C is locally an \mathbb{R} -tree⁸ with branching precisely at the type-2 points. The type-1 points should be thought of as infinitely far away from every point of Val_C .

The local \mathbb{R} -tree structure arises in the following way. If \mathcal{C} is an R -model for C and Z is a reduced and irreducible component of the special fiber of \mathcal{C} , then Z corresponds to a type-2 point y_Z of \mathcal{C} . Blowing up a nonsingular closed point of Z (with respect to some choice of a uniformizer $\varpi \in \mathfrak{m}_R$) gives a new point $y_{Z'}$ of C^{an} corresponding to the exceptional divisor Z' of the blowup. We can then blow up a nonsingular closed point on the exceptional divisor Z' to obtain a new point of C^{an} , and so forth. The resulting constellation of points obtained by all such sequences of blowups, and varying over all possible choices of ϖ , is an \mathbb{R} -tree T_Z rooted at y_Z , as pictured in Fig. 10. The distance between the points y_Z and $y_{Z'}$ is $\text{val}(\varpi)$.

A *semistable vertex set* is a finite set of type-2 points whose complement is a disjoint union of finitely many open annuli and infinitely many open balls. There is a one-to-one correspondence between semistable vertex sets and semistable models of C . More specifically, the normalized components of the central fiber of this semistable model are precisely the curves C_y for y in the semistable vertex set, and the preimages of the nodes under specialization are the annuli. The annulus corresponding to a node where C_y meets $C_{y'}$ contains a unique open segment with endpoints y and y' , and its length (with respect to the natural metric on Val_C) is the logarithmic modulus of the annulus. The union of these open segments together with the semistable vertex set is a closed connected metric graph Γ contained in C^{an} , called the *skeleton* of the semistable model \mathcal{C} . If C has genus at least 2, then there is a unique minimal semistable vertex set in C^{an} and a corresponding minimal skeleton.

Fix a semistable model \mathcal{C} of C and a corresponding skeleton $\Gamma = \Gamma_{\mathcal{C}}$. Each connected component of $C^{\text{an}} \setminus \Gamma$ has a unique boundary point in Γ , and there is a canonical retraction to the skeleton


$$\tau : C^{\text{an}} \rightarrow \Gamma$$

taking a connected component of $C^{\text{an}} \setminus \Gamma$ to its boundary point. There is a natural homeomorphism of topological spaces $C^{\text{an}} \cong \varprojlim \Gamma_{\mathcal{C}}$, where the inverse limit is taken over all semistable models \mathcal{C} (cf. [25, Theorem 5.2]).

Example 6.2. Figure 10 depicts the Berkovich analytification of an elliptic curve E/K with non-integral j -invariant j_E . In this case, the skeleton Γ associated with a minimal proper semistable model of R is isometric to a circle with circumference $-\text{val}(j_E)$. There are an infinite number of infinitely branched \mathbb{R} -trees emanating from the circle at each type-2 point of the skeleton. The retraction map takes a point

⁸See [21, Appendix B] for an introduction to the theory of \mathbb{R} -trees. For our purposes, what is most important about \mathbb{R} -trees is that there is a unique path between any two points.

Fig. 10 The skeleton of an elliptic curve with non-integral j -invariant



$x \in E^{\text{an}}$ to the endpoint in Γ of the unique path from x to Γ . The points of $E(K)$ lie “out at infinity” in the picture: they are ends of the \mathbb{R} -trees.

6.3 Tropicalization of Divisors and Functions on Curves

Restricting the retraction map τ to $C(K)$ and extending linearly give the tropicalization map on divisors

$$\text{Trop} : \text{Div}(C) \rightarrow \text{Div}(\Gamma).$$

If K_0 is a discretely valued subfield of K over which C is defined and has semistable reduction, and if D is a divisor on C whose support consists of K_0 -rational points, then the divisor $\text{mdeg}(D)$ on G (identified with a divisor on Γ via the natural inclusion) coincides with the tropicalization $\text{Trop}(D)$ defined via retraction to the skeleton.

Example 6.3. Returning to Example 3.1, in which the metric graph Γ is a closed line segment of length 1, by considering the divisor cut out by $y^a = x^b z^{a-b}$ for positive integers $a > b$ we see that a divisor on $C_{\bar{K}}$ can tropicalize to any rational point on Γ .

Given a rational function $f \in K(C)^*$, we write $\text{trop}(f)$ for the real valued function on the skeleton Γ given by $y \mapsto \text{val}_y(f)$. The function $\text{trop}(f)$ is piecewise linear with integer slopes, and thus we obtain a map

$$\text{trop} : K(C)^* \rightarrow \text{PL}(\Gamma).$$

Moreover, this map respects linear equivalence of divisors, in the sense that if $D \sim D'$ on C then $\text{trop}(D) \sim \text{trop}(D')$ on Γ . In particular, the tropicalization map on divisors descends to a natural map on Picard groups

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(\Gamma).$$

One can refine this observation as follows. Let x be a type-2 point in C^{an} . Given a nonzero rational function f on C , one can define its *normalized reduction* \tilde{f}_x with respect to x as follows. Choose $c \in K^*$ such that $|f|_x = |c|$. Define $\tilde{f}_x \in \kappa(C_x)^*$ to be the image of $c^{-1}f$ in the residue field of $K(C)$ with respect to val_x , which by definition is isomorphic to $\kappa(C_x)$. Although f_x is only well-defined up to multiplication by an element of κ^* , its divisor $\text{div}(f_x)$ is completely well-defined. We define the normalized reduction of the zero function to be zero.

Given an $(r + 1)$ -dimensional K -vector space $H \subset K(C)$, its reduction $\tilde{H}_x = \{\tilde{f}_x \mid f \in H\}$ is an $(r + 1)$ -dimensional vector space over κ . Given \tilde{f} in the function field of C_x and a closed point ν of C_x , we let $s^\nu(\tilde{f}) := \text{ord}_\nu(\tilde{f})$ be the order of vanishing of \tilde{f} at ν . If $\tilde{f} = \tilde{f}_x$ for $f \in K(C)^*$, then $s^\nu(\tilde{f})$ is equal to the slope of $\text{trop}(f)$ in the tangent direction at x corresponding to ν . This is a consequence of the non-Archimedean Poincaré–Lelong formula, due to Thuillier [116]. Using this observation, one deduces the following important result (cf. [25, Theorem 5.15]):

Theorem 6.4 (Slope Formula). *For any nonzero rational function $f \in K(C)$,*

$$\text{Trop}(\text{div}(f)) = \text{div}(\text{trop}(f)).$$

6.4 Skeletons of Higher-Dimensional Berkovich Spaces

The construction of the skeleton of a semistable model of a curve given in Sect. 6.2 can be generalized in various ways to higher dimensions. For brevity, we mention just three such generalizations. In what follows, X will denote a proper variety of dimension n over K .

1. *Semistable models.* Suppose \mathcal{X} is a strictly semistable model of X over R . Then the geometric realization of the dual complex $\Delta(\overline{\mathcal{X}})$ of the special fiber embeds naturally in the Berkovich analytic space X^{an} , and as in the case of curves there is a strong deformation retraction of X^{an} onto $\Delta(\overline{\mathcal{X}})$. These facts are special cases of results due to Berkovich; see [104] for a lucid explanation of the constructions in the special case of strictly semistable models, and [70] for a generalization to “extended skeleta.”
2. *Toroidal embeddings.* A *toroidal embedding* is, roughly speaking, something which looks étale-locally like a toric variety together with its dense big open torus. When K is trivially valued, Thuillier [117] associates a skeleton $\Sigma(X)$ of U and an extended skeleton $\overline{\Sigma}(X)$ of X to any toroidal embedding $U \subset X$ (see also [1]). As in the case of semistable models, the skeleton $\overline{\Sigma}(X)$ embeds naturally into X^{an} and there is a strong deformation retract $X^{\text{an}} \rightarrow \overline{\Sigma}(X)$.
3. *Abelian varieties.* If E/K is an elliptic curve with non-integral j -invariant, the skeleton associated with a minimal proper semistable model of R can also be constructed using Tate’s non-Archimedean uniformization theory. In this case,

the skeleton of E^{an} is the quotient of the skeleton of \mathbf{G}_m (which is isomorphic to \mathbb{R} and consists of the unique path from 0 to ∞ in $(\mathbb{P}^1)^{\text{an}}$) by the map $x \mapsto x - \text{val}(j_E)$. Using Mumford’s higher-dimensional generalization of Tate’s theory [102], one can define a skeleton associated with any totally degenerate abelian variety A ; it is a real torus of dimension $\dim(A)$. This can be generalized further using Raynaud’s uniformization theory to define a canonical notion of skeleton for an arbitrary abelian variety A/K (see, e.g., [68]). If A is principally polarized, there is an induced *tropical principal polarization* on the skeleton of A , see [20, Sect. 3.7] for a definition. It is shown in [20] that the skeleton of the Jacobian of a curve C is isomorphic to the Jacobian of the skeleton as principally polarized real tori:

Theorem 6.5 ([20]). *Let C be a curve over an algebraically closed field, complete with respect to a nontrivial valuation, such that the minimal skeleton of the Berkovich analytic space C^{an} is isometric to Γ . Then there is a canonical isomorphism of principally polarized real tori $\text{Jac}(\Gamma) \cong \Sigma(\text{Jac}(C)^{\text{an}})$ making the following diagram commute.*

$$\begin{array}{ccc}
 C^{\text{an}} & \xrightarrow{AJ} & \text{Jac}(C)^{\text{an}} \\
 \downarrow & & \searrow \\
 \Gamma & \xrightarrow{AJ} & \text{Jac}(\Gamma) \xrightarrow{\sim} \Sigma(\text{Jac}(C)^{\text{an}}).
 \end{array}$$

Remark 6.6. Theorem 6.5 has the following interpretation in terms of tropical moduli spaces, which we discuss in greater detail in Sect. 7. There is a map

$$\text{trop} : M_g \rightarrow M_g^{\text{trop}}$$

from the moduli space of genus $g \geq 2$ curves to the moduli space of tropical curves of genus g which takes a curve C to its minimal skeleton, considered as a vertex-weighted metric graph. There is also a map (of sets, for example)

$$\text{trop} : A_g \rightarrow A_g^{\text{trop}}$$

from the moduli space of principally polarized abelian varieties of dimension g to the moduli space of “principally polarized tropical abelian varieties” of dimension g , taking an abelian variety to its skeleton in the sense of Berkovich. Finally, there are Torelli maps $M_g \rightarrow A_g$ (resp., $M_g^{\text{trop}} \rightarrow A_g^{\text{trop}}$) which take a curve (resp., metric graph) to its Jacobian [31]. Theorem 6.5 implies that the following square commutes:

$$\begin{array}{ccc}
 M_g & \xrightarrow{\text{trop}} & M_g^{\text{trop}} \\
 \downarrow & & \downarrow \\
 A_g & \xrightarrow{\text{trop}} & A_g^{\text{trop}}
 \end{array}$$

This is also proved, with slightly different hypotheses, in [118, Theorem A].

7 Moduli Spaces

7.1 Moduli of Tropical Curves

The moduli space of tropical curves M_g^{trop} has been constructed by numerous authors [1, 33, 64, 87]. In this section, we give a brief description of this object, with an emphasis on applications to classical algebraic geometry.

Given a finite vertex-weighted graph $\mathbf{G} = (G, \omega)$ in the sense of Sect. 4.4, the set of all vertex-weighted metric graphs (Γ, ω) with underlying finite graph G is naturally identified with

$$M_{\mathbf{G}}^{\text{trop}} := \mathbb{R}_{>0}^{|E(G)|} / \text{Aut}(\mathbf{G})$$

with the Euclidean topology. If \mathbf{G}' is obtained from \mathbf{G} by contracting an edge, then we may think of a metric graph in $M_{\mathbf{G}'}^{\text{trop}}$ as a limit of graphs in $M_{\mathbf{G}}^{\text{trop}}$ in which the length of the given edge approaches zero. Similarly, if \mathbf{G}' is obtained from \mathbf{G} by contracting a cycle to a vertex v and augmenting the weight of v by one, we may think of a metric graph in $M_{\mathbf{G}'}^{\text{trop}}$ as a limit of graphs in $M_{\mathbf{G}}^{\text{trop}}$. In this way, we may construct the *moduli space of tropical curves*

$$M_g^{\text{trop}} := \bigsqcup M_{\mathbf{G}}^{\text{trop}},$$

where the union is over all *stable*⁹ vertex-weighted graphs \mathbf{G} of genus g , and the topology is induced by gluing $M_{\mathbf{G}'}^{\text{trop}}$ to the boundary of $M_{\mathbf{G}}^{\text{trop}}$ whenever \mathbf{G}' is a contraction of \mathbf{G} in one of the two senses above.

We note that the moduli space M_g^{trop} is not compact, since edge lengths in a metric graph must be finite and thus there is no limit if we let some edge length tend to infinity. There exists a compactification $\overline{M}_g^{\text{trop}}$, known as the moduli space of *extended tropical curves*, which parameterizes vertex-weighted metric graphs with possibly infinite edges; we refer to [1] for details.

⁹A vertex-weighted finite graph (G, ω) is called *stable* if every vertex of weight zero has valence at least 3.

Let M_g be the (coarse) moduli space of genus g curves and \overline{M}_g its Deligne–Mumford compactification, considered as varieties over \mathbb{C} endowed with the trivial valuation. Points of M_g^{an} can be identified with equivalence classes of points of $M_g(L)$, where L is a complete non-Archimedean field extension of \mathbb{C} (with possibly nontrivial valuation). There is a natural map $\text{Trop} : M_g^{\text{an}} \rightarrow M_g^{\text{trop}}$ which on the level of L -points takes a smooth proper genus g curve C/L to the minimal skeleton of its Berkovich analytification C^{an} . This map extends naturally to a map $\text{Trop} : \overline{M}_g^{\text{an}} \rightarrow \overline{M}_g^{\text{trop}}$.

Let $\Sigma(\overline{M}_g)$ (resp., $\overline{\Sigma}(\overline{M}_g)$) denote the skeleton, in the sense of Thuillier, of M_g^{an} (resp., $\overline{M}_g^{\text{an}}$) with respect to the natural toroidal structure coming from the boundary strata of $\overline{M}_g \setminus M_g$. According to the main result of Abramovich et al. [1], there is a very close connection between the moduli space of tropical curves M_g^{trop} and the Thuillier skeleton $\Sigma(\overline{M}_g)$:

Theorem 7.1 ([1]). *There is a canonical homeomorphism¹⁰*

$$\Phi : \Sigma(\overline{M}_g^{\text{an}}) \rightarrow M_g^{\text{trop}}$$

which extends uniquely to a map

$$\overline{\Phi} : \overline{\Sigma}(\overline{M}_g^{\text{an}}) \rightarrow \overline{M}_g^{\text{trop}}$$

of compactifications in such a way that

$$\begin{array}{ccc} \overline{M}_g^{\text{an}} & \xrightarrow{\overline{P}} & \overline{\Sigma}(\overline{M}_g^{\text{an}}) \\ & \searrow \text{Trop} & \downarrow \overline{\Phi} \\ & & \overline{M}_g^{\text{trop}} \end{array}$$

commutes, where $\overline{P} : \overline{M}_g^{\text{an}} \rightarrow \overline{\Sigma}(\overline{M}_g^{\text{an}})$ is the canonical deformation retraction.

