

Algebra & Number Theory

Volume 2
2008
No. 6

**Specialization of linear systems from curves to
graphs**

Matthew Baker
Appendix by Brian Conrad



mathematical sciences publishers

Specialization of linear systems from curves to graphs

Matthew Baker
Appendix by Brian Conrad

We investigate the interplay between linear systems on curves and graphs in the context of specialization of divisors on an arithmetic surface. We also provide some applications of our results to graph theory, arithmetic geometry, and tropical geometry.

1. Introduction

1A. Notation and terminology. We set the notation which will be used throughout the paper unless otherwise noted.

- G a graph, by which we will mean a finite, unweighted, connected multigraph without loop edges. We let $V(G)$ (respectively, $E(G)$) denote the set of vertices (respectively, edges) of G .
- Γ a metric graph (see Section 1D for the definition).
- $\Gamma_{\mathbb{Q}}$ the set of “rational points” of Γ (see Section 1D).
- R a complete discrete valuation ring with field of fractions K and algebraically closed residue field k .
- \overline{K} a fixed algebraic closure of K .
- X a smooth, proper, geometrically connected curve over K .
- \mathfrak{X} a proper model for X over R . For simplicity, we assume unless otherwise stated that \mathfrak{X} is regular, that the irreducible components of \mathfrak{X}_k are all smooth, and that all singularities of \mathfrak{X}_k are ordinary double points.

Unless otherwise specified, by a *smooth curve* we will always mean a smooth, proper, geometrically connected curve over a field, and by an *arithmetic surface* we will always mean a proper flat scheme \mathfrak{X} over a discrete valuation ring such that the generic fiber of \mathfrak{X} is a smooth curve. We will usually, but not always, be

MSC2000: primary 14H25, 05C99; secondary 14H51, 14H55.

Keywords: arithmetic surface, specialization, dual graph, reduction graph, Brill–Noether theory, Weierstrass point, tropical geometry.

Baker’s work was supported in part by NSF grant DMS-0600027. Conrad’s work was partially supported by NSF grant DMS-0600919.

working with *regular* arithmetic surfaces. A *model* for a smooth curve X/K is an arithmetic surface \mathfrak{X}/R whose generic fiber is X . An arithmetic surface \mathfrak{X} is called *semistable* if its special fiber \mathfrak{X}_k is reduced and has only ordinary double points as singularities. If in addition the irreducible components of \mathfrak{X}_k are all smooth (so that there are no components with self-crossings), we will say that \mathfrak{X} is *strongly semistable*.

1B. Overview. In this paper, we show that there is a close connection between linear systems on a curve X/K and linear systems, in the sense of [Baker and Norine 2007b], on the dual graph G of a regular semistable model \mathfrak{X}/R for X . A brief outline of the paper is as follows. In Section 2, we prove a basic inequality — the “Specialization Lemma” — which says that the dimension of a linear system can only go up under specialization from curves to graphs (see Lemma 2.8 for the precise statement). In the rest of the paper, we explore various applications of this fact, illustrating along the way a fruitful interaction between divisors on graphs and curves. The interplay works in both directions: for example, in Section 3 we use Brill–Noether theory for curves to prove and/or conjecture some new results about graphs (compare Theorem 3.12 and Conjecture 3.9) and, in the other direction, in Section 4 we use the notion of Weierstrass points on graphs to gain new insight into Weierstrass points on curves (see Corollaries 4.9 and 4.10).

Another fruitful interaction which emerges from our approach is a “machine” for transporting certain theorems about curves from classical to tropical algebraic geometry.¹ The connection goes through the theory of arithmetic surfaces, by way of the deformation-theoretic result proved in Appendix B, and uses the approximation method introduced in [Gathmann and Kerber 2008] to pass from \mathbb{Q} -graphs to arbitrary metric graphs, and finally to tropical curves. As an illustration of this machine, we prove an analogue for tropical curves (see Theorem 3.20 below) of the classical fact that if g , r and d are nonnegative integers for which the Brill–Noether number

$$\rho(g, r, d) := g - (r + 1)(g - d + r)$$

is nonnegative, then on every smooth curve X/\mathbb{C} of genus g there is a divisor D with

$$\dim |D| = r \quad \text{and} \quad \deg(D) \leq d.$$

We also prove, just as in the classical case of algebraic curves, that there exist Weierstrass points on every tropical curve of genus $g \geq 2$ (see Theorem 4.13).

We conclude the paper with two appendices which can be read independently of the rest of the paper. In Appendix A, we provide a reformulation of certain parts of Raynaud’s theory [1970] of “specialization of the Picard functor” in terms of linear

¹See, for example, [Mikhalkin 2006] for an introduction to tropical geometry.

systems on graphs. We also point out some useful consequences of Raynaud’s results for which we do not know any references. Although we do not actually use Raynaud’s results in the body of the paper, it should be useful for future work on the interplay between curves and graphs to highlight the compatibility between Raynaud’s theory and our notion of linear equivalence on graphs. Appendix B, written by Brian Conrad, discusses a result from the deformation theory of stable marked curves, which implies that every finite graph occurs as the dual graph of a regular semistable model for some smooth curve X/K with totally degenerate special fiber. This result, which seems known to the experts but for which we could not find a suitable reference, is used several times throughout the main body of the paper.

1C. Divisors and linear systems on graphs. By a *graph*, we will always mean a finite, connected multigraph without loop edges.

Let G be a graph, and let $V(G)$ (respectively, $E(G)$) denote the set of vertices (respectively, edges) of G . We let $\text{Div}(G)$ denote the free abelian group on $V(G)$, and refer to elements of $\text{Div}(G)$ as *divisors* on G . We can write each divisor on G as

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

with $a_v \in \mathbb{Z}$, and we will say that $D \geq 0$ if $a_v \geq 0$ for all $v \in V(G)$. We define

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

to be the *degree* of D . We let

$$\text{Div}_+(G) = \{E \in \text{Div}(G) : E \geq 0\}$$

denote the set of *effective* divisors on G , and we let $\text{Div}^0(G)$ denote the set of divisors of degree zero on G . Finally, we let $\text{Div}_+^d(G)$ denote the set of effective divisors of degree d on G .

Let $\mathcal{M}(G)$ be the group of \mathbb{Z} -valued functions on $V(G)$, and define the Laplacian operator $\Delta : \mathcal{M}(G) \rightarrow \text{Div}^0(G)$ by

$$\Delta(\varphi) = \sum_{v \in V(G)} \sum_{e=vw \in E(G)} (\varphi(v) - \varphi(w))(v).$$

We let

$$\text{Prin}(G) = \Delta(\mathcal{M}(G)) \subseteq \text{Div}^0(G)$$

be the subgroup of $\text{Div}^0(G)$ consisting of *principal divisors*.

If $D, D' \in \text{Div}(G)$, we write $D \sim D'$ if $D - D' \in \text{Prin}(G)$, and set

$$|D| = \{E \in \text{Div}(G) : E \geq 0 \text{ and } E \sim D\}.$$

We refer to $|D|$ as the (complete) linear system associated to D , and call divisors D and D' with $D \sim D'$ linearly equivalent.

Given a divisor D on G , define $r(D) = -1$ if $|D| = \emptyset$, and otherwise set

$$r(D) = \max\{k \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all } E \in \text{Div}_+^k(G)\}.$$

Note that $r(D)$ depends only on the linear equivalence class of D , and therefore can be thought of as an invariant of the complete linear system $|D|$.

When we wish to emphasize the underlying graph G , we will sometimes write $r_G(D)$ instead of $r(D)$.

We define the canonical divisor on G to be

$$K_G = \sum_{v \in V(G)} (\deg(v) - 2)(v).$$

We have $\deg(K_G) = 2g - 2$, where $g = |E(G)| - |V(G)| + 1$ is the genus (or cyclomatic number) of G .

Theorem 1.1 (Riemann–Roch for graphs [Baker and Norine 2007b, Theorem 1.12]). *Let D be a divisor on a graph G . Then*

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

For each linear ordering $<$ of the vertices of G , we define a corresponding divisor $v \in \text{Div}(G)$ of degree $g - 1$ by the formula

$$v = \sum_{v \in V(G)} (|\{e = vw \in E(G) : w < v\}| - 1)(v).$$

One of the main ingredients in the proof of Theorem 1.1, which is also quite useful for computing $r(D)$ in specific examples, is this:

Theorem 1.2 [Baker and Norine 2007b, Theorem 3.3]. *For every $D \in \text{Div}(G)$, exactly one of the following holds:*

- (1) $r(D) \geq 0$, or
- (2) $r(v - D) \geq 0$ for some divisor v associated to a linear ordering $<$ of $V(G)$.

In particular, note that $r(v) = -1$ for any divisor v associated to a linear ordering $<$ of $V(G)$.

1D. Subdivisions, metric graphs, and \mathbb{Q} -graphs. By a weighted graph, we will mean a graph in which each edge is assigned a positive real number called the length of the edge. Following the terminology of [Baker and Faber 2006], a metric graph (or metrized graph) is a compact, connected metric space Γ which arises by viewing the edges of a weighted graph G as line segments. Somewhat more formally, a metric graph should be thought of as corresponding to an equivalence

class of weighted graphs, where two weighted graphs G and G' are *equivalent* if they admit a common refinement. (A *refinement* of G is any weighted graph obtained by subdividing the edges of G in a length-preserving fashion.) A weighted graph G in the equivalence class corresponding to Γ is called a *model* for Γ . Under the correspondence between equivalence classes of weighted graphs and metric graphs, after choosing an orientation, each edge e in the model G can be identified with the real interval $[0, \ell(e)] \subseteq \Gamma$.

We let $\text{Div}(\Gamma)$ denote the free abelian group on the points of the metric space Γ , and refer to elements of $\text{Div}(\Gamma)$ as *divisors* on Γ . We can write an element $D \in \text{Div}(\Gamma)$ as

$$D = \sum_{P \in \Gamma} a_P(P)$$

with $a_P \in \mathbb{Z}$ for all P and $a_P = 0$ for all but finitely many P . We will say that $D \geq 0$ if $a_P \geq 0$ for all $P \in \Gamma$. We let

$$\deg(D) = \sum_{P \in \Gamma} a_P$$

be the *degree* of D , we let

$$\text{Div}_+(\Gamma) = \{E \in \text{Div}(\Gamma) : E \geq 0\}$$

denote the set of *effective* divisors on Γ , and we let $\text{Div}^0(\Gamma)$ denote the subgroup of divisors of degree zero on Γ . Finally, we let $\text{Div}_+^d(\Gamma)$ denote the set of effective divisors of degree d on Γ .

Following [Gathmann and Kerber 2008], a \mathbb{Q} -*graph* is a metric graph Γ having a model G whose edge lengths are rational numbers. We call such a model a *rational model* for Γ . An ordinary unweighted graph G can be thought of as a \mathbb{Q} -graph whose edge lengths are all 1. We denote by $\Gamma_{\mathbb{Q}}$ the set of points of Γ whose distance from every vertex of G is rational; we call elements of $\Gamma_{\mathbb{Q}}$ *rational points* of Γ . It is immediate that the set $\Gamma_{\mathbb{Q}}$ does not depend on the choice of a rational model G for Γ . We let $\text{Div}_{\mathbb{Q}}(\Gamma)$ be the free abelian group on $\Gamma_{\mathbb{Q}}$, and refer to elements of $\text{Div}_{\mathbb{Q}}(\Gamma)$ as \mathbb{Q} -*rational divisors* on Γ .

A *rational function* on a metric graph Γ is a continuous, piecewise affine function $f : \Gamma \rightarrow \mathbb{R}$, all of whose slopes are integers. We let $\mathcal{M}(\Gamma)$ denote the space of rational functions on Γ . The *divisor* of a rational function $f \in \mathcal{M}(\Gamma)$ is defined² as

$$(f) = - \sum_{P \in \Gamma} \sigma_P(f)(P),$$

²Here we follow the sign conventions from [Baker and Faber 2006]. In [Gathmann and Kerber 2008], the divisor of f is defined to be the negative of the one we define here.

where $\sigma_P(f)$ is the sum of the slopes of Γ in all directions emanating from P . We let

$$\text{Prin}(\Gamma) = \{(f) : f \in \mathcal{M}(\Gamma)\}$$

be the subgroup of $\text{Div}(\Gamma)$ consisting of *principal divisors*. It follows from [Baker and Faber 2006, Corollary 1] that (f) has degree zero for all $f \in \mathcal{M}(\Gamma)$, that is,

$$\text{Prin}(\Gamma) \subseteq \text{Div}^0(\Gamma).$$

If Γ is a \mathbb{Q} -graph, we denote by $\text{Prin}_{\mathbb{Q}}(\Gamma)$ the group of principal divisors supported on $\Gamma_{\mathbb{Q}}$.

Remark 1.3. As explained in [Baker and Faber 2006], if we identify a rational function $f \in \mathcal{M}(\Gamma)$ with its restriction to the vertices of any model G for which f is affine along each edge of G , then (f) can be naturally identified with the *weighted combinatorial Laplacian* $\Delta(f)$ of f on G .

If $D, D' \in \text{Div}(\Gamma)$, we write $D \sim D'$ if $D - D' \in \text{Prin}(\Gamma)$, and set

$$|D|_{\mathbb{Q}} = \{E \in \text{Div}_{\mathbb{Q}}(\Gamma) : E \geq 0 \text{ and } E \sim D\}$$

and

$$|D| = \{E \in \text{Div}(\Gamma) : E \geq 0 \text{ and } E \sim D\}.$$

It is straightforward using Remark 1.3 to show that if G is a graph and Γ is the corresponding \mathbb{Q} -graph all of whose edge lengths are 1, then two divisors $D, D' \in \text{Div}(G)$ are equivalent on G (in the sense of Section 1C) if and only if they are equivalent on Γ in the sense just defined.

Given a \mathbb{Q} -graph Γ and a \mathbb{Q} -rational divisor D on Γ , define $r_{\mathbb{Q}}(D) = -1$ if $|D|_{\mathbb{Q}} = \emptyset$, and otherwise set

$$r_{\mathbb{Q}}(D) = \max\{k \in \mathbb{Z} : |D - E|_{\mathbb{Q}} \neq \emptyset \text{ for all } E \in \text{Div}_{\mathbb{Q}}(\Gamma) \text{ with } E \geq 0, \text{deg}(E) = k\}.$$

Similarly, given an arbitrary metric graph Γ and a divisor D on Γ , we define $r_{\Gamma}(D) = -1$ if $|D| = \emptyset$, and otherwise set

$$r_{\Gamma}(D) = \max\{k \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all } E \in \text{Div}_{+}^k(\Gamma)\}.$$

Let k be a positive integer, and let $\sigma_k(G)$ be the graph obtained from the (ordinary unweighted) graph G by subdividing each edge of G into k edges. We call $\sigma_k(G)$ the k -th *regular subdivision* of G . A divisor D on G can also be thought of as a divisor on $\sigma_k(G)$ for all $k \geq 1$ in the obvious way. The following recent combinatorial result, which had been conjectured by the author, relates the quantities $r(D)$ on G and $\sigma_k(G)$:

Theorem 1.4 [Hladky–Král–Norine 2007]. *Let G be a graph. If $D \in \text{Div}(G)$, then for every integer $k \geq 1$, we have*

$$r_G(D) = r_{\sigma_k(G)}(D).$$

When working with metric graphs, if Γ is the metric graph corresponding to G (in which every edge of G has length 1), then we will usually think of each edge of $\sigma_k(G)$ as having length $1/k$; in this way, each finite graph $\sigma_k(G)$ can be viewed as a model for the same underlying metric \mathbb{Q} -graph Γ .

