

LINEAR SERIES ON METRIZED COMPLEXES OF ALGEBRAIC CURVES

OMID AMINI AND MATTHEW BAKER

ABSTRACT. A *metrized complex of algebraic curves* over an algebraically closed field κ is, roughly speaking, a finite metric graph Γ together with a collection of marked complete nonsingular algebraic curves C_v over κ , one for each vertex v of Γ ; the marked points on C_v are in bijection with the edges of Γ incident to v . We define linear equivalence of divisors and establish a Riemann-Roch theorem for metrized complexes of curves which combines the classical Riemann-Roch theorem over κ with its graph-theoretic and tropical analogues from [AC, BN, GK, MZ], providing a common generalization of all of these results. For a complete nonsingular curve X defined over a non-Archimedean field \mathbb{K} , together with a strongly semistable model \mathfrak{X} for X over the valuation ring R of \mathbb{K} , we define a corresponding metrized complex $\mathfrak{C}\mathfrak{X}$ of curves over the residue field κ of \mathbb{K} and a canonical specialization map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$ which preserves degrees and linear equivalence. We then establish generalizations of the specialization lemma from [B] and its weighted graph analogue from [AC], showing that the rank of a divisor cannot go down under specialization from X to $\mathfrak{C}\mathfrak{X}$. As an application, we establish a concrete link between specialization of divisors from curves to metrized complexes and the Eisenbud-Harris theory [EH] of limit linear series. Using this link, we formulate a generalization of the notion of limit linear series to curves which are not necessarily of compact type and prove, among other things, that any degeneration of a \mathfrak{g}_d^r in a regular family of semistable curves is a limit \mathfrak{g}_d^r on the special fiber.

CONTENTS

1. Introduction	2
2. Metrized complexes of algebraic curves	11
3. A Riemann-Roch theorem for metrized complexes of curves	13
4. The specialization map and specialization inequality	21
5. Limit linear series	30
Appendix A. Rank-determining sets for metrized complexes	41
References	45

Acknowledgments: The authors would like to thank Vladimir Berkovich, Lucia Caporaso, Ethan Cotterill, Eric Katz, Johannes Nicaise, Joe Rabinoff, Frank-Olaf Schreyer, David Zureick-Brown, and the referees for helpful discussions and remarks. The second author was supported in part by NSF grant DMS-0901487.

1. INTRODUCTION

We begin by quoting the introduction from the groundbreaking paper [EH] of Eisenbud and Harris:

“One of the most potent methods in the theory of (complex projective algebraic) curves and their linear systems since the work of Castelnuovo has been degeneration to singular curves... Most problems of interest about curves are, or can be, formulated in terms of (families of) linear series. Thus, in order to use degenerations to reducible curves for studying smooth curves, it is necessary to understand what happens to linear series in the course of such a degeneration, and in particular, to understand what structure on a reducible curve plays the part of a linear series.”

This has been very successful in practice and has led to important advances in the study of algebraic curves (see [HM98] for an overview); we mention for example Brill-Noether theory [EH, Bi], the geometry of the moduli space of curves [HM82, EH2], and Weierstrass points and their monodromy [EH3, EM, EH4]. In view of the Deligne-Mumford compactification of \mathcal{M}_g and the stable reduction theorem [DM], it is particularly natural to consider degeneration of smooth curves to stable (or more generally semistable) curves. Eisenbud and Harris were able to settle a number of longstanding open problems with their theory of limit linear series on such curves, and the theory has undergone numerous developments and improvements in the last 25 years. However, the theory of limit linear series only applies, for the most part, to a rather restricted class of reducible curves, namely those of *compact type* (i.e., nodal curves whose dual graph is a tree).

In [B], the second author introduced a new framework for degenerating linear series on curves, degenerating linear series on a regular semistable family of curves to a linear series on the *dual graph* of the special fiber. This theory, which is closely related to tropical geometry and also to the theory of Berkovich analytic spaces, is more or less orthogonal to the Eisenbud-Harris theory, in that it works best for special fibers which are *maximally degenerate*, meaning that the dual graph has first Betti number equal to the genus of the generic fiber. Specializing linear series to the dual graph provides no information whatsoever when the special fiber is of compact type. Intriguingly, both the Eisenbud-Harris theory and the second author’s theory from [B] lead to simple proofs of the celebrated Brill-Noether theorem of Griffiths-Harris (see [EH, CDPR]).

The aim of the present paper is to introduce a theoretical framework suitable for generalizing both the Eisenbud-Harris theory of limit linear series and the second author’s theory of specialization from curves to graphs. The main new object of study is what we call a *metrized complex of algebraic curves* over an algebraically closed field κ ; this is, roughly speaking, a finite metric graph Γ together with a collection of marked complete nonsingular algebraic curves C_v over κ , one for each vertex v (with a marked point on C_v for each edge of Γ incident to v). We define divisors, linear equivalence of divisors, and rank of divisors on metrized complexes of curves and establish a Riemann-Roch theorem for metrized complexes of curves which combines the classical Riemann-Roch theorem over κ with its graph-theoretic and tropical analogues proved in [AC, BN, GK, MZ]. For a curve X defined over a non-Archimedean field \mathbb{K} , together with a strongly semistable model \mathfrak{X} for X over the valuation ring R of \mathbb{K} , we define a corresponding metrized complex $\mathfrak{C}\mathfrak{X}$ of curves over the residue field κ of \mathbb{K} and a specialization map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$ preserving degrees and linear equivalence. We then establish generalizations of the

specialization lemma from [B] and its weighted graph analogue from [AC], showing that the rank of a divisor cannot go down under specialization from X to $\mathfrak{C}\mathfrak{X}$. As an application, we establish a concrete link between specialization of divisors from curves to metrized complexes and the Eisenbud-Harris theory of limit linear series. Using this link, we formulate a generalization of the notion of limit linear series to curves which are not necessarily of compact type and prove, among other things, that any degeneration of a \mathfrak{g}_d^r in a regular family of semistable curves is a limit \mathfrak{g}_d^r on the special fiber.

As mentioned above, the specialization theory from [B] works best in the case where the family is maximally degenerate; in other cases, the dual graph of the family forgets too much information. This loss of information was partially remedied in [AC] by looking at vertex-weighted graphs, where the weight attached to a vertex is the genus of the corresponding irreducible component of the special fiber. The theory of metrized complexes of curves offers a richer solution in which one keeps track of the (normalization of the) irreducible components C_v of the special fiber themselves, and not just their genera. From the point of view of non-Archimedean geometry, this corresponds to retracting divisors on X to the skeleton $\Gamma_{\mathfrak{X}}$, a metric graph canonically embedded in the Berkovich analytic space X^{an} , but also keeping track, for points retracting to a vertex of $\Gamma_{\mathfrak{X}}$, of the tangent direction in which the vertex is approached. This is quite natural from the valuation-theoretic point of view: the metric graph $\Gamma_{\mathfrak{X}}$ associated to a strongly semistable model \mathfrak{X} is a canonical subset of the Berkovich analytic space X^{an} , and the curve C_v over κ is naturally identified (for each $v \in V$) with the fiber over v of the canonical retraction map $X^{\text{ad}} \rightarrow X^{\text{an}}$ (which sends a continuous valuation of arbitrary rank in the Huber adic space X^{ad} to its canonical rank-1 generalization), c.f. [Te, Remark 2.6]. Though we could possibly (with quite a bit of additional effort) have formulated many of our theorems and proofs without mentioning Berkovich spaces, it would have resulted in a significant loss of elegance and clarity, and in any case Berkovich's theory seems ideally suited for the point of view taken here. We do not, however, require any non-trivial facts from Huber's theory of adic spaces, so we will not mention them again even though they are certainly lurking in the background.

There are a number of connections between the ideas in the present paper and tropical geometry; these will be explored in more detail in future work. For example, the theory of morphisms between metrized complexes of curves studied in [ABBR1, ABBR2] sheds interesting new light on the question of which morphisms between tropical curves are liftable. The theory developed in [BPR] shows that there is a close connection between metrized complexes of curves and "exploded tropicalizations" in the sense of [Pay, Definition 2.9]; we plan to say more about this in the future as well. As noted by Payne in [Pay], exploded tropicalizations can be thought of as an algebraic analogue of the "exploded torus fibrations" studied by Parker from a symplectic viewpoint [Par]. It could be interesting to explore connections between *loc. cit.* and the present work.

The ideas in the present paper also have Diophantine applications to the study of rational points on curves over number fields (specifically, to the method of Coleman-Chabauty). Indeed, Eric Katz and David Zureick-Brown have recently proved a result similar to Theorem 3.4 below (Clifford's theorem for metrized complexes), and they use this result to answer a question of M. Stoll. A special case of the main result in [KZB] is the following. Let X be a smooth projective geometrically irreducible curve of genus $g \geq 2$ over \mathbf{Q} and assume that the Mordell-Weil rank r of the Jacobian of X is less than g . Fix a prime number $p > 2r + 2$ and let \mathfrak{X} be a proper (not necessarily semistable) regular model for X

over \mathbf{Z}_p . Then (letting $\bar{\mathfrak{X}}^{\text{sm}}$ denote the smooth locus of $\bar{\mathfrak{X}}$)

$$\#X(\mathbf{Q}) \leq \bar{\mathfrak{X}}^{\text{sm}}(\mathbf{F}_p) + 2r.$$

We discuss the theorem of Katz and Zureick-Brown in more detail in Section 5.5 below, explaining how limit linear series on metrized complexes of curves can be used to illuminate their proof and put it into a broader context.

Finally, we would like to mention that the ideas presented in this paper have been used to prove the two following recent results, see [AB] and [A14].

The first one [AB] is the multiplicity conjecture of [CDPR]. In [AB], we apply the theory developed in this paper to a generic chain of genus-one curves, and provide an analysis of the Brill-Noether theory of the corresponding metrized complexes. The description of ranks of divisors on such metrized complexes turn out to be quite similar to the example considered [CDPR] of a generic chain of genus-one metric graphs. In particular, we are able to provide an interpretation of the lingering paths of [CDPR] in terms of the η -function we introduce in Section 4.6. In addition, we show that every divisor of degree d and rank r on such a metrized complex can be completed to a limit \mathfrak{g}_d^r . In the case $\rho = 0$, these results, combined with a geometric argument, can be used to prove the conjecture of [CDPR], namely that each divisor D of degree d and rank r on the metric graph considered in [CDPR] can be lifted to a \mathfrak{g}_d^r on the corresponding curve.