It follows from Theorem 7.1 that the map $\text{Trop} : \overline{M}_g^{\text{an}} \rightarrow \overline{M}_g^{\text{trop}}$ is continuous, proper, and surjective. From this, one easily deduces:

Corollary 7.2. *Let K be a complete and algebraically closed non-Archimedean field with value group \mathbb{R} , and let Γ be a stable metric graph of genus at least 2. Then there exists a curve C over K such that the minimal skeleton of the Berkovich analytic space C^{an} is isometric to Γ .*

¹⁰This homeomorphism is in fact an isomorphism of “generalized cone complexes with integral structure” in the sense of Abramovich et al. [1].

Remark 7.3. A more direct proof of Corollary 7.2, which in fact proves a stronger statement by replacing Γ with an arbitrary metrized complex of curves, and the field K with any complete and algebraically closed non-Archimedean field whose value group contains all edge lengths in some model for Γ , can be found in Theorem 3.24 of [8]. The proof uses formal and rigid geometry. A variant of Corollary 7.2 for discretely valued fields, proved using deformation theory, can be found in Appendix B of [17].

Remark 7.4. Let R be a complete DVR with field of fractions K and infinite residue field κ . The argument in Appendix B of [17] shows that for any finite connected graph G , there exists a regular, proper, flat curve \mathcal{C} over R whose generic fiber is smooth and whose special fiber is a maximally degenerate semistable curve with dual graph G . One should note the assumption here that κ has infinite residue field. In the case where the residue field is finite—for example, when $K = \mathbb{Q}_p$ —the question of which graphs arise in this way remains an open problem. The significance of this problem is its relation to the rational points of the moduli space of curves. For example, the existence of Brill–Noether general curves defined over \mathbb{Q} for large g is a well-known open question. Lang’s conjecture predicts that, for large g , such curves should be contained in a proper closed subset of the moduli space of curves. One suggested candidate for this closed subset is the stable base locus of the canonical bundle, which is known to contain only Brill–Noether special curves.

7.2 Brill–Noether Rank

The motivating problem of Brill–Noether theory is to describe the variety $W_d^r(C)$ parameterizing divisors of a given degree and rank on a curve C . A first step in such a description should be to compute numerical invariants of $W_d^r(C)$, such as its dimension. Our goal is to use the combinatorics of the dual graph Γ to describe $W_d^r(\Gamma)$. Combining this combinatorial description with the Specialization Theorem, we can then hope to understand the Brill–Noether locus of our original curve. One might be tempted to think that the tropical analogue of $\dim W_d^r(C)$ should be $\dim W_d^r(\Gamma)$, but as in the case of linear series, the dimension is not a well-behaved tropical invariant. We note one example of such poor behavior.

Example 7.5. In [93, Theorem 1.1], it is shown that the function that takes a metric graph Γ to $\dim W_d^r(\Gamma)$ is not upper semicontinuous on M_g^{trop} . To see this, the authors construct the following example. Let Γ be the loop of loops of genus 4 depicted in Fig. 11, with edges of length $\ell_1 < \ell_2 < \ell_3$ as pictured. Suppose that $\ell_1 + \ell_2 > \ell_3$. Then $\dim W_3^1(\Gamma) = 1$. If, however, we consider the limiting metric graph Γ_0 as ℓ_1, ℓ_2 , and ℓ_3 approach zero, then on this graph the only divisor of degree 3 and rank 1 is the sum of the three vertices, hence $\dim W_3^1(\Gamma_0) = 0$.




Fig. 11 The metric graph Γ from [93]

The solution to this problem has a very similar flavor to the definition of rank recorded in Definition 4.2. Specifically, given a curve C , consider the incidence correspondence

$$\Phi = \{(p_1, \dots, p_d, D) \in C^d \times W_d^r(C) \mid p_1 + \dots + p_d \in |D|\}.$$

The forgetful map to $W_d^r(C)$ has fibers of dimension r , so $\dim \Phi = r + \dim W_d^r(C)$, and hence the image of Φ in C^d has the same dimension. This suggests the following surrogate for the dimension of $W_d^r(C)$.

Definition 7.6. Let Γ be a metric graph, and suppose that $W_d^r(\Gamma)$ is nonempty. The Brill–Noether rank $w_d^r(\Gamma)$ is the largest integer k such that, for every effective divisor E of degree $r + k$, there exists a divisor $D \in W_d^r(\Gamma)$ such that $|D - E| \neq \emptyset$.

Example 7.7. Note that, in the previous example, although $\dim W_3^1(\Gamma) = 1$, the Brill–Noether rank $w_3^1(\Gamma) = 0$. To see this, it suffices to find a pair of points such that no divisor of degree 3 and rank 1 passes through both points simultaneously. Indeed, it is shown in [93, Theorem 1.9] that no divisor of rank 1 and degree 3 contains $v_3 + w_3$.

The Brill–Noether rank is much better behaved than the dimension of the Brill–Noether locus; for example (cf. [93, Theorem 1.6] and [90, Theorem 5.4]):

Theorem 7.8. *The Brill–Noether rank is upper semicontinuous on the moduli space of tropical curves.*

The Brill–Noether rank also satisfies the following analogue of the Specialization Theorem (cf. [93, Theorem 1.7] and [90, Theorem 5.7]):

Theorem 7.9. *If C is a curve over an algebraically closed field K with nontrivial valuation, and the skeleton of the Berkovich analytic space C^{an} is isometric to Γ , then*

$$\dim W_d^r(C) \leq w_d^r(\Gamma).$$

We note the following generalization of Theorem 4.6.

Corollary 7.10. *Let Γ be a metric graph of genus g . Then $w_d^r(\Gamma) \geq \rho := g - (r + 1)(g - d + r)$.*

Proof. The general theory of determinantal varieties shows that, if $W_d^r(C)$ is nonempty, then its dimension is at least ρ . The result then follows from [84, 86] and Theorem 7.9.

Remark 7.11. It is unknown whether $W_d^r(\Gamma)$ must have local dimension at least ρ . Note, however, that this must hold in a neighborhood of a divisor $D \in \text{Trop } W_d^r(C)$. Hence, if $W_d^r(\Gamma)$ has smaller than the expected local dimension in a neighborhood of some divisor D , then D does not lift to a divisor of rank r on a curve C having Γ as its tropicalization.

8 Metrized Complexes of Curves and Limit Linear Series

In this section we describe the work of Amini and Baker [5] on the Riemann–Roch and Specialization Theorems for divisors on metrized complexes of curves, along with applications to the theory of limit linear series.

8.1 Metrized Complexes of Curves

Metrized complexes of curves can be thought of, loosely, as objects which interpolate between classical and tropical algebraic geometry. More precisely, a *metrized complex of algebraic curves* over an algebraically closed field κ is a finite metric graph Γ together with a fixed model G and a collection of marked complete nonsingular algebraic curves C_v over κ , one for each vertex v of G ; the set \mathcal{A}_v of marked points on C_v is in bijection with the edges of G incident to v . A metrized complex over \mathbb{C} can be visualized as a collection of compact Riemann surfaces connected together via real line segments, as in Fig. 12.

The *geometric realization* $|\mathfrak{C}|$ of a metrized complex of curves \mathfrak{C} is defined as the topological space given by the union of the edges of G and the collection of curves C_v , with each endpoint v of an edge e identified with the corresponding marked point x_v^e (as suggested by Fig. 12). The *genus* of a metrized complex of curves \mathfrak{C} , denoted $g(\mathfrak{C})$, is by definition $g(\mathfrak{C}) = g(\Gamma) + \sum_{v \in V} g_v$, where g_v is the genus of C_v and $g(\Gamma)$ is the first Betti number of Γ .

A *divisor* on a metrized complex of curves \mathcal{C} is an element \mathcal{D} of the free abelian group on $|\mathcal{C}|$. Thus a divisor on \mathcal{C} can be written uniquely as $\mathcal{D} = \sum_{x \in |\mathcal{C}|} a_x x$ where $a_x \in \mathbb{Z}$, all but finitely many of the a_x are zero, and the sum is over all points of $\Gamma \setminus V$ as well as $C_v(\kappa)$ for $v \in V$. The *degree* of \mathcal{D} is defined to be $\sum a_x$.

A *nonzero rational function* f on a metrized complex of curves \mathcal{C} is the data of a rational function $f_\Gamma \in \text{PL}(\Gamma)$ and nonzero rational functions f_v on C_v for each $v \in V$. We call f_Γ the Γ -*part* of f and f_v the C_v -*part* of f . The *divisor* of a nonzero rational function f on \mathcal{C} is defined to be

$$\text{div}(f) := \sum_{x \in |\mathcal{C}|} \text{ord}_x(f) x,$$


where $\text{ord}_x(f)$ is defined as follows¹¹:

- If $x \in \Gamma \setminus V$, then $\text{ord}_x(f) = \text{ord}_x(f_\Gamma)$, where $\text{ord}_x(f_\Gamma)$ is the sum of the slopes of f_Γ in all tangent directions emanating from x .
- If $x \in C_v(\kappa) \setminus \mathcal{A}_v$, then $\text{ord}_x(f) = \text{ord}_x(f_v)$.
- If $x = x_v^e \in \mathcal{A}_v$, then $\text{ord}_x(f) = \text{ord}_x(f_v) + \text{slp}_e(f_\Gamma)$, where $\text{slp}_e(f_\Gamma)$ is the outgoing slope of f_Γ at v in the direction of e .

Divisors of the form $\text{div}(f)$ are called *principal*, and the principal divisors form a subgroup of $\text{Div}^0(\mathcal{C})$, the group of divisors of degree zero on \mathcal{C} . Two divisors in $\text{Div}(\mathcal{C})$ are called *linearly equivalent* if they differ by a principal divisor. Linear equivalence of divisors on \mathcal{C} can be understood rather intuitively in terms of “chip-firing moves” on \mathcal{C} . We refer the reader to Sect. 1.2 of [5] for details.

A divisor $\mathcal{E} = \sum_{x \in |\mathcal{C}|} a_x(x)$ on \mathcal{C} is called *effective* if $a_x \geq 0$ for all x . The *rank* $r_{\mathcal{C}}$ of a divisor $\mathcal{D} \in \text{Div}(\mathcal{C})$ is defined to be the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors \mathcal{E} of degree k on \mathcal{C} (so in particular $r_{\mathcal{C}}(\mathcal{D}) \geq 0$ if and only if \mathcal{D} is linearly equivalent to an effective divisor, and otherwise $r_{\mathcal{C}}(\mathcal{D}) = -1$).

Fig. 12 An example of a metrized complex



¹¹Note that our sign convention here for the divisor of a rational function on Γ , which coincides with the one used in [5], is the opposite of the sign convention used in Sect. 3.2, which is also used in a number of other papers in the subject. This should not cause any confusion, but it is good for the reader to be aware of this variability when perusing the literature.

The theory of divisors, linear equivalence, and ranks on metrized complexes of curves generalizes both the classical theory for algebraic curves and the corresponding theory for metric graphs. The former corresponds to the case where G consists of a single vertex v and no edges and $C = C_v$ is an arbitrary smooth curve. The latter corresponds to the case where the curves C_v have genus zero for all $v \in V$. Since any two points on a curve of genus zero are linearly equivalent, it is easy to see that the divisor theories and rank functions on \mathcal{C} and Γ are essentially the same.

The *canonical divisor* on \mathcal{C} is

$$\mathcal{K} = \sum_{v \in V} (K_v + \sum_{w \in \mathcal{A}_v} w),$$

where K_v is a canonical divisor on C_v .

The following result generalizes both the classical Riemann–Roch theorem for algebraic curves and the Riemann–Roch theorem for metric graphs:

Riemann–Roch for Metrized Complexes. *Let \mathcal{C} be a metrized complex of algebraic curves over κ . For any divisor $\mathcal{D} \in \text{Div}(\mathcal{C})$, we have*

$$r_{\mathcal{C}}(\mathcal{D}) - r_{\mathcal{C}}(\mathcal{K} - \mathcal{D}) = \text{deg}(\mathcal{D}) - g(\mathcal{C}) + 1.$$

As with the tropical Riemann–Roch theorem, the proof of this theorem makes use of a suitable notion of *reduced divisors* for metrized complexes of curves. We note that the proof of the Riemann–Roch theorem for metrized complexes uses the Riemann–Roch theorem for algebraic curves and does not furnish a new proof of that result.

8.2 Specialization of Divisors from Curves to Metrized Complexes

Let K be a complete and algebraically closed non-Archimedean field with valuation ring R and residue field κ , and let C be a smooth proper curve over K . As in Sect. 6, there is a metrized complex \mathcal{C} canonically associated with any strongly semistable model \mathcal{C} of C over R . The specialization map Trop defined in Sects. 2 and 3 can be enhanced in a canonical way to a map from divisors on C to divisors on \mathcal{C} . The enhanced specialization map, which by abuse of terminology we continue to denote by Trop, is obtained by linearly extending a map $\tau : C(K) \rightarrow |\mathcal{C}|$. The map τ is defined as follows:

- For $P \in C(K)$ reducing to a smooth point $\text{red}(P)$ of the special fiber C_0 of \mathcal{C} , $\tau(P)$ is just the point $\text{red}(P)$.
- For $P \in C(K)$ reducing to a singular point, $\tau(P)$ is the point $\text{Trop}(P)$ in the relative interior of the corresponding edge of the skeleton Γ of \mathcal{C} .

The motivation for the definitions of $\text{Trop} : C(K) \rightarrow |\mathcal{C}|$ and $\text{div}(f)$ comes in part from the following extension of the Slope Formula (Theorem 6.4):

Proposition 8.1. *Let f be a nonzero rational function on C and let \mathfrak{f} be the corresponding nonzero rational function on \mathfrak{C} , where f_Γ is the restriction to Γ of the piecewise linear function $\log |f|$ on C^{an} and $f_v \in \kappa(C_v)$ for $v \in V$ is the normalized reduction of f to C_v (cf. Sect. 6.1). Then*

$$\text{Trop}(\text{div}(f)) = \text{div}(\mathfrak{f}).$$

In particular, we have $\text{Trop}(\text{Prin}(C)) \subseteq \text{Prin}(\mathfrak{C})$.