It is evident from the definitions that $r_{\mathbb{Q}}(D)$ and $r_{\Gamma}(D)$ do not change if the length of every edge in (some model for) Γ is multiplied by a positive integer k . Using this observation, together with [Gathmann and Kerber 2008, Proposition 2.4], one deduces from Theorem 1.4 this result:

Corollary 1.5. *If G is a graph and Γ is the corresponding metric \mathbb{Q} -graph in which every edge of G has length 1, then for every divisor D on G we have*

$$r_G(D) = r_{\mathbb{Q}}(D) = r_{\Gamma}(D). \tag{1.6}$$

By Corollary 1.5, we may unambiguously write $r(D)$ to refer to any of the three quantities appearing in (1.6).

Remark 1.7. Our only use of Theorem 1.4 in this paper, other than the notational convenience of not having to worry about the distinction between $r_{\mathbb{Q}}(D)$ and $r_G(D)$, will be in Remark 4.1. In practice, however, Theorem 1.4 is quite useful, since without it there is no obvious way to calculate the quantity $r_{\mathbb{Q}}(D)$ for a divisor D on a metric \mathbb{Q} -graph Γ .

On the other hand, we *will* make use of the equality $r_{\mathbb{Q}}(D) = r_{\Gamma}(D)$ given by [Gathmann and Kerber 2008, Proposition 2.4] when we develop a machine for deducing theorems about metric graphs from the corresponding results for \mathbb{Q} -graphs (see Section 3D).

Finally, we recall the statement of the Riemann–Roch theorem for metric graphs. Define the *canonical divisor* on Γ to be

$$K_{\Gamma} = \sum_{v \in V(G)} (\deg(v) - 2)(v)$$

for any model G of Γ . It is easy to see that K_{Γ} is independent of the choice of a model G , and that

$$\deg(K_{\Gamma}) = 2g - 2,$$

where $g = |E(G)| - |V(G)| + 1$ is the *genus* (or *cyclomatic number*) of Γ .

The following result is proved in [Gathmann and Kerber 2008] and [Mikhalkin and Zharkov 2007]:

Theorem 1.8 (Riemann–Roch for metric graphs). *Let D be a divisor on a metric graph Γ . Then*

$$r_{\Gamma}(D) - r_{\Gamma}(K_{\Gamma} - D) = \deg(D) + 1 - g. \quad (1.9)$$

By Corollary 1.5, there is a natural “compatibility” between Theorems 1.1 and 1.8.

1E. Tropical curves. Tropical geometry is a relatively recent and highly active area of research, and in dimension one it is closely connected with the theory of metric graphs as discussed in the previous section. For the sake of brevity, we adopt a rather minimalist view of tropical curves in this paper; the interested reader should see [Gathmann and Kerber 2008; Mikhalkin 2006; Mikhalkin and Zharkov 2007] for motivation and a more extensive discussion.

Following [Gathmann and Kerber 2008, §1], we define a *tropical curve* to be a “metric graph with possibly unbounded ends”. More concretely, a tropical curve $\tilde{\Gamma}$ can be thought of as the geometric realization of a pair (G, ℓ) , where G is a graph and

$$\ell : E(G) \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$$

is a length function; each edge e of G having finite length is identified with the real interval $[0, \ell(e)]$, and each edge of length ∞ is identified with the extended real interval $[0, +\infty]$ in such a way that the ∞ endpoint of the edge has valence 1. The main difference between a tropical curve $\tilde{\Gamma}$ and a metric graph Γ is that we allow finitely many edges of $\tilde{\Gamma}$ to have infinite length.³ In particular, every metric graph is also a tropical curve.

One can define divisors, rational functions, and linear equivalence for tropical curves exactly as we have done in Section 1D for metric graphs; see [Gathmann and Kerber 2008, §1 and Definition 3.2] for details. (The only real difference is that one must allow a rational function to take the values $\pm\infty$ at the unbounded ends of $\tilde{\Gamma}$.) Using the tropical notion of linear equivalence, one defines $r_{\tilde{\Gamma}}(D)$ for divisors on tropical curves just as we defined $r_{\Gamma}(D)$ in Section 1D for divisors on metric graphs. With these definitions in place, the Riemann–Roch formula (1.9) holds in the context of tropical curves, a result which can be deduced easily from Theorem 1.8 (see [Gathmann and Kerber 2008, §3] for details).

2. The specialization lemma

In this section, we investigate the behavior of the quantity $r(D)$ under specialization from curves to graphs. In order to make this precise, we first need to introduce some notation and background facts concerning divisors on arithmetic surfaces.

³Unlike some other definitions in the literature, the definition of a tropical curve from [Gathmann and Kerber 2008] allows vertices of valence 1 and 2, and requires that there is a “point at infinity” at the end of each unbounded edge.

2A. The specialization map. Let R be a complete discrete valuation ring with field of fractions K and algebraically closed residue field k . Let X be a smooth curve over K , and let \mathfrak{X}/R be a strongly semistable regular model for X with special fiber \mathfrak{X}_k (see Section 1A). We let

$$\mathcal{C} = \{C_1, \dots, C_n\}$$

be the set of irreducible components of \mathfrak{X}_k .

Let G be the *dual graph* of \mathfrak{X}_k , that is, G is the finite graph whose vertices v_i correspond to the irreducible components C_i of \mathfrak{X}_k , and whose edges correspond to intersections between these components (so that there is one edge between v_i and v_j for each point of intersection between C_i and C_j). The assumption that \mathfrak{X} is strongly semistable implies that G is well-defined and has no loop edges.

We let $\text{Div}(X)$ (respectively, $\text{Div}(\mathfrak{X})$) be the group of Cartier divisors on X (respectively, on \mathfrak{X}); since X is smooth and \mathfrak{X} is regular, Cartier divisors on X (respectively, \mathfrak{X}) are the same as Weil divisors. Recall that when K is perfect, $\text{Div}(X)$ can be identified with the group of $\text{Gal}(\bar{K}/K)$ -invariant elements of the free abelian group $\text{Div}(X(\bar{K}))$ on $X(\bar{K})$.

We also let $\text{Prin}(X)$ (respectively, $\text{Prin}(\mathfrak{X})$) denote the group of principal Cartier divisors on X (respectively, \mathfrak{X}).

There is a natural inclusion $\mathcal{C} \subset \text{Div}(\mathfrak{X})$, and an intersection pairing

$$\mathcal{C} \times \text{Div}(\mathfrak{X}) \rightarrow \mathbb{Z}, \quad (C_i, \mathcal{D}) \mapsto (C_i \cdot \mathcal{D}),$$

where

$$(C_i \cdot \mathcal{D}) = \text{deg}(\mathcal{O}_{\mathfrak{X}}(\mathcal{D})|_{C_i}).$$

The intersection pairing gives rise to a homomorphism $\rho : \text{Div}(\mathfrak{X}) \rightarrow \text{Div}(G)$ by the formula

$$\rho(\mathcal{D}) = \sum_{v_i \in V(G)} (C_i \cdot \mathcal{D})(v_i).$$

(If we wish to emphasize the dependence on the ground field K , we will sometimes write $\rho_K(\mathcal{D})$ instead of $\rho(\mathcal{D})$.) We call the homomorphism ρ the *specialization map*. By intersection theory, the group $\text{Prin}(\mathfrak{X})$ is contained in the kernel of ρ .

The Zariski closure in \mathfrak{X} of an effective divisor on X is a Cartier divisor. Extending by linearity, we can associate to each $D \in \text{Div}(X)$ a Cartier divisor $\text{cl}(D)$ on \mathfrak{X} , which we refer to as the *Zariski closure* of D . By abuse of terminology, we will also denote by ρ the composition of $\rho : \text{Div}(\mathfrak{X}) \rightarrow \text{Div}(G)$ with the map cl . By construction, if $D \in \text{Div}(X)$ is *effective*, then $\rho(D)$ is an effective divisor on G . Furthermore, $\rho : \text{Div}(X) \rightarrow \text{Div}(G)$ is a degree-preserving homomorphism.

A divisor $\mathcal{D} \in \text{Div}(\mathfrak{X})$ is called *vertical* if it is supported on \mathfrak{X}_k , and *horizontal* if it is the Zariski closure of a divisor on X . If \mathcal{D} is a vertical divisor, then

$$\rho(\mathcal{D}) \in \text{Prin}(G),$$

which follows from the fact that $\rho(C_i)$ is the negative Laplacian of the characteristic function of the vertex v_i . Since every divisor $\mathcal{D} \in \text{Div}(\mathfrak{X})$ can be written uniquely as

$$\mathcal{D}_h + \mathcal{D}_v$$

with \mathcal{D}_h horizontal and \mathcal{D}_v vertical, it follows that $\rho(\mathcal{D})$ and $\rho(\mathcal{D}_h)$ are linearly equivalent divisors on G .

Consequently, if $D \in \text{Prin}(X)$, then although the horizontal divisor $\mathcal{D} := \text{cl}(D)$ may not belong to $\text{Prin}(\mathfrak{X})$, it differs from a principal divisor $\mathcal{D}' \in \text{Prin}(\mathfrak{X})$ by a vertical divisor $\mathcal{F} \in \text{Div}(\mathfrak{X})$ for which $\rho(\mathcal{F}) \in \text{Prin}(G)$. Thus we deduce a basic fact:

Lemma 2.1. *If $D \in \text{Prin}(X)$, then $\rho(D) \in \text{Prin}(G)$.*

When D corresponds to a Weil divisor

$$\sum_{P \in X(K)} n_P(P)$$

supported on $X(K)$, there is another, more concrete, description of the map ρ . Since \mathfrak{X} is regular, each point $P \in X(K) = \mathfrak{X}(R)$ specializes to a nonsingular point of \mathfrak{X}_k , and hence to a well-defined irreducible component $c(P) \in \mathcal{C}$, which we may identify with a vertex $v(P) \in V(G)$. Then by [Liu 2002, Corollary 9.1.32], we have

$$\rho(D) = \sum_P n_P(v(P)). \tag{2.2}$$

Remark 2.3. Since the natural map from $X(K) = \mathfrak{X}(R)$ to the smooth locus of $\mathfrak{X}_k(k)$ is surjective (see, for example, [Liu 2002, Proposition 10.1.40(b)]), it follows from (2.2) that $\rho : \text{Div}(X) \rightarrow \text{Div}(G)$ is surjective. In fact, this implies the stronger fact that the restriction of ρ to $\text{Div}(X(K))$ (the free abelian group on $X(K)$) is surjective.

2B. Behavior of $r(D)$ under specialization. Let $D \in \text{Div}(X)$, and let $\bar{D} = \rho(D) \in \text{Div}(G)$ be its specialization to G . We want to compare the dimension of the complete linear system $|D|$ on X (in the sense of classical algebraic geometry) with the quantity $r(\bar{D})$ defined in Section 1C. In order to do this, we first need some simple facts about linear systems on curves.

We temporarily suspend our convention that K denotes the field of fractions of a discrete valuation ring R , and allow K to be an arbitrary field, with X still denoting

a smooth curve over K . We let $\text{Div}(X(K))$ denote the free abelian group on $X(K)$, which we can view in a natural way as a subgroup of $\text{Div}(X)$. For $D \in \text{Div}(X)$, let

$$|D| = \{E \in \text{Div}(X) : E \geq 0, E \sim D\}.$$

Set $r(D) = -1$ if $|D| = \emptyset$, and otherwise put

$$r_{X(K)}(D) := \max\{k \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all } E \in \text{Div}_+^k(X(K))\}.$$

Lemma 2.4. *Let X be a smooth curve over a field K , and assume that $X(K)$ is infinite. Then for $D \in \text{Div}(X)$, we have*

$$r_{X(K)}(D) = \dim_K L(D) - 1,$$

where $L(D) = \{f \in K(X) : (f) + D \geq 0\} \cup \{0\}$.

Proof. It is well known that $\dim L(D - P) \geq \dim L(D) - 1$ for all $P \in X(K)$. If $\dim L(D) \geq k + 1$, it follows that for any points $P_1, \dots, P_k \in X(K)$ we have $\dim L(D - P_1 - \dots - P_k) \geq 1$, so

$$L(D - P_1 - \dots - P_k) \neq (0) \quad \text{and} \quad |D - P_1 - \dots - P_k| \neq \emptyset.$$

Conversely, we prove by induction on k that if $\dim L(D) = k$, then there exist $P_1, \dots, P_k \in X(K)$ such that $L(D - P_1 - \dots - P_k) = (0)$, that is, $|D - P_1 - \dots - P_k| = \emptyset$. This is clearly true for the base case $k = 0$. Suppose $\dim L(D) = k \geq 1$, and choose a nonzero rational function $f \in L(D)$. Since f has only finitely many zeros and $X(K)$ is infinite, there exists $P = P_1 \in X(K)$ for which $f(P) \neq 0$. It follows that $L(D - P) \subsetneq L(D)$, so $\dim L(D - P) = k - 1$. By induction, there exist $P_2, \dots, P_k \in X(K)$ such that $|D - P - P_2 - \dots - P_k| = \emptyset$, which proves what we want. \square

Since $\dim L(D)$ remains constant under base change by an arbitrary field extension K'/K , we conclude:

Corollary 2.5. *Let X be a smooth curve over a field K , and assume that $X(K)$ is infinite. Let K' be any extension field. Then for $D \in \text{Div}(X)$, we have*

$$r_{X(K)}(D) = r_{X(K')}(D).$$

In view of Lemma 2.4 and Corollary 2.5, we will simply write $r_X(D)$, or even just $r(D)$, to denote the quantity $r_{X(K)}(D) = \dim_K L(D) - 1$.

Lemma 2.6. *Let X be a smooth curve over a field K , and assume that $X(K)$ is infinite. If $D \in \text{Div}(X)$, then $r(D - P) \geq r(D) - 1$ for all $P \in X(K)$, and if $r(D) \geq 0$, then $r(D - P) = r(D) - 1$ for some $P \in X(K)$.*

Proof. Let $k = r(D)$. The result is clear for $r(D) \leq 0$, so we may assume that $k \geq 1$. If $P = P_1, P_2, \dots, P_k \in X(K)$ are arbitrary, then since $r(D) \geq k$, we have

$$|D - P - P_2 - \dots - P_k| \neq \emptyset,$$

and therefore $r(D - P) \geq k - 1$. Also, since $r(D) = k$, it follows that there exist $P = P_1, P_2, \dots, P_{k+1} \in X(K)$ such that

$$|D - P - P_2 - \dots - P_{k+1}| = \emptyset,$$

and therefore $r(D - P) \leq k - 1$ for this particular choice of P . □

The same proof shows that an analogous result holds in the context of graphs:

Lemma 2.7. *Let G be a graph, and let $D \in \text{Div}(G)$. Then $r(D - P) \geq r(D) - 1$ for all $P \in V(G)$, and if $r(D) \geq 0$, then $r(D - P) = r(D) - 1$ for some $P \in V(G)$.*

We now come to the main result of this section. Returning to our conventional notation, we let R be a complete discrete valuation ring with field of fractions K and algebraically closed residue field k . We let X be a smooth curve over K , and let \mathfrak{X}/R be a strongly semistable regular model for X with special fiber \mathfrak{X}_k and dual graph G .