The second application is the non-Archimedean analogue of Mumford-Neeman theorem on distribution of Weierstrass points. A theorem of Mumford and Neeman [N] states that for a compact Riemann surface S , and for a line bundle L of positive degree on S , the discrete measures μ_n on S supported on Weierstrass points of $L^{\otimes n}$ converge weakly to the Arakelov (canonical) measure on S when n goes to infinity. The first author proves the following non-Archimedean analogue of Mumford-Neeman theorem in [A14].

Theorem 1.1 (Amini [A14]). *Let X be a smooth proper curve of genus at least one over an algebraically closed and complete non-Archimedean field of equicharacteristic zero. Let X^{an} be the Berkovich analytification of X , and Γ the minimal skeleton of X^{an} (which is a metric graph). Let L be a line bundle of positive degree on X , and H_n the multiset of Weierstrass points of $L^{\otimes n}$. Consider the discrete measure $\mu_n := \frac{1}{|H_n|} \sum_{x \in H_n} \delta_{\tau(x)}$ on Γ , supported on the set of all points $\tau(x)$, reduction of a point $x \in H_n$ on Γ . The measures μ_n converge weakly to the canonical admissible measure μ_{ad} on Γ when n goes to infinity.*

The admissible measure μ_{ad} was constructed by Zhang. For any divisor D on the metric graph Γ of degree different from -2 , by extending the results of Chinburg and Rumely [CR] in the case $D = 0$, Zhang associates in [Z] the unique measure μ_D such that for any point $x \in \Gamma$, one has $g_{\mu_D}(D, x) + g_{\mu_D}(x, x) = c$ for a unique constant $c = c(D)$. Here $g_{\mu}(\cdot, \cdot)$ is the Green function on $\Gamma \times \Gamma$ associated to the measure μ . The canonical admissible measure is the measure μ_K , where K is the canonical divisor of the weighted metric graph Γ [Z].

1.1. Notation. We set the following notation, which will be used throughout this paper.

- G a finite edge-weighted graph with vertex set V and edge set E
- Γ a compact metric graph (geometric realization of an edge-weighted graph G)
- \mathbb{K} a complete and algebraically closed non-trivially valued non-Archimedean field
- R the valuation ring of \mathbb{K}
- κ the residue field of \mathbb{K}

X	a smooth, proper, connected algebraic curve defined over \mathbb{K}
X^{an}	the Berkovich analytification of X/\mathbb{K}
\mathfrak{X}	a proper and flat semistable R -model for X , usually assumed to be strongly semistable
C_v	a smooth, proper, connected algebraic curve over κ (one for each vertex $v \in V$)
\mathfrak{C}	a metrized complex of algebraic curves over κ
$ \mathfrak{C} $	the geometric realization of \mathfrak{C}
$\mathfrak{C}\mathfrak{X}$	the metrized complex associated to a strongly semistable R -model \mathfrak{X} of X
X_0	a nodal curve defined over κ
X_v	an irreducible component of X_0
$\mathfrak{C}X_0$	the regularization of the nodal curve X_0
$\text{Div}(\mathfrak{C})$	the group of divisors on the metrized complex \mathfrak{C}
\mathcal{D}	a divisor on \mathfrak{C}
\mathcal{E}	an effective divisor on \mathfrak{C}
f	a rational function on \mathfrak{C}
D_v	a divisor on C_v
E_v	an effective divisor on C_v
f_Γ	a rational function on Γ , usually the Γ -part of a rational function f on \mathfrak{C}
f_v	a rational function on C_v , usually the C_v -part of a rational function f on \mathfrak{C}
$\text{slp}_e(f_\Gamma)$	the outgoing slope of f_Γ along an edge (or tangent direction) e
$r_{\mathfrak{C}}$	the rank function on $\text{Div}(\mathfrak{C})$
$r_{\mathfrak{C},\mathcal{F}}$	the restricted rank function with respect to a collection $\mathcal{F} = \{F_v\}_{v \in V}$ of subspaces of $\kappa(C_v)$
τ	the retraction map from X^{an} to the skeleton $\Gamma_{\mathfrak{X}}$ of \mathfrak{X}
$\tau_*^{\mathfrak{C}\mathfrak{X}}$	the specialization map from $\text{Div}(X)$ to $\text{Div}(\mathfrak{C}\mathfrak{X})$
\mathcal{A}_v	the marked points of C_v in the metrized complex \mathfrak{C}
x_v^e	the point of \mathcal{A}_v corresponding to the edge (or tangent direction) e
A_v	the sum in $\text{Div}(C_v)$ of the points of \mathcal{A}_v

1.2. Overview. We now discuss the contents of this paper in more detail.

An *edge-weighted graph* G is a connected multigraph, possibly with loop edges, having vertex set V and edge set E , and endowed with a weight (or length) function $\ell : E \rightarrow \mathbb{R}_{>0}$. A *metric graph* Γ is the geometric realization of an edge-weighted graph G in which each edge e of G is identified with a line segment of length $\ell(e)$. We call G a *model* for Γ . Subdividing an edge of G in a length-preserving fashion changes the model but not the underlying metric graph.

Let κ be an algebraically closed field.¹ A *metrized complex* \mathfrak{C} of κ -curves consists of the following data:

- A connected finite graph G with vertex set V and edge set E .
- A metric graph Γ having G as a model (or equivalently, a length function $\ell : E \rightarrow \mathbb{R}_{>0}$).
- For each vertex v of G , a complete, nonsingular, irreducible curve C_v over κ .
- For each vertex v of G , a bijection $e \mapsto x_v^e$ between the edges of G incident to v (with loop edges counted twice) and a subset $\mathcal{A}_v = \{x_v^e\}_{e \ni v}$ of $C_v(\kappa)$.

¹To simplify the presentation, we restrict to algebraically closed fields. In §2.3 of the January 2013 arXiv version of this paper, we indicate how the theory can be developed over an arbitrary field κ .




FIGURE 1. The geometric realization of a metrized complex of genus four.

The *geometric realization* $|\mathfrak{C}|$ of \mathfrak{C} is defined to be the union of the edges of G and the collection of curves C_v , with each endpoint v of an edge e identified with the corresponding marked point x_v^e . See Figure 1. When we think of $|\mathfrak{C}|$ as a set, we identify it with the disjoint union of $\Gamma \setminus V$ and $\bigcup_{v \in V} C_v(\kappa)$. Thus, when we write $x \in |\mathfrak{C}|$, we mean that x is either a non-vertex point of Γ (a *graphical point* of \mathfrak{C}) or a point of $C_v(\kappa)$ for some $v \in V$ (a *geometric point* of \mathfrak{C}).

The *genus* of a metrized complex of curves \mathfrak{C} , denoted $g(\mathfrak{C})$, is by definition $g(\mathfrak{C}) = g(\Gamma) + \sum_{v \in V} g_v$, where g_v is the genus of C_v and $g(\Gamma)$ is the first Betti number of Γ .

Given a metrized complex \mathfrak{C} of κ -curves, there is an associated semistable curve X_0 over κ obtained by gluing the curves C_v along the points x_v^e (one intersection for each edge e of G) and forgetting the metric structure on Γ . Conversely, given a semistable curve X_0 over κ together with a positive real number for each node (which we call a “length function”), one obtains an associated metrized complex of κ -curves by letting G be the dual graph of X_0 , Γ the metric graph associated to G and the given length function, C_v the normalization of the irreducible component X_v of X_0 corresponding to v , and \mathcal{A}_v the preimage in C_v of the set of nodes of X_0 belonging to X_v .

Let \mathbb{K} be a complete and algebraically closed non-Archimedean field with residue field κ and let X/\mathbb{K} be a smooth, proper, connected algebraic curve. There is a metrized complex $\mathfrak{C} = \mathfrak{C}\mathfrak{X}$ canonically associated to any strongly semistable model \mathfrak{X} of X over the valuation ring R of \mathbb{K} : the special fiber $\bar{\mathfrak{X}}$ of \mathfrak{X} is a semistable curve over κ , and one defines the length of an edge e of the dual graph of $\bar{\mathfrak{X}}$ to be the *modulus* of the open annulus $\text{red}^{-1}(z)$, where z is the singular point of $\bar{\mathfrak{X}}$ corresponding to e and $\text{red} : X(\mathbb{K}) \rightarrow \bar{\mathfrak{X}}(\kappa)$ is the canonical reduction map. Equivalently, there is a local analytic equation for z of the form $xy = \varpi$ with ϖ in the maximal ideal of R , and the length of e is $\text{val}(\varpi)$. The corresponding metric graph $\Gamma = \Gamma_{\mathfrak{X}}$ is called the *skeleton* of \mathfrak{X} , and there is a canonical retraction map $\tau = \tau_{\mathfrak{X}} : X^{\text{an}} \rightarrow \Gamma_{\mathfrak{X}}$, where X^{an} denotes the Berkovich analytification of X . By restricting to $X(\mathbb{K})$ and then extending by linearity, we obtain a specialization map $\tau_* : \text{Div}(X) \rightarrow \text{Div}(\Gamma_{\mathfrak{X}})$.

If R is a discrete valuation ring and the fibered surface \mathfrak{X}/R is regular, then one can use intersection-theoretic methods to define and study specialization of divisors from curves to graphs. For example, in [B], using intersection theory, a specialization homomorphism $\rho : \text{Div}(X) \rightarrow \text{Div}(G)$ was defined, where G is the dual graph of $\bar{\mathfrak{X}}$, with the property that $\rho(D)$ and $\tau_*(D)$ are linearly equivalent as divisors on $\Gamma_{\mathfrak{X}}$ for all $D \in \text{Div}(X)$. In the present paper, we focus exclusively on τ_* and its generalization to metrized complexes of curves, though in principle all of our results can be translated (when R is a discrete valuation ring and \mathfrak{X} is regular) into results about this alternative specialization map.