The Specialization Theorem from Sect. 4.2 generalizes to metrized complexes as follows:

Specialization Theorem for Metrized Complexes. *For all $D \in \text{Div}(C)$, we have*

$$r_C(D) \leq r_{\mathfrak{C}}(\text{trop}(D)).$$

Since $r_{\mathfrak{C}}(\text{trop}(D)) \leq r_\Gamma(\text{trop}(D))$, the specialization theorem for metrized complexes is a strengthening of the analogous specialization result for metrized graphs. In conjunction with a simple combinatorial argument, this theorem also refines the Specialization Theorem for vertex-weighted graphs.

A simple consequence of the Riemann–Roch and Specialization Theorems for metrized complexes is that for any canonical divisor K_C on C , the divisor $\text{trop}(K_C)$ belongs to the canonical class on \mathfrak{C} . Indeed, the Specialization Theorem shows that $r_{\mathfrak{C}}(\text{trop}(K_C)) \geq g - 1$, while Riemann–Roch shows that a divisor of degree $2g - 2$ and rank at least $g - 1$ must be equivalent to \mathcal{K} .

There is also a version of specialization in which one has *equality* rather than just an inequality. One can naturally associate to a rank r divisor D on C a collection $\mathcal{H} = \{H_v\}_{v \in V}$ of $(r + 1)$ -dimensional subspaces of $\kappa(C_v)$, where H_v is the normalized reduction of $\mathcal{L}(D)$ to C_v (cf. Sect. 6). If $\mathcal{F} = \{F_v\}_{v \in V}$, where F_v is any κ -subspace of the function field $\kappa(C_v)$, then for $\mathcal{D} \in \text{Div}(\mathfrak{C})$ we define the \mathcal{F} -restricted rank of \mathcal{D} , denoted $r_{\mathfrak{C}, \mathcal{F}}(\mathcal{D})$, to be the largest integer k such that for any effective divisor \mathcal{E} of degree k on \mathfrak{C} , there is a rational function \mathfrak{f} on \mathfrak{C} whose C_v -parts f_v belong to F_v for all $v \in V$, and such that $\mathcal{D} - \mathcal{E} + \text{div}(\mathfrak{f}) \geq 0$.

Theorem 8.2 (Specialization Theorem for Restricted Ranks). *With notation as above, the \mathcal{H} -restricted rank of the specialization of D is equal to the rank of D , i.e., $r_{\mathfrak{C}, \mathcal{H}}(\text{trop}(D)) = r$.*

8.3 Connections with the Theory of Limit Linear Series

The theory of linear series on metrized complexes of curves has close connections with the Eisenbud–Harris theory of limit linear series for strongly semistable curves of compact type, and allows one to generalize the basic definitions in the Eisenbud–Harris theory to more general semistable curves. The Eisenbud–Harris theory, which they used to settle a number of longstanding open problems in the theory of

algebraic curves, only applies to a rather restricted class of reducible curves, namely, those of *compact type* (i.e., nodal curves whose dual graph is a tree). It has been an open problem for some time to generalize their theory to more general semistable curves.¹²

Recall that the *vanishing sequence* of a linear series $L = (L, W)$ at $p \in C$, where $W \subset H^0(C, L)$ is the ordered sequence

$$a_0^L(p) < \dots < a_r^L(p)$$

of integers k with the property that there exists some $s \in W$ vanishing to order exactly k at p . For strongly semistable curves of compact type, Eisenbud and Harris define a notion of *crude limit* $\mathfrak{g}_d^r L$ on C_0 , which is the data of a (not necessarily complete) degree d and rank r linear series L_v on C_v for each vertex $v \in V$ with the following property: if two components C_u and C_v of C_0 meet at a node p , then for any $0 \leq i \leq r$,

$$a_i^{L_v}(p) + a_{r-i}^{L_u}(p) \geq d.$$

We can canonically associate to a proper strongly semistable curve C_0 a metrized complex \mathfrak{C} of κ -curves, called the *regularization* of C_0 , by assigning a length of 1 to each edge of G . This is the metrized complex associated with any regular smoothing \mathfrak{C} of C_0 over any discrete valuation ring R with residue field κ .

Theorem 8.3. *Let \mathfrak{C} be the metrized complex of curves associated with a strongly semistable curve C_0/κ of compact type. Then there is a bijective correspondence between the following:*

1. *Crude limit \mathfrak{g}_d^r 's on C_0 in the sense of Eisenbud and Harris.*
2. *Equivalence classes of pairs $(\mathcal{H}, \mathcal{D})$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r + 1)$ -dimensional subspace of $\kappa(C_v)$ for each $v \in V$, and \mathcal{D} is a divisor of degree d supported on the vertices of \mathfrak{C} with $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) = r$. Here we say that $(\mathcal{H}, \mathcal{D}) \sim (\mathcal{H}', \mathcal{D}')$ if there is a rational function \mathfrak{f} on \mathfrak{C} such that $\mathcal{D}' = \mathcal{D} + \text{div}(\mathfrak{f})$ and $H_v = H'_v \cdot f_v$ for all $v \in V$, where f_v denotes the C_v -part of \mathfrak{f} .*

Theorem 8.3, combined with the Riemann–Roch theorem for metrized complexes of curves, provides a new proof of the fact, originally established in [58], that limit linear series satisfy analogues of the classical theorems of Riemann and Clifford. The point is that $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$ for all $\mathcal{D} \in \text{Div}(\mathfrak{C})$, and therefore upper bounds on $r_{\mathfrak{C}}(\mathcal{D})$ which follow from Riemann–Roch imply corresponding upper bounds on the restricted rank $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D})$.

Motivated by Theorem 8.3, Amiri and Baker propose the following definition.

Definition 8.4. Let C_0 be a strongly semistable (but not necessarily compact type) curve over κ with regularization \mathfrak{C} . A *limit \mathfrak{g}_d^r* on C_0 is an equivalence class of pairs

¹²Brian Osserman [107] has recently proposed a different framework for doing this.

$(\{H_v\}, \mathcal{D})$ as above, where H_v is an $(r + 1)$ -dimensional subspace of $\kappa(C_v)$ for each $v \in V$, and \mathcal{D} is a degree d divisor on \mathcal{C} with $r_{\mathcal{C}, \mathcal{H}}(\mathcal{D}) = r$.

As partial additional justification for Definition 8.4, Amini and Baker prove, using specialization, that a \mathfrak{g}_d^r on the smooth general fiber C of a semistable family \mathcal{C} gives rise in a natural way to a crude limit \mathfrak{g}_d^r on the central fiber.

Part 3: Applications

In this part, we discuss several recent applications of tropical Brill–Noether theory to problems in algebraic and arithmetic geometry. These sections are largely independent of each other, so the reader should be able to peruse them according to his or her interest.

9 Applications of Tropical Linear Series to Classical Brill–Noether Theory

Recent years have witnessed several applications of tropical Brill–Noether theory to problems in classical algebraic geometry. In this section, we survey the major recent developments in the field.

9.1 The Brill–Noether Theorem

The Brill–Noether theorem predicts the dimension of the space $W_d^r(C)$ parameterizing divisor classes (or, equivalently, complete linear series) of a given degree and rank on a general curve C .

Brill–Noether Theorem ([66]). *Let C be a general curve of genus g over \mathbb{C} . Then $W_d^r(C)$ has pure dimension $\rho(g, r, d) = g - (r + 1)(g - d + r)$, if this is nonnegative, and is empty otherwise.*

The original proof of the Brill–Noether theorem, due to Griffiths and Harris, involves a subtle degeneration argument [66]. The later development of limit linear series by Eisenbud and Harris led to a simpler proof of this theorem [57, 58]. The literature contains several other proofs, some of which work in any characteristic. One that is often referenced is due to Lazarsfeld, because rather than using degenerations, Lazarsfeld’s argument involves vector bundles on K3 surfaces [89].

The first significant application of tropical Brill–Noether theory was the new proof of the Brill–Noether theorem by Cools, Draisma, Payne, and Robeva [48], which successfully realized the program laid out in [17]. In [48], the authors




Fig. 13 The graph Γ

consider the family of graphs pictured in Fig. 13, colloquially known as the chain of loops.¹³ The edge lengths are further assumed to be generic, which in this case means that, if ℓ_i, m_i are the lengths of the bottom and top edges of the i th loop, then ℓ_i/m_i is not equal to the ratio of two positive integers whose sum is less than or equal to $2g - 2$.

Using Theorem 4.6 as the only input from algebraic geometry, the authors of [48] employ an intricate combinatorial argument to prove the following:

Theorem 9.1 ([48]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . Then $\dim W_d^r(C) = \rho(g, r, d)$ if this number is nonnegative, and $W_d^r(C) = \emptyset$ otherwise.*

We note that such a curve C exists by Corollary 7.2. The Brill–Noether theorem (over an arbitrary algebraically closed field) then follows from Theorem 9.1 using the theory of Brill–Noether rank discussed in Sect. 7.2.

In fact, [48] proves more. Theorem 4.6 of [48] completely describes $W_d^r(\Gamma)$, explicitly classifying all divisors of given degree and rank on a generic chain of loops. Indeed, it is shown that $W_d^r(\Gamma)$ is a union of ρ -dimensional tori. The set of tori is in bijection with the so-called lingering lattice paths, which in turn are in bijection with standard Young tableaux on a rectangle with $r + 1$ columns and $g - d + r$ rows containing the numbers $1, \dots, g$. From this, one can compute the number of tori to be

$$\binom{g}{\rho} (g - \rho)! \prod_{i=0}^r \frac{i!}{(g - d + r + i)!}$$

if $\rho \geq 0$, and 0 if $\rho < 0$.

We briefly discuss the argument here. Given an effective divisor D , we may assume that D is v_1 -reduced. The divisor D then has some number d_1 of chips at

¹³In fact, they consider the graph in which the lengths of the bridge edges between the loops are all zero. There is, however, a natural rank-preserving isomorphism between the Jacobian of a metric graph with a bridge and the Jacobian of the graph in which that bridge has been contracted, so their argument works equally well in this case. We consider the graph with bridges because of its use in [78, 79].




Fig. 14 A decomposition of Γ

v_1 , and by Dhar’s burning algorithm D has at most 1 chip on each of the half-open loops γ_k pictured in Fig. 14, and no chips on the half-open bridges br_k .

The associated *lingering lattice path* is a sequence of vectors $p_i \in \mathbb{Z}^r$, starting at $p_1 = (d_1, d_1 - 1, \dots, d_1 - r + 1)$, with the i th step given by

$$p_{i+1} - p_i = \left\{ \begin{array}{ll} (-1, -1, \dots, -1) & \text{if } D \text{ has no chip on } \gamma_i \\ e_j & \text{if } D \text{ has a chip on } \gamma_i, \text{ the distance from } v_i \\ & \text{to this chip is precisely } (p_i(j) + 1)m_{i+1}, \\ & \text{and both } p_i \text{ and } p_i + e_j \text{ are in } \mathcal{W} \\ 0 & \text{otherwise} \end{array} \right\}$$

Here, the distance from v_i is in the counterclockwise direction. Since the chip lies on a circle of circumference $m_i + \ell_i$, this distance should be understood to be modulo $(m_i + \ell_i)$. The symbols e_0, \dots, e_{r-1} represent the standard basis vectors in \mathbb{Z}^r and \mathcal{W} is the open Weyl chamber

$$\mathcal{W} = \{y \in \mathbb{Z}^r | y_0 > y_1 > \dots > y_{r-1} > 0\}.$$

The steps where $p_{i+1} = p_i$ are known as *lingering steps*. The basic idea of the lingering lattice path is as follows. By Theorem 5.7, the set $\{v_1, v_2, \dots, v_g, w_g\}$ is rank-determining. Hence, if D fails to have rank r , there is an effective divisor E of degree r , supported on these vertices, such that $|D - E| = \emptyset$. Starting with the v_1 -reduced divisor D , we move chips to the right and record the v_i -degree of the equivalent v_i -reduced divisor. The number $p_i(j)$ is then the minimum, over all effective divisors E of degree j supported at v_1, \dots, v_i , of the v_i -degree of the v_i -reduced divisor equivalent to $D - E$. From this it follows that, when D has rank at least r , we must have $p_i(j) \geq r - j$, so the corresponding lingering lattice path must lie in the open Weyl chamber \mathcal{W} .

The corresponding tableau is constructed by placing the moves in the direction e_j in the i th column of the rectangle, and the moves in the direction $(-1, \dots, -1)$ in the last column. If the k th step is lingering, then the integer k does not appear in the tableau. Given this description, we see that each tableau determines the existence and position of the chip on the half-open loop γ_k if and only if the integer k appears in the tableau. Otherwise, the chip on the k th loop is allowed to move freely. The number of chips that are allowed to move freely is therefore $\rho = g - (r + 1)(g - d + r)$. Indeed, we see that not only is the Brill–Noether rank $w'_d(\Gamma)$ equal to ρ , but in fact

$\dim W_d^r(\Gamma) = \rho$ as well. Theorem 9.1 then follows from the specialization result for Brill–Noether rank, Theorem 7.9.

9.2 The Gieseker–Petri Theorem

Assume that $\rho(g, r, d) \geq 0$. The variety $W_d^r(C)$ is singular along $W_d^{r+1}(C)$. Blowing up along this subvariety yields the variety $\mathcal{G}_d^r(C)$ parameterizing (not necessarily complete) linear series of degree d and rank r on C . A natural generalization of the Brill–Noether theorem is the following:

Gieseker–Petri Theorem ([65]). *Let C be a general curve of genus g . If $\rho(g, r, d) \geq 0$, then $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$.*

It is a standard result, following [12, Sect. IV.4], that the Zariski cotangent space to $\mathcal{G}_d^r(C)$ at a point corresponding to a complete linear series $\mathcal{L}(D)$ is naturally isomorphic to the cokernel of the adjoint multiplication map

$$\mu_D : \mathcal{L}(D) \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C).$$

Thus the cotangent space has dimension $\rho(g, r, d) + \dim \ker \mu_D$, and in particular, $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$ at such a point if and only if the multiplication map μ_D is injective. More generally, if $P \in \mathcal{G}_d^r(C)$ corresponds to a possibly incomplete linear series $W \subset \mathcal{L}(D)$, then $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$ at P if and only if the multiplication map $W \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C)$ is injective. One deduces that the Gieseker–Petri theorem is equivalent to the assertion that if C is a general curve of genus g , then μ_D is injective for all divisors D on C .

A recent application of tropical Brill–Noether theory is the following result [78], which yields a new proof of the Gieseker–Petri theorem:

Theorem 9.2 ([78]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . Then the multiplication map*

$$\mu_D : \mathcal{L}(D) \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C)$$

is injective for all divisors D on C .