Lemma 2.8 (Specialization Lemma). *For all $D \in \text{Div}(X)$, we have*

$$r_G(\rho(D)) \geq r_X(D).$$

Proof. Let $\bar{D} := \rho(D)$. We prove by induction on k that if $r_X(D) \geq k$, then $r_G(\bar{D}) \geq k$ as well. The base case $k = -1$ is obvious. Now suppose $k = 0$, so $r_X(D) \geq 0$. Then there exists an effective divisor $E \in \text{Div}(X)$ with $D \sim E$, so that $D - E \in \text{Prin}(X)$. Since ρ is a homomorphism and takes principal (respectively, effective) divisors on X to principal (respectively, effective) divisors on G , we have

$$\bar{D} = \rho(D) \sim \rho(E) \geq 0,$$

so $r_G(\bar{D}) \geq 0$ as well.

We may therefore assume that $k \geq 1$. Let $\bar{P} \in V(G)$ be arbitrary. By Remark 2.3, there exists $P \in X(K)$ such that $\rho(P) = \bar{P}$. Then $r_X(D - P) \geq k - 1$, so by induction we have

$$r_G(\bar{D} - \bar{P}) \geq k - 1$$

as well (and in particular, $r_G(\bar{D}) \geq 0$). Since this is true for all $\bar{P} \in V(G)$, it follows from Lemma 2.7 that $r_G(\bar{D}) \geq k$ as desired. □

Remark 2.9. In the situation of Lemma 2.8, it can certainly happen that

$$r_G(\rho(D)) > r_X(D).$$

For example, on the dual graph of the modular curve $X_0(73)$ there is an effective divisor \bar{D} of degree 2 with $r(\bar{D}) = 1$, but since $X_0(73)$ is not hyperelliptic, $r(D) = 0$ for every effective divisor D of degree 2 on X with $\rho(D) = \bar{D}$ (see Example 3.6).

2C. Compatibility with base change. Let K'/K be a finite extension, let R' be the valuation ring of K' , and let

$$X_{K'} := X \times_K K'.$$

It is known that there is a unique *relatively minimal* regular semistable model \mathfrak{X}'/R' which dominates $\mathfrak{X} \times_R R'$, and the dual graph G' of the special fiber of \mathfrak{X}' is isomorphic to $\sigma_e(G)$, where e is the ramification index of K'/K . If we assign a length of $\frac{1}{e}$ to each edge of G' , then the corresponding *metric graph* is the same for all finite extensions K'/K . In other words, G and G' are different models for the same metric \mathbb{Q} -graph Γ , which we call the *reduction graph* associated to the model \mathfrak{X}/R (see [Chinburg and Rumely 1993] for further discussion).

The discussion in [Chinburg and Rumely 1993], together with (2.2), shows that there is a unique surjective map $\tau : X(\bar{K}) \rightarrow \Gamma_{\mathbb{Q}}$ for which the induced homomorphism

$$\tau_* : \text{Div}(X_{\bar{K}}) \cong \text{Div}(X(\bar{K})) \rightarrow \text{Div}_{\mathbb{Q}}(\Gamma)$$

is compatible with ρ , in the sense that for $D \in \text{Div}(X(K'))$, we have

$$\tau_*(D) = \rho_{K'}(D).$$

Concretely, if K'/K is a finite extension and $P \in X(K')$, then $\tau(P)$ is the point of $\Gamma_{\mathbb{Q}}$ corresponding to the irreducible component of the special fiber of \mathfrak{X}' to which P specializes.

Remark 2.10. In general, for $D \in \text{Div}(X)$ we will not always have $\rho(D) = \tau_*(D)$, but $\rho(D)$ and $\tau_*(D)$ will at least be linearly equivalent as divisors on the metric graph Γ . (This is a consequence of standard facts from arithmetic intersection theory, see, for example [Liu 2002, Propositions 9.2.15 and 9.2.23].)

From the discussion in Section 1D, we deduce from the proof of Lemma 2.8:

Corollary 2.11. *Let $D \in \text{Div}(X_{\bar{K}})$. Then*

$$r_{\mathbb{Q}}(\tau_*(D)) \geq r_X(D).$$

Remark 2.12. For the reader familiar with Berkovich’s theory of analytic spaces [1990], it may be helpful to remark that the metric \mathbb{Q} -graph Γ can be identified with the *skeleton* of the formal model associated to \mathfrak{X} , and the map $\tau : X(\bar{K}) \rightarrow \Gamma_{\mathbb{Q}}$ can be identified with the restriction to $X(\bar{K}) \subset X^{\text{an}}$ of the natural deformation retraction $X^{\text{an}} \twoheadrightarrow \Gamma$, where X^{an} denotes the Berkovich K -analytic space associated to X .

3. Some applications of the specialization lemma

3A. Specialization of g_d^r 's. Recall that a *complete* g_d^r on X/K is defined to be a complete linear system $|D|$ with

$$D \in \text{Div}(X_{\bar{K}}), \quad \deg(D) = d, \quad \text{and} \quad r(D) = r.$$

For simplicity, we will omit the word “complete” and just refer to such a linear system as a g_d^r . A g_d^r is called *K-rational* if we can choose the divisor D to lie in $\text{Div}(X)$.

By analogy, we define a g_d^r on a graph G (respectively, a metric graph Γ) to be a complete linear system $|D|$ with $D \in \text{Div}(G)$ (respectively, $D \in \text{Div}(\Gamma)$) such that $\deg(D) = d$ and $r(D) = r$. Also, we will denote by $g_{\leq d}^r$ (respectively, $g_d^{\geq r}$) a complete linear system $|D|$ with $\deg(D) \leq d$ and $r(D) = r$ (respectively, $\deg(D) = d$ and $r(D) \geq r$).

As an immediate consequence of Lemma 2.8 and Corollary 2.11, we obtain:

Corollary 3.1. *Let X be a smooth curve over K , and let \mathfrak{X}/R be a strongly semistable regular model for X with special fiber \mathfrak{X}_k . Let G be the dual graph of \mathfrak{X}_k , and let Γ be the corresponding metric graph. If there exists a K -rational g_d^r on X , then there exists a $g_d^{\geq r}$ and a $g_{\leq d}^r$ on G . Similarly, if there exists a g_d^r on X , then there exists a $g_d^{\geq r}$ and a $g_{\leq d}^r$ on Γ .*

This result places restrictions on the possible graphs which can appear as the dual graph of some regular model of a given curve X/K .

In the particular case $r = 1$, we refer to the smallest positive integer d for which there exists a g_d^1 (respectively, a K -rational g_d^1) on X as the *gonality* (respectively, *K-gonality*) of X .

Similarly, we define the *gonality* of a graph G (or a metric graph Γ) to be the smallest positive integer d for which there exists a g_d^1 on G (or Γ).

As a special case of Corollary 3.1, we have:

Corollary 3.2. *The gonality of G (respectively, Γ) is at most the K -gonality (respectively, gonality) of X .*

Example 3.3. Let K_n denote the complete graph on $n \geq 2$ vertices.

Claim. *The gonality of K_n is equal to $n - 1$.*

Indeed, let $D = \sum a_v(v)$ be an effective divisor of degree at most $n - 2$ on K_n , and label the vertices v_1, \dots, v_n of G so that $a_{v_1} \leq \dots \leq a_{v_n}$. Then it is easy to see that $a_{v_1} = 0$ and $a_{v_i} \leq i - 2$ for all $2 \leq i \leq n$. If ν is the divisor associated to the linear ordering $v_1 < \dots < v_n$ of $V(K_n)$, it follows that $D - (v_1) \leq \nu$, so $r(D - (v_1)) = -1$ by Theorem 1.2. In particular, we have $r(D) \leq 0$, and thus the

gonality of K_n is at least $n - 1$. On the other hand, for any vertex $v_0 \in V(K_n)$, the divisor

$$D = \sum_{v \in V(K_n) \setminus \{v_0\}} (v)$$

has degree $n - 1$ and rank at least 1, since

$$D - (v_0) \sim (n - 2)(v_0).$$

It follows from Corollary 3.2 that if X/K has K -gonality at most $n - 2$, then no regular model \mathfrak{X}/R for X can have K_n as its dual graph. For example, K_4 cannot be the dual graph of any regular model of a hyperelliptic curve X/K .

3B. Hyperelliptic graphs. Focusing now on the special case $d = 2$, we recall from [Hartshorne 1977, §IV.5] that a smooth curve X/K of genus g is called *hyperelliptic* if $g \geq 2$ and there exists a g_2^1 on X . If such a g_2^1 exists, it is automatically unique and K -rational.

Similarly, we say that a graph G (or a metric graph Γ) of genus g is *hyperelliptic* if $g \geq 2$ and there exists a g_2^1 on G (or Γ).

Remark 3.4. One can show that if such a g_2^1 exists, it is automatically unique. Also, if G is 2-edge-connected of genus at least 2, then G is hyperelliptic if and only if there is an involution h on G for which the quotient graph $G/\langle h \rangle$ is a tree. These and other matters are discussed in [Baker and Norine 2007a].

By Clifford’s theorem for graphs [Baker and Norine 2007b, Corollary 3.5], if $g \geq 2$ and D is a divisor of degree 2 on G with $r(D) \geq 1$, then in fact $r(D) = 1$, and thus G is hyperelliptic. Combining this observation with Corollary 3.1, we find:

Corollary 3.5. *If X is hyperelliptic and G has genus at least 2, then G is hyperelliptic as well.*

The converse is false, as the following example shows.

Example 3.6. (1) Let $G = B_n$ be the “banana graph” of genus $n - 1$ consisting of 2 vertices Q_1, Q_2 connected by $n \geq 3$ edges. Then the divisor $D = (Q_1) + (Q_2)$ on G has degree 2 and $r(D) = 1$, so G is hyperelliptic. On the other hand, there are certainly nonhyperelliptic curves X possessing a regular strongly semistable model with dual graph G . For example, let $p \equiv 1 \pmod{12}$ be prime, and let $K = \mathbb{Q}_p^{\text{nr}}$ be the completion of the maximal unramified extension of \mathbb{Q}_p . Then the modular curve $X_0(p)$ has a regular semistable model over K (the “Deligne–Rapoport model” [1973]) whose dual graph is isomorphic to B_n with $n = \frac{p-1}{12}$. However, by a result of Ogg [1974], $X_0(p)$ is never hyperelliptic when $p > 71$.

(2) More generally, let $G = B(\ell_1, \dots, \ell_n)$ be the graph obtained by subdividing the i -th edge of B_n into ℓ_i edges for $1 \leq i \leq n$ (so that $B(1, 1, \dots, 1) = B_n$). Then one

easily checks that $|(Q_1) + (Q_2)|$ is still a g_2^1 , so G is hyperelliptic. The dual graph of \mathfrak{X} is always of this type when the special fiber of \mathfrak{X}_k consists of two projective lines over k intersecting transversely at n points. For example, the modular curve $X_0(p)$ with $p \geq 23$ prime has a regular model whose dual graph G is of this type. For all primes $p > 71$, G is hyperelliptic even though $X_0(p)$ is not.

Remark 3.7. Every graph of genus 2 is hyperelliptic, since by the Riemann–Roch theorem for graphs, the canonical divisor K_G has degree 2 and dimension 1. It is not hard to prove that for every integer $g \geq 3$, there are both hyperelliptic and nonhyperelliptic graphs of genus g .

3C. Brill–Noether theory for graphs. Classically, it is known by *Brill–Noether theory* that every smooth curve of genus g over the complex numbers has gonality at most $\lfloor \frac{1}{2}(g+3) \rfloor$, and this bound is tight: for every $g \geq 0$, the general curve of genus g has gonality exactly equal to $\lfloor \frac{1}{2}(g+3) \rfloor$. More generally:

Theorem 3.8 (Classical Brill–Noether theory). *Fix integers $g, r, d \geq 0$, and define the Brill–Noether number*

$$\rho(g, r, d) = g - (r+1)(g-d+r).$$

Then

- (1) *if $\rho(g, r, d) \geq 0$, then every smooth curve X/\mathbb{C} of genus g has a divisor D with $r(D) = r$ and $\deg(D) \leq d$;*
- (2) *if $\rho(g, r, d) < 0$, then on a general smooth curve of genus g , there is no divisor D with $r(D) = r$ and $\deg(D) \leq d$.*

Based on extensive computer calculations by Adam Tart (an undergraduate at Georgia Tech), we conjecture that similar results hold in the purely combinatorial setting of finite graphs:

Conjecture 3.9 (Brill–Noether conjecture for graphs). *Fix integers $g, r, d \geq 0$, and set*

$$\rho(g, r, d) = g - (r+1)(g-d+r).$$

Then

- (1) *if $\rho(g, r, d) \geq 0$, then every graph of genus g has a divisor D with $r(D) = r$ and $\deg(D) \leq d$;*
- (2) *if $\rho(g, r, d) < 0$, there exists a graph of genus g for which there is no divisor D with $r(D) = r$ and $\deg(D) \leq d$.*

In the special case $r = 1$, Conjecture 3.9 can be reformulated as follows:

Conjecture 3.10 (Gonality conjecture for graphs). *For each integer $g \geq 0$,*

- (1) *the gonality of any graph of genus g is at most $\lfloor \frac{1}{2}(g+3) \rfloor$ and*

(2) *there exists a graph of genus g with gonality exactly $\lfloor \frac{1}{2}(g+3) \rfloor$.*

Adam Tart has verified Conjecture 3.10 (2) for $g \leq 12$, and Conjecture 3.9 (2) for $2 \leq r \leq 4$ and $g \leq 10$. He has also verified that Conjecture 3.10 (1) holds for approximately 1000 randomly generated graphs of genus at most 10, and has similarly verified Conjecture 3.9 (1) for around 100 graphs in the case $2 \leq r \leq 4$.