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

To a divisor on \mathfrak{C} , we can naturally associate a divisor D_Γ on Γ , called the Γ -*part* of \mathcal{D} , as well as, for each $v \in V$, a divisor D_v on C_v (called the C_v -*part* of \mathcal{D}). The divisor D_v is simply the restriction of \mathcal{D} to C_v , i.e. $D_v = \sum_{x \in C_v(\kappa)} \mathcal{D}(x)(x)$, and D_Γ is defined as

$$D_\Gamma = \sum_{x \in \Gamma \setminus V} \mathcal{D}(x)(x) + \sum_{v \in V} \deg(D_v)(v),$$

where $\mathcal{D}(x)$ denotes the coefficient of x in \mathcal{D} .

In particular, the degree of \mathcal{D} equals the degree of D_Γ . One could equivalently define a divisor on \mathfrak{C} to be an element of the form $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ of $\text{Div}(\Gamma) \oplus (\oplus_v \text{Div}(C_v))$ such that $\deg(D_v) = D_\Gamma(v)$ for all v in V .

The specialization map τ_* can be enhanced in a canonical way to a map from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$. The fiber $\tau^{-1}(v)$ of the retraction map $\tau : X(\mathbb{K}) \rightarrow \Gamma$ over a vertex $v \in V$ can be canonically identified with $\text{red}^{-1}(C_v^{\text{sm}}(\kappa))$, the fiber of the reduction map $\text{red} : X(\mathbb{K}) \rightarrow \bar{\mathfrak{X}}(\kappa)$ over the smooth locus $C_v(\kappa) \setminus \mathcal{A}_v$ of C_v . We define $\tau^{\mathfrak{C}\mathfrak{X}} : X(\mathbb{K}) \rightarrow \mathfrak{C}\mathfrak{X}$ by setting

$$\tau^{\mathfrak{C}\mathfrak{X}}(P) = \begin{cases} \tau(P) & \tau(P) \notin V \\ \text{red}(P) & \tau(P) \in V \end{cases}$$

and we extend this by linearity to a map $\tau_*^{\mathfrak{C}\mathfrak{X}} : \text{Div}(X) \rightarrow \text{Div}(\mathfrak{C}\mathfrak{X})$.

Intuitively speaking, for points P of $X(\mathbb{K})$ which reduce to smooth points of $\bar{\mathfrak{X}}$, the map $\tau^{\mathfrak{C}\mathfrak{X}}$ keeps track of the reduced point $\text{red}(P) \in \bar{\mathfrak{X}}^{\text{sm}}(\kappa)$, while if P reduces to a singular point of $\bar{\mathfrak{X}}$ then $\tau^{\mathfrak{C}\mathfrak{X}}$ instead keeps track of the retraction of P to the skeleton of the open annulus $\text{red}^{-1}(P)$, which is canonically identified with the relative interior of the corresponding edge of $\Gamma_{\bar{\mathfrak{X}}}$.

A *nonzero rational function* \mathfrak{f} on a metrized complex of curves \mathfrak{C} is the data of a rational function f_Γ on Γ and nonzero rational functions f_v on C_v for each $v \in V$. (We do *not* impose any compatibility conditions on the rational functions f_Γ and f_v .) We call f_Γ the Γ -*part* of \mathfrak{f} and f_v the C_v -*part* of \mathfrak{f} .

The *divisor* of a nonzero rational function \mathfrak{f} on \mathfrak{C} is defined to be

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

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

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

If D_Γ (resp. D_v) denotes the Γ -part (resp. the C_v -part) of $\text{div}(\mathfrak{f})$, then

$$D_\Gamma = \text{div}(f_\Gamma) = \sum_{u \in \Gamma} \text{ord}_u(f_\Gamma)(u) \quad \text{and} \quad D_v = \text{div}(f_v) + \text{div}_v(f_\Gamma), \quad \text{where}$$

$$(1) \quad \operatorname{div}_v(f_\Gamma) := \sum_{e \ni v} \operatorname{slp}_e(f_\Gamma)(x_v^e).$$

Divisors of the form $\operatorname{div}(f)$ are called *principal*, and the principal divisors form a subgroup of $\operatorname{Div}^0(\mathfrak{C})$, the group of divisors of degree zero on \mathfrak{C} . Two divisors in $\operatorname{Div}(\mathfrak{C})$ are called *linearly equivalent* if they differ by a principal divisor.

Linear equivalence of divisors on \mathfrak{C} can be understood quite intuitively in terms of “chip-firing moves”. To state this formally, for each vertex v of G , let A_v be the sum of the $\deg_G(v)$ points in \mathcal{A}_v and let ℓ_v be the minimum length of an edge e incident to v . Also, for an edge e having v as an endpoint and a positive real number $\varepsilon < \ell(e)$, let $p_{v,\varepsilon}^e$ be the unique point of e at distance ε from v . Then it is not hard to show that linear equivalence of divisors on \mathfrak{C} is the equivalence relation generated by the following three kinds of “moves”:

- (1) (Firing on C_v) Choose a vertex v and replace the divisor D_v with a linearly equivalent divisor D'_v on C_v .
- (2) (Firing a vertex) Choose a vertex v and a positive real number $\varepsilon < \ell_v$ and replace D_v with $D_v - A_v$ and D_Γ with $D_\Gamma - \deg_G(v)(v) + \sum_{e \ni v} (p_\varepsilon^e)$.
- (3) (Firing a non-vertex) Choose a non-vertex $p \in \Gamma$ and a positive real number ε less than the distance from p to the nearest vertex and replace D_Γ by the divisor $D_\Gamma - 2(p) + (p_\varepsilon^+) + (p_\varepsilon^-)$, where p_ε^\pm are the two points at distance ε from p .

The motivation for our definitions of $\tau_*^{\mathfrak{C}\mathfrak{X}}$ and $\operatorname{div}(f)$ come from the following fundamental relation, a consequence of the non-Archimedean Poincaré-Lelong formula due to Thuillier (see §5 of [BPR]). Let f be a nonzero rational function on X and let \mathfrak{f} be the corresponding nonzero rational function on $\mathfrak{C}\mathfrak{X}$, where f_Γ is the restriction to Γ of the piecewise linear function $\log|f|$ on X^{an} and $f_v \in \kappa(C_v)$ for $v \in V$ is the normalized reduction of f to C_v (cf. §4.4). Then

$$\tau_*^{\mathfrak{C}\mathfrak{X}}(\operatorname{div}(f)) = \operatorname{div}(\mathfrak{f}).$$

In particular, we have $\tau_*^{\mathfrak{C}\mathfrak{X}}(\operatorname{Prin}(X)) \subseteq \operatorname{Prin}(\mathfrak{C}\mathfrak{X})$.

A divisor $\mathcal{E} = \sum_{x \in |\mathfrak{C}|} a_x(x)$ on \mathfrak{C} is called *effective* if $a_x \geq 0$ for all x (or, equivalently, if the Γ -part E_Γ and the C_v -parts E_v of \mathcal{E} are effective for every $v \in V$). The *rank* $r_{\mathfrak{C}}$ of a divisor $\mathcal{D} \in \operatorname{Div}(\mathfrak{C})$ is defined to be the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors \mathcal{E} of degree k on \mathfrak{C} (so in particular $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ if and only if \mathcal{D} is linearly equivalent to an effective divisor, and otherwise $r_{\mathfrak{C}}(\mathcal{D}) = -1$).

If $r_X(D)$ denotes the usual rank function $r_X(D) = \dim |D| = h^0(D) - 1$ on $\operatorname{Div}(X)$, there is an important semicontinuity result relating r_X to $r_{\mathfrak{C}\mathfrak{X}}$:

Theorem 1.2 (Specialization Theorem). *For all $D \in \operatorname{Div}(X)$, we have*

$$r_X(D) \leq r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)).$$

Since $r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) \leq r_\Gamma(\tau_*(D))$, Theorem 1.2 is a strengthening of the analogous specialization result from [B]. In conjunction with a simple combinatorial argument, Theorem 1.2 also refines the specialization lemma for vertex-weighted graphs from [AC].

We also state and prove a version of Theorem 1.2 in which one has *equality* rather than just an inequality. One can naturally associate to a rank r divisor D on X not only a divisor $\tau_*^{\mathfrak{C}\mathfrak{X}}(D)$ on $\mathfrak{C}\mathfrak{X}$, but also a collection $\mathcal{H} = \{H_v\}_{v \in V}$ of $(r+1)$ -dimensional subspaces of $\kappa(C_v)$, where H_v is the reduction to C_v of all functions in $H^0(X, \mathcal{L}(D))$ (c.f. Section 4.4

for a discussion of reduction of rational functions). If $\mathcal{F} = \{F_v\}_{v \in V}$, where F_v is any κ -subspace of the function field $\kappa(C_v)$, then for $\mathcal{D} \in \text{Div}(\mathfrak{C})$ we define the \mathcal{F} -restricted rank of \mathcal{D} , denoted $r_{\mathfrak{C}, \mathcal{F}}(\mathcal{D})$, to be the largest integer k such that for any effective divisor \mathcal{E} of degree k on \mathfrak{C} , there is a rational function f on \mathfrak{C} whose C_v -parts f_v belong to F_v for all $v \in V$, and such that $\mathcal{D} - \mathcal{E} + \text{div}(f) \geq 0$. (In terms of the ‘‘chip-firing moves’’ described above, one has to be able to make $\mathcal{D} - \mathcal{E}$ effective using only type (1) moves corresponding to functions belonging to one of the spaces F_v , together with moves of type (2) or (3).) The following variant of Theorem 1.2 will be proved as Theorem 5.9 below:

Theorem 1.3 (Specialization Theorem for Restricted Ranks). *With notation as above, the \mathcal{H} -restricted rank of the specialization of D is equal to the rank of D , i.e., $r_{\mathfrak{C}, \mathcal{H}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) = r_{\mathfrak{C}}(D)$.*

The theory of divisors and linear equivalence on metrized complexes of curves generalizes both the classical theory for algebraic curves and the corresponding theory for metric graphs and tropical curves found in [GK, MZ]. The former corresponds to the case where G consists of a single vertex v and no edges and $C = C_v$ is an arbitrary smooth curve. The latter corresponds to the case where the curves C_v have genus zero for all $v \in V$. Since any two points on a curve of genus zero are linearly equivalent, it is easy to see that the divisor theories and rank functions on \mathfrak{C} and Γ are essentially equivalent. More precisely, (i) two divisors \mathcal{D} and \mathcal{D}' on \mathfrak{C} are linearly equivalent if and only if their Γ -parts D_Γ and D'_Γ are linearly equivalent on Γ , and (ii) for every $\mathcal{D} \in \text{Div}(\mathfrak{C})$ we have $r_{\mathfrak{C}}(\mathcal{D}) = r_\Gamma(D_\Gamma)$. In the presence of higher genus curves among the C_v , the divisor theories on Γ and \mathfrak{C} can be very different. In addition, different choices of \mathcal{A}_v can drastically change both the linear equivalence relation and the rank function.