The argument has much in common with the tropical proof of the Brill–Noether theorem, using the same metric graph with the same genericity conditions on edge lengths. The new ingredient is the idea of tropical independence, as defined in Sect. 5.3. Given a divisor $D \in W_d^r(C)$, the goal is to find functions

$$f_0, \dots, f_r \in \text{trop}(\mathcal{L}(D))$$

$$g_0, \dots, g_{g-d+r-1} \in \text{trop}(\mathcal{L}(K_C - D))$$

such that $\{f_i + g_j\}_{i,j}$ is tropically independent.

There is a dense open subset of $W_d^r(\Gamma)$ consisting of divisors D with the following property: given an integer $0 \leq i \leq r$, there exists a unique divisor $D_i \sim D$ such that

$$D_i - iw_g - (r - i)v_1 \geq 0.$$

These are the divisors referred to as *vertex-avoiding* in [41].

We first describe the proof of Theorem 9.2 in the case that D is vertex-avoiding. If D is the specialization of a divisor $\mathcal{D} \in W_d^r(C)$, and $p_1, p_g \in C$ are points specializing to v_1, w_g , respectively, then there exists a divisor $\mathcal{D}_i \sim \mathcal{D}$ such that

$$\mathcal{D}_i - ip_g - (r - i)p_1 \geq 0,$$

and, by the uniqueness of D_i, \mathcal{D}_i must specialize to D_i . It follows that there is a function $f_i \in \text{trop}(\mathcal{L}(\mathcal{D}))$ such that $\text{div}(f_i) = D_i - D$.

For this open subset of divisors, the argument then proceeds as follows. By the classification in [48], the divisor D_i fails to have a chip on the k th loop if and only if the integer k appears in the i th column of the corresponding tableau. The adjoint divisor $E = K_\Gamma - D$ corresponds to the transpose tableau [2, Theorem 39], so the divisor $D_i + E_j$ fails to have a chip on the k th loop if and only if k appears in the (i, j) position of the tableau. Since for each k at most one of these divisors fails to have a chip on the k th loop, we see that if

$$\theta = \min\{f_i + g_j + b_{i,j}\}$$

occurs at least twice at every point of Γ , then the divisor

$$\Delta = \text{div}(\theta) + K_\Gamma$$

must have a chip on the k th loop for all k .

To see that this is impossible, let p_k be a point of Δ in γ_k , and let

$$D' = p_1 + \cdots + p_g.$$

Then by construction $K_\Gamma - D'$ is equivalent to an effective divisor, so by the tropical Riemann–Roch theorem we see that $r(D') \geq 1$. On the other hand, Dhar’s burning algorithm shows that D' is universally reduced, so by Proposition 5.16 we have $r(D') = 0$, a contradiction.

It is interesting to note that this obstruction is, at heart, combinatorial. Unlike the earlier proofs via limit linear series, which arrive at a contradiction by constructing a canonical divisor of impossible *degree* (larger than $2g - 2$), this argument arrives at a contradiction by constructing a canonical divisor of impossible *shape*.

The major obstacle to extending this argument to the case where D is not vertex-avoiding is that the containment $\text{trop}(\mathcal{L}(D)) \subseteq R(\text{Trop}(D))$ is often strict. Given an arbitrary divisor $D \in W_d^r(C)$ and function $f \in R(\text{Trop}(D))$, it is difficult to determine whether f is the specialization of a function in $\mathcal{L}(D)$. To avoid this issue, the authors

make use of a patching construction, gluing together tropicalizations of different rational functions in a fixed algebraic linear series on different parts of the graph, to arrive at a piecewise linear function in $R(K_\Gamma)$ that may not be in $\text{trop}(\mathcal{L}(K_C))$. Once this piecewise linear function is constructed, the argument proceeds very similarly to the vertex-avoiding case.

9.3 The Maximal Rank Conjecture

One of the most well-known open problems in Brill–Noether theory is the Maximal Rank Conjecture, which predicts the Hilbert function for sufficiently general embeddings of sufficiently general curves. This conjecture is attributed to Noether in [11, p. 4] (see [105, Sect. 8] and [43, pp. 172–173] for details), was studied classically by Severi [112, Sect. 10], and popularized by Harris [71, p. 79].

Maximal Rank Conjecture. *Fix nonnegative integers g, r, d , let C be a general curve of genus g , and let $V \subset \mathcal{L}(D)$ be a general linear series of rank r and degree d on C . Then the multiplication maps*

$$\mu_m : \text{Sym}^m V \rightarrow \mathcal{L}(mD)$$

have maximal rank for all m . That is, each μ_m is either injective or surjective.

While the Maximal Rank Conjecture remains open in general, several important cases are known [27, 61, 115, 119]. For example, it is shown in [27] that the Maximal Rank Conjecture holds in the non-special range $d \geq g + r$. When $d < g + r$, the general linear series of degree d and rank r on a general curve is complete, and for this reason, most of the work in the subject focuses on the case where $V = \mathcal{L}(D)$. We note that the arguments of Ballico and Ellia [27] and Farkas [61] involve degenerations to unions of two curves that meet in more than one point. Since such curves are not of compact type, the arguments do not make use of limit linear series.

In [79], tropical Brill–Noether theory is used to prove the $m = 2$ case of the Maximal Rank Conjecture.

Theorem 9.3 ([79]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . For a given r and d , let D be a general divisor of rank r and degree d on C . Then the multiplication map*

$$\mu_2 : \text{Sym}^2 \mathcal{L}(D) \rightarrow \mathcal{L}(2D)$$

has maximal rank.

The genericity conditions placed on the edge lengths of Γ in Theorem 9.3 are stricter than those appearing in the tropical proofs of the Brill–Noether and

Gieseker–Petri theorems. First, the bridges between the loops are assumed to be much longer than the loops themselves, and second, one must assume that certain integer linear combinations of the edge lengths do not vanish.

A simplifying aspect of the Maximal Rank Conjecture is that it concerns a general, rather than arbitrary, divisor. It therefore suffices to prove that the maximal rank condition holds for a single divisor of the given degree and rank on C . The main result of Cartwright et al. [41] is that every divisor on the generic chain of loops is the specialization of a divisor of the same rank on C . We are therefore free to choose whatever divisor we wish to work with, and in particular we may choose one of the vertex-avoiding divisors described in the previous section. Recall that, if $D \in W_d^r(\Gamma)$ is vertex-avoiding, then we have an explicit set of piecewise linear functions $f_i \in R(D)$ that are tropicalizations of a basis for the linear series on the curve C . The goal, in the case where the multiplication map is supposed to be injective, is to show that the set $\{f_i + f_j\}_{i \leq j}$ is tropically independent. In the surjective case, we must choose a subset of the appropriate size, and then show that this subset is tropically independent.

The basic idea of the argument is as follows. Assume that

$$\theta = \min\{f_i + f_j + b_{i,j}\}$$

occurs at least twice at every point of Γ , and consider the divisor

$$\Delta = \text{div}(\theta) + 2D.$$

To arrive at a contradiction, one studies the degree distribution of the divisor Δ across the loops of Γ . More precisely, one defines

$$\delta_k := \text{deg}(\Delta|_{\gamma_k}).$$

The first step is to show that $\delta_k \geq 2$ for all k . One then identifies intervals $[a, b]$ for which this inequality must be strict for at least one $k \in [a, b]$. As one proceeds from left to right across the graph, one encounters such intervals sufficiently many times to obtain $\text{deg } \Delta > 2 \text{ deg } D$, a contradiction.

10 Lifting Problems for Divisors on Metric Graphs

In this section we discuss the lifting problem in tropical Brill–Noether theory: given a divisor of rank r on a metric graph Γ , when is it the tropicalization of a rank r divisor on a smooth curve C ? There are essentially two formulations of this problem, one in which the curve C is fixed, and one in which it is not.

Throughout this section, we let K be a complete and algebraically closed nontrivially valued non-Archimedean field.

Question 10.1. Given a metric graph Γ and a divisor D on Γ , under what conditions do there exist a curve C/K (together with a semistable model R) and a divisor of the same rank as D tropicalizing to Γ and D , respectively?

Question 10.2. Given a curve C/K (together with a semistable model R) tropicalizing to a metric graph Γ , and given a divisor D on Γ , under what conditions does there exist a divisor on C of the same rank as D tropicalizing to D ?

These are very difficult questions. Even for the earlier theory of limit linear series on curves of compact type, the analogous questions remain open. A partial answer in that setting is given by the Regeneration Theorem of Eisenbud and Harris [58], which says that if the space of limit linear series has local dimension equal to the Brill–Noether number ρ , then the given limit linear series lifts in any one-parameter smoothing. At the time of writing, there is no corresponding theorem in the tropical setting.¹⁴

10.1 Specialization of Hyperelliptic Curves

One of the first results concerning lifting of divisors is the classification of vertex-weighted metric graphs that are the specialization of a hyperelliptic curve. Recall from Sect. 4.4 that given a curve C/K and a semistable model \mathcal{C}/R for C , there is a natural way to associate to \mathcal{C} a vertex-weighted metric graph (Γ, ω) . We call such a pair *minimal* if there is no vertex v with $\text{val}(v) = 1$ and $\omega(v) = 0$.

Theorem 10.3 ([9, 34]). *Let (Γ, ω) be a minimal vertex-weighted metric graph. There is a smooth projective hyperelliptic curve over a discretely valued field with a regular, strongly semistable model whose special fiber has dual graph Γ if and only if the following conditions hold:*


- (HYP1) *there exists an involution s on Γ such that the quotient Γ/s is a tree and $s(v) = v$ for all $v \in \Gamma$ with $\omega(v) > 0$ and*
- (HYP2) *for every point $v \in \Gamma$, the number of bridge edges adjacent to v is at most $2\omega(v) + 2$.*

Kawaguchi and Yamaki show moreover that, when Γ satisfies these conditions, there is a smoothing C for which every divisor on Γ lifts to a divisor of the same rank on C [83].

We outline the necessity of the conditions above in the special case where $\omega = 0$, which is equivalent to requiring that $g(C) = g(\Gamma)$. Note that if C is a hyperelliptic curve, then by the Specialization Theorem any divisor of degree 2 and rank 1 on C specializes to a divisor D of rank at least 1 on Γ , and by tropical Clifford’s theorem D must have rank exactly 1. Now, if $P \in \Gamma$ is not contained in a bridge, then $|P| = \{P\}$.

¹⁴Amini has apparently made substantial progress in this direction.

Fig. 15 A hyperelliptic metric graph of genus 3 that is not the skeleton of any hyperelliptic curve of genus 3



On the other hand, if $P \in \Gamma$ is contained in a bridge, a simple analysis reveals that $D \sim 2P$. In this way we obtain an involution s on Γ mapping each point P to the P -reduced divisor equivalent to $D - P$.

To see why (HYP2) holds, note that for each type-2 point $v \in \Gamma$, the linear series of degree 2 and rank 1 on C specializes to a linear series of degree 2 and rank 1 on the corresponding curve C_v . Each of the bridges adjacent to v corresponds to ramification points of this linear series, but such a linear series has only $2g(C_v) + 2 = 2$ ramification points.

To see that the conditions (HYP1) and (HYP2) are sufficient requires significantly more work.

Example 10.4. Consider the metric graph Γ pictured in Fig. 15, consisting of a tree with a loop attached to each leaf, with all edge lengths being arbitrary. Then Γ is hyperelliptic, because any divisor of degree 2 supported on the tree has rank 1. On the other hand, this graph is not the dual graph of the limit of any family of genus 3 hyperelliptic curves, because the vertex of valence 3 in the tree is adjacent to more than 2 bridges.

For metric graphs of higher gonality, the lifting problem is significantly harder. In [97], Luo and Manjunath describe an algorithm for smoothability of rank one generalized limit linear series on metrized complexes.

10.2 Lifting Divisors on the Chain of Loops

For some specific families of graphs, such as the chain of loops discussed Sect. 9, one can show that the lifting problem is unobstructed.

Theorem 10.5 ([41]). *Let C/K be a smooth projective curve of genus g . If the dual graph of the central fiber of some regular model of C is isometric to a generic chain of loops Γ of genus g , then every divisor class on Γ that is rational over the value group of K lifts to a divisor class of the same rank on C .*

The general strategy for proving Theorem 10.5 is to study the Brill–Noether loci as subschemes

$$W_d^r(C) \subset \text{Jac}(C).$$

Since C is maximally degenerate, the universal cover of $\text{Jac}(C)^{\text{an}}$ gives a uniformization

$$T^{\text{an}} \rightarrow \text{Jac}(C)^{\text{an}}$$

by an algebraic torus T of dimension g . The tropicalization of this torus is the universal cover of the skeleton of $\text{Jac}(C)$, which as discussed in Sect. 6 is canonically identified with the tropical Jacobian of Γ [20].

A key tool in the proof of Theorem 10.5 is Rabinoff’s lifting theorem [110], which can be applied to the analytic preimages in T of algebraic subschemes of $\text{Jac}(C)$. This lifting theorem says that isolated points in complete intersections of tropicalizations of analytic hypersurfaces lift to points in the analytic intersection with appropriate multiplicities. This theorem can be applied to translates of the preimage of the theta divisor $\Theta_\Gamma = W_{g-1}^0(\Gamma)$, as follows.

When Γ is the generic chain of loops, one can use the explicit description of $W_d^r(\Gamma)$ from [48] to produce explicit translates of Θ_C whose tropicalizations intersect transversally and locally cut out $W_d^r(\Gamma)$. By intersecting with ρ additional translates of Θ_C , one obtains an isolated point in a tropical complete intersection, to which we may apply Rabinoff’s lifting theorem. This complete intersection is typically larger than $W_d^r(\Gamma)$, but the argument shows that the tropicalization map from a 0-dimensional slice of $W_d^r(C)$ to the corresponding slice of $W_d^r(\Gamma)$ is injective. Using again the explicit description of $W_d^r(\Gamma)$, one then shows that the two finite sets have the same cardinality, and hence the map is bijective.

Remark 10.6. As mentioned in the section on the Maximal Rank Conjecture, Theorem 10.5 is one of the key ingredients in the proof of Theorem 9.3 (the Maximal Rank Conjecture for quadrics). In particular, in order to show that the maximal rank condition holds for a generic line bundle of a given degree and rank, it suffices to show that it holds for a single line bundle. Since every divisor of a given rank on the chain of loops lifts to a line bundle on C of the same rank, one is free to work with any divisor of this rank on the chain of loops.

10.3 Examples of Divisors That Do Not Lift

Among the results on lifting divisors, there is a plethora of examples of divisors that do not lift. For example, in [53], Coppens defines a base-point free divisor on a metric graph Γ to be a divisor D such that $r(D - p) < r(D)$ for all $p \in \Gamma$. He then shows that the Clifford and Riemann–Roch bounds are the only obstructions to the existence of base-point free divisors on metric graphs of arbitrary genus. This is in contrast to the case of algebraic curves, where, for example, a curve of genus greater than 6 cannot have a base-point free divisor of degree 5 and rank 2.

Another example of divisors that do not lift comes from the theory of matroids.