Although we do not know how to handle the general case, it is easy to prove that Conjecture 3.10 (1) holds for small values of g :

Lemma 3.11. *Conjecture 3.10 (1) is true for $g \leq 3$.*

Proof. For $g \leq 2$, this is a straightforward consequence of Riemann–Roch for graphs. For $g = 3$, we argue as follows. The canonical divisor K_G on any genus 3 graph G has degree 4 and $r(K_G) = 2$. By Lemma 2.7, there exists a vertex $P \in V(G)$ for which the degree 3 divisor $K_G - P$ has rank 1. (In fact, it is not hard to see that $r(K_G - P) = 1$ for every vertex P .) Therefore G has a g_3^1 , so the gonality of G is at most 3, proving the lemma. \square

For *metric \mathbb{Q} -graphs*, we can prove the analogue of Conjecture 3.9 (1) using the Specialization Lemma (Lemma 2.8) and classical Brill–Noether theory (Theorem 3.8):

Theorem 3.12. *Fix nonnegative integers g, r , and d for which $g \geq (r+1)(g-d+r)$. Then every metric \mathbb{Q} -graph Γ of genus g has a divisor D with $r(D) = r$ and $\deg(D) \leq d$.*

Proof. By uniformly rescaling the edges of a suitable model for Γ so that they all have integer lengths, then adding vertices of valence 2 as necessary, we may assume that Γ is the metric graph associated to a graph G (and that every edge of G has length 1). By Theorem B.2, there exists a strongly semistable regular arithmetic surface \mathfrak{X}/R whose generic fiber is smooth of genus g and whose special fiber has dual graph G . By classical Brill–Noether theory, there exists a $g_{d'}^r$ on X for some $d' \leq d$, so according to (the proof of) Corollary 3.1, there is a \mathbb{Q} -rational divisor D on Γ with $\deg(D) \leq d$ and $r(D) = r$. \square

In Section 3D, we will generalize Theorem 3.12 to arbitrary metric graphs, and then to tropical curves, using ideas from [Gathmann and Kerber 2008].

Remark 3.13. It would be very interesting to give a direct combinatorial proof of Theorem 3.12. In any case, we view the above proof of Theorem 3.12 as an example of how one can use the Specialization Lemma, in conjunction with known theorems about algebraic curves, to prove nontrivial results about graphs (or more precisely, in this case, *metric graphs*).

For a given graph G (or metric graph Γ) and an integer $r \geq 1$, let $D(G, r)$ (or $D(\Gamma, r)$) be the minimal degree d of a g_d^r on G (or Γ).

Conjecture 3.14. *Let G be a graph, and let Γ be the associated \mathbb{Q} -graph. Then for every $r \geq 1$, we have*

- (1) $D(G, r) = D(\sigma_k(G), r)$ for all $k \geq 1$ and
- (2) $D(G, r) = D(\Gamma, r)$.

Adam Tart has verified Conjecture 3.14 (1) for 100 different graphs with $1 \leq r, k \leq 4$, and for 1000 randomly generated graphs of genus up to 10 in the special case where $r = 1$ and $k = 2$ or 3.

Note that Conjecture 3.14, in conjunction with Theorem 3.12, would imply Conjecture 3.9 (1).

Finally, we have the analogue of Conjecture 3.9 (2) for metric graphs:

Conjecture 3.15 (Brill–Noether conjecture for metric graphs). *Fix integers $g, r, d \geq 0$, and set*

$$\rho(g, r, d) = g - (r + 1)(g - d + r).$$

If $\rho(g, r, d) < 0$, then there exists a metric graph of genus g for which there is no divisor D with $r(D) = r$ and $\deg(D) \leq d$.

Note that Conjecture 3.15 would follow from Conjecture 3.14 and Conjecture 3.9 (2).

Remark 3.16. By a simple argument based on Theorem B.2 and Corollary 3.2, a direct combinatorial proof of Conjecture 3.15 in the special case $r = 1$ would yield a new proof of the classical fact that for every $g \geq 0$, there exists a smooth curve X of genus g over an algebraically closed field of characteristic zero having gonality at least $\lfloor \frac{1}{2}(g + 3) \rfloor$.

3D. A tropical Brill–Noether theorem. In this section, we show how the ideas from [Gathmann and Kerber 2008] can be used to generalize Theorem 3.12 from metric \mathbb{Q} -graphs to arbitrary metric graphs and tropical curves.

The key result is the following “Semicontinuity Lemma”, which allows one to transfer certain results about divisors on \mathbb{Q} -graphs to arbitrary metric graphs. For the statement, fix a metric graph Γ and a positive real number ϵ smaller than all edge lengths in some fixed model G for Γ . We denote by $A_\epsilon(\Gamma)$ the “moduli space” of all metric graphs that are of the same combinatorial type as Γ , and whose edge lengths are within ϵ of the corresponding edge lengths in Γ . Then

$$A_\epsilon(\Gamma) \cong \prod_{e \in E(G)} [\ell(e) - \epsilon, \ell(e) + \epsilon]$$

can naturally be viewed as a product of closed intervals. In particular, there is a well-defined notion of convergence in $A_\epsilon(\Gamma)$.

Similarly, for each positive integer d , we define

$$M = M_\epsilon^d(\Gamma)$$

to be the compact polyhedral complex whose underlying point set is

$$M := \{(\Gamma', D') : \Gamma' \in A_\epsilon(\Gamma), D' \in \text{Div}_+^d(\Gamma')\}.$$

Lemma 3.17 (Semicontinuity Lemma). *The function $r : M \rightarrow \mathbb{Z}$ given by*

$$r(\Gamma', D') = r_{\Gamma'}(D')$$

is upper semicontinuous, that is, the set $\{(\Gamma', D') : r_{\Gamma'}(D') \geq i\}$ is closed for all i .

Proof. Following the general strategy of [Gathmann and Kerber 2008, Proof of Proposition 3.1], but with some slight variations in notation, we set

$$S := \{(\Gamma', D', f, P_1, \dots, P_d) : \Gamma' \in A_\epsilon(\Gamma), D' \in \text{Div}_+^d(\Gamma'), f \in \mathcal{M}(\Gamma'), (f) + D' = P_1 + \dots + P_d\}.$$

Also, for each $i = 0, \dots, d$, set

$$M_i := \{(\Gamma', D', P_1, \dots, P_i) : \Gamma' \in A_\epsilon(\Gamma), D' \in \text{Div}_+^d(\Gamma'), P_1, \dots, P_i \in \Gamma'\}.$$

As in [Gathmann and Kerber 2008, Lemma 1.9 and Proposition 3.1], one can endow each of the spaces S and M_i ($0 \leq i \leq d$) with the structure of a polyhedral complex.

The obvious “forgetful morphisms”

$$\pi_i : S \rightarrow M_i, (\Gamma', D', f, P_1, \dots, P_d) \mapsto (\Gamma', D', P_1, \dots, P_i)$$

and

$$p_i : M_i \rightarrow M, (\Gamma', D', P_1, \dots, P_i) \mapsto (\Gamma', D')$$

are morphisms of polyhedral complexes, and in particular they are continuous maps between topological spaces. Following [Gathmann and Kerber 2008, Proof of Proposition 3.1], we make some observations:

- (1) p_i is an *open map* for all i (since it is locally just a linear projection).
- (2) $M_i \setminus \pi_i(S)$ is a union of open polyhedra, and in particular, is an open subset of M_i .
- (3) For $(\Gamma', D') \in M$, we have $r_{\Gamma'}(D') \geq i$ if and only if $(\Gamma', D') \notin p_i(M_i \setminus \pi_i(S))$.

From (1) and (2), it follows that $p_i(M_i \setminus \pi_i(S))$ is open in M . So by (3), we see that the subset $\{(\Gamma', D') : r_{\Gamma'}(D') \geq i\}$ is closed in M , as desired. \square

The following corollary shows that the condition for a metric graph to have a $g_{\leq d}^r$ is closed:

Corollary 3.18. *Suppose Γ_n is a sequence of metric graphs in $A_\epsilon(\Gamma)$ converging to Γ . If there exists a $g_{\leq d}^r$ on Γ_n for all n , then there exists a $g_{\leq d}^r$ on Γ as well.*

Proof. Without loss of generality, we may assume that $r \geq 0$. Passing to a subsequence and replacing d by some $d' \leq d$ if necessary, we may assume that for each n , there exists an effective divisor $D_n \in \text{Div}_+(\Gamma_n)$ with $\deg(D_n) = d$ and $r(D_n) = r$. Since M is compact, $\{(\Gamma_n, D_n)\}$ has a convergent subsequence; by passing to this subsequence, we may assume that $(\Gamma_n, D_n) \rightarrow (\Gamma, D)$ for some divisor $D \in \text{Div}(\Gamma)$. By Lemma 3.17, we have $r(D) \geq r$. Subtracting points from D if necessary, we find that there is an effective divisor $D' \in \text{Div}(\Gamma)$ with $\deg(D') \leq d$ and $r(D') = r$, as desired. \square

Corollary 3.19. *Fix nonnegative integers g, r , and d . If there exists a $g_{\leq d}^r$ on every \mathbb{Q} -graph of genus g , then there exists a $g_{\leq d}^r$ on every metric graph of genus g .*

Proof. We can approximate a metric graph Γ by a sequence of \mathbb{Q} -graphs in $A_\epsilon(\Gamma)$ for some $\epsilon > 0$, so the result follows directly from Corollary 3.18. \square

Finally, we give our promised application of the Semicontinuity Lemma to Brill–Noether theory for tropical curves:

Theorem 3.20. *Fix integers $g, r, d \geq 0$ such that*

$$\rho(g, r, d) = g - (r + 1)(g - d + r) \geq 0.$$

Then every tropical curve of genus g has a divisor D with $r(D) = r$ and $\deg D \leq d$.

Proof. By Theorem 3.12, there exists a $g_{\leq d}^r$ on every metric \mathbb{Q} -graph, so it follows from Corollary 3.19 that the same is true for *all* metric graphs. By [Gathmann and Kerber 2008, Remark 3.6], if $\tilde{\Gamma}$ is a tropical curve and Γ is the metric graph obtained from $\tilde{\Gamma}$ by removing all unbounded edges, then for every $D \in \text{Div}(\Gamma)$ we have $r_{\tilde{\Gamma}}(D) = r_\Gamma(D)$. Therefore the existence of a $g_{\leq d}^r$ on Γ implies the existence of a $g_{\leq d}^r$ on $\tilde{\Gamma}$. \square

4. Weierstrass points on curves and graphs

As another illustration of the Specialization Lemma in action, in this section we will explore the relationship between Weierstrass points on curves and (a suitable notion of) Weierstrass points on graphs. As an application, we will generalize and place into a more conceptual framework a well-known result of Ogg concerning Weierstrass points on the modular curve $X_0(p)$. We will also prove the existence of Weierstrass points on tropical curves of genus $g \geq 2$.

4A. Weierstrass points on graphs. Let G be a graph of genus g . By analogy with the theory of algebraic curves, we say that $P \in V(G)$ is a *Weierstrass point* if $r(g(P)) \geq 1$. We define Weierstrass points on *metric graphs* and *tropical curves* in exactly the same way.

Remark 4.1. By Corollary 1.5, if Γ is the \mathbb{Q} -graph corresponding to G (so that every edge of G has length 1), then $P \in V(G)$ is a Weierstrass point on G if and only if P is a Weierstrass point on Γ .

Let $P \in V(G)$. An integer $k \geq 1$ is called a *Weierstrass gap* for P if

$$r(k(P)) = r((k - 1)(P)).$$

The Riemann–Roch theorem for graphs, together with the usual arguments from the theory of algebraic curves, yields the following result, whose proof we leave to the reader:

Lemma 4.2. *The following are equivalent:*

- (1) P is a Weierstrass point.
- (2) There exists a positive integer $k \leq g$ which is a Weierstrass gap for P .
- (3) $r(K_G - g(P)) \geq 0$.

Remark 4.3. Unlike the situation for algebraic curves, there exist graphs of genus at least 2 with no Weierstrass points. For example, consider the graph $G = B_n$ of genus $g = n - 1$ introduced in Example 3.6. We claim that B_n has no Weierstrass points if $n \geq 3$. Indeed, the canonical divisor K_G is $(g - 1)(Q_1) + (g - 1)(Q_2)$, and by symmetry it suffices to show that

$$r((g - 1)(Q_2) - (Q_1)) = -1.$$

This follows directly from Theorem 1.2, since

$$(g - 1)(Q_2) - (Q_1) \leq \nu := g(Q_2) - (Q_1)$$

and ν is the divisor associated to the linear ordering $Q_1 < Q_2$ of $V(G)$.

More generally, let $G = B(\ell_1, \dots, \ell_n)$ be the graph of genus $g = n - 1$ obtained by subdividing the i -th edge of B_n into ℓ_i edges. Let R_{ij} for $1 \leq j \leq \ell_i - 1$ denote the vertices strictly between Q_1 and Q_2 lying on the i -th edge (in sequential order). Then Q_1 and Q_2 are not Weierstrass points of G . Indeed, by symmetry it again suffices to show that $r((g - 1)(Q_2) - (Q_1)) = -1$, and this follows from Theorem 1.2 by considering the linear ordering

$$Q_1 < R_{11} < R_{12} < \dots < R_{1(\ell_1-1)} < R_{21} < \dots < R_{n(\ell_n-1)} < Q_2.$$

Other examples of families of graphs with no Weierstrass points are given in [Baker and Norine 2007a]. The graphs in these examples are all hyperelliptic.

More recently, S. Norine and P. Whalen have discovered examples of nonhyperelliptic graphs of genus 3 and 4 without Weierstrass points. It remains an interesting open problem to classify all graphs without Weierstrass points.

By the proof of Theorem 4.13 below, given any graph G of genus at least 2, there exists a positive integer k for which the regular subdivision $\sigma_k(G)$ (see Section 1D) has at least one Weierstrass point. In particular, there are always Weierstrass points on the metric graph associated to G . From the point of view of arithmetic geometry, this is related to the fact that Weierstrass points on an algebraic curve X/K of genus at least 2 always exist, but in general they are not K -rational. So just as one sometimes needs to pass to a finite extension L/K in order to see the Weierstrass points on a curve, one needs in general to pass to a regular subdivision of G in order to find Weierstrass points.

Example 4.4. On the complete graph $G = K_n$ on $n \geq 4$ vertices, every vertex is a Weierstrass point. Indeed, if $P, Q \in V(G)$ are arbitrary, then $g(P) - (Q)$ is equivalent to the effective divisor

$$(g - (n - 1))(P) - (Q) + \sum_{v \in V(G)} (v),$$

and thus $r(g(P)) \geq 1$.

The following example, due to Serguei Norine, shows that there exist metric \mathbb{Q} -graphs with infinitely many Weierstrass points:

Example 4.5. Let Γ be the metric \mathbb{Q} -graph associated to the banana graph B_n for some $n \geq 4$. Then Γ has infinitely many Weierstrass points.

Indeed, label the edges of Γ as e_1, \dots, e_n , and identify each e_i with the segment $[0, 1]$, where Q_1 corresponds to 0, say, and Q_2 corresponds to 1. We write $x(P)$ for the element of $[0, 1]$ corresponding to the point $P \in e_i$ under this parametrization. Then for each i and each $P \in e_i$ with $x(P) \in [\frac{1}{3}, \frac{2}{3}]$, we claim that $r(3(P)) \geq 1$, and hence P is a Weierstrass point on Γ .