As with the corresponding specialization theory from [B], the main utility of Theorem 1.2 is that the rank function $r_{\mathfrak{C}}$ on a metrized complex of curves is surprisingly well-behaved; for example, it satisfies an analogue of the Riemann-Roch formula. In order to state this result, we need to define the canonical class. A *canonical divisor* on \mathfrak{C} , denoted \mathcal{K} , is defined to be any divisor linearly equivalent to $\sum_{v \in V} (K_v + A_v)$, where K_v is a canonical divisor on C_v and A_v is the sum of the $\deg_G(v)$ points in \mathcal{A}_v . The Γ -part of \mathcal{K} is $K^\# = \sum_v (\deg_G(v) + 2g_v - 2)(v)$, and the C_v -part of \mathcal{K} is $K_v + A_v$.

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

Theorem 1.4 (Riemann-Roch for metrized complexes of algebraic curves). *Let \mathfrak{C} be a metrized complex of algebraic curves over κ and \mathcal{K} a divisor in the canonical class of \mathfrak{C} . For any divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})$, we have*

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

The proof of Theorem 1.4 makes use of a suitable notion of *reduced divisors* for metrized complexes of curves.

Taken together, Theorems 1.2 and 1.4 have some interesting consequences. For example, the specialization theorem implies that the specialization to $\mathfrak{C}\mathfrak{X}$ of any canonical divisor on X has degree $2g - 2$ and rank at least $g - 1$ (where $g = g(X) = g(\mathfrak{C}\mathfrak{X})$), while the Riemann-Roch theorem easily implies that any such divisor on $\mathfrak{C}\mathfrak{X}$ belongs to the canonical divisor class. Thus we have:

Corollary 1.5. *For any canonical divisor K_X on X , its specialization $\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X)$ belongs to the canonical divisor class on $\mathfrak{C}\mathfrak{X}$.*

In particular, Corollary 1.5 implies that $\tau_*(K_X)$ is a canonical divisor on Γ . This fact was proved in [B], for discretely valued R , by a completely different argument based on the adjunction formula for arithmetic surfaces. It does not seem easy to prove the general (not necessarily discretely valued) case using arithmetic intersection theory, since there does not seem to be a satisfactory theory of relative dualizing sheaves and adjunction in the non-Noetherian setting.

Our theory of linear series on metrized complexes of curves has close connections with the Eisenbud-Harris theory of limit linear series for strongly semistable curves of compact type, and suggests a way to generalize the Eisenbud-Harris theory to more general semistable curves. A proper nodal curve X_0 over κ is of *compact type* if its dual graph G is a tree. For such curves, Eisenbud and Harris define a notion of (crude) *limit* \mathfrak{g}_d^r , which we discuss in detail in Section 5.2 below. A *crude limit* \mathfrak{g}_d^r L on X_0 is the data of a (not necessarily complete) degree d and rank r linear series L_v on X_v for each vertex $v \in V$ such that if two components X_u and X_v of X_0 meet at a node p , then for any $0 \leq i \leq r$,

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

where $a_i^L(p)$ denotes the i^{th} term in the vanishing sequence of a linear series L at p . A crude limit series is *refined* if all the inequalities in the above definition are equalities. For simplicity, all limit linear series in the remainder of this introduction will be crude.

We can canonically associate to a proper strongly curve X_0 a metrized complex $\mathfrak{C}X_0$ of κ -curves, called the *regularization* of X_0 , by assigning a length of 1 to each edge of G , and we write X_v for the irreducible component of X_0 corresponding to a vertex $v \in V$. (This is the metrized complex associated to any regular smoothing \mathfrak{X} of X_0 over any discrete valuation ring R with residue field κ .)

Theorem 1.6. *Let $\mathfrak{C}X_0$ be the metrized complex of curves associated to a proper strongly semistable curve X_0/κ of compact type. Then there is a bijective correspondence between the following:*

- *Crude limit \mathfrak{g}_d^r 's on X_0 .*
- *Equivalence classes of pairs $(\mathcal{H}, \mathcal{D})$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r+1)$ -dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and \mathcal{D} is a divisor of degree d supported on the vertices of $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$. Here we say that $(\mathcal{H}, \mathcal{D}) \sim (\mathcal{H}', \mathcal{D}')$ if there is a rational function \mathfrak{f} on $\mathfrak{C}X_0$ such that $\mathcal{D}' = \mathcal{D} + \text{div}(\mathfrak{f})$ and $H_v = H'_v \cdot \mathfrak{f}_v$ for all $v \in V$, where \mathfrak{f}_v denotes the C_v -part of \mathfrak{f} .*

Theorem 1.6, combined with our Riemann-Roch theorem for metrized complexes of curves, provides an arguably more conceptual proof of the fact (originally established in [EH]) that limit linear series satisfy analogues of the classical theorems of Riemann and Clifford. The point is that $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}X_0}(\mathcal{D})$ for all $\mathcal{D} \in \text{Div}(\mathfrak{C}X_0)$ and therefore upper bounds on $r_{\mathfrak{C}X_0}(\mathcal{D})$ which follow from Riemann-Roch imply corresponding upper bounds on the restricted rank $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D})$.

Motivated by Theorem 1.6, we propose the following definition. Let X_0 be a proper strongly semistable (but not necessarily compact type) curve over κ with regularization $\mathfrak{C}X_0$. A *limit* \mathfrak{g}_d^r on X_0 is an equivalence class of pairs $(\mathcal{H} = \{H_v\}, \mathcal{D})$ as above, where H_v

is an $(r + 1)$ -dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and \mathcal{D} is a degree d divisor on $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$.

As a partial justification for this definition, we prove that if R is a discrete valuation ring with residue field κ and \mathfrak{X}/R is a regular arithmetic surface whose generic fiber X is smooth and whose special fiber X_0 is strongly semistable, then for any divisor D on X with $r_X(D) = r$ and $\deg(D) = d$, our specialization machine gives rise in a natural way to a limit \mathfrak{g}_d^r on X_0 (see Theorem 5.10 below for a precise statement, and Theorem 5.9 for a more general statement).

2. METRIZED COMPLEXES OF ALGEBRAIC CURVES

Let κ be an algebraically closed field of arbitrary characteristic and let \mathfrak{C} be a metrized complex of κ -curves, as defined in Section 1.2.

2.1. Divisors and their rank. If \mathcal{G} is a subgroup of \mathbb{R} containing the lengths of all edges of G , we write $\text{Div}(\mathfrak{C})_{\mathcal{G}}$ for the subgroup of $\text{Div}(\mathfrak{C})$ consisting of all divisors \mathcal{D} such that the support of the Γ -part of \mathcal{D} consists entirely of \mathcal{G} -rational points of Γ (i.e., points of Γ whose distances to the vertices of G are in \mathcal{G}).

Recall from Section 1.2 that the *rank* $r_{\mathfrak{C}}$ of a divisor \mathcal{D} on \mathfrak{C} is the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors \mathcal{E} of degree k on \mathfrak{C} . If \mathcal{G} is any subgroup of \mathbb{R} containing the lengths of all edges of G , one can define an analogous rank function $r_{\mathfrak{C}, \mathcal{G}}$ for divisors in $\text{Div}(\mathfrak{C})_{\mathcal{G}}$ by restricting the effective divisors \mathcal{E} in the above definition to lie in $\text{Div}(\mathfrak{C})_{\mathcal{G}}$. By Corollary A.5, $r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D}) = r_{\mathfrak{C}}(\mathcal{D})$ for any divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})_{\mathcal{G}}$, so the restriction in the definition of $r_{\mathfrak{C}, \mathcal{G}}$ does not affect the rank of divisors.

2.2. Regularization of nodal curves. Let X_0 be a (reduced) connected projective nodal curve over κ with irreducible components X_{v_1}, \dots, X_{v_k} . We can associate to X_0 a metrized complex $\mathfrak{C}X_0$, called the *regularization* of X_0 , as follows. The underlying metric graph Γ_0 has model the dual graph $G_0 = (V_0, E_0)$ of X_0 . Recall that the vertex set V_0 of G_0 consists of vertices v_1, \dots, v_k (in bijection with the irreducible components of X_0) and the edges of G_0 are in bijection with singular points of X_0 . (Note that G_0 might have loop edges.) The length of all the edges in Γ_0 are equal to one. The κ -curve C_{v_i} in $\mathfrak{C}X_0$ is the normalization of X_i for $i = 1, \dots, k$. The collection \mathcal{A}_{v_i} is the set of all κ -points of C_{v_i} which lie over a singular point of X_0 in X_{v_i} . By the definition of the dual graph, these points are in bijection with the edges adjacent to v_i in G_0 .

If X_0 is *strongly semistable*, so that G_0 does not have any loop edge and $C_{v_i} = X_{v_i}$, the rank function $r_{\mathfrak{C}X_0, \mathbb{Z}}$ on $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ can be reformulated as follows. Let $\text{Pic}(X_0)$ be the Picard group of X_0 , and consider the restriction map $\pi : \text{Pic}(X_0) \rightarrow \bigoplus_i \text{Pic}(X_{v_i})$. For any line bundle \mathcal{L} on X_0 , and $v \in V_0$, denote by \mathcal{L}_v the restriction of \mathcal{L} to X_v .