Theorem 10.7 ([40]). *Let M be any rank 3 matroid. Then there exist a graph G_M and a rank 2 divisor D_M on G_M such that, for any infinite field k , there are a curve C over $k((t))$ (together with a semistable model \mathcal{C} of C over $k[[t]]$) and a rank 2 divisor on C tropicalizing to G_M and D_M , respectively, if and only if M is realizable over k .*

Combining this with the scheme-theoretic analogue of Mnëv universality, due to Lafforgue [88], one obtains the following.

Corollary 10.8 ([40]). *Let X be a scheme of finite type over $\text{Spec } \mathbb{Z}$. Then there exist a graph G and a rank 2 divisor D on G such that, for any infinite field k , G , and D are the tropicalizations of a curve $C/k((t))$ and a rank 2 divisor on C if and only if X has a k -point.*

In other words, the obstructions to lifting over a valued field of the form $k((t))$ are essentially as general as possible.

Cartwright’s construction is as follows. Recall that a rank 3 simple matroid on a finite set E consists of a collection of subsets of E , called *flats*, such that every pair of elements is contained in exactly one flat. (Here we are abusing language, using the word flat to refer to the maximal, or rank 2, flats.) The bipartite graph G_M is the *Levi graph* of the matroid M , where the vertices correspond to elements and flats, and there is an edge between two vertices if the corresponding element is contained in the corresponding flat. The divisor D_M is simply the sum of the vertices corresponding to elements of E . A combinatorial argument then shows that the rank of D_M is precisely 2.

If M is realizable over k , then by definition, there exists a configuration of lines in \mathbb{P}_k^2 where the lines correspond to the elements of E , and the flats correspond to points where two or more of the lines intersect. If we blow up the plane at the intersection points, the dual graph of the resulting configuration is the Levi graph G_M , and the pullback of the hyperplane class specializes to the divisor D_M . After some technical deformation arguments, one then sees that the pair (G_M, D_M) admits a lifting when M is realizable over k .

For the converse, one must essentially show that the above construction is the only possibility. That is, if \mathcal{C} is a regular semistable curve over $k[[t]]$, the dual graph of the central fiber is G_M and the divisor D_M is the specialization of a rank 2 divisor on \mathcal{C} , then in fact the image of the central fiber under the corresponding linear series must provide a realization of the matroid M in \mathbb{P}_k^2 .

11 Bounding the Number of Rational Points on Curves

By Faltings’ theorem (née the Mordell conjecture), if C is a curve of genus $g \geq 2$ over a number field K , then the set $C(K)$ of rational points on C is finite. Shortly after Faltings proved this theorem, Vojta published a new proof which furnishes an effective upper bound on the number of points in $C(K)$. However, the Vojta bound is completely theoretical—to our knowledge no one has ever written down the bound

explicitly (and the bound is surely quite far from optimal). None of the existing proofs of the Mordell conjecture gives an algorithm—even in theory!—to compute the set $C(K)$. And in practice the situation is even worse—it seems safe to say that no one has ever used the Faltings or Vojta proofs of the Mordell conjecture to compute $C(K)$ in a single nontrivial example.

11.1 *The Katz–Zureick-Brown Refinement of Coleman’s Bound*

One of the first significant results in the direction of the Mordell conjecture was Chabauty’s theorem that $C(K)$ is finite provided that the rank of the finitely generated abelian group $\text{Jac}(C)(K)$ is less than g . Much later, Coleman used his theory of p -adic integration to give an effective upper bound on $C(K)$ in this situation. Coleman’s bound has the advantage of being sharp in certain cases, and the method of proof can be used to compute $C(K)$ in a wide range of concrete examples. For simplicity, we state the results in this section for $K = \mathbb{Q}$ only, but everything extends with minor modifications to curves over a number field K . Coleman’s theorem is as follows.

Theorem 11.1 ([50]). *Let C be a curve of genus g over \mathbb{Q} , and suppose that the Mordell–Weil rank r of $\text{Jac}(C)(\mathbb{Q})$ is strictly less than the genus g . Then for every prime $p > 2g$ of good reduction for C , we have*

$$\#C(\mathbb{Q}) \leq \#C(\mathbb{F}_p) + 2g - 2. \quad (2)$$

Coleman’s theorem was subsequently strengthened in different ways. In [95], Lorenzini and Tucker (see also McCallum–Poonen [99]) generalized Theorem 11.1 to primes of bad reduction, replacing $C(\mathbb{F}_p)$ in (2) by the smooth \mathbb{F}_p -points of the special fiber of the minimal proper regular model for C over \mathbb{Z}_p . Stoll replaced the quantity $2g - 2$ in (2) by $2r$ when C has good reduction at p , and asked if this improvement could be established in the bad reduction case as well. Stoll’s question was answered affirmatively by Katz and Zureick-Brown in [80] by supplementing Stoll’s method with results from the theory of linear series on tropical curves:

Theorem 11.2 ([80]). *Let C be a curve of genus g over \mathbb{Q} and suppose that the rank r of $\text{Jac}(C)(\mathbb{Q})$ is less than g . Then for every prime $p > 2r + 2$, we have*

$$\#C(\mathbb{Q}) \leq \#C^{\text{sm}}(\mathbb{F}_p) + 2r,$$

where C denotes the minimal proper regular model of C over \mathbb{Z}_p .

In order to explain the relevance of linear series on tropical curves to such a result, we need to briefly explain the basic ideas underlying the previous work of Coleman et al. Let us first outline a proof of Theorem 11.1. Fix a rational point $P \in C(\mathbb{Q})$

(if no such point exists, then the theorem is vacuously true) and let $\iota : C \hookrightarrow J$ be the corresponding Abel–Jacobi embedding. Coleman’s theory of p -adic integration of 1-forms associates to each $\omega \in H^0(C, \Omega_C^1)$ and $Q \in C(\mathbb{Q}_p)$ a (definite) p -adic integral $\int_p^Q \omega \in \mathbb{Q}_p$, obtained by pulling back a corresponding p -adic integral on J via the map ι . Locally on C , such p -adic integrals can be computed by formally integrating a power series expansion $f_\omega(T)$ for ω with respect to a local parameter T on some residue disc U . One can show fairly easily that the p -adic closure $\overline{J(\mathbb{Q})}$ of $J(\mathbb{Q})$ in $J(\mathbb{Q}_p)$ has dimension at most r as a p -adic manifold. The formalism of Coleman’s theory implies that forcing the p -adic integral of a 1-form on J to vanish identically on $\overline{J(\mathbb{Q})}$ imposes at most r linear conditions on $H^0(J, \Omega_J^1)$. The functoriality of Coleman integration implies that the \mathbb{Q}_p -vector space V_{chab} of all $\omega \in H^0(C, \Omega_C^1)$ such that $\int_p^Q \omega = 0$ for all $Q \in C(\mathbb{Q})$ has dimension at least $g - r > 0$.

The condition $p > 2g$ implies, by a p -adic analogue of Rolle’s theorem which can be proved in an elementary way with Newton polygons, that if $f_\omega(T)$ has n zeroes on U then $\int f_\omega(T) dT$ has at most $n + 1$ zeroes on U . Using this observation, Coleman deduces, by summing over all residue classes, that if ω is a nonzero 1-form in $H^0(C, \Omega_C^1)$ vanishing on all of $C(\mathbb{Q})$ then

$$\#C(\mathbb{Q}) \leq \sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} (1 + \text{ord}_{\overline{Q}} \overline{\omega}),$$

where $\overline{\omega}$ denotes the reduction of ω to \overline{C} . Since the 1-form $\overline{\omega}$ on \overline{C} has a total of $2g - 2$ zeros counting multiplicity, we have

$$\sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} \text{ord}_{\overline{Q}} \overline{\omega} \leq 2g - 2,$$

which yields Coleman’s bound.

Stoll observed in [113] that one could do better than this by adapting the differential ω to the point \overline{Q} rather than using the same differential ω for all residue classes. Define the *Chabauty divisor*

$$D_{\text{chab}} = \sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} n_{\overline{Q}}(\overline{Q}),$$

where $n_{\overline{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\text{ord}_{\overline{Q}} \overline{\omega}$, and let d be the degree of D_{chab} . Since D_{chab} and $K_{\overline{C}} - D_{\text{chab}}$ are both equivalent to effective divisors, Clifford’s inequality (applied to the smooth proper curve \overline{C}) implies that

$$h^0(K_{\overline{C}} - D_{\text{chab}}) - 1 \leq \frac{1}{2}(2g - 2 - d).$$

On the other hand, the semicontinuity of $h^0(D) = r(D) + 1$ under specialization shows that $h^0(D_{\text{chab}}) \geq \dim V_{\text{chab}} \geq g - r$. Combining these inequalities gives

$$g - r - 1 \leq \frac{1}{2}(2g - 2 - d)$$

and thus $d \leq 2r$, giving Stoll’s refinement of Coleman’s bound.

Lorenzini and Tucker [95] had shown earlier that one can generalize Coleman’s bound to the case of bad reduction as follows. Since points of $C(\mathbb{Q})$ specialize to the set $\bar{C}^{\text{sm}}(\mathbb{F}_p)$ of smooth \mathbb{F}_p -points on the special fiber of C under the reduction map, one obtains by an argument similar to the one above the bound

$$\#C(\mathbb{Q}) \leq \sum_{\bar{Q} \in \bar{C}^{\text{sm}}(\mathbb{F}_p)} (1 + n_{\bar{Q}}), \tag{3}$$

where $\bar{\omega}$ denotes the reduction of ω to the unique irreducible component of the special fiber of C containing \bar{Q} . Choosing a nonzero $\omega \in V_{\text{chab}}$ as in Coleman’s bound, the fact that the relative dualizing sheaf for C has degree $2g - 2$ gives the Lorenzini–Tucker bound. A similar argument was found independently by McCallum and Poonen [99].

We now explain where the subtlety occurs when one tries to combine the bounds of Stoll and Lorenzini–Tucker. As above, we define the Chabauty divisor

$$D_{\text{chab}} = \sum_{\bar{Q} \in \bar{C}^{\text{sm}}(\mathbb{F}_p)} n_{\bar{Q}}(\bar{Q})$$

and we let d be its degree. As in the case where C has good reduction, the goal is to show that $d \leq 2r$. When C has good reduction, Stoll proves this by combining the semicontinuity of h^0 and Clifford’s inequality. For singular curves, one can still define h^0 of a line bundle and it satisfies the desired semicontinuity theorem. However, even when C has semistable reduction, it is well-known that Clifford’s inequality does not hold in the form needed here. Katz and Zureick-Brown replace the use of Clifford’s inequality in Stoll’s argument by a hybrid between the classical Clifford inequality and Clifford’s inequality for linear series on tropical curves. In this way, they are able to obtain the desired bound $d \leq 2r$.

We briefly highlight the main steps in the argument, following the reformulation in terms of metrized complexes given in [5].

1. As noted by Katz and Zureick-Brown, if one makes a base change from \mathbb{Q}_p to an extension field over which there is a regular semistable model C' for C dominating the base change of C , the corresponding Chabauty divisors satisfy $D'_{\text{chab}} \geq D_{\text{chab}}$. We may therefore assume that C is a regular semistable model for C .
2. Let $s = \dim V_{\text{chab}} - 1$. We can identify V_{chab} with an $(s + 1)$ -dimensional space W of rational functions on C in the usual way by identifying $H^0(C, \Omega_C^1)$ with

$\mathcal{L}(K_C)$. The divisor D_{chab} on $\bar{\mathcal{C}}^{\text{sm}}$ defines in a natural way a divisor \mathcal{D} of degree d on the metrized complex \mathcal{C} associated with \mathcal{C} . We can promote the divisor $K_{\mathcal{C}} - \mathcal{D}$ to a limit linear series $(K_{\mathcal{C}} - \mathcal{D}, \{H_v\})$ by defining H_v to be the reduction of W to C_v for each $v \in V(G)$. By the definition of D_{chab} , each element of H_v vanishes to order at least $n_{\bar{Q}}$ at each point \bar{Q} in $\text{supp}(D_{\text{chab}}) \cap C_v$. The Specialization Theorem for limit linear series on metrized complexes then shows that

$$r_{\mathcal{C}}(K_{\mathcal{C}} - \mathcal{D}) \geq s \geq g - r - 1.$$

3. On the other hand, Clifford’s inequality for metrized complexes implies that

$$r_{\mathcal{C}}(K_{\mathcal{C}} - \mathcal{D}) \leq \frac{1}{2}(2g - 2 - d).$$

Combining these inequalities gives $d \leq 2r$ as desired.

11.2 The Uniformity Theorems of Katz–Rabinoff–Zureick-Brown

Together with Rabinoff, Katz, and Zureick-Brown have recently used linear series on tropical curves to refine another result due to Stoll. In [37], Caporaso, Harris, and Mazur proved that if one assumes the Bombieri–Lang conjecture then there is a uniform bound $M(g, K)$ depending only on g and the number field K such that $|C(K)| \leq M(g, K)$ for every curve C of genus $g \geq 2$ over K . The Bombieri–Lang conjecture, which asserts that the set of rational points on a variety of general type over a number field is not Zariski dense, remains wide open, and until recently little progress had been made in the direction of unconditional proofs of the Caporaso–Harris–Mazur result. In [114], Stoll proved that a uniform bound $M(g, K)$ exists for hyperelliptic curves provided that one assumes in addition that the Mordell–Weil rank of $\text{Jac}(C)(K)$ is at most $g - 3$. Katz, Rabinoff, and Zureick-Brown succeeded in removing the hypothesis in Stoll’s theorem that C is hyperelliptic, obtaining the following result.

Theorem 11.3 ([81]). *There is an explicit bound $N(g, d)$ such that if C is a curve of genus $g \geq 3$ defined over a number field K of degree d over \mathbb{Q} and having Mordell–Weil rank $r \leq g - 3$, then*

$$\#C(K) \leq N(g, d).$$

When $K = \mathbb{Q}$, one can take $N(g, 1) = 76g^2 - 82g + 22$.

Note that the bound in Theorem 11.2 is not uniform, because the quantity $|C^{\text{sm}}(\mathbb{F}_p)|$ can be arbitrarily large for a given prime p of bad reduction, and the smallest prime p of good reduction can be arbitrarily large as a function of g . Stoll’s main new idea was to apply the Chabauty–Coleman method on residue

annuli instead of just on discs. Stoll’s proof exploits the concrete description of differentials on a hyperelliptic curve as $f(x)dx/y$; the restriction of such a differential to an annulus has a bounded numerator, and Stoll is able to analyze the zeroes of the resulting p -adic integral via explicit computations with Newton polygons.