To see this, we will show explicitly that for every $Q \in \Gamma$ we have

$$|3(P) - (Q)| \neq \emptyset.$$

For this, it suffices to construct a function $f \in \mathcal{M}(\Gamma)$ for which

$$\Delta(f) \geq -3(P) + (Q).$$

This is easy if $P = Q$. Otherwise we have:

Case 1(a): If $Q \in e_i$ and $x(P) < x(Q)$, let $y = \frac{1}{2}(3x(P) - x(Q))$ and take f to be constant on e_j for $j \neq i$, and on e_i to have slope -2 on $[y, x(P)]$, slope 1 on $[x(P), x(Q)]$, and slope 0 elsewhere.

Case 1(b): If $Q \in e_i$ and $x(Q) < x(P)$, we again let $y = \frac{1}{2}(3x(P) - x(Q))$ and take f to be constant on e_j for $j \neq i$, and on e_i to have slope -1 on $[x(Q), x(P)]$, slope 2 on $[x(P), y]$, and slope 0 elsewhere.

If $Q \in e_j$ for some $j \neq i$, take f to be constant on e_k for $k \neq i, j$. On e_j , let $z = \min(x(Q), 1 - x(Q))$, and take f to have slope 1 on $[0, z]$, slope 0 on $[z, 1 - z]$, and slope -1 on $[1 - z, 1]$. Finally, along e_i , we have two cases:

Case 2(a): If $x(P) \in [\frac{1}{3}, \frac{1}{2}]$, let $y = 3x(P) - 1$, and define f on e_i to have slope -1 on $[0, y]$, slope -2 on $[y, x(P)]$, and slope 1 on $[x(P), 1]$.

Case 2(b): If $x(P) \in [\frac{1}{2}, \frac{2}{3}]$, we again let $y = 3x(P) - 1$, and define f on e_i to have slope -1 on $[0, x(P)]$, slope 2 on $[x(P), y]$, and slope 1 on $[y, 1]$.

Remark 4.6. Similarly, one can show that for each integer $m \geq 2$, if $x(P) \in [\frac{1}{m}, \frac{m-1}{m}]$ then $r(m(P)) \geq 1$, and thus P is a Weierstrass point on Γ as long as the genus of Γ is at least m .

We close this section with a result which generalizes Remark 4.3, and which can be used in practice to identify non-Weierstrass points on certain graphs.

Lemma 4.7. *Let v be a vertex of a graph G of genus $g \geq 2$, and let G' be the graph obtained by deleting the vertex v and all edges incident to v . If G' is a tree, then v is not a Weierstrass point.*

Proof. Since G' is a tree, there is a linear ordering

$$v_1 < \dots < v_{n-1}$$

of the vertices of G' such that each vertex other than the first one has exactly one neighbor preceding it in the order. Extend $<$ to a linear ordering of $V(G)$ by letting v be the last element in the order. Since the corresponding divisor ν is equal to $g(v) - (v_1)$, it follows from Theorem 1.2 that $|g(v) - (v_1)| = \emptyset$, and therefore $r(g(v)) = 0$. Thus v is not a Weierstrass point. □

Remark 4.8. It is easy to see that if (G, ν) satisfies the hypothesis of Lemma 4.7, then so does (\tilde{G}, ν) , where \tilde{G} is obtained by subdividing each edge e_i of G into m_i edges for some positive integer m_i .

4B. Specialization of Weierstrass points on totally degenerate curves. We say that an arithmetic surface \mathfrak{X}/R is *totally degenerate* if the genus of its dual graph G is the same as the genus of X . Under our hypotheses on \mathfrak{X} , the genus of X is the sum of the genus of G and the genera of all irreducible components of \mathfrak{X}_k , so \mathfrak{X} is totally degenerate if and only if all irreducible components of \mathfrak{X}_k have genus 0 .

Applying the Specialization Lemma and the definition of a Weierstrass point, we immediately obtain:

Corollary 4.9. *If \mathfrak{X} is a strongly semistable, regular, and totally degenerate arithmetic surface, then for every K -rational Weierstrass point $P \in X(K)$, $\rho(P)$ is a Weierstrass point of the dual graph G of \mathfrak{X} . More generally, for every Weierstrass point $P \in X(\overline{K})$, $\tau_*(P)$ is a Weierstrass point of the reduction graph Γ of \mathfrak{X} .*

As a sample consequence, we have this concrete result:

Corollary 4.10. (1) *Let \mathfrak{X}/R be a strongly semistable, regular, totally degenerate arithmetic surface whose special fiber has a dual graph with no Weierstrass points (for example, the graph B_n for some $n \geq 3$). Then X does not possess any K -rational Weierstrass points.*

(2) *Let \mathfrak{X}/R be a (not necessarily regular) arithmetic surface whose special fiber consists of two genus 0 curves intersecting transversely at 3 or more points. Then every Weierstrass point of $X(K)$ specializes to a singular point of \mathfrak{X}_k .*

(3) *More generally, let \mathfrak{X}/R be a (not necessarily regular) strongly semistable and totally degenerate arithmetic surface whose dual graph G contains a vertex v for which $G' := G \setminus \{v\}$ is a tree. Then there are no K -rational Weierstrass points on X specializing to the component C of \mathfrak{X}_k corresponding to v .*

Proof. Part (1) follows from what we have already said. For (2), it suffices to note that X has a strongly semistable regular model \mathfrak{X}' whose dual graph G' is isomorphic to $B(\ell_1, \dots, \ell_g)$ for some positive integers ℓ_i , and a point of $X(K)$ which specializes to a nonsingular point of \mathfrak{X}_k will specialize to either Q_1 or Q_2 in G' , neither of which is a Weierstrass point by Remark 4.3. Finally, (3) follows easily from Lemma 4.7 and Remark 4.8. \square

We view Corollary 4.10 as a generalization of Ogg's argument [1978] showing that the cusp ∞ is never a Weierstrass point on $X_0(p)$ for $p \geq 23$ prime, since $X_0(p)/\mathbb{Q}_p^{\text{nr}}$ has a model $\mathfrak{X}_0(p)$ of the type described in Corollary 4.10 (2) (the Deligne–Rapoport model), and ∞ specializes to a nonsingular point on the special fiber of $\mathfrak{X}_0(p)$. More generally, Corollary 4.10 shows (as does Ogg's original argument) that all Weierstrass points of $X_0(p)$ specialize to supersingular points in the Deligne–Rapoport model. Corollaries 4.9 and 4.10 give a recipe for extending Ogg's result to a much broader class of curves with totally degenerate reduction.

Remark 4.11. Corollary 4.10 has strong implications concerning the arithmetic of Weierstrass points on curves. For example, in the special case where $K = \mathbb{Q}_p^{\text{nr}}$, the conclusion of Corollary 4.10 (1) implies that every Weierstrass point on X is ramified at p .

Example 4.12. The hypothesis that \mathfrak{X} is totally degenerate is necessary in the statement of Corollary 4.9. For example, it follows from [Atkin 1967, Theorem 1] that the cusp ∞ is a \mathbb{Q} -rational Weierstrass point on $X_0(180)$. The mod 5 reduction of the Deligne–Rapoport model of $X_0(180)$ consists of two copies of $X_0(36)_{\mathbb{F}_5}$

intersecting transversely at the supersingular points, and the cusp ∞ specializes to a nonsingular point on one of these components. (This does not contradict Corollary 4.10 because $X_0(36)$ does not have genus 0.)

We conclude this section with another application of algebraic geometry to tropical geometry: we use the classical fact that Weierstrass points exist on every smooth curve of genus at least 2 to show that there exist Weierstrass points on every tropical curve of genus at least 2.

Theorem 4.13. *Let $\tilde{\Gamma}$ be a tropical curve of genus $g \geq 2$. Then there exists at least one Weierstrass point on $\tilde{\Gamma}$.*

Proof. We first consider the case of a \mathbb{Q} -graph Γ . By rescaling if necessary, we may assume that Γ is the metric graph associated to a finite graph G , with every edge of G having length one. By Theorem B.2, there exists a strongly semistable, regular, totally degenerate arithmetic surface \mathfrak{X}/R whose generic fiber is smooth of genus g and whose special fiber has reduction graph Γ . Let $P \in X(\bar{K})$ be a Weierstrass point, which exists by classical algebraic geometry since $g \geq 2$. By Corollary 4.9, $\tau_*(P)$ is a Weierstrass point of Γ . We have thus shown that every metric \mathbb{Q} -graph of genus at least 2 has a Weierstrass point.

Now let Γ be an arbitrary metric graph. As in Section 3D, for $\epsilon > 0$ sufficiently small, we can approximate Γ by a sequence Γ_n of \mathbb{Q} -graphs within the space $A_\epsilon(\Gamma)$. Let $P_n \in \Gamma_n$ be a Weierstrass point. Passing to a subsequence if necessary, we may assume without loss of generality that

$$(\Gamma_n, P_n) \rightarrow (\Gamma, P)$$

in $M_\epsilon^1(\Gamma)$ for some point $P \in \Gamma$. Since $r_{\Gamma_n}(gP_n) \geq 1$ for all n and

$$(\Gamma_n, gP_n) \rightarrow (\Gamma, gP)$$

in $M_\epsilon^g(\Gamma)$, we conclude from the Semicontinuity Lemma that $r_\Gamma(gP) \geq 1$, that is, P is a Weierstrass point on Γ .

Finally, suppose that $\tilde{\Gamma}$ is a tropical curve, and let Γ be the metric graph obtained from $\tilde{\Gamma}$ by removing all unbounded edges. It follows from [Gathmann and Kerber 2008, Remark 3.6] that every Weierstrass point on Γ is also a Weierstrass point on $\tilde{\Gamma}$. Therefore the existence of Weierstrass points on Γ implies the existence of Weierstrass points on $\tilde{\Gamma}$. □

Remark 4.14. We do not know a direct combinatorial proof of Theorem 4.13, but it would certainly be interesting to give such a proof. Also, in light of Example 4.5, it is not clear if there is an analogue for metric graphs of the classical fact that the total weight of all Weierstrass points on a smooth curve of genus $g \geq 2$ is $g^3 - g$.

4C. Specialization of a canonical divisor. Since $P \in X(\bar{K})$ is a Weierstrass point of X if and only if $r(K_X - g(P)) \geq 0$, where K_X denotes a canonical divisor on X , and since $P \in V(G)$ is a Weierstrass point of G if and only if $r(K_G - g(P)) \geq 0$, Corollary 4.9 suggests a relationship between the canonical divisor of G and the specialization of K_X when \mathfrak{X} is totally degenerate. We investigate this relationship in this section.

Let \mathfrak{X} be a strongly semistable, regular, totally degenerate arithmetic surface. Let $\omega_{\mathfrak{X}/R}$ be the *canonical sheaf* for \mathfrak{X}/R , and let $K_{\mathfrak{X}}$ be a Cartier divisor such that

$$\mathcal{O}_{\mathfrak{X}}(K_{\mathfrak{X}}) \cong \omega_{\mathfrak{X}/R}.$$

We call any such $K_{\mathfrak{X}}$ a *canonical divisor*.

Lemma 4.15. $\rho(K_{\mathfrak{X}}) = K_G$.

Proof. This is a consequence of the *adjunction formula* for arithmetic surfaces (see [Liu 2002, Theorem 9.1.37]), which tells us that

$$(K_{\mathfrak{X}} \cdot C_i) = 2g(C_i) - 2 - (C_i \cdot C_i) = -2 - (C_i \cdot C_i) \tag{4.16}$$

for all i . Since $(C_i \cdot \sum_j C_j) = 0$ for all i , we have

$$(C_i \cdot C_i) = - \sum_{j \neq i} (C_i \cdot C_j) = - \deg(v_i). \tag{4.17}$$

Combining (4.16) and (4.17) gives

$$(K_{\mathfrak{X}} \cdot C_i) = \deg(v_i) - 2$$

for all i , as desired. □

Remark 4.18. More generally, if we do not assume that X is totally degenerate, then the above proof shows that

$$\rho(K_{\mathfrak{X}}) = K_G + 2 \sum_{i=1}^n g(C_i)(v_i),$$

where C_i is the irreducible component of \mathfrak{X}_k corresponding to the vertex v_i of G .

Remark 4.19. Lemma 4.15 helps explain why there is in fact a canonical *divisor* on a graph G , rather than just a canonical *divisor class*, and also explains the connection between the canonical divisor on a graph and the canonical divisor class in algebraic geometry. This connection is implicit in the earlier work of S. Zhang [1993].

Lemma 4.20. *Let $K_X \in \text{Div}(X)$ be a canonical divisor. Then $\rho(K_X)$ is linearly equivalent to K_G .*

Proof. Since the Zariski closure of K_X differs from a canonical divisor $K_{\mathfrak{X}}$ on \mathfrak{X} by a vertical divisor, this follows from Lemma 4.15 and the remarks preceding Lemma 2.1. \square

Remark 4.21. By a general *moving lemma* (see, for example, [Liu 2002, Corollary 9.1.10 or Proposition 9.1.11]), there exists a horizontal canonical divisor $K_{\mathfrak{X}}$ on \mathfrak{X} . Since $K_{\mathfrak{X}}$ is the Zariski closure of K_X in this case, it follows that there exists a canonical divisor $K_X \in \text{Div}(X)$ for which $\rho(K_X)$ is equal to K_G (and not just linearly equivalent to it).

4D. An example. We conclude with an explicit example which illustrates many of the concepts that have been discussed in this paper.

Let p be an odd prime, and let X be the smooth plane quartic curve over \mathbb{Q}_p given by $F(x, y, z) = 0$, where

$$F(x, y, z) = (x^2 - 2y^2 + z^2)(x^2 - z^2) + py^3z. \tag{4.22}$$

By classical algebraic geometry, we have $g(X) = 3$, and the $\overline{\mathbb{Q}_p}$ -gonality of X is also 3. We let $K = \mathbb{Q}_p^{\text{nr}}$, and consider X as an algebraic curve over K .

Let \mathfrak{X}' be the model for X over the valuation ring R of K given by (4.22). According to [Liu 2002, Exercise 10.3.10], the special fiber of \mathfrak{X}' is semistable and consists of two projective lines ℓ_1 and ℓ_2 with equations $x = z$ and $x = -z$, respectively, which intersect transversely at the point $(0 : 1 : 0)$, together with the conic C defined by $x^2 - 2y^2 + z^2 = 0$, which intersects each of ℓ_1 and ℓ_2 transversely at 2 points. The model \mathfrak{X}' is not regular, but a regular model \mathfrak{X} can be obtained from \mathfrak{X}' by blowing up the point $(0 : 1 : 0)$ of the special fiber of \mathfrak{X}' , which produces an exceptional divisor E in \mathfrak{X}_k isomorphic to \mathbb{P}_k^1 , and which intersects each of ℓ_1 and ℓ_2 transversely at a single point (see [Liu 2002, Corollary 10.3.25]). The special fiber \mathfrak{X}_k of \mathfrak{X} and the dual graph G of \mathfrak{X}_k are depicted in Figure 1.

In the diagram on the right, the vertex P corresponds to the conic C , Q_i corresponds to the line ℓ_i ($i = 1, 2$), and P' corresponds to the exceptional divisor E of




Figure 1. Left: The special fiber \mathfrak{X}_k . Right: The dual graph G of \mathfrak{X}_k .

the blowup. Note that G is a graph of genus 3, and that \mathfrak{X} has totally degenerate strongly semistable reduction.