Two line bundles \mathcal{L} and \mathcal{L}' in $\text{Pic}(X_0)$ are said to be *combinatorially equivalent* if there exists a function $f : V_0 \rightarrow \mathbb{Z}$ such that $\mathcal{L}'_v = \mathcal{L}_v(\text{div}_v(f))$ in $\text{Pic}(X_v)$ for any vertex v of G_0 . (In a regular smoothing of X_0 , c.f. Section 5 for the definition, the role of f is to specify a particular twist of \mathcal{L} by a divisor supported on the irreducible components of X_0 .) In particular, two line bundles \mathcal{L} and \mathcal{L}' in $\text{Pic}(X_0)$ with $\pi(\mathcal{L}) = \pi(\mathcal{L}')$ are combinatorially equivalent. In the above definition, div_v is defined in analogy with (1) as follows:

$$\text{div}_v(f) = \sum_{u: \{u, v\} \in E_0} (f(u) - f(v))(x_v^{\{u, v\}})$$

where $x_v^{\{u,v\}}$ is the point of C_v labelled with the edge $\{u,v\}$. Note that the function $f : V_0 \rightarrow \mathbb{Z}$ in the definition of combinatorially equivalent line bundles above is unique up to an additive constant. For such an f , we denote by \mathcal{L}^f any line bundle \mathcal{L}' in $\text{Pic}(X_0)$ such that $\mathcal{L}'_v = \mathcal{L}_v(\text{div}_v(f))$ for all $v \in V_0$.

For a line bundle \mathcal{L} on X_0 , define the *combinatorial rank* of \mathcal{L} , denoted $r_c(\mathcal{L})$, to be the maximum integer r such that for any effective divisor $E = \sum_{v \in V_0} E(v)(v)$ on G_0 of degree r , there is a line bundle \mathcal{L}' combinatorially equivalent to \mathcal{L} with

$$\dim_{\kappa} H^0(X_v, \mathcal{L}'_v) \geq E(v) + 1$$

for all $v \in V$. In particular, if no non-negative integer r with the above property exists, the combinatorial rank of \mathcal{L} will be equal to -1 . The combinatorial rank of \mathcal{L} clearly depends only on the combinatorial equivalence class of \mathcal{L} .

Let \mathcal{D} be a divisor in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ with X_v -part D_v . We denote by $\mathcal{L}(\mathcal{D})$ any line bundle on X_0 whose restriction to X_v is $\mathcal{L}(D_v)$ for each vertex v of G_0 .

Proposition 2.1. *Let X_0 be a strongly semistable curve over κ and let $\mathfrak{C}X_0$ be the corresponding metrized complex. Then:*

- (1) *Two divisors \mathcal{D} and \mathcal{D}' in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ are linearly equivalent iff $\mathcal{L}(\mathcal{D})$ and $\mathcal{L}(\mathcal{D}')$ are combinatorially equivalent.*
- (2) *For any divisor $\mathcal{D} \in \text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$, the combinatorial rank of $\mathcal{L}(\mathcal{D})$ is equal to $r_{\mathfrak{C}X_0}(\mathcal{D})$.*

Proof. The first part follows by definition. To see the second part, let r be the combinatorial rank of $\mathcal{L}(\mathcal{D})$. It follows from Corollary A.6 that $r_{\mathfrak{C}X_0}(\mathcal{D}) = r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D})$. Thus, it will be enough to show that $r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D}) = r$.

We first prove that $r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D}) \geq r$. Let \mathcal{E} be an effective divisor in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ with Γ -part E_{Γ} and X_v -part E_v . By the definition of the combinatorial rank, there exists a function $f : V_0 \rightarrow \mathbb{Z}$ such that

$$\dim_{\kappa} H^0(X_v, \mathcal{L}_v^f) \geq E_{\Gamma}(v) + 1 \quad \text{for any vertex } v \text{ in } G_0.$$

Since $\deg(E_v) = E_{\Gamma}(v)$, this implies the existence of a global section f_v of \mathcal{L}_v^f such that f_v has order of vanishing at least $E_v(x)$ at any point x of X_v .

Let f_{Γ} be the rational function defined by f on Γ . (Recall that $f_{\Gamma}|_{V_0} = f$ and f is integer affine with slope $f(u) - f(v)$ on each edge $\{u,v\}$.) For the rational function \mathfrak{f} on $\mathfrak{C}X_0$ defined by f_{Γ} and $\{f_v\}_{v \in V_0}$, one has $\text{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} \geq 0$. This shows that $r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D}) \geq r$.

We now prove that $r \geq r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D})$. Let $\mathcal{L} = \mathcal{L}(\mathcal{D})$ be a line bundle on X_0 defined by \mathcal{D} . Let E be an effective divisor of degree $r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D})$ on G_0 , and let \mathcal{E} be any effective divisor in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ with Γ -part equal to E . Then there exists a rational function \mathfrak{f} with $\text{div}(\mathfrak{f}) \in \text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ such that $\text{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} \geq 0$. In particular, the Γ -part f_{Γ} of \mathfrak{f} is linear on each edge of G_0 and $\text{div}(f_{\Gamma}) + D_{\Gamma} \geq 0$, which shows that up to an additive constant, the set F of all functions f_{Γ} with these properties is finite. Let S_{Γ} denote the finite set $\{\text{div}(f_{\Gamma}) \mid f_{\Gamma} \in F\}$. Let $d_v = E(v)$, and for $D_0 \in S_{\Gamma}$, let S_{D_0} be the subset of $\prod_{v \in V(G_0)} X_v^{(d_v)}$ defined by all collections of effective divisors E_v of degree d_v such that there exists a rational function \mathfrak{f} on $\mathfrak{C}X_0$ with Γ -component f_{Γ} satisfying $\mathcal{D} - \mathcal{E} + \text{div}(\mathfrak{f}) \geq 0$. (Here $X^{(n)}$ denotes the n^{th} symmetric product of a curve X .) Now let S be the union over all $D_0 \in S_{\Gamma}$ of S_{D_0} . Then one sees easily that:

- (i) Each S_{D_0} is Zariski-closed in $\prod_v X_v^{(d_v)}$. (Use the fact that S_{D_0} is the projection of a Zariski closed subset of a product of proper curves.)
- (ii) $\bigcup_{D_0 \in S_\Gamma} S_{D_0} = \prod_v X_v^{(d_v)}$.

Therefore there exists $D_0 \in S_\Gamma$ such that $S_{D_0} = \prod_v X_v^{(d_v)}$. In other words, we can suppose that the Γ -part f_Γ of all the rational functions f giving $\mathcal{D} - \mathcal{E} + \text{div}(f) \geq 0$, when the effective divisors E_v vary, is constant (and we have $D_0 = \text{div}(f_\Gamma)$). If we fix a corresponding twist \mathcal{L}^{f_Γ} , then

$$\dim_\kappa H^0(X_v, \mathcal{L}^{f_\Gamma}) \geq E(v) + 1.$$

Since E was an arbitrary effective divisor of degree $r_{\mathfrak{C}X_0, \mathbb{Z}}(\mathcal{D})$ on G_0 , we obtain the desired inequality. □

3. A RIEMANN-ROCH THEOREM FOR METRIZED COMPLEXES OF CURVES

Our aim in this section is to formulate and prove a Riemann-Roch theorem for metrized complexes of κ -curves.

Let \mathfrak{C} be a metrized complex of algebraic curves over κ . For each vertex $v \in V$, let g_v be the genus of C_v and let $g(\mathfrak{C}) = g(\Gamma) + \sum_v g_v$ be the genus of \mathfrak{C} . For each curve C_v , let K_v denote a divisor of degree $2g_v - 2$ in the canonical class of C_v . Denote by A_v the divisor in C_v consisting of the sum of the $\deg_G(v)$ points in \mathcal{A}_v . The *canonical class* of \mathfrak{C} is defined to be the linear equivalence class of the divisor $\mathcal{K} = \sum_{v \in V} (K_v + A_v)$.

Remark 3.1. Let X_0 be a strongly semistable curve. The definition of the canonical class \mathcal{K} in $\mathfrak{C}X_0$ is compatible in a natural sense with the definition of the dualizing sheaf (or sheaf of logarithmic 1-forms) $\omega_{X_0/\kappa}$. Indeed, if $\alpha : X \rightarrow X_0$ denotes the normalization of X_0 , and y^e, z^e denote the points of X above the singular point x^e of X_0 for $e \in E(G_0)$, then sections of the sheaf $\omega_{X_0/\kappa}$ on an open set U consist of all the meromorphic 1-forms δ on X which are regular everywhere on U except for possible simple poles at y^e, z^e , with $\text{Res}_{y^e}(\delta) + \text{Res}_{z^e}(\delta) = 0$ for all e . From this, it is easy to see that the restriction of $\omega_{X_0/\kappa}$ to each component X_v of X_0 is the invertible sheaf $\omega_{X_v}(A_v)$ corresponding to $K_v + A_v$ for any canonical divisor K_v on X_v .

For a divisor \mathcal{D} on \mathfrak{C} , let $r_{\mathfrak{C}}(\mathcal{D})$ denote the rank of \mathcal{D} as defined in Section 1.2.

Theorem 3.2 (Riemann-Roch for metrized complexes of algebraic curves). *Let \mathfrak{C} be a metrized complex of κ -curves and \mathcal{K} a divisor in the canonical class of \mathfrak{C} . For every divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})$, we have*

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

Let X_0 be a strongly semistable curve of genus g over κ . For any $\mathcal{L} \in \text{Pic}(X_0)$, let $r_c(\mathcal{L})$ be the combinatorial rank of \mathcal{L} as defined in Section 2. As a corollary of Theorem 3.2, Remark 3.1, and Proposition 2.1, we have:

Corollary 3.3 (Riemann-Roch for strongly semistable curves). *For any $\mathcal{L} \in \text{Pic}(X_0)$,*

$$r_c(\mathcal{L}) - r_c(\omega_{X_0/\kappa} \otimes \mathcal{L}^{-1}) = \deg(\mathcal{L}) - g(X_0) + 1.$$

The Riemann-Roch theorem has a number of well-known formal consequences which can be transported to our situation. For example, we obtain the following analogue of Clifford's bound for the rank of a special divisor. (A divisor \mathcal{D} on a metrized complex \mathfrak{C} is called *special* if $K_{\mathfrak{C}} - \mathcal{D}$ is linearly equivalent to an effective divisor.)

Theorem 3.4 (Clifford's theorem for metrized complexes). *For any special divisor \mathcal{D} on a metrized complex \mathfrak{C} , $r_{\mathfrak{C}}(\mathcal{D}) \leq \deg(\mathcal{D})/2$.*

Remark 3.5. Our Riemann-Roch theorem seems to be rather different from the classical Riemann-Roch theorem on semistable curves. We refer to [C1] for a discussion of Riemann's and Clifford's theorems for linear series on semistable curves, and for examples showing the failure of Clifford's theorem for $h^0(\mathcal{L})$ with $\mathcal{L} \in \text{Pic}(X_0)$ (even for line bundles in the compactified Picard scheme of X_0).