For general curves, such an explicit description of differentials and the Newton polygons of their p -adic integrals is not possible. This is where the theory of linear systems on metric graphs becomes useful. To circumvent the difficulty posed by not having an explicit description of differentials on C , Katz, Rabinoff, and Zureick-Brown generalize the Slope Formula (Theorem 6.4) to sections of a metrized line bundle. For a differential ω , the associated tropical function $F = \log |\omega|$ on the skeleton Γ of C belongs to the space $R(K_\Gamma^\#)$ of tropical rational functions G with $K_\Gamma^\# + \text{div}(G) \geq 0$. (The absolute value here comes from a natural formal metric on the canonical bundle.) Belonging to $R(K_\Gamma^\#)$ gives strong constraints on the slopes of F , and hence on the number of zeroes of the p -adic integral of ω . The Slope Formula thus replaces the Newton polygons in Stoll’s arguments, and estimates on the slopes of the Newton polygon are replaced by properties of the tropical linear series $|K_\Gamma^\#|$.

A major issue one faces in trying to establish Theorem 11.3 (which also shows up in the earlier work of Stoll) is that when C has bad reduction at p , there are two different kinds of p -adic integrals which need to be considered. On the one hand, there are the p -adic abelian integrals studied by Colmez, Zarhin, and Vologodsky, which have no periods and are obtained by pulling back the logarithm map on the p -adic Lie group $\text{Jac}(C)(\mathbb{Q}_p)$ to C . These are the integrals for which one knows that $\dim(V_{\text{chab}}) \geq g - r$. On the other hand, there are the p -adic integrals of Berkovich and Coleman–de Shalit which do have periods but also have better functoriality properties. These are the integrals which are given locally on residue annuli of a semistable model \mathcal{C} by formally integrating a local Laurent series expansion of $\omega \in H^0(C, \Omega^1)$. In order to prove Theorem 11.3, one needs to study the difference between the two kinds of p -adic integrals. One of the new discoveries of Katz, Rabinoff, and Zureick-Brown is that the difference can be understood quite concretely using tropical geometry by combining Theorem 6.5 with Raynaud’s uniformization theory.

The methods used by Katz–Rabinoff–Zureick-Brown in [81] also provide new results in the direction of a “uniform Manin-Mumford conjecture.” The Manin–Mumford conjecture, proved by Raynaud, asserts that if C is a curve of genus at least 2 embedded in its Jacobian via an Abel–Jacobi map $\iota : C \rightarrow \text{Jac}(C)$, then $\iota(C) \cap \text{Jac}(C)(\bar{K})_{\text{tors}}$ is finite. One can ask whether there is a uniform bound on the size of this intersection as one varies over all curves of a fixed genus g . The following uniform result for the number of K -rational points on C which are torsion on J is proved in [81]:

Theorem 11.4. *There is an explicit bound $N(g, d)_{\text{tors}}$ (which one can equal to the bound $N(g, d)$ above) such that if C is a curve of genus $g \geq 3$ defined over a number field K of degree d over \mathbb{Q} and $\iota : C \rightarrow \text{Jac}(C)$ is an Abel–Jacobi embedding defined over K , then*

$$\#\iota(C) \cap \text{Jac}(C)(K)_{\text{tors}} \leq N(g, d)_{\text{tors}}.$$

Note that in Theorem 11.4 there is no restriction on the Mordell–Weil rank of $\text{Jac}(C)$. It has been conjectured that $\#A(K)_{\text{tors}}$ is bounded uniformly in terms of $[K : \mathbb{Q}]$ and g for all abelian varieties of dimension g over K , which would of course imply Theorem 11.4 as a special case, but this is known only for $g = 1$ [100] and the general case seems far out of reach at present.

Katz, Rabinoff, and Zureick-Brown also prove a uniformity result concerning the number of *geometric* (\bar{K} -rational) torsion points lying on C , under a technical assumption about the structure of the stable model at some prime p . We refer to [81] for the precise statement.

12 Limiting Behavior of Weierstrass Points in Degenerating Families

The theory discussed in this paper has interesting applications to the behavior of Weierstrass points under specialization. To motivate this kind of question, we begin with a seemingly unrelated classical result due to Andrew Ogg [106].

12.1 Weierstrass Points on Modular Curves

Let N be a positive integer. The finite-dimensional space $S = S_2(\Gamma_0(N))$ of weight 2 cusp forms for the congruence subgroup $\Gamma_0(N)$ of $\text{SL}_2(\mathbb{Z})$ is an important object in number theory. An element $f \in S$ has a *q-expansion* of the form $f = \sum_{n=1}^{\infty} a_n q^n$ with $a_n \in \mathbb{C}$, which uniquely determines f . For $f \neq 0$ in S , define

$$\text{ord}(f) = \inf\{n \mid a_n \neq 0\} - 1. \tag{4}$$

If $g = g_0(N) = \dim(S)$, then by Gaussian elimination there exists an element $f \in S$ with $\text{ord}(f) \geq g - 1$. Is there any unexpected cancellation? Under certain restrictions on the level N , the answer is no:

Theorem 12.1 ([106]). *If $N = pM$ with p prime, $p \nmid M$, and $g_0(M) = 0$, then there is no nonzero element f of $S_2(\Gamma_0(N))$ with $\text{ord}(f) \geq g$. (In particular, this holds if $N = p$ is prime.)*

One can give an enlightening proof of Ogg’s theorem using specialization of divisors from curves to metric graphs; the following argument is taken from [17].

First of all, Ogg’s theorem can be recast in the following purely geometric way, which is in fact how Ogg formulated and proved the result in [106]:

Theorem 12.2. *If $N = pM$ with p prime, $p \nmid M$, and $g_0(M) = 0$, then the cusp ∞ is not a Weierstrass point on the modular curve $X_0(N)$.*

Recall that a point P on a genus g curve X is called a *Weierstrass point* if there exists a holomorphic differential $\omega \in H^0(X, \Omega_X^1)$ vanishing to order at least g at P . To see the equivalence between Theorems 12.1 and 12.2, recall that q is an analytic local parameter on $X_0(N)$ at the cusp ∞ and the map $f \mapsto f(q)\frac{dq}{q}$ gives an isomorphism between $S_2(\Gamma_0(N))$ and the space of holomorphic differentials on $X_0(N)$. Under this isomorphism, the function ord defined in (4) becomes the order of vanishing of the corresponding differential at ∞ . So there is a nonzero element f of $S_2(\Gamma_0(N))$ with $\text{ord}(f) \geq g$ if and only if there is a nonzero holomorphic differential ω vanishing to order at least g at ∞ .


The reduction of $X = X_0(N)$ modulo p when p exactly divides $N = pM$ is well-understood; the special fiber of the so-called *Deligne–Rapoport model* for $X_0(N)$ over \mathbb{Z}_p consists of two copies of $X_0(M)$ intersecting transversely at the supersingular points in characteristic p . This model is always semistable but is not in general regular. (It is very easy to describe the minimal regular model, but we will not need this here.) In any case, the skeleton Γ of $X_0(N)$ over \mathbb{Q}_p is a “metric banana graph” consisting of two vertices connected by a number of edges, as pictured in Fig. 16, and the cusp ∞ specializes to one of the two vertices, call it P . Under the hypotheses of Theorem 12.1, each $X_0(M)$ is a rational curve and so the genus of Γ is equal to the genus of $X_0(N)$. That is, there are $g + 1$ edges. By the Specialization Theorem, if there is a nonzero global section of K_X vanishing to order at least g at ∞ , then $r(K_\Gamma - gP) \geq 0$. However, since $K_\Gamma = (g - 1)P + (g - 1)Q$, where Q is the other vertex, we have $K_\Gamma - gP = (g - 1)Q - P$, which is P -reduced by Dhar’s algorithm and non-effective. Therefore $r(K_\Gamma - gP) = -1$, and Ogg’s theorem is proved.

12.2 Specialization of Weierstrass Points

The essence of the above argument is that if C is a *totally degenerate* curve, meaning that the genus of its minimal skeleton Γ equals the genus of C , then the Weierstrass points on C must specialize to Weierstrass points on Γ , where a Weierstrass point on Γ is a point P such that $r(K_\Gamma - gP) \geq 0$. It follows from the Specialization Theorem and the corresponding fact from algebraic geometry that if Γ is a metric graph of genus $g \geq 2$ then the set of Weierstrass points on Γ is nonempty. A purely combinatorial proof of this fact was given by Amini [3].

The specialization of Weierstrass points is also a natural thing to study from the purely algebro-geometric point of view, where one is asking about the limiting

Fig. 16 The “Banana” graph of genus 2



behavior of the Weierstrass points in a semistable one-parameter family of curves. This subject, which was previously studied by Eisenbud–Harris [59], Esteves–Medeiros [60], and several other authors, has seen important recent advances by Amini [4]. We now summarize the main results proved in Amini’s paper.

Let L be a line bundle of degree d and rank $r \geq 0$ on a curve C of genus g over an algebraically closed field k of characteristic zero. Given a point $P \in C(k)$, we define the *vanishing set* $S_P(L)$ of L at P to be the set of orders of vanishing of global sections of L at P . We have $|S_P(L)| = r + 1$ for all $P \in C(k)$, and for all but finitely many $P \in C(k)$ the vanishing set is $[r] := \{0, 1, \dots, r\}$. A point $P \in C(k)$ whose vanishing set is not $[r]$ is called a *Weierstrass point* for L . Equivalently, P is a Weierstrass point for L if there exists a global section of L vanishing to order at least $r + 1$ at P . A Weierstrass point of C is by definition a Weierstrass point for the canonical bundle K_C .

The L -weight of a point $P \in C(k)$ is

$$\text{wt}_P(L) = \left(\sum_{m \in S_P(L)} m \right) - \binom{r+1}{2} = \sum_{m \in S_P(L)} m - \sum_{i \in [r]} i.$$

Thus $\text{wt}_P(L) \geq 0$ for all $P \in C(k)$ and $\text{wt}_P(L) > 0$ if and only if P is a Weierstrass point for L . The *Weierstrass divisor* for L is $\mathcal{W} = \mathcal{W}(L) = \sum_{P \in C(k)} \text{wt}_P(L)(P)$. If we fix a basis \mathcal{F} for $H^0(C, L)$, the corresponding *Wronskian* $\text{Wr}_{\mathcal{F}}$ is a nonzero global section of $L^{\otimes(r+1)} \otimes K_C^{\otimes \frac{r(r+1)}{2}}$ whose divisor is precisely $\mathcal{W}(L)$. In particular, the degree of $\mathcal{W}(L)$ (i.e., the total number of Weierstrass points counted according to their weights) is $W(L) := d(r + 1) + (g - 1)r(r + 1)$.

We seek an explicit formula for $\text{Trop}(\mathcal{W})$. For this, it is convenient to fix a divisor with $L = L(D)$, and to define as usual $\mathcal{L}(D) = \{f \in k(C)^* \mid \text{div}(f) + D \geq 0\}$. Let $D_\Gamma = \sum_{x \in \Gamma} d_x(x)$ be the specialization $\text{Trop}(D)$ of D to Γ . Let $K_\Gamma^\#$ be the canonical divisor of Γ considered as a vertex-weighted metric graph, as in Sect. 4.4. Concretely, we have $K_\Gamma^\# = \sum_{x \in \Gamma} (2g_x - 2 + \text{val}(x))x$.

For a tangent direction ν at x , define $S^\nu(D)$ to be the set of integers occurring as $s^\nu(f)$ for some $f \in \mathcal{L}(D)$, where $s^\nu(f)$ is defined as in Sect. 6 to be the slope of $\text{trop}(f)$ in the tangent direction ν . Since $s^\nu(f)$ coincides with the order of vanishing of the normalized reduction \tilde{f}_x at the point of C_x corresponding to ν , one sees easily that $|S^\nu(D)| = r + 1$.

For $x \in \Gamma$, let

$$S_x(D) = \begin{cases} \sum_{\nu \in T_x(\Gamma)} \left(\sum_{s \in S^\nu(D)} s \right) & \text{if } x \text{ is of type-2} \\ 0 & \text{otherwise,} \end{cases}$$

where $T_x(\Gamma)$ denotes the set of tangent directions at x in Γ , and let

$$S(D) = \sum_{x \in \Gamma} S_x(D) x.$$

Note that $\deg(S(D)) = 0$, since if $f \in k(C)^*$ then the slope of $F = -\log |f|$ along an oriented edge \vec{e} of Γ is the negative of the slope of F along the same edge with the orientation reversed.

The following formula is due to Amini. When Γ is the skeleton of a semistable R -model \mathcal{C} for C , the formula shows how the Weierstrass points of the generic fiber C specialize to the various components of the special fiber of \mathcal{C} , providing a simple and satisfying answer to a question of Eisenbud and Harris.

Theorem 12.3 ([3]). *Let $\text{Trop} : \text{Div}(C_{\bar{k}}) \rightarrow \text{Div}(\Gamma)$ be the natural map. Then*

$$\text{Trop}(\mathcal{W}(L)) = (r + 1) \text{Trop}(D) + \binom{r + 1}{2} K_{\Gamma}^{\#} - S(D). \tag{5}$$

Note that since $\deg(S(D)) = 0$, the degree of the right-hand side of (5) is $W(L) = \deg(\mathcal{W}(L))$ as expected. Amini also proves an analogue of (5) when the residue field of k has positive characteristic. As this is more technical to state, we will not discuss this here.

Remark 12.4. A metric graph can have infinitely many Weierstrass points; this happens, for example, with the banana graphs of genus $g \geq 3$ discussed above (see [17]). In general, the set of Weierstrass points on a metric graph Γ is a finite disjoint union of closed connected sets. It is an open problem to determine whether there are intrinsic multiplicities $m(A)$ attached to each connected component A of the Weierstrass locus on a metric graph Γ such that for any curve C having Γ as a skeleton, exactly $m(A)$ Weierstrass points of C tropicalize to A .

12.3 Distribution of Weierstrass Points

Amini uses formula (5) to prove a non-Archimedean analogue of the Mumford–Neeman equidistribution theorem, previously conjectured by Baker. We first recall the statement of the latter result, and then present Amini’s analogous theorem.

Let C be a compact Riemann surface of genus at least 1. There is a natural volume form ω_{Ar} on C , called the *Arakelov form*, which can be defined as follows. Let $\omega_1, \dots, \omega_g$ be an orthonormal basis of $\mathcal{L}(K_C)$ with respect to the Hermitian inner product

$$\langle \omega, \nu \rangle = \frac{i}{2} \int_C \omega \wedge \bar{\nu}.$$

Then the $(1, 1)$ -form $\omega_C = \frac{i}{2} \sum_{j=1}^g \omega_j \wedge \bar{\omega}_j$ does not depend on the choice of $\omega_1, \dots, \omega_g$ and has total mass g . We define

$$\omega_{Ar} := \frac{1}{g} \omega_C.$$

Geometrically, the curvature form of ω_C is the pullback of the curvature form of the flat metric on the Jacobian J of C with respect to any Abel–Jacobi map $C \rightarrow J$. Since the flat metric on J is translation-invariant, the pullback in question is independent of the choice of base point in the definition of the Abel–Jacobi map.