Claim. Q_1 and Q_2 are Weierstrass points of G , while P and P' are not.

Indeed, since $3(Q_1) \sim 3(Q_2) \sim 2(P) + (P')$, it follows easily that $r(3(Q_1)) = r(3(Q_2)) \geq 1$, and therefore Q_1 and Q_2 are Weierstrass points. On the other hand, P is not a Weierstrass point of G by Lemma 4.7, and P' is not a Weierstrass point either, since $3(P') - (P)$ is equivalent to $(Q_1) + (Q_2) + (P') - (P)$, which is the divisor associated to the linear ordering $P < Q_1 < Q_2 < P'$ of $V(G)$.

Claim. The gonality of G is 3.

We have already seen that $r(3(Q_1)) \geq 1$, so the gonality of G is at most 3. It remains to show that G is not hyperelliptic, that is, there is no g_2^1 on G . By symmetry, and using the fact that $(Q_1) + (Q_2) \sim 2(P')$, it suffices to show that $r(D) = 0$ for $D = (P) + (P')$ and for each of the divisors $(Q_1) + (X)$ with $X \in V(G)$. For this, it suffices to show that $|D| = \emptyset$ for each of the divisors $(Q_1) + 2(Q_2) - (P)$, $(P) + 2(Q_1) - (P')$, and $(P) + (P') + (Q_1) - (Q_2)$. But these are the v -divisors associated to the linear orderings $P < Q_1 < P' < Q_2$, $P' < Q_2 < P < Q_1$, and $Q_1 < P < Q_2 < P'$ of $V(G)$, respectively. The claim therefore follows from Theorem 1.2.

The canonical divisor K_G on G is $2(P) + (Q_1) + (Q_2)$. We now compute the specializations of various canonical divisors in $\text{Div}(X(K))$. Since a canonical divisor on X is just a hyperplane section, the following divisors are all canonical:

$$\begin{aligned} K_1 : (x = z) \cap X &= (0 : 1 : 0) + 3(1 : 0 : 1), \\ K_2 : (x = -z) \cap X &= (0 : 1 : 0) + 3(1 : 0 : -1), \\ K_3 : (z = 0) \cap X &= 2(0 : 1 : 0) + (\sqrt{2} : 1 : 0) + (-\sqrt{2} : 1 : 0), \\ K_4 : (y = 0) \cap X &= (1 : 0 : 1) + (1 : 0 : -1) + (1 : 0 : \sqrt{-1}) + (1 : 0 : -\sqrt{-1}). \end{aligned}$$

The specializations of these divisors under ρ (or equivalently, under τ_*) are

$$\begin{aligned} \rho(K_1) &: (P') + 3(Q_1), \\ \rho(K_2) &: (P') + 3(Q_2), \\ \rho(K_3) &: 2(P') + 2(P), \\ \rho(K_4) &: (Q_1) + (Q_2) + 2(P). \end{aligned}$$

It is straightforward to check that each of these divisors is linearly equivalent to $K_G = (Q_1) + (Q_2) + 2(P)$, in agreement with Lemma 4.20.

Finally, note that (as follows from the above calculations) $(1 : 0 : 1)$ and $(1 : 0 : -1)$ are Weierstrass points on X , and they specialize to Q_1 and Q_2 , respectively. As we have seen, these are both Weierstrass points of G , as predicted by Corollary 4.9.

On the metric graph Γ associated to G , there are additional Weierstrass points. A somewhat lengthy case-by-case analysis shows that the Weierstrass points of Γ are the four points at distance $\frac{1}{3}$ from P , together with the two intervals $[Q_i, R_i]$ of length $\frac{1}{3}$, where R_1, R_2 are the points at distance $\frac{2}{3}$ from P' . It would be interesting to compute the specializations to Γ of the remaining Weierstrass points in $X(\bar{K})$.

Appendix A. A reformulation of Raynaud's description of the Néron model of a Jacobian

In this appendix, we reinterpret in the language of divisors on graphs some results of Raynaud concerning the relation between a proper regular model for a curve and the Néron model of its Jacobian. The main result here is that the diagrams (A.5) and (A.6) below are exact and commutative. This may not be a new observation, but since we could not find a reference, we will attempt to explain how it follows in a straightforward way from Raynaud's work.

In order to keep this appendix self-contained, we have repeated certain definitions which appear in the main body of the text. Some references for the results described here are [Bertolini and Darmon 1997, Appendix; Bosch et al. 1990; Edixhoven 1998; Raynaud 1970].

Raynaud's description. Let R be a complete discrete valuation ring with field of fractions K and algebraically closed residue field k . Let X be a smooth, proper, geometrically connected curve over K , and let \mathfrak{X}/R be a proper model for X with reduced special fiber \mathfrak{X}_k . For simplicity, we assume throughout that \mathfrak{X} is regular, that the irreducible components of \mathfrak{X}_k are all smooth, and that all singularities of \mathfrak{X}_k are ordinary double points. We let $\mathcal{C} = \{C_1, \dots, C_n\}$ be the set of irreducible components of \mathfrak{X}_k .

Let J be the Jacobian of X over K , let \mathcal{J} be the Néron model of J/R , and let \mathcal{J}^0 be the connected component of the identity in \mathcal{J} . We denote by

$$\Phi = \mathcal{J}_k / \mathcal{J}_k^0$$

the group of connected components of the special fiber \mathcal{J}_k of \mathcal{J} .

Let $\text{Div}(X)$ (respectively, $\text{Div}(\mathfrak{X})$) be the group of Cartier divisors on X (respectively, on \mathfrak{X}); since X is smooth and \mathfrak{X} is regular, Cartier divisors on X (respectively, \mathfrak{X}) are the same as Weil divisors.

The Zariski closure in \mathfrak{X} of an effective divisor on X is a Cartier divisor. Extending by linearity, we can associate to each $D \in \text{Div}(X)$ a Cartier divisor \mathcal{D} on \mathfrak{X} , which we refer to as the *Zariski closure* of D .

Let $\text{Div}^0(X)$ denote the subgroup of Cartier divisors of degree zero on X . In addition, let $\text{Div}^{(0)}(\mathfrak{X})$ denote the subgroup of $\text{Div}(\mathfrak{X})$ consisting of those Cartier divisors \mathcal{D} for which the restriction of the associated line bundle $\mathcal{O}_{\mathfrak{X}}(\mathcal{D})$ to each

irreducible component of \mathfrak{X}_k has degree zero, that is, for which

$$\deg(\mathbb{O}_{\mathfrak{X}}(\mathcal{D})|_{C_i}) = 0 \text{ for all } C_i \in \mathcal{C}.$$

Finally, let

$$\text{Div}^{(0)}(X) = \{D \in \text{Div}^0(X) : \mathcal{D} \in \text{Div}^{(0)}(\mathfrak{X})\},$$

where \mathcal{D} is the Zariski closure of D .

Let $\text{Prin}(X)$ (respectively, $\text{Prin}(\mathfrak{X})$) denote the group of principal Cartier divisors on X (respectively, \mathfrak{X}). There is a well-known isomorphism

$$J(K) = \mathcal{F}(R) \cong \text{Div}^0(X) / \text{Prin}(X),$$

and according to Raynaud, there is an isomorphism

$$J^0(K) := \mathcal{F}^0(R) \cong \text{Div}^{(0)}(X) / \text{Prin}^{(0)}(X), \tag{A.1}$$

where

$$\text{Prin}^{(0)}(X) := \text{Div}^{(0)}(X) \cap \text{Prin}(X).$$

The isomorphism in (A.1) comes from the fact that $\mathcal{F}^0 = \text{Pic}_{\mathfrak{X}/R}^0$ represents the functor “isomorphism classes of line bundles whose restriction to each element of \mathcal{C} has degree zero”. (Recall that there is a canonical isomorphism between isomorphism classes of line bundles on \mathfrak{X} and the Cartier class group of \mathfrak{X} .)

Remark A.2. In particular, it follows from the above discussion that every element $P \in J^0(K)$ can be represented as the class of D for some $D \in \text{Div}^{(0)}(X)$.

There is a natural inclusion $\mathcal{C} \subset \text{Div}(\mathfrak{X})$, and an intersection pairing

$$\mathcal{C} \times \text{Div}(\mathfrak{X}) \rightarrow \mathbb{Z}, \quad (C_i, \mathcal{D}) \mapsto (C_i \cdot \mathcal{D}),$$

where $(C_i \cdot \mathcal{D}) = \deg(\mathbb{O}_{\mathfrak{X}}(\mathcal{D})|_{C_i})$.

The intersection pairing gives rise to a map

$$\begin{aligned} \alpha : \mathbb{Z}^{\mathcal{C}} &\rightarrow \mathbb{Z}^{\mathcal{C}}, \\ f &\mapsto \left(C_i \mapsto \sum_{C_j \in \mathcal{C}} (C_i \cdot C_j) f(C_j) \right). \end{aligned}$$

Since k is algebraically closed and the canonical map $J(K) \rightarrow \mathcal{F}_k(k)$ is surjective by [Liu 2002, Proposition 10.1.40(b)], there is a canonical isomorphism $J(K)/J^0(K) \cong \Phi$. According to Raynaud, the component group Φ is canonically isomorphic to the homology of the complex

$$\mathbb{Z}^{\mathcal{C}} \xrightarrow{\alpha} \mathbb{Z}^{\mathcal{C}} \xrightarrow{\text{deg}} \mathbb{Z}, \tag{A.3}$$

where

$$\text{deg} : f \mapsto \sum_{C_i} f(C_i).$$

The isomorphism

$$\phi : J(K)/J^0(K) \cong \text{Ker}(\text{deg})/\text{Im}(\alpha)$$

can be described in the following way. Let $P \in J(K)$, and choose $D \in \text{Div}^0(X)$ such that $P = [D]$. Let $\mathcal{D} \in \text{Div}(\mathfrak{X})$ be the Zariski closure of D . Then

$$\phi(P) = [C_i \mapsto (C_i \cdot \mathcal{D})].$$

When D corresponds to a Weil divisor supported on $X(K)$, we have another description of the map ϕ . Write

$$D = \sum_{P \in X(K)} n_P(P)$$

with $\sum n_P = 0$. Since \mathfrak{X} is regular, each point $P \in X(K) = \mathfrak{X}(R)$ specializes to a well-defined element $c(P)$ of \mathcal{C} . Identifying a formal sum $\sum_{C_i \in \mathcal{C}} a_i C_i$ with the function $C_i \mapsto a_i \in \mathbb{Z}^{\mathcal{C}}$, we have

$$\phi([D]) = \left[\sum_P n_P c(P) \right]. \tag{A.4}$$

The quantities appearing in (A.3) can be interpreted in a more suggestive fashion using the language of graphs. Let G be the *dual graph* of \mathfrak{X}_k , that is, G is the finite graph whose vertices v_i correspond to the irreducible components C_i of \mathfrak{X}_k , and whose edges correspond to intersections between these components (so that there is one edge between v_i and v_j for each point of intersection between C_i and C_j). We let $\text{Div}(G)$ denote the free abelian group on the set of vertices of G , and define $\text{Div}^0(G)$ to be the kernel of the natural map $\text{deg} : \text{Div}(G) \rightarrow \mathbb{Z}$ given by

$$\text{deg}\left(\sum a_i(v_i)\right) = \sum a_i.$$

In particular, the set $V(G)$ of vertices of G is in bijection with \mathcal{C} , and the group $\text{Div}(G)$ is isomorphic to $\mathbb{Z}^{\mathcal{C}}$, with $\text{Div}^0(G)$ corresponding to $\text{Ker}(\text{deg})$.

Let $\mathcal{M}(G) = \mathbb{Z}^{V(G)}$ be the set of \mathbb{Z} -linear functions on $V(G)$, and define the Laplacian operator $\Delta : \mathcal{M}(G) \rightarrow \text{Div}^0(G)$ by

$$\Delta(\varphi) = \sum_{v \in V(G)} \sum_{e=vw} (\varphi(v) - \varphi(w)) (v),$$

where the inner sum is over all edges e of G having v as an endpoint. Finally, define

$$\text{Prin}(G) = \Delta(\mathcal{M}(G)) \subseteq \text{Div}^0(G),$$

and let

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

be the *Jacobian* of G . It is a consequence of Kirchhoff’s Matrix-Tree Theorem that $\text{Jac}(G)$ is a finite abelian group whose order is equal to the number of spanning trees of G .

Since the graph G is connected, one knows that $\text{Ker}(\Delta)$ consists precisely of the constant functions, and it follows from (A.3) that there is a canonical exact sequence

$$0 \longrightarrow \text{Prin}(G) \xrightarrow{\gamma_1} \text{Div}^0(G) \xrightarrow{\gamma_2} \Phi \longrightarrow 0.$$

In other words, the component group Φ is canonically isomorphic to the Jacobian group of the graph G .

We can summarize much of the preceding discussion by saying that the following diagram is commutative and exact:

$$\begin{array}{ccccccc}
 & & 0 & & 0 & & 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \text{Prin}^{(0)}(X) & \xrightarrow{\alpha_1} & \text{Div}^{(0)}(X) & \xrightarrow{\alpha_2} & J^0(K) \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \text{Prin}(X) & \xrightarrow{\beta_1} & \text{Div}^0(X) & \xrightarrow{\beta_2} & J(K) \longrightarrow 0 \quad (\text{A.5}) \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \text{Prin}(G) & \xrightarrow{\gamma_1} & \text{Div}^0(G) & \xrightarrow{\gamma_2} & \text{Jac}(G) \cong \Phi \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

A few remarks are in order about the exactness of the rows and columns in (A.5). It is well known that the natural map from $X(K) = \mathfrak{X}(R)$ to the smooth locus of $\mathfrak{X}_k(k)$ is surjective (see for example [Liu 2002, Proposition 10.1.40(b)]); by (A.4), this implies that the natural maps

$$\text{Div}(X) \rightarrow \text{Div}(G) \quad \text{and} \quad \text{Div}^0(X) \rightarrow \text{Div}^0(G)$$

are surjective. The surjectivity of the horizontal map $\alpha_2 : \text{Div}^{(0)}(X) \rightarrow J^0(K)$ follows from Remark A.2. Using this, we see from the Snake Lemma that since the vertical map $\text{Div}^0(X) \rightarrow \text{Div}^0(G)$ is surjective, the vertical map $\text{Prin}(X) \rightarrow \text{Prin}(G)$ is also surjective. All of the other claims about the commutativity and exactness of (A.5) follow in a straightforward way from the definitions.