The rest of this section is devoted to the proof of Theorem 3.2. Our proof follows and extends the original arguments of [BN] in the proof of Riemann-Roch theorem for graphs. In particular, we are going to first extend the notion of reduced divisors to the context of metrized complexes of algebraic curves, and then study the minimal non-special divisors.

3.1. Reduced divisors: existence and uniqueness. Let v_0 be a fixed base point of Γ . We introduce the notion of v_0 -reduced divisors and show that each equivalence class of divisors on \mathfrak{C} contains a quasi-unique v_0 -reduced divisor, in a precise sense to be defined immediately preceding Theorem 3.7 below.

For a closed connected subset S of Γ and a point $v \in \partial S$ (the topological boundary of S), the number of "outgoing" tangent directions at v (i.e., tangent directions emanating outward from S) is denoted by $\text{outdeg}_S(v)$; this is also the maximum size of a collection of internally disjoint segments in $\Gamma \setminus (S - \{v\})$ with one end at v . If in addition we have $v \in V$, we denote by $\text{div}_v(\partial S)$ the divisor in C_v associated to the outgoing edges at v ; this is by definition the sum of all the points x_v^e of C_v indexed by the edges e leaving S at v . In what follows we refer to a closed connected subset S of Γ as a *cut* in Γ .

Let \mathcal{D} be a divisor on \mathfrak{C} with Γ -part D_{Γ} and C_v -part D_v , and let S be a cut in Γ . A boundary point $x \in \partial S$ is called *saturated* with respect to \mathcal{D} and S if

- $x \notin V$ and $\text{outdeg}_S(x) \leq D_{\Gamma}(x)$; or
- $x = v$ for some $v \in V$ and $D_v - \text{div}_v(\partial S)$ is equivalent to an effective divisor on C_v .

Otherwise, $x \in \partial S$ is called *non-saturated*. A cut S is said to be saturated if all its boundary points are saturated. (When talking about saturated and non-saturated points, we will sometimes omit the divisor \mathcal{D} or the set S if they are clear from the context.)

The divisor \mathcal{D} is said to be *v_0 -reduced* if the following three properties are satisfied:

- (i) For all points $x \neq v_0$ of Γ , $D_{\Gamma}(x) \geq 0$, i.e., all the coefficients of D_{Γ} are non-negative except possibly at the base point v_0 .
- (ii) For all points $v \in V \setminus \{v_0\}$ ($= V$ if v_0 does not belong to V), there exists an effective divisor E_v linearly equivalent to D_v on C_v .
- (iii) For every cut S of Γ which does not contain v_0 , there exists a non-saturated point $x \in \partial S$.

Remark 3.6. There is an efficient "burning algorithm" for checking whether or not a given divisor \mathcal{D} on \mathfrak{C} which satisfies (i) is v_0 -reduced. The algorithm can be described informally as follows (compare with [L, §2] or [BS, §5.1]). Imagine that Γ is made of a flammable

material, and light a fire at v_0 . The fire begins spreading across Γ in a continuous manner and can only be blocked at a point $x \in \Gamma$ if one of the following holds:

- $x \notin V$ and the number of burnt directions coming into x exceeds $D_\Gamma(x)$. (One should imagine $D_\Gamma(x)$ firefighters standing at x , each of whom can block fire in one incoming direction.)
- $x = v$ for some $v \in V$ and D_v minus the sum of the marked points corresponding to burnt directions is not equivalent to any effective divisor on C_v .

It is straightforward to check, following the same ideas as the proof of [L, Algorithm 2.5], that \mathcal{D} is v_0 -reduced if and only if the fire eventually burns through all of Γ . We omit the details since we will not need this result in the sequel.

We now show that every divisor \mathcal{D} on \mathfrak{C} is linearly equivalent to a quasi-unique v_0 -reduced divisor \mathcal{D}^{v_0} . The quasi-uniqueness is understood in the following sense: the Γ -part $D_\Gamma^{v_0}$ of \mathcal{D}^{v_0} is unique, and for all $v \in V$, the divisor class $[D_v^{v_0}]$ on C_v defined by the v -part of \mathcal{D}^{v_0} is unique.

Theorem 3.7. *Let \mathfrak{C} be a metrized complex of κ -curves and v_0 a base point of Γ . For every divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})$, there exists a quasi-unique v_0 -reduced divisor \mathcal{D}^{v_0} such that $\mathcal{D}^{v_0} \sim \mathcal{D}$.*

Proof. Let $\mathcal{T}_\mathcal{D}$ be the set of all divisors \mathcal{D}' linearly equivalent to \mathcal{D} such that:

- All the coefficients of the Γ -part D'_Γ of \mathcal{D}' are non-negative at every point of Γ except possibly at v_0 .
- For each point $v \in V \setminus \{v_0\}$, the v -part D'_v of \mathcal{D}' has non-negative rank on C_v .
- The coefficient of v_0 in D'_Γ is maximal with respect to the two properties above.

One shows as in [A, Proof of Theorem 2] or [BN, Proof of Proposition 3.1] that there is a divisor \mathcal{D}' linearly equivalent to \mathcal{D} on \mathfrak{C} with the property that the coefficient of every point $x \in \Gamma \setminus (V \cup \{v_0\})$ is non-negative and the coefficient of D'_Γ at each point $v \in V$ (which coincides with the degree of D'_v) is at least g_v . (The intuitive idea is to repeatedly fire v_0 and then use further chip-firing moves to spread chips wherever needed on Γ , assuring that each $v \in V$ gets at least g_v chips and that no $p \in \Gamma$ except for v_0 remains in debt.) By the (classical) Riemann-Roch theorem, each D'_v has non-negative rank on C_v , which shows that $\mathcal{T}_\mathcal{D}$ is non-empty.

Define $T_\mathcal{D}$ as the set of all $D'_\Gamma \in \text{Div}(\Gamma)$ for which there exists $\mathcal{D}' \in \mathcal{T}_\mathcal{D}$ whose Γ -part is D'_Γ . Each divisor $D'_\Gamma \in T_\mathcal{D}$ has degree $\deg(\mathcal{D})$, and all divisors in $T_\mathcal{D}$ have the same coefficient at v_0 . Thus the number $N := \sum_{x \in \Gamma \setminus \{v_0\}} D'_\Gamma(x)$ is independent of the choice of a divisor D'_Γ in $T_\mathcal{D}$. As a consequence, $T_\mathcal{D}$ inherits a natural topology from the topology of Γ , since it can be embedded as a subset of $\text{Sym}^N \Gamma$, and is compact. Let \mathcal{A} be the subset of \mathbb{R}^N defined by

$$\mathcal{A} := \left\{ (x_1, \dots, x_N) \mid 0 \leq x_1 \leq x_2 \leq \dots \leq x_N \right\},$$

equipped with the total order defined by the lexicographical rule $(x_i) \leq (y_i)$ iff $x_1 = y_1, \dots, x_i = y_i$ and $x_{i+1} < y_{i+1}$. One considers a map $F : T_\mathcal{D} \rightarrow \mathcal{A}$ defined as follows. For each divisor $D'_\Gamma \in T_\mathcal{D}$, consider the multiset $A(D'_\Gamma)$ of points in $\Gamma \setminus \{v_0\}$ where each point $v \neq v_0$ appears in this multiset exactly $D'_\Gamma(v)$ times. Define $F(D'_\Gamma)$ to be the point of \mathcal{A} defined by the multiset of distances $\text{dist}_\Gamma(v, v_0)$ for $v \in A(D'_\Gamma)$, ordered in an increasing way. It is straightforward to verify that the map F is continuous. Since $T_\mathcal{D}$ is compact and F is continuous, there exists a divisor $D_\Gamma^{v_0}$ in $T_\mathcal{D}$ such that F takes its minimum value at $D_\Gamma^{v_0}$, i.e., $F(D_\Gamma^{v_0}) = \min_{D'_\Gamma \in T_\mathcal{D}} F(D'_\Gamma)$. Let \mathcal{D}^{v_0} be an element of $\mathcal{T}_\mathcal{D}$ whose Γ -part is $D_\Gamma^{v_0}$.

Claim: The divisor $\mathcal{D}^{v_0} \sim \mathcal{D}$ is v_0 -reduced.

Properties (i) and (ii) above are clearly satisfied. The only point one needs to check is that every cut S which does not contain v_0 has a non-saturated point on its boundary. For the sake of contradiction, suppose this is not the case, and let S be a closed connected set violating this condition. This means that the following hold:

- For all $x \in \partial S \setminus V$, $\text{outdeg}_S(x) \leq D_\Gamma^{v_0}(v)$.
- For each $v \in \partial S \cap V$, $D_v^{v_0} - \text{div}_v(\partial S)$ has non-negative rank on C_v , i.e., there exists a rational function f_v on C_v such that $D_v^{v_0} - \text{div}_v(\partial S) + \text{div}(f_v) \geq 0$.

By the definition of outdeg_S , there exists an $\epsilon > 0$ such that for each vertex $x \in \partial S$, there are closed segments $I_1^x, \dots, I_{\text{outdeg}_S(x)}^x$ emanating from x with the following properties:

- For $x \in \partial S$ and $1 \leq j \leq \text{outdeg}_S(x)$, the half-open segments $I_j^x \setminus x$ are disjoint from S and from each other and do not contain v_0 .
- Each segment I_j^x has length ϵ , for $x \in \partial S$ and $1 \leq j \leq \text{outdeg}_S(x)$.

These data give rise to a rational function $f_\Gamma : \Gamma \rightarrow \mathbb{R}$ which is identically zero on S , is linear of slope -1 on each interval I_j^x for $x \in \partial S$ and $1 \leq j \leq \text{outdeg}_S(x)$, and takes the constant value $-\epsilon$ at all points of $\Gamma \setminus (S \cup \bigcup_{x,j} I_j^x)$. Consider the divisor $\mathcal{D}^* = \mathcal{D}^{v_0} + \text{div}(\mathfrak{f})$ on \mathfrak{C} , where \mathfrak{f} is the rational function on \mathfrak{C} consisting of the rational function f_Γ on Γ and the rational functions f_v on C_v . Letting D_Γ^* be the Γ -part of \mathcal{D}^* , one verifies that \mathcal{D}^* and D_Γ^* lie in $\mathcal{T}_\mathcal{D}$ and $T_\mathcal{D}$, respectively, and that $F(D_\Gamma^*) < F(D_\Gamma^{v_0})$, contradicting the choice of $D_\Gamma^{v_0}$. This proves the claim and hence the existence part of the theorem.