The Mumford–Neeman theorem [103] asserts that for any ample line bundle L on C , the Weierstrass points of $L^{\otimes n}$ become equidistributed with respect to ω_{Ar} as n tends to infinity:

Theorem 12.5. *Let C be a compact Riemann surface of genus at least 1 and let L be an ample line bundle on C . Let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in C} \text{wt}_P(L^{\otimes n}) \delta_P$$

be the probability measure supported equally on the Weierstrass points of $L^{\otimes n}$. Then as n tends to infinity, the measures δ_n converge weakly¹⁵ to the Arakelov metric ω_{Ar} .

In order to state Amini’s non-Archimedean analogue of Theorem 12.5, we will first define the analogue of the Zhang measure on vertex-weighted metric graphs/Berkovich curves following [18].

Let Γ be a metric graph of genus g . We fix a weighted graph model G of Γ and for each edge e of G let $\ell(e)$ denote the length of e . For each spanning tree T of G , let e_1, \dots, e_g denote the edges of G not belonging to T , and let

$$\mu_T = \sum_{j=1}^g \lambda(e_j)$$

where $\lambda(e)$ is Lebesgue measure along e , normalized to have total mass 1 (so that μ_T has total mass g). We also let $w(T) = \prod_{j=1}^g \ell(e_j)$, and let

$$w(G) = \sum_T w(T)$$

be the sum of $w(T)$ over all spanning trees T of G . Then the measure

$$\mu_\Gamma = \sum_T \frac{w(T)}{w(G)} \mu_T$$

is a weighted average of the measures μ_T over all spanning trees T , and in particular has total mass g .

¹⁵This means that for every continuous function $f : C \rightarrow \mathbb{R}$, we have $\int_C f \delta_n = \int_C f \omega_{Ar}$.

In other words, a random point in the complement of a random spanning tree of G is distributed according to the probability measure $\frac{1}{g}\mu_\Gamma$.

Now let (Γ, ω) be a vertex-weighted metric graph, in the sense of Sect. 4.4, of genus $g = g(\Gamma) + \sum_{x \in \Gamma} \omega(x)$. Then the measure

$$\mu_{(\Gamma, \omega)} := \mu_\Gamma + \sum_{x \in \Gamma} \omega_x \delta_x$$

has total mass g . If $g \geq 1$, we define the *Zhang measure* on (Γ, ω) to be the probability measure

$$\mu_{\text{Zh}} := \frac{1}{g} \mu_{(\Gamma, \omega)}.$$

Theorem 12.6 ([4]). *Let C be an algebraic curve of genus at least 1 over the non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Let \mathcal{C} be a strongly semistable model of C over the valuation ring of k , let (Γ, ω) be the weighted graph associated with \mathcal{C} in the sense of Sect. 4.4, and let μ_{Zh} be the Zhang measure associated with (Γ, ω) . Finally, let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in \mathcal{C}} \text{wt}_P(L^{\otimes n}) \delta_{\text{Trop}(P)}$$

be the probability measure on Γ supported equally on the tropicalizations of the Weierstrass points of $L^{\otimes n}$ (taken with multiplicities). Then as n tends to infinity, the measures δ_n converge weakly on Γ to μ_{Zh} .

A new and concrete consequence of Theorem 12.6 is the following:

Corollary 12.7. *Let C be an algebraic curve of genus at least 1 over a non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Fix a strongly semistable model \mathcal{C} for C over the valuation ring of k , let Z be an irreducible component of the special fiber of \mathcal{C} , and let g_Z be the genus of Z . Let $\mathcal{W}_Z(L^{\otimes n})$ be the set of Weierstrass points of $L^{\otimes n}$ specializing to a nonsingular point of Z . Then*

$$\lim_{n \rightarrow \infty} \frac{|\mathcal{W}_Z(L^{\otimes n})|}{|\mathcal{W}(L^{\otimes n})|} = \frac{g_Z}{g}.$$

It is convenient and enlightening to rephrase Theorem 12.6 in terms of the Berkovich analytic space C^{an} . If Γ is any skeleton of C^{an} and ω is the corresponding weight function defined by $\omega(x) = g_x$, we define the *Zhang measure* on C^{an} to be the probability measure

$$\mu_{\text{Zh}} := \frac{1}{g} \iota_* \mu_{(\Gamma, \omega)}$$

with respect to the natural inclusion $\iota : \Gamma \rightarrow C^{\text{an}}$. This measure on C^{an} is easily seen to be independent of the choice of Γ , and has total mass equal to the genus g of C . Using the fact that $C^{\text{an}} \cong \varprojlim \Gamma_C$ (cf. Sect. 6.2), one deduces using standard results from real analysis that Theorem 12.6 is equivalent to the following reformulation, which more closely resembles Theorem 12.5:

Theorem 12.8 ([4]). *Let C be an algebraic curve of genus at least 1 over the non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in C} \text{wt}_P(L^{\otimes n}) \delta_P$$

be the probability measure on C^{an} supported equally on the Weierstrass points of $L^{\otimes n}$. Then as n tends to infinity, the measures δ_n converge weakly on C^{an} to the Zhang measure μ_{Zh} .

Remark 12.9. The measure μ_{Zh} , which was first introduced in the context of vertex-weighted metric graphs by Shouwu Zhang in [120], plays the role in the non-Archimedean setting of the Arakelov volume form. Using a result of Heinz [74] and the recent work of Chambert-Loir–Ducros [44] and Gubler–Kunemann [69] on non-Archimedean Arakelov theory, one can show that, as in the Archimedean case, μ_C is obtained by pulling back the curvature form of a canonical translation-invariant metric on J via an Abel–Jacobi map. There is also evidently a close connection between the measure μ_Γ and the polyhedral decomposition $\{C_T\}$ of $\text{Pic}^g(\Gamma)$ associated with G (cf. Sect. 5.4) which is deserving of further study.

The proof of Theorem 12.6 (and its equivalent formulation Theorem 12.8) is based on formula (5) together with the theory of *Okounkov bodies*. The rough idea is that fixing a type-2 point x of C^{an} and a tangent direction v at x , as well as a divisor D with $L = L(D)$, the rational numbers $\frac{1}{n}S^v(nD)$ defined above become equidistributed in a real interval of length $d = \deg(L)$ as $n \rightarrow \infty$. Combining this “local” equidistribution result with (5) and a careful analysis of the variation of the minimum slope along edges of Γ give the desired result.

13 Further Reading

There are many topics closely related to the contents of this paper which we have not had space to discuss. Here is a brief and non-exhaustive list of some related topics and papers which we recommend to the interested reader:

1. *Harmonic morphisms.* In algebraic geometry, a base-point free linear series of rank r on a curve C is more or less the same thing as a morphism $C \rightarrow \mathbb{P}^r$. In tropical geometry, the situation is much more subtle, and no satisfactory analogue of this correspondence is known. For $r = 1$, there is a close relationship

(although not a precise correspondence) between tropical g_d^1 's on a metric graph Γ and degree d harmonic morphisms from Γ to a metric tree. The theory of harmonic morphisms of metric graphs and metrized complexes of curves is explored in detail in the papers [8, 9, 45, 97], among others.

2. *Spectral bounds for gonality.* In [49], Cornelissen et. al. establish a spectral lower bound for the *stable gonality* (in the sense of harmonic morphisms) of a graph G in terms of the smallest nonzero eigenvalue of the Laplacian of G . This is a tropical analogue of the Li–Yau inequality for Riemann surfaces. They give applications of their tropical Li–Yau inequality to uniform boundedness of torsion points on rank two Drinfeld modules, as well as to lower bounds on the modular degree of elliptic curves over function fields. The spectral bound from [49] was subsequently refined by Amini and Kool in [7] to a spectral lower bound for the *divisorial gonality* (i.e., the minimal degree of a rank 1 divisor) of a metric graph Γ . In [7], as well as in the related paper [54], this circle of ideas is applied to show that the expected gonality of a random graph is asymptotic to the number of vertices.
3. *Tropical complexes.* In [39], Cartwright formulates a higher-dimensional analogue of the basic theory of linear series on graphs, including a Specialization Theorem for the rank function. He calls the objects on which his higher-dimensional linear series live *tropical complexes*. A generalization of the Slope Formula to the context of non-Archimedean varieties and tropical complexes is proved in [70].
4. *Abstract versus embedded tropical curves.* In this paper we have dealt exclusively with linear series on abstract tropical curves (thought of as metric graphs) and have eschewed the more traditional perspective of tropical varieties as non-Archimedean amoebas associated with subvarieties of tori. The two approaches are closely related, however: see, for example, [24, 42, 70]. The theory of linear series on abstract tropical curves has concrete consequences for embedded tropical curves, e.g., with respect to the theory of bitangents and theta characteristics as in [26, 46].
5. *Algebraic rank.* In [35], Caporaso introduces a notion of rank for divisors on graphs known as the *algebraic rank*, which is defined geometrically by varying over all curves with the given dual graph and all line bundles with the given specialization. The algebraic rank differs in general from the combinatorial rank [38], but the two invariants agree for hyperelliptic graphs and graphs of genus 3 [82]. Many of the results we have discussed also hold for the algebraic rank. For example, there are specialization, Riemann–Roch, and Clifford's theorems for algebraic rank [35], and Mnëv universality holds for obstructions to the lifting problem for algebraic rank [91].

Acknowledgements The authors would like to thank Omid Amini and Sam Payne for enlightening discussions, and Eric Katz and Joe Rabinoff for helpful feedback on our summary of Katz et al. [81]. They also thank Spencer Backman, Dustin Cartwright, Melody Chan, Yoav Len, and Sam Payne for their comments on an early draft of this paper.

References

1. Abramovich, D., Caporaso, L., Payne, S.: The tropicalization of the moduli space of curves. *Ann. Sci. Éc. Norm. Supér.* **48**(4), 765–809 (2015)
2. Agrawal, R., Musiker, G., Sotirov, V., Wei, F.: Involutions on standard Young tableaux and divisors on metric graphs. *Electron. J. Comb.* **20**(3) (2013)
3. Amini, O.: Reduced divisors and embeddings of tropical curves. *Trans. Am. Math. Soc.* **365**(9), 4851–4880 (2013)
4. Amini, O.: Equidistribution of Weierstrass points on curves over non-Archimedean fields. Preprint, arXiv:1412.0926v1 (2014)
5. Amini, O., Baker, M.: Linear series on metrized complexes of algebraic curves. *Math. Ann.* **362**(1–2), 55–106 (2015)
6. Amini, O., Caporaso, L.: Riemann–Roch theory for weighted graphs and tropical curves. *Adv. Math.* **240**, 1–23 (2013)
7. Amini, O., Kool, F.: A spectral lower bound for the divisorial gonality of metric graphs. Preprint, arXiv:1407.5614 (2014)
8. Amini, O., Baker, M., Brugallé, E., Rabinoff, J.: Lifting harmonic morphisms I: metrized complexes and Berkovich skeleta. *Res. Math. Sci.* **2**, Art. 7, 67 (2015)
9. Amini, O., Baker, M., Brugallé, E., Rabinoff, J.: Lifting harmonic morphisms II: tropical curves and metrized complexes. *Algebra Number Theory.* **9**(2), 267–315 (2015)
10. An, Y., Baker, M., Kuperberg, G., Shokrieh, F.: Canonical representatives for divisor classes on tropical curves and the matrix–tree theorem. *Forum Math. Sigma* **2**, e24, 25 (2014)
11. Arbarello, E., Ciliberto, C.: Adjoint hypersurfaces to curves in \mathbf{P}^r following Petri. In: *Commutative Algebra (Trento, 1981)*. Lecture Notes in Pure and Applied Mathematics, vol. 84, pp. 1–21. Dekker, New York (1983)
12. Arbarello, E., Cornalba, M., Griffiths, P.A., Harris, J.: *Geometry of Algebraic Curves. Volume I*. Grundlehren der Mathematischen Wissenschaften, vol. 267. Springer, New York (1985)
13. Babbage, D.W.: A note on the quadrics through a canonical curve. *J. Lond. Math. Soc.* **14**, 310–315 (1939)
14. Backman, S.: Infinite reduction of divisors on metric graphs. *European J. Comb.* **35**, 67–74 (2014)
15. Backman, S.: Riemann–Roch theory for graph orientations. Preprint, arXiv:1401.3309 (2014)
16. Baker, M.: An introduction to Berkovich analytic spaces and non-Archimedean potential theory on curves. In: *p -Adic Geometry*. University Lecture Series, vol. 45, pp. 123–174. American Mathematical Society, Providence, RI (2008)
17. Baker, M.: Specialization of linear systems from curves to graphs. *Algebra Number Theory* **2**(6), 613–653 (2008)
18. Baker, M., Faber, X.: Metric properties of the tropical Abel–Jacobi map. *J. Algebraic Combin.* **33**(3), 349–381 (2011)
19. Baker, M., Norine, S.: Riemann–Roch and Abel–Jacobi theory on a finite graph. *Adv. Math.* **215**(2), 766–788 (2007)
20. Baker, M., Rabinoff, J.: The skeleton of the Jacobian, the Jacobian of the skeleton, and lifting meromorphic functions from tropical to algebraic curves. *Int. Math. Res. Not.* **2015**(16), 7436–7472 (2015)
21. Baker, M., Rumely, R.: Potential theory and dynamics on the Berkovich projective line. *Mathematical Surveys and Monographs*, vol. 159. American Mathematical Society, Providence, RI (2010)
22. Baker, M., Shokrieh, F.: Chip-firing games, potential theory on graphs, and spanning trees. *J. Comb. Theory Ser. A* **120**(1), 164–182 (2013)
23. Baker, M., Wang, Y.: The Bernardi process and torsor structures on spanning trees. Preprint, arXiv:1406.1584 (2014)
24. Baker, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization, and metrics on curves. *Algebraic Geom.* Preprint, arXiv:1104.0320v3 (2011, to appear)