Passage to the limit. If K'/K is a finite extension of degree m with ramification index $e \mid m$ and valuation ring R' , then by a sequence of blow-ups we can obtain a regular model \mathfrak{X}'/R' for X whose corresponding dual graph G' is the graph $\sigma_e(G)$ obtained by subdividing each edge of G into e edges. If we think of G as an unweighted graph and of $\sigma_e(G)$ as a weighted graph in which every edge has length $\frac{1}{e}$, then G and $\sigma_e(G)$ are different models for the same metric \mathbb{Q} -graph Γ , which one calls the *reduction graph* of \mathfrak{X}/R . The discussion in [Chinburg and Rumely 1993] shows that the various maps

$$c_{K'} : X(K') \rightarrow G'$$

are compatible, in the sense that they give rise to a specialization map $\tau : X(\bar{K}) \rightarrow \Gamma$ which takes $X(\bar{K})$ surjectively onto $\Gamma_{\mathbb{Q}}$.

It is straightforward to check that the diagram (A.5) behaves functorially with respect to finite extensions, and therefore that there is a commutative and exact diagram

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 0 & \longrightarrow & \text{Prin}^{(0)}(X(\bar{K})) & \xrightarrow{\alpha_1} & \text{Div}^{(0)}(X(\bar{K})) & \xrightarrow{\alpha_2} & J^0(\bar{K}) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 0 & \longrightarrow & \text{Prin}(X(\bar{K})) & \xrightarrow{\beta_1} & \text{Div}^0(X(\bar{K})) & \xrightarrow{\beta_2} & J(\bar{K}) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 0 & \longrightarrow & \text{Prin}_{\mathbb{Q}}(\Gamma) & \xrightarrow{\gamma_1} & \text{Div}_{\mathbb{Q}}^0(\Gamma) & \xrightarrow{\gamma_2} & \text{Jac}_{\mathbb{Q}}(\Gamma) & \longrightarrow & 0. \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & & \text{(A.6)}
 \end{array}$$

Remark A.7. Let \tilde{K} be the completion of an algebraic closure \bar{K} of K , so that \tilde{K} is a complete and algebraically closed field equipped with a valuation

$$v : \tilde{K} \rightarrow \mathbb{Q} \cup \{+\infty\},$$

and \bar{K} is dense in \tilde{K} . By continuity, one can extend τ to a map $\tau : X(\tilde{K}) \rightarrow \Gamma$ and replace \bar{K} by \tilde{K} everywhere in the diagram (A.6).

A few explanations are in order concerning the definitions of the various groups and group homomorphisms which appear in (A.6). Since \bar{K} is algebraically closed, we may identify the group $\text{Div}(X_{\bar{K}})$ of Cartier (or Weil) divisors on $X_{\bar{K}}$ with $\text{Div}(X(\bar{K}))$, the free abelian group on the set $X(\bar{K})$. We define $\text{Prin}(X(\bar{K}))$ to be the subgroup of $\text{Div}(X(\bar{K}))$ consisting of principal divisors. The group $\text{Div}(X(\bar{K}))$

(respectively, $\text{Prin}(X(\bar{K}))$) can be identified with the direct limit of $\text{Div}(X_{K'})$ (respectively, $\text{Prin}(X_{K'})$) over all finite extensions K'/K . Accordingly, we define the group $J^0(\bar{K})$ to be the direct limit of the groups $J^0(K')$ over all finite extensions K'/K , and we define $\text{Div}^{(0)}(X(\bar{K}))$ and $\text{Prin}^{(0)}(X(\bar{K}))$ similarly. Finally, we define $\text{Jac}_{\mathbb{Q}}(\Gamma)$ to be the quotient $\text{Div}_{\mathbb{Q}}^0(\Gamma)/\text{Prin}_{\mathbb{Q}}(\Gamma)$.

That $\text{Prin}_{\mathbb{Q}}(\Gamma)$, as defined in Section 1, coincides with the direct limit over all finite extensions K'/K of the groups $\text{Prin}(G')$ follows easily from Remark 1.3.

With these definitions in place, it is straightforward to check using (A.5) that the diagram (A.6) is both commutative and exact.

We note the following consequence of the exactness of (A.5) and (A.6):

Corollary A.8. *The canonical maps*

$$\text{Prin}(X) \rightarrow \text{Prin}(G) \quad \text{and} \quad \text{Prin}(X(\bar{K})) \rightarrow \text{Prin}_{\mathbb{Q}}(\Gamma)$$

are surjective.

Remark A.9. It follows from Corollary A.8 that if G is a graph, the group $\text{Prin}(G)$ can be characterized as the image of $\text{Prin}(X)$ under the specialization map from $\text{Div}(X)$ to $\text{Div}(G)$ for any strongly semistable regular arithmetic surface \mathfrak{X}/R whose special fiber has dual graph isomorphic to G . (Such an \mathfrak{X} always exists by Corollary B.3 below.)

Remark A.10. Another consequence of (A.6) is that there is a canonical isomorphism

$$J(\bar{K})/J^0(\bar{K}) \cong \text{Jac}_{\mathbb{Q}}(\Gamma),$$

so that the group $\text{Jac}_{\mathbb{Q}}(\Gamma)$ plays the role of the component group of the Néron model in this situation, even though there is not a well-defined Néron model for J over \bar{K} or \tilde{K} , since the valuations on these fields are not discrete. One can show using elementary methods that $\text{Jac}_{\mathbb{Q}}(\Gamma)$ is (noncanonically) isomorphic to $(\mathbb{Q}/\mathbb{Z})^g$ (compare with the discussion in [Grothendieck 1972, Exposé IX, §11.8]).

Appendix B. A result from the deformation theory of stable marked curves by Brian Conrad

Recall from Section 1A that by an *arithmetic surface*, we mean an integral scheme that is proper and flat over a discrete valuation ring such that its generic fiber is a smooth and geometrically connected curve. By Stein factorization, the special fiber of an arithmetic surface is automatically geometrically connected.

In this appendix, we describe how one can realize an arbitrary graph G as the dual graph of the special fiber of some regular arithmetic surface whose special fiber C is a *totally degenerate semistable curve* (or *Mumford curve*), meaning that

C is semistable and connected, every irreducible component of C is isomorphic to the projective line over the residue field k , and all singularities of C are k -rational.

Lemma B.1. *Let G be a connected graph, and let k be an infinite field. Then there exists a totally degenerate semistable curve C/k whose dual graph is isomorphic to G .*

The proof is left to the reader.

The crux of the matter is the following theorem, whose proof is a standard application of the deformation theory of stable marked curves.

Theorem B.2. *Let C be a proper and geometrically connected semistable curve over a field k , and let R be a complete discrete valuation ring with residue field k . Then there exists an arithmetic surface \mathfrak{X} over R with special fiber C such that \mathfrak{X} is a regular scheme.*

Combining these two results, we obtain:

Corollary B.3. *Let R be a complete discrete valuation ring with field of fractions K and infinite residue field k . For any connected graph G , there exists a regular arithmetic surface \mathfrak{X}/R whose generic fiber is a smooth, proper, and geometrically connected curve X/K , and whose special fiber is a totally degenerate semistable curve with dual graph isomorphic to G .*

Remark B.4. By [Liu 2002, Lemma 10.3.18], the genus of X coincides with the genus $g = |E(G)| - |V(G)| + 1$ of the graph G .

Proof of Theorem B.2. Let $g = \dim H^1(C, \mathbb{O}_C)$ denote the arithmetic genus of C . The structure theorem for ordinary double points [Freitag and Kiehl 1988, III, §2] ensures that C^{sing} splits over a separable extension of k , so we can choose a finite Galois extension k'/k so that the locus $C_{k'}^{\text{sing}}$ of nonsmooth points in $C_{k'}$ consists entirely of k' -rational points and every irreducible component of $C_{k'}$ is geometrically irreducible and has a k' -rational point. (If C is a Mumford curve then we can take $k' = k$.) In particular, each smooth component in $C_{k'}$ of arithmetic genus 0 is isomorphic to $\mathbf{P}_{k'}^1$ and so admits at least three k' -rational points. We can then construct a $\text{Gal}(k'/k)$ -stable étale divisor $D' \subseteq C_{k'}^{\text{sm}}$ whose support consists entirely of k' -rational points in the smooth locus such that for each component X' of $C_{k'}$ isomorphic to $\mathbf{P}_{k'}^1$ we have

$$\#(X' \cap C_{k'}^{\text{sing}}) + \#(X' \cap D') \geq 3.$$

In particular, if we choose an enumeration of $D'(k')$ then the pair $(C_{k'}, D')$ is a stable n -pointed genus- g curve, where $n = \#D'(k')$ and $2g - 2 + n > 0$. Let $D \subseteq C^{\text{sm}}$ be the étale divisor that descends D' . We let R' be the local finite étale R -algebra with residue field k'/k .

The stack $\mathcal{M}_{g,n}$ classifying stable n -pointed genus- g curves for any $g, n \geq 0$ such that $2g - 2 + n > 0$ is a proper smooth Deligne–Mumford stack over $\mathrm{Spec} \mathbb{Z}$.⁴ The existence of $\mathcal{M}_{g,n}$ as a smooth Deligne–Mumford stack ensures that $(C_{k'}, D')$ admits a universal formal deformation $(\widehat{\mathcal{C}'}, \widehat{\mathcal{D}'})$ over a complete local noetherian R' -algebra A' with residue field k' , and that A' is a formal power series ring over R' . Moreover, there is a relatively ample line bundle canonically associated to any stable n -pointed genus- g curve $(X, \{\sigma_1, \dots, \sigma_n\}) \rightarrow S$ over a scheme S , namely the twist

$$\omega_{X/S}(\sum \sigma_i)$$

of the relative dualizing sheaf of the curve by the étale divisor defined by the marked points. Hence, the universal formal deformation uniquely algebraizes to a pair $(\mathcal{C}', \mathcal{D}')$ over $\mathrm{Spec} A'$.

Since $(C_{k'}, D') = k' \otimes_k (C, D)$, universality provides an action of the Galois group $\Gamma = \mathrm{Gal}(k'/k)$ on A' and on $(\mathcal{C}', \mathcal{D}')$ covering the natural Γ -action on $(C_{k'}, D')$ and on R' . The action by Γ on \mathcal{C}' is compatible with one on the canonically associated ample line bundle $\omega_{\mathcal{C}'/A'}(\mathcal{D}')$. Since $R \rightarrow R'$ is finite étale with Galois group Γ , the Γ -action on everything in sight (including the relatively ample line bundle) defines effective descent data: $A = (A')^\Gamma$ is a complete local noetherian R -algebra with residue field k such that $R' \otimes_R A \rightarrow A'$ is an isomorphism, and $(\mathcal{C}', \mathcal{D}')$ canonically descends to a deformation $(\mathcal{C}, \mathcal{D})$ of (C, D) over A . This can likewise be shown to be a universal deformation of (C, D) in the category of complete local noetherian R -algebras with residue field k , but we do not need this fact.

What matters for our purposes is structural information about A and the proper flat A -curve \mathcal{C} . By the functorial characterization of formal smoothness, A is formally smooth over R since the same holds for A' over R' , so A is a power series ring over R in finitely many variables. We claim that the Zariski-open locus of smooth curves in $\mathcal{M}_{g,n}$ is dense, which is to say that every geometric point has a smooth deformation. Indeed, since Knudsen's contraction and stabilization operations do nothing to smooth curves, this claim immediately reduces to the special cases $n = 0$ with $g \geq 2$, $g = n = 1$, and $g = 0, n = 3$. The final two cases are obvious and the first case was proved by Deligne and Mumford [1969, 1.9] via deformation theory. It follows that the generic fiber of \mathcal{C}' over A' is a smooth curve, so the same holds for \mathcal{C} over A . In other words, the Zariski-closed locus in \mathcal{C} where $\Omega_{\mathcal{C}/A}^1$ is not invertible has its closed image in $\mathrm{Spec}(A)$ given by a closed subset $\mathrm{Spec}(A/I)$ for

⁴ This is a standard fact: it can be found in [Deligne and Mumford 1969, 5.2] if $n = 0$ (so $g \geq 2$), in [Deligne and Rapoport 1973, IV, 2.2] if $g = n = 1$, and is trivial if $g = 0, n = 3$ (in which case $(\mathbf{P}^1, \{0, 1, \infty\})$ is the only such object, so the stack is $\mathrm{Spec} \mathbb{Z}$). The general case follows from these cases by realizing $\mathcal{M}_{g,n}$ as the universal curve over $\mathcal{M}_{g,n-1}$ (due to Knudsen's contraction and stabilization operations [1983, 2.7]).

a unique nonzero radical ideal $I \subseteq A$. Since A is a power series ring over R and I is a nonzero ideal, we can certainly find a local R -algebra map $\phi : A \rightarrow R$ in which I has nonzero image. The pullback \mathfrak{X}_ϕ of \mathcal{C} along ϕ is a proper flat semistable curve over R deforming C such that the generic fiber is smooth (as otherwise the map $\text{Spec}(R) \rightarrow \text{Spec}(A)$ would factor through $\text{Spec}(A/I)$, a contradiction).

The only remaining problem is to show that if ϕ is chosen more carefully then \mathfrak{X}_ϕ is also regular. If C is k -smooth then \mathcal{C} is A -smooth, so \mathfrak{X}_ϕ is R -smooth (and in particular regular). Thus, we may assume that C is not k -smooth. The main point is to use an understanding of the structure of \mathcal{C} near each singular point $c \in C - C^{\text{sm}}$. We noted at the outset that each finite extension $k(c)/k$ is separable, and if R_c is the corresponding local finite étale extension of R then the structure theory of ordinary double points [Freitag and Kiehl 1988, III, §2] provides an R_c -algebra isomorphism of henselizations

$$\mathcal{O}_{\mathcal{C},c}^h \simeq A_c[u, v]_{(\mathfrak{m}_c, u, v)}^h / (uv - a_c)$$

for some nonzero nonunit $a_c \in A_c$, where (A_c, \mathfrak{m}_c) is the local finite étale extension of A with residue field $k(c)/k$; we have $a_c \neq 0$ since \mathcal{C} has smooth generic fiber. For

$$B = \mathbb{Z}[t, u, v] / (uv - t),$$

the B -module $\bigwedge^2(\Omega_{B/\mathbb{Z}[t]}^1)$ is $B/(u, v) = B/(t)$. Thus, by a formal computation at each c we see that the annihilator ideal of $\bigwedge^2(\Omega_{\mathcal{C}/A}^1)$ on \mathcal{C} cuts out an A -finite closed subscheme of the nonsmooth locus in \mathcal{C} that is a pullback of a unique closed subscheme $\text{Spec}(A/J) \subseteq \text{Spec}(A)$. We made the initial choice of k'/k large enough so that it splits each $k(c)/k$. Hence, the method of proof of [Deligne and Mumford 1969, 1.5] and the discussion following that result show that for each singularity c' in $C_{k'}$, the corresponding element $a_{c'} \in A'$ may be chosen so that the $a_{c'}$'s are part of a system of variables for A' as a formal power series ring over R' . The ideal JA' is the intersection of the ideals $(a_{c'})$. To summarize, A is a formal power series ring over R and we can choose the variables for A' over R' such that JA' is generated by a product of such variables, one for each singularity on $C_{k'}$. In particular, the local interpretation of each $a_{c'}$ on \mathcal{C}' shows that for a local R' -algebra map $\phi' : A' \rightarrow R'$, the pullback $\mathfrak{X}'_{\phi'}$ of \mathcal{C}' along ϕ' is regular if and only if $\phi'(a_{c'}) \in R'$ is a uniformizer for each c' . This condition on ϕ' is equivalent to saying that $\phi'(JA') \subseteq R'$ is a proper nonzero ideal with multiplicity equal to the number ν of geometric singularities on C .