It remains to prove the quasi-uniqueness. Assume for the sake of contradiction that there are linearly equivalent v_0 -reduced divisors \mathcal{D} and \mathcal{D}' on \mathfrak{C} such that either $D_\Gamma \neq D'_\Gamma$ or D_v and D'_v are not linearly equivalent on C_v for some $v \in V$. Then there exists a non-constant rational function \mathfrak{f} on \mathfrak{C} such that $\mathcal{D}' = \mathcal{D} + \text{div}(\mathfrak{f})$. If f_Γ is constant then we obtain an immediate contradiction, so we may assume that f_Γ is non-constant. Without loss of generality, we may assume that f_Γ does not take its maximum at v_0 (otherwise, we can interchange the role of \mathcal{D} and \mathcal{D}'). Let S be a connected component of the set of all points where f_Γ takes its maximum. Note that $v_0 \notin S$. For all points $x \in \partial S$, the slope of f_Γ at any outgoing segment I_j^x emanating from x is at most -1 . Since \mathcal{D} is v_0 -reduced, there exists a point $x \in \partial S$ such that either $x \notin V$ and $D_\Gamma(x) < \text{outdeg}_S(x)$, or $x = v$ for some $v \in V$ and the divisor $D_v - \text{div}(\partial S)$ has negative rank. In the first case, $D'_\Gamma(x) \leq D_\Gamma(x) - \text{outdeg}_S(x) < 0$, contradicting the assumption that the coefficient of D'_Γ is non-negative at $x \neq v_0$. In the second case, $D'_v \sim D_v + \sum_{e \in E: e \sim v} \text{slp}_e(f_\Gamma)(x_v^e) \leq D_v - \text{div}_v(\partial S)$, which implies that D'_v has negative rank, a contradiction. \square

Remark 3.8. One can presumably give an “iterated Dhar algorithm” for algorithmically finding a v_0 -reduced divisor equivalent to a given divisor \mathcal{D} by combining the ideas in the proof of Theorem 3.7 with the ideas behind [L, Algorithm 2.12].

3.2. Description of minimal non-special divisors. Let \mathfrak{C} be a metrized complex of algebraic curves. For each $v \in V$, let \mathfrak{N}_v be the set of all *minimal non-special* divisors on C_v :

$$\mathfrak{N}_v = \{D \in \text{Div}(C_v) : \deg(D) = g(C_v) - 1 \text{ and } |D| = \emptyset\}.$$

Remark 3.9. Recall from [H77, IV, Example 1.3.4] that a divisor D on a smooth projective curve C of genus g over κ is called *special* if $r(K_C - D) \geq 0$. By Riemann-Roch, every

divisor of degree at most $g - 2$ on C is special, and if $\deg(D) = g - 1$ then D is special if and only if $r(D) \geq 0$. Moreover, every non-special divisor dominates a minimal non-special divisor. This explains the term “minimal non-special divisor”.

Similarly, define the set of minimal non-special divisors on \mathfrak{C} as

$$\mathfrak{N} = \{\mathcal{D} \in \text{Div}(\mathfrak{C}) : \deg(\mathcal{D}) = g(\mathfrak{C}) - 1 \text{ and } |\mathcal{D}| = \emptyset\}.$$

In this section we provide an explicit description of the minimal non-special divisors on metrized complexes of algebraic curves, generalizing the corresponding description of minimal non-special divisors on graphs and metric graphs from [BN, MZ].

An *acyclic orientation* on Γ is an acyclic orientation on some model of Γ , i.e., a decomposition of Γ into closed directed edges with disjoint interiors such that no directed cycles are formed. Given an acyclic orientation π on Γ , we denote by $\deg_{\pi}^{+}(x)$ the number of tangent directions emanating from x which are compatible with π . Note that for all but finitely many points of Γ , $\deg_{\pi}^{+}(x) = 1$. For $v \in V$, we denote by $E^{+}(v)$ the set of edges incident to v which (locally near v) are oriented outward from v . A point x with $\deg_{\pi}^{+}(x) = 0$ is called a *sink*. It is well known and easy to prove that every acyclic orientation contains at least one sink (start at any point and follow the orientation until you get stuck, which must eventually happen by acyclicity).

Given a collection of minimal non-special divisors $D_v \in \mathfrak{N}_v$ for each $v \in V$, together with an acyclic orientation π of Γ , define a corresponding divisor $\mathcal{D}^{\pi, \{D_v\}}$ by the formula

$$\mathcal{D}^{\pi, \{D_v\}} = \sum_{x \in \Gamma \setminus V} (\deg_{\pi}^{+}(x) - 1)(x) + \sum_{v \in V} (A_v^{\pi} + D_v),$$

where A_v^{π} is the sum of all points $x_v^e \in \mathcal{A}_v$ for which the π -orientation on e points away from v .

The Γ -part of $\mathcal{D}^{\pi, \{D_v\}}$ is

$$D_{\Gamma}^{\pi} := \sum_{x \in \Gamma} (\deg_{\pi}^{+}(x) + g_x - 1)(x)$$

and the C_v -part of $\mathcal{D}^{\pi, \{D_v\}}$ is $D_v^{\pi} := A_v^{\pi} + D_v$. If $\tilde{G} = (\tilde{V}, \tilde{E})$ is a (loopless) model of Γ such that the orientation π of Γ is induced by an acyclic orientation \tilde{G}^{π} on the edges of \tilde{G} , then the degree of $\mathcal{D}^{\pi, \{D_v\}}$ is given by

$$\begin{aligned} \deg(\mathcal{D}^{\pi, \{D_v\}}) &= \sum_{v \in \tilde{V}} (\deg_{\pi}^{+}(v) + g_v - 1) \\ &= \left[\sum_{v \in \tilde{V}} (\deg_{\pi}^{+}(v) - 1) \right] + \sum_{v \in V} g_v = g(\Gamma) - 1 + \sum_v g_v = g(\mathfrak{C}) - 1. \end{aligned}$$

A divisor $\mathcal{M} \in \text{Div}(\mathfrak{C})$ of the form $\mathcal{D}^{\pi, \{D_v\}}$ is called a *moderator* on \mathfrak{C} . (This terminology comes from [MZ].) Given a moderator $\mathcal{M} = \mathcal{D}^{\pi, \{D_v\}}$, the *dual moderator* $\bar{\mathcal{M}}$ is defined to be $\mathcal{D}^{\bar{\pi}, \{K_v - D_v\}}$, where $\bar{\pi}$ is obtained from π by reversing the orientation of every oriented segment. It is easy to see that $\mathcal{M} + \bar{\mathcal{M}}$ belongs to the canonical class on \mathfrak{C} .

The following two lemmas are essentially what we need for the proof of Theorem 3.2.

Lemma 3.10. *For any acyclic orientation π of Γ and any collection $D_v \in \mathfrak{N}_v$ of minimal non-special divisors on C_v , the moderator $\mathcal{D}^{\pi, \{D_v\}}$ is a minimal non-special divisor on \mathfrak{C} .*

Proof. Suppose for the sake of contradiction that there is a rational function f on \mathfrak{C} such that $\mathcal{D}^{\pi, \{D_v\}} + \text{div}(f) \geq 0$. Let $\tilde{G} = (\tilde{V}, \tilde{E})$ be a (loopless) model of Γ such that the orientation π of Γ is induced by an acyclic orientation \tilde{G}^π on the edges of \tilde{G} . Let S be the closed subset of Γ consisting of all the points of Γ where the Γ -part f_Γ of f takes its maximum value. It is easy to see that $\partial S \subseteq \tilde{V}$, since at a hypothetical point $x \in \partial S \setminus \tilde{V}$ we would necessarily have $D_\Gamma^\pi(x) + \text{div}_x(f_\Gamma) < 0$, contradicting the assumption on f .

Note that since $\partial S \subseteq \tilde{V}$, for any edge e of \tilde{G} either S contains e entirely or S does not intersect the interior of e . Consider now the restriction of the orientation π to the induced subgraph $\tilde{G}[\tilde{V} \cap S]$. The resulting directed graph $\tilde{G}^\pi[\tilde{V} \cap S]$ is acyclic and therefore contains a sink vertex u .

Claim: $u \in \partial S$.

It suffices to prove that no sink vertex of the directed graph \tilde{G}^π belongs to S , so assume for the sake of contradiction that $w \in S$ is a sink vertex of \tilde{G}^π . There are two cases to consider:

- If $w \in V$, then $D_w + \text{div}(f_w) \geq D_w^\pi + \text{div}_w(f_\Gamma) + \text{div}(f_w) \geq 0$ (since f_Γ achieves its maximum value at w and we supposed that $\mathcal{D}^{\pi, \{D_v\}} + \text{div}(f) \geq 0$), contradicting the assumption that D_w is a minimal non-special divisor on C_w .
- If $w \notin V$, then by definition the coefficient of w in D_Γ^π is at most -1 , from which it follows that the coefficient of w in $\mathcal{D}^{\pi, \{D_v\}} + \text{div}(f)$ is negative, a contradiction.

This proves the claim.

By our choice of $u \in \partial S$ and the definition of S , f_Γ has strictly negative slope along all outgoing edges at u , and it has slope zero along all other edges incident to u . This shows in particular that $\text{ord}_u(f_\Gamma) \leq -\text{outdeg}_S(u) \leq -\text{deg}_\pi^+(u)$. There are now two different cases to consider, both of which will lead to a contradiction (and hence complete our proof of the lemma):

- If $u \in \tilde{V} \setminus V$, then $D_\Gamma^\pi(u) + \text{ord}_u(f_\Gamma) \leq \text{deg}_\pi^+(u) - 1 - \text{deg}_\pi^+(u) < 0$, contradicting the choice of the rational function f on \mathfrak{C} .
- If $u \in V$, then for each edge e incident to u we have $\text{slp}_e(f_\Gamma) \leq 0$, with strict inequality if e is an outgoing edge at u in \tilde{G}^π (because u is a sink of the oriented graph $\tilde{G}^\pi[S \cap \tilde{V}]$). By the definition of D_u^π , we infer that $D_u^\pi + \sum_{e \in \tilde{E}: e \sim u} \text{slp}_e(f_\Gamma)(x_u^e) \leq D_u$, which shows that $D_u^\pi + \sum_{e \in \tilde{E}: e \sim u} \text{slp}_e(f_\Gamma)(x_u^e) + \text{div}(f_u)$ is not effective, since D_u is minimal non-special. This contradicts our assumptions on f .