25. Baker, M., Payne, S., Rabinoff, J.: On the structure of non-Archimedean analytic curves. In: Tropical and Non-Archimedean Geometry. Contemporary Mathematics, vol. 605, pp. 93–121. American Mathematical Society, Providence, RI (2013)
26. Baker, M., Len, Y., Morrison, R., Pflueger, N., Ren, Q.: Bitangents of tropical plane quartic curves. Preprint, arXiv:1404.7568 (2014)
27. Ballico, E., Ellia, P.: The maximal rank conjecture for nonspecial curves in \mathbf{P}^3 . *Invent. Math.* **79**(3), 541–555 (1985)
28. Bayer, D., Eisenbud, D.: Graph curves. *Adv. Math.* **86**(1), 1–40 (1991) With an appendix by Sung Won Park.
29. Berkovich, V.: Spectral Theory and Analytic Geometry over Non-Archimedean Fields. Mathematical Surveys and Monographs, vol. 33. American Mathematical Society, Providence, RI (1990)
30. Biggs, N.L.: Chip-firing and the critical group of a graph. *J. Algebraic Combin.* **9**(1), 25–45 (1999)
31. Brannetti, S., Melo, M., Viviani, F.: On the tropical Torelli map. *Adv. Math.* **226**(3), 2546–2586 (2011)
32. Caporaso, L.: Néron models and compactified Picard schemes over the moduli stack of stable curves. *Am. J. Math.* **130**(1), 1–47 (2008)
33. Caporaso, L.: Algebraic and tropical curves: comparing their moduli spaces. In: Farkas and Morrison (eds.) Handbook of Moduli. Adv. Lect. Math. (ALM), vol. 24, pp. 119–160. Int. Press, Somerville (2013)
34. Caporaso, L.: Algebraic and combinatorial Brill-Noether theory. In: Compact Moduli Spaces and Vector Bundles. Contemporary Mathematics, vol. 564, pp. 69–85. American Mathematical Society, Providence, RI (2012)
35. Caporaso, L.: Rank of divisors on graphs: an algebro-geometric analysis. In: A Celebration of Algebraic Geometry. Clay Mathematics Proceedings, vol. 18, pp. 45–64. American Mathematical Society, Providence, RI (2013)
36. Caporaso, L., Viviani, F.: Torelli theorem for graphs and tropical curves. *Duke Math. J.* **153**(1), 129–171 (2010)
37. Caporaso, L., Harris, J., Mazur, B.: Uniformity of rational points. *J. Am. Math. Soc.* **10**(1), 1–35 (1997)
38. Caporaso, L., Len, Y., Melo, M.: Algebraic and combinatorial rank of divisors on finite graphs. Preprint, arXiv:1401:5730v2 (2014)
39. Cartwright, D.: Tropical complexes. Preprint, arXiv:1308.3813 (2013)
40. Cartwright, D.: Lifting matroid divisors on tropical curves. Preprint, arXiv:1502.03759 (2015)
41. Cartwright, D., Jensen, D., Payne, S.: Lifting divisors on a generic chain of loops. *Canad. Math. Bull.* **58**(2), 250–262 (2015)
42. Cartwright, D., Dudzik, A., Manjunath, M., Yao, Y.: Embeddings and immersions of tropical curves. Preprint, arXiv:1409.7372v1 (2014)
43. Castelnuovo, G., Enriques, F., Severi, F.: Max Noether. *Math. Ann.* **93**(1), 161–181 (1925)
44. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. Preprint, arXiv:1204.6277 (2012)
45. Chan, M.: Tropical hyperelliptic curves. *J. Algebraic Combin.* **37**(2), 331–359 (2013)
46. Chan, M., Jiradilok, P.: Theta characteristics of tropical K_4 curves. Preprint, arXiv:1503:05776v1 (2015)
47. Ciliberto, C., Harris, J., Miranda, R.: On the surjectivity of the Wahl map. *Duke Math. J.* **57**(3), 829–858 (1988)
48. Cools, F., Draisma, J., Payne, S., Robeva, E.: A tropical proof of the Brill-Noether theorem. *Adv. Math.* **230**(2), 759–776 (2012)
49. Cornelissen, G., Fumiharu, K., Kool, J.: A combinatorial Li–Yau inequality and rational points on curves. *Math. Ann.* **361**(1–2), 211–258 (2015)
50. Coleman, R.F.: Effective Chabauty. *Duke Math. J.* **52**(3), 765–770 (1985)
51. Conrad, B.: Several approaches to non-Archimedean geometry. In: p -Adic Geometry. University Lecture Series, vol. 45, pp. 9–63. American Mathematical Society, Providence, RI (2008)

52. Coppens, M.: Clifford's theorem for graphs. Preprint, arXiv:1304.6101 (2014)
53. Coppens, M.: Free divisors on metric graphs. Preprint, arXiv:1410.3114 (2014)
54. Deveau, A., Jensen, D., Kainic, J., Mitropolsky, D.: Gonality of random graphs. Preprint, arXiv:1409.5688 (2014)
55. Dhar, D.: Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.* **64**(14), 1613–1616 (1990)
56. Dhar, D., Ruelle, P., Sen, S., Verma, D.-N.: Algebraic aspects of Abelian sandpile models. *J. Phys. A* **28**(4), 805–831 (1995)
57. Eisenbud, D., Harris, J.: A simpler proof of the Gieseker-Petri theorem on special divisors. *Invent. Math.* **74**(2), 269–280 (1983)
58. Eisenbud, D., Harris, J.: Limit linear series: basic theory. *Invent. Math.* **85**(2), 337–371 (1986)
59. Eisenbud, D., Harris, J.: Existence, decomposition, and limits of certain Weierstrass points. *Invent. Math.* **87**(3), 495–515 (1987)
60. Esteves, E., Medeiros, N.: Limits of Weierstrass points in regular smoothings of curves with two components. *C. R. Acad. Sci. Paris Sér. I Math.* **330**(10), 873–878 (2000)
61. Farkas, G.: Koszul divisors on moduli spaces of curves. *Am. J. Math.* **131**(3), 819–867 (2009)
62. Foster, T., Gross, P., Payne, S.: Limits of tropicalizations. *Israel J. Math.* **201**(2), 835–846 (2014)
63. Gathmann, A., Kerber, M.: A Riemann-Roch theorem in tropical geometry. *Math. Z.* **259**(1), 217–230 (2008)
64. Gathmann, A., Kerber, M., Markwig, H.: Tropical fans and the moduli spaces of tropical curves. *Compos. Math.* **145**(1), 173–195 (2009)
65. Gieseker, D.: Stable curves and special divisors: Petri's conjecture. *Invent. Math.* **66**(2), 251–275 (1982)
66. Griffiths, P., Harris, J.: On the variety of special linear systems on a general algebraic curve. *Duke Math. J.* **47**(1), 233–272 (1980)
67. Gubler, W.: Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.* **169**(2), 321–376 (2007)
68. Gubler, W.: Non-Archimedean canonical measures on abelian varieties. *Compos. Math.* **146**(3), 683–730 (2010)
69. Gubler, W., Künnemann, K.: A tropical approach to non-Archimedean Arakelov theory. Preprint, arXiv:1406.7637 (2014)
70. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalizations. Preprint, arXiv:1404.7044 (2014)
71. Harris, J.: *Curves in Projective Space*. Séminaire de Mathématiques Supérieures, vol. 85 Presses de l'Université de Montréal, Montreal, Que (1982) With the collaboration of David Eisenbud
72. Harris, J., Morrison, I.: *Moduli of Curves*. Graduate Texts in Mathematics, vol. 187. Springer, New York (1998)
73. Haase, C., Musiker, G., Yu, J.: Linear systems on tropical curves. *Math. Z.* **270**(3–4), 1111–1140 (2012)
74. Heinz, N.: Admissible metrics for line bundles on curves and abelian varieties over non-Archimedean local fields. *Arch. Math. (Basel)* **82**(2), 128–139 (2004)
75. Hladký, J., Král', D., Norine, S.: Rank of divisors on tropical curves. *J. Combin. Theory Ser. A* **120**(7), 1521–1538 (2013)
76. Holroyd, A., Levine, L., Mészáros, K., Peres, Y., Propp, J., Wilson, D.: Chip-firing and rotor-routing on directed graphs. In: *In and Out of Equilibrium 2*. Progress in Probability, vol. 60, pp. 331–364. Birkhäuser, Basel (2008)
77. Jensen, D.: The locus of Brill-Noether general graphs is not dense. Preprint, arXiv:1405.6338 (2014)
78. Jensen, D., Payne, S.: Tropical independence I: Shapes of divisors and a proof of the Gieseker-Petri theorem. *Algebra Number Theory* **8**(9), 2043–2066 (2014)
79. Jensen, D., Payne, S.: Tropical independence II: the maximal rank conjecture for quadrics. *Algebra Number Theory*. <http://arxiv.org/abs/1505.05460>

80. Katz, E., Zureick-Brown, D.: The Chabauty-Coleman bound at a prime of bad reduction and Clifford bounds for geometric rank functions. *Compos. Math.* **149**(11), 1818–1838 (2013)
81. Katz, E., Rabinoff, J., Zureick-Brown, D.: Uniform bounds for the number of rational points on curves of small Mordell-Weil rank. *Duke Math J.* arXiv:1504.00694 (2015)
82. Kawaguchi, S., Yamaki, K.: Algebraic rank of hyperelliptic graphs and graphs of genus 3. Preprint, arXiv:1401:3935 (2014)
83. Kawaguchi, S., Yamaki, K.: Rank of divisors on hyperelliptic curves and graphs under specialization. Preprint, arXiv:1304.6979v3 (2014)
84. Kempf, G.: *Schubert Methods with an Application to Algebraic Curves*. Publ. Math. Centrum (1971)
85. Kiss, V., Tóthmérész, L.: Chip-firing games on Eulerian digraphs and NP-hardness of computing the rank of a divisor on a graph. Preprint, arXiv:1407.6598 (2014)
86. Kleiman, S.L., Laksov, D.: Another proof of the existence of special divisors. *Acta Math.* **132**, 163–176 (1974)
87. Kozlov, D.: The topology of moduli spaces of tropical curves with marked points. *Asian J. Math.* **13**(3), 385–403 (2009)
88. Lafforgue, L.: *Chirurgie des grassmanniennes*. CRM Monograph Series. vol. 19. American Mathematical Society, Providence, RI (2003)
89. Lazarsfeld, R.: Brill-Noether-Petri without degenerations. *J. Differential Geom.* **23**(3), 299–307 (1986)
90. Len, Y.: The Brill-Noether rank of a tropical curve. Preprint, arXiv:1209.6309v1 (2012)
91. Len, Y.L. A note on algebraic rank, matroids, and metrized complexes. Preprint, arXiv:1401:8156 (2014)
92. Levine, L., Propp, J.: What is . . . a sandpile? *Not. Am. Math. Soc.* **57**(8), 976–979 (2010)
93. Lim, C.-M., Payne, S., Potashnik, N.: A note on Brill-Noether theory and rank-determining sets for metric graphs. *Int. Math. Res. Not. IMRN* **23**:5484–5504 (2012)
94. Liu, Q.: *Algebraic Geometry and Arithmetic Curves*. Oxford Graduate Texts in Mathematics, vol. 6. Oxford University Press, Oxford (2002) Translated from the French by Reinie Ern e, Oxford Science Publications
95. Lorenzini, D., Tucker, T.J.: Thue equations and the method of Chabauty-Coleman. *Invent. Math.* **148**(1), 47–77 (2002)
96. Luo, Y.: Rank-determining sets of metric graphs. *J. Combin. Theory Ser. A* **118**, 1775–1793 (2011)
97. Luo, Y., Manjunath, M.: Smoothing of limit linear series of rank one on saturated metrized complexes of algebraic curves. Preprint, arXiv:1411.2325 (2014)
98. Maclagan, D., Sturmfels, B.: *Introduction to Tropical Geometry*. Graduate Studies in Mathematics, vol. 161, pp. xii+363. American Mathematical Society, Providence (2015)
99. McCallum, W., Poonen, B.: The method of Chabauty and Coleman. In: *Explicit Methods in Number Theory*. Panorama Synth eses, vol. 36 pp. 99–117. Soci t  Math matique de France, Paris (2012)
100. Merel, L.: Bornes pour la torsion des courbes elliptiques sur les corps de nombres. *Invent. Math.* **124**(1–3), 437–449 (1996)
101. Mikhalkin, G., Zharkov, I.: Tropical curves, their Jacobians and theta functions. In: *Curves and Abelian Varieties*. Contemporary Mathematics, pp. 203–230, vol. 465. American Mathematical Society, Providence, RI (2008)
102. Mumford, D.: An analytic construction of degenerating abelian varieties over complete rings. *Compositio Math.* **24**, 239–272 (1972)
103. Neeman, A.: The distribution of Weierstrass points on a compact Riemann surface. *Ann. Math. (2)* **120**(2), 317–328 (1984)
104. Nicaise, J.: Singular cohomology of the analytic Milnor fiber, and mixed Hodge structure on the nearby cohomology. *J. Algebraic Geom.* **20**(2), 199–237 (2011)
105. Noether, M.: Zur Grundlegung der Theorie der algebraische raumcurven. *J. Reine Angew. Math.* **93**, 271–318 (1882)
106. Ogg, A.P.: On the Weierstrass points of $X_0(N)$. III. *J. Math.* **22**(1), 31–35 (1978)

107. Osserman, B.: Limit linear series for curves not of compact type. Preprint, arXiv:1406.6699 (2014)
108. Payne, S.: Analytification is the limit of all tropicalizations. *Math. Res. Lett.* **16**(3), 543–556 (2009)
109. Perkinson, D., Perlman, J., Wilmes, J.: Primer for the algebraic geometry of sand-piles. In: *Tropical and Non-Archimedean Geometry. Contemporary Mathematics*, vol. 605, pp. 211–256. American Mathematical Society, Providence, RI (2013)
110. Rabinoff, J.: Tropical analytic geometry, Newton polygons, and tropical intersections. *Adv. Math.* **229**(6), 3192–3255 (2012)
111. Raynaud, M.: Spécialisation du foncteur de Picard. *Inst. Hautes Études Sci. Publ. Math.* **(38)**, 27–76 (1970)
112. Severi, F.: Sulla classificazione delle curve algebriche e sul teorema di esistenza di Riemann. *Rend. R. Acc. Naz. Lincei* **24**(5), 887–888 (1915)
113. Stoll, M.: Independence of rational points on twists of a given curve. *Compos. Math.* **142**(5), 1201–1214 (2006)
114. Stoll, M.: Uniform bounds for the number of rational points on hyperelliptic curves of small Mordell-Weil rank. *J. Eur. Math. Soc.* arXiv:1307.1773v4 (2013)
115. Teixidor i Bigas, M.: Injectivity of the symmetric map for line bundles. *Manuscripta Math.* **112**(4), 511–517 (2003)
116. Thuillier, A.: Théorie du potentiel sur les courbes en géométrie analytique non archimédienne. Applications à la théorie d’Arakelov. Ph.D. thesis, University of Rennes (2005)
117. Thuillier, A.: Géométrie toroïdale et géométrie analytique non archimédienne. *Manuscripta Math.* **123**(4), 381–451 (2007)
118. Viviani, F.: Tropicalizing vs. compactifying the Torelli morphism. In: *Tropical and Non-Archimedean Geometry. Contemporary Mathematics*, vol. 605, pp. 181–210. American Mathematical Society, Providence, RI (2013)
119. Voisin, C.: Sur l’application de Wahl des courbes satisfaisant la condition de Brill-Noether-Petri. *Acta Math.* **168**(3–4), 249–272 (1992)
120. Zhang, S.: Admissible pairing on a curve. *Invent. Math.* **112**(1), 171–193 (1993)