Since $JA' = J \otimes_A A'$ is a principal nonzero proper ideal, and hence it is invertible as an A' -module, it follows that J is principal as well, say $J = (\alpha)$ for some nonzero nonunit $\alpha \in A$. We seek an R -algebra map $\phi : A \rightarrow R$ such that $\phi(I) \neq 0$ and $\phi(\alpha) \in R$ is nonzero with order ν , for then $\phi^*(\mathcal{C})$ will have smooth generic fiber (by

our earlier discussion) and will become regular over R' , and so it will be regular since R' is finite étale over R . The information we have about $\alpha \in A$ is that in the formal power series ring A' over R' we can choose the variables so that α is a product of ν of the variables. We are now reduced to the following problem in commutative algebra. Let $R \rightarrow R'$ be a local finite étale extension of discrete valuation rings, $A = R[[x_1, \dots, x_N]]$, $I \subseteq A$ a nonzero ideal, and $\alpha \in A$ an element such that in $A' = R' \otimes_R A$ we can write $\alpha = x'_1 \cdots x'_\nu$ for some $1 \leq \nu \leq N$ and choice of R' -algebra isomorphism $A' \simeq R'[[x'_1, \dots, x'_N]]$. Then we claim that there is an R -algebra map $\phi : A \rightarrow R$ such that $\text{ord}_R(\phi(\alpha)) = \nu$ and $\phi(I) \neq 0$.

If we can choose $k' = k$, such as in the case when C is a Mumford curve, then such a ϕ obviously exists. Hence, for the intended application to Corollary B.3, we are done. To prove the claim in general, consider the expansions

$$x'_j = a'_{0j} + \sum_{i=1}^N a'_{ij}x_i + \cdots$$

where $a'_{0j} \in \mathfrak{m}_{R'}$ and $(a'_{ij})_{1 \leq i, j \leq N}$ is invertible over R' . Let $\pi \in \mathfrak{m}_R$ be a uniformizer. We seek ϕ of the form $\phi(x_i) = t_i\pi$ for $t_i \in R$. The requirement on the t_i 's is that

$$(a'_{0j}/\pi) + \sum a'_{ij}t_i \in R'^{\times}$$

for $1 \leq j \leq \nu$ and that

$$h(t_1\pi, \dots, t_N\pi) \neq 0$$

for some fixed nonzero power series $h \in I$. The unit conditions only depend on $t_i \pmod{\mathfrak{m}_R}$. Thus, once we find t_i that satisfy these unit conditions, the remaining nonvanishing condition on the nonzero power series h is trivial to satisfy by modifying the higher-order parts of the t_i 's appropriately. It remains to consider the unit conditions, which is a consequence of the lemma below. □

Lemma B.5. *Let k'/k be a finite extension of fields. For $1 \leq \nu \leq N$ let $\{H'_1, \dots, H'_\nu\}$ be a collection of independent hyperplanes in k'^N . For any $v'_1, \dots, v'_\nu \in k'^N$, the union of the affine-linear hyperplanes $v'_i + H'_i$ in k'^N cannot contain k'^N .*

Proof. The case of infinite k is trivial, but to handle finite k we have to do more work. Throughout the argument the ground field k may be arbitrary. Assuming we are in a case with $k^N \subseteq \bigcup (v'_i + H'_i)$, we seek a contradiction. Observe that $\nu \leq N$ since the H'_i are linearly independent hyperplanes in k'^N . Each overlap $k^N \cap (v'_i + H'_i)$ is either empty or a translate of $V_i = k^N \cap H'_i$ by some $v_i \in V = k^N$. For $1 \leq d \leq N$, any d -fold intersection

$$V_{i_1} \cap \cdots \cap V_{i_d} = k^N \cap (H'_{i_1} \cap \cdots \cap H'_{i_d})$$

has k -dimension at most $N - d$. Indeed, otherwise it contains $N - d + 1$ linearly independent vectors in k^N , and these may also be viewed as k' -linearly independent vectors in the overlap $H'_{i_1} \cap \dots \cap H'_{i_d}$ of d linearly independent hyperplanes in k'^N . This contradicts the linear independence of such hyperplanes. It now remains to prove the following claim that does not involve k' and concerns subspace arrangements over k : if V is a vector space of dimension $N \geq 1$ over a field k and if V_1, \dots, V_m are linear subspaces with $1 \leq m \leq N$ such that V is a union of the translates $v_i + V_i$ for $v_1, \dots, v_m \in V$ then there is $1 \leq d \leq m$ such that some d -fold intersection $V_{i_1} \cap \dots \cap V_{i_d}$ has dimension at least $N - d + 1$. This claim was suggested by T. Tao, and the following inductive proof of it was provided by S. Norine.

The claim is trivial for $N = 1$, and in general we induct on N so we may assume $N > 1$. The case $m = 1$ is trivial, so with $N > 1$ fixed we can assume $m > 1$. The case $V_i = V$ for all i is also trivial, so we may assume V_m is contained in a hyperplane H . Let $v = 0$ if $v_m \notin H$ and $v \in V - H$ if $v_m \in H$, so $(v_m + V_m) \cap (v + H)$ is empty. Since the $v_i + V_i$ cover V , it follows that $v + H$ is covered by the $v_i + V_i$ for $1 \leq i \leq m - 1$, so H is covered by the $(v_i - v) + V_i$ for $1 \leq i \leq m - 1 \leq N - 1$. Each overlap $H \cap ((v_i - v) + V_i)$ for such i is either empty or a translate $w_i + W_i$ of $W_i = V_i \cap H$. Setting $w_i = 0$ if $H \cap ((v_i - v) + V_i)$ is empty, we can relabel the V_i 's such that H is covered by $w_i + W_i$ for $1 \leq i \leq m - 1$. By induction there is $1 \leq d \leq N - 1$ so that after relabeling we have $\dim(W_1 \cap \dots \cap W_d) \geq \dim(H) - d + 1 = N - d$. Hence, for $W = V_1 \cap \dots \cap V_d$ we have that $\dim(H \cap W) = N - d$. In particular, $\dim(W) \geq N - d$. If $\dim(W) \geq N - d + 1$ then $\{V_1, \dots, V_d\}$ works as required in the claim we are aiming to prove, so we can assume $\dim(W) = N - d > 0$. Let $W' \subseteq V$ be a complementary subspace to W , and consider $W'_i = W' \cap V_i$ for $1 \leq i \leq d$. Obviously $\dim W' = d$ and

$$V_i = W \oplus W'_i$$

for such i . If some collection of translates $w'_i + W'_i$ for $1 \leq i \leq d$ and $w'_i \in W'$ covers W' then by induction there is $1 \leq d' \leq d$ such that (after relabeling) $\dim(W'_1 \cap \dots \cap W'_{d'}) \geq d - d' + 1$, so

$$\begin{aligned} \dim(V_1 \cap \dots \cap V_{d'}) &= \dim(W) + \dim(W'_1 \cap \dots \cap W'_{d'}) \\ &= (N - d) + d - d' + 1 = N - d' + 1, \end{aligned}$$

as required.

Hence, we can now assume that no collection of translates $w'_i + W'_i$ in W' for $1 \leq i \leq d$ can cover W' . In particular, under the decomposition $V = W \oplus W'$, we may write $v_i = (w_i, w'_i)$ for $1 \leq i \leq N$ and these $w'_i + W'_i$ for $1 \leq i \leq d$ do not cover W' . Since each of V_1, \dots, V_d contains W , it follows that $v_1 + V_1, \dots, v_d + V_d$ do not cover W' . Thus, we can choose $w' \in W'$ not in any of these d translates, so

$w' + W$ is disjoint from all of them (since the projections into W' are disjoint as well). By the initial covering hypothesis we therefore have that $w' + W$ is covered by $v_j + V_j$ for $d + 1 \leq j \leq m$, so W is covered by translates of $V_j \cap W$ for such j . The number of such j 's is $m - d \leq N - d$, so since $0 < \dim W = N - d < N$ we can conclude by induction that there is some $1 \leq d'' \leq m - d$ so that (after relabeling) the intersection

$$(V_{d+1} \cap W) \cap \cdots \cap (V_{d+d''} \cap W) = V_1 \cap \cdots \cap V_d \cap V_{d+1} \cap \cdots \cap V_{d+d''}$$

has dimension at least $(N - d) - d'' + 1 = N - (d + d'') + 1$. This completes the induction and so proves the claim. \square

Acknowledgments

I am very grateful to Marc Coppens for pointing out an error in an earlier version of this paper; his comments led to numerous improvements in the present version. I would also like to thank Serguei Norine for many enlightening conversations, my student Adam Tart for performing some very useful computer calculations, and Michael Kerber for stimulating discussions concerning tropical curves. I am also grateful to Kevin Buzzard, Filip Cools, Bas Edixhoven, Xander Faber, and Ezra Miller for their comments on earlier versions of this paper. Finally, I would like to express my gratitude to Brian Conrad for contributing Appendix B.

References

- [Atkin 1967] A. O. L. Atkin, “Weierstrass points at cusps of $\Gamma_o(n)$ ”, *Ann. of Math.* (2) **85** (1967), 42–45. MR 36 #1646 Zbl 0145.31904
- [Baker and Faber 2006] M. Baker and X. Faber, “Metriized graphs, Laplacian operators, and electrical networks”, pp. 15–33 in *Quantum graphs and their applications* (Snowbird, UT, 2005), edited by G. Berkolaiko et al., Contemp. Math. **415**, Amer. Math. Soc., Providence, RI, 2006. MR 2007j:05136 Zbl 1114.94025
- [Baker and Norine 2007a] M. Baker and S. Norine, “Harmonic morphisms and hyperelliptic graphs”, preprint, 2007. arXiv 0707.1309
- [Baker and Norine 2007b] M. Baker and S. Norine, “Riemann–Roch and Abel–Jacobi theory on a finite graph”, *Adv. Math.* **215**:2 (2007), 766–788. MR 2355607 Zbl 1124.05049
- [Berkovich 1990] V. G. Berkovich, *Spectral theory and analytic geometry over non-Archimedean fields*, Mathematical Surveys and Monographs **33**, American Mathematical Society, Providence, RI, 1990. MR 91k:32038 Zbl 0715.14013
- [Bertolini and Darmon 1997] M. Bertolini and H. Darmon, “A rigid analytic Gross–Zagier formula and arithmetic applications”, *Ann. of Math.* (2) **146**:1 (1997), 111–147. MR 99f:11079 Zbl 1029.11027
- [Bosch et al. 1990] S. Bosch, W. Lütkebohmert, and M. Raynaud, *Néron models*, Ergebnisse der Mathematik (3) **21**, Springer, Berlin, 1990. MR 91i:14034 Zbl 0705.14001

- [Chinburg and Rumely 1993] T. Chinburg and R. Rumely, “The capacity pairing”, *J. Reine Angew. Math.* **434** (1993), 1–44. MR 94b:14019 Zbl 0756.14013
- [Deligne and Mumford 1969] P. Deligne and D. Mumford, “The irreducibility of the space of curves of given genus”, *Inst. Hautes Études Sci. Publ. Math.* 36 (1969), 75–109. MR 41 #6850 Zbl 0181.48803
- [Deligne and Rapoport 1973] P. Deligne and M. Rapoport, “Les schémas de modules de courbes elliptiques”, pp. 143–316 in *Modular functions of one variable, II* (Antwerp, 1972), Lecture Notes in Math. **349**, Springer, Berlin, 1973. MR 49 #2762 Zbl 0281.14010
- [Edixhoven 1998] B. Edixhoven, “On Néron models, divisors and modular curves”, *J. Ramanujan Math. Soc.* **13**:2 (1998), 157–194. MR 2000a:11094 Zbl 0931.11021
- [Freitag and Kiehl 1988] E. Freitag and R. Kiehl, *Étale cohomology and the Weil conjecture*, Ergebnisse der Mathematik (3) **13**, Springer, Berlin, 1988. MR 89f:14017 Zbl 0643.14012
- [Gathmann and Kerber 2008] A. Gathmann and M. Kerber, “A Riemann–Roch theorem in tropical geometry”, *Math. Z.* **259**:1 (2008), 217–230. MR 2377750 Zbl 05261613
- [Grothendieck 1972] A. Grothendieck (editor), *Groupes de monodromie en géométrie algébrique, I* (SGA 7 I), Lecture Notes in Mathematics **288**, Springer, Berlin, 1972. MR 50 #7134 Zbl 0237.00013
- [Hartshorne 1977] R. Hartshorne, *Algebraic geometry*, Graduate Texts in Mathematics **52**, Springer, New York, 1977. MR 57 #3116 Zbl 0367.14001
- [Hladky et al. 2007] J. Hladky, D. Král, and S. Norine, “Rank of divisors on tropical curves”, preprint, 2007. arXiv 0709.4485
- [Knudsen 1983] F. F. Knudsen, “The projectivity of the moduli space of stable curves, II: The stacks $M_{g,n}$ ”, *Math. Scand.* **52**:2 (1983), 161–199. MR 85d:14038a Zbl 0544.14020
- [Liu 2002] Q. Liu, *Algebraic geometry and arithmetic curves*, Oxford Graduate Texts in Mathematics **6**, Oxford University Press, Oxford, 2002. MR 2003g:14001 Zbl 0996.14005
- [Mikhalkin 2006] G. Mikhalkin, “Tropical geometry and its applications”, pp. 827–852 in *International Congress of Mathematicians* (Madrid, 2006), vol. II, edited by M. Sanz-Solé et al., Eur. Math. Soc., Zürich, 2006. MR 2008c:14077
- [Mikhalkin and Zharkov 2007] G. Mikhalkin and I. Zharkov, “Tropical curves, their Jacobians and Theta functions”, preprint, 2007. arXiv math.AG/0612267v2
- [Ogg 1974] A. P. Ogg, “Hyperelliptic modular curves”, *Bull. Soc. Math. France* **102** (1974), 449–462. MR 51 #514 Zbl 0314.10018
- [Ogg 1978] A. P. Ogg, “On the Weierstrass points of $X_0(N)$ ”, *Illinois J. Math.* **22**:1 (1978), 31–35. MR 57 #3136 Zbl 0374.14005
- [Raynaud 1970] M. Raynaud, “Spécialisation du foncteur de Picard”, *Inst. Hautes Études Sci. Publ. Math.* 38 (1970), 27–76. MR 44 #227 Zbl 0207.51602
- [Zhang 1993] S. Zhang, “Admissible pairing on a curve”, *Invent. Math.* **112**:1 (1993), 171–193. MR 94h:14023 Zbl 0795.14015

Communicated by Bjorn Poonen

Received 2008-02-21

Revised 2008-05-13

Accepted 2008-06-12

mbaker@math.gatech.edu

*School of Mathematics, Georgia Institute of Technology,
Atlanta, GA 30332-0160, United States*