□

For $v_0 \in \Gamma$, we denote by $\mathcal{AO}_{v_0}(\Gamma)$ the set of all acyclic orientations of Γ with a unique sink at v_0 , i.e., such that v_0 has out-degree zero and all other points of Γ have out-degree at least one.

Lemma 3.11. *Let \mathcal{D} be a v_0 -reduced divisor on \mathfrak{C} . Then $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ if and only if the following two conditions hold:*

- *The coefficient $D_\Gamma(v_0)$ of D_Γ at v_0 is non-negative.*
- *If $v_0 \in V$, the divisor D_{v_0} on C_{v_0} has non-negative rank.*

More precisely, if the above conditions do not both hold, then there exists a moderator $\mathcal{M} = \mathcal{D}^{\pi, \{D_v^\}}$ with $\pi \in \mathcal{AO}_{v_0}(\Gamma)$ such that $\mathcal{D} \leq \mathcal{M}$.*

Proof. The intuitive idea is that one constructs the moderator $\mathcal{M} = \mathcal{D}^{\pi, \{D_v^*\}}$ using the burning algorithm described in Remark 3.6. Beginning with a v_0 -reduced divisor \mathcal{D} , burn through Γ following the algorithm and keep track of the direction in which the fire spreads; reversing all the arrows defines an acyclic orientation $\pi \in \mathcal{AO}_{v_0}(\Gamma)$. If \mathcal{M} is the corresponding moderator, one checks that $\mathcal{D} \leq \mathcal{M}$ outside v_0 , and if \mathcal{D} is not effective then $\mathcal{D} \leq \mathcal{M}$ everywhere. A more rigorous version of this argument is as follows.

First suppose that $D_\Gamma(v_0) \geq 0$, and in the case $v_0 \in V$ that there exists a rational function f_{v_0} on C_{v_0} such that $D_{v_0} + \text{div}(f_{v_0}) \geq 0$. We infer that $D_\Gamma \geq 0$, and since each D_v for $v \in V \setminus \{v_0\}$ is of non-negative rank, there exists a rational function f_v on C_v such that $D_v + \text{div}(f_v) \geq 0$. Let f be a rational function on \mathfrak{C} consisting of a constant rational function on Γ and f_v on C_v for $v \in V$. We obviously have $\mathcal{D} + \text{div}(f) \geq 0$, which implies that $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$.

For the other direction, assume that $D_\Gamma(v_0) < 0$ if $v_0 \notin V$, and that D_{v_0} has negative rank if $v_0 \in V$. To show that $r_{\mathfrak{C}}(\mathcal{D}) = -1$, by Lemma 3.10 it will be enough to show the existence of an acyclic orientation $\pi \in \mathcal{AO}_{v_0}(\Gamma)$ and a set of minimal non-special divisors $D_v^* \in \mathfrak{N}_v$ such that $\mathcal{D} \leq \mathcal{D}^{\pi, \{D_v^*\}}$.

Let \tilde{V} be the union of V and all the points in the support of D_Γ . Note that in both cases above, $v_0 \in \tilde{V}$. Let $\tilde{G} = (\tilde{V}, \tilde{E})$ be the corresponding model of Γ . We are going to recursively define an orientation π of \tilde{G} and the collection $\{D_v^*\}$ by handling at each step the orientation of all the edges incident to a vertex $v \in \tilde{V}$ and the minimal non-special divisor D_v^* in the case $v \in V$. We start by considering the vertex $v_0 \in \tilde{V}$. The orientations of all the edges incident to v_0 are defined so that $\text{deg}_\pi^+(v_0) = 0$; in other words, all these oriented edges are incoming at v_0 . Note that in the case $v_0 \notin V$, we have $D_\Gamma(v_0) \leq -1 = D_\Gamma^\pi(v_0)$. In the case $v_0 \in V$, since $r(D_{v_0}) < 0$ it follows from the Riemann-Roch theorem for C_{v_0} that there exists a minimal non-special divisor $D_{v_0}^*$ such that $D_{v_0} \leq D_{v_0}^*$.

Suppose that the orientation of all edges adjacent to vertices $v_0, \dots, v_i \in \tilde{V}$ has been defined, and that for all $v_j \in V$ with $j \leq i$, a minimal non-special divisor D_v^* on C_v has been given. Let S_i be a connected component of the induced subgraph $\tilde{G}[\tilde{V} \setminus \{v_0, \dots, v_i\}]$. Since \mathcal{D} is v_0 -reduced and S_i is a cut not containing v_0 , there exists a point v_{i+1} on the boundary of S_i which is non-saturated. (Note that v_{i+1} also lies in \tilde{V} .) This means that either:

- (1) $v_{i+1} \in \Gamma \setminus V$ and $\text{outdeg}_{S_i}(v_{i+1}) > D_\Gamma(v_{i+1})$; or
- (2) $v_{i+1} \in V$ and $D_v - \text{div}_v(\partial S_i)$ has negative rank.

All outgoing edges from S_i adjacent to v_{i+1} have already been oriented (and are outgoing from v_{i+1} by the definition of the orientation π). Orient all other edges incident to v_{i+1} in such a way that they are all incoming at v_{i+1} . Note that $\text{deg}_\pi^+(v_{i+1}) = \text{outdeg}_{S_i}(v_{i+1})$, and so in Case (1),

$$D_\Gamma(v_{i+1}) \leq \text{deg}_\pi^+(v_{i+1}) - 1.$$

In Case (2), since $r(D_v - \text{div}_v(\partial S_i)) = -1$, it follows from the Riemann-Roch theorem for C_v that there exists a minimal non-special divisor D_v^* such that

$$D_v - \text{div}_v(\partial S_i) \leq D_v^*.$$

This shows that $D_v \leq \text{div}_v(\partial S_i) + D_v^*$. By the definition of the orientation π we must have $D_{v+1}^{*\pi} = \text{div}_{v+1}(\partial S_i) + D_v^*$. Therefore, by the definition of the orientation π at v_{i+1} and the choice of $D_{v_{i+1}}^*$, we have $D_{v_{i+1}} \leq D_{v_{i+1}}^{*\pi}$.

Let π be the orientation of Γ just constructed, and let $\{D_v^*\}_{v \in V}$ be the collection of minimal non-special divisors on C_v defined above. By the definition of π , the only vertex of \tilde{V} with out-degree zero is v_0 and thus π belongs to $\mathcal{AO}_{v_0}(\Gamma)$. We clearly have $\mathcal{D} \leq \mathcal{D}^{\pi, \{D_v^*\}}$ as well, which completes the proof. \square

As a corollary of the above lemma, we obtain:

Corollary 3.12. *Any minimal non-special divisor \mathcal{D} on \mathfrak{C} is linearly equivalent to a moderator.*

Proof. Fix $v_0 \in \Gamma$. Let \mathcal{D} be a minimal non-special divisor on \mathfrak{C} and let \mathcal{D}^{v_0} be the v_0 -reduced divisor linearly equivalent to \mathcal{D} . By Lemma 3.11, there exists a moderator \mathcal{M} such that $\mathcal{D}^{v_0} \leq \mathcal{M}$. The effective divisor $\mathcal{M} - \mathcal{D}^{v_0}$ has degree zero, and so must be equal to zero. \square

3.3. Proof of Theorem 3.2. The proof of Theorem 3.2 can now be completed using (a slight simplification of) the idea behind the proof of Riemann-Roch theorem for graphs given in [BN]. The key is the following formula, which provides a useful description of the rank of \mathcal{D} in terms of minimal non-special divisors. We denote by $\text{deg}^+(\mathcal{D})$ (resp. $\text{deg}^-(\mathcal{D})$) the sum of non-negative (resp. non-positive) coefficients in \mathcal{D} . (Note that $\text{deg}^+(\mathcal{D}) + \text{deg}^-(\mathcal{D}) = \text{deg}(\mathcal{D})$).

Proposition 3.13. *For any divisor \mathcal{D} on \mathfrak{C} , we have*

$$r_{\mathfrak{C}}(\mathcal{D}) = \min_{\mathcal{N} \in \mathfrak{N}} \text{deg}^+(\mathcal{D} - \mathcal{N}) - 1.$$

Proof. Fix $\mathcal{N} \in \mathfrak{N}$, and let \mathcal{E} be the non-negative part of $\mathcal{D} - \mathcal{N}$, which is of degree $\text{deg}^+(\mathcal{D} - \mathcal{N})$. Then $r_{\mathfrak{C}}(\mathcal{D}) \leq \text{deg}^+(\mathcal{D} - \mathcal{N})$, since $\mathcal{D} - \mathcal{E} \leq \mathcal{N}$ and $r(\mathcal{N}) = -1$ imply that $|\mathcal{D} - \mathcal{E}| = \emptyset$. This shows that $r_{\mathfrak{C}}(\mathcal{D}) \leq \min_{\mathcal{N} \in \mathfrak{N}} \text{deg}^+(\mathcal{D} - \mathcal{N}) - 1$.

To prove the opposite inequality, let \mathcal{E} be an effective divisor of degree $r_{\mathfrak{C}}(\mathcal{D}) + 1$ such that $|\mathcal{D} - \mathcal{E}| = \emptyset$. By Lemma 3.11, there exists $\mathcal{N} \in \mathfrak{N}$ such that $\mathcal{D} - \mathcal{E} \leq \mathcal{N}$, or equivalently, $\mathcal{D} - \mathcal{N} \leq \mathcal{E}$. In particular, $\text{deg}^+(\mathcal{D} - \mathcal{N}) \leq \text{deg}(\mathcal{E}) = r_{\mathfrak{C}}(\mathcal{D}) + 1$, which proves the proposition. \square

To finish the proof of Theorem 3.2, note that

$$\text{deg}^+(\mathcal{D} - \mathcal{N}) = \text{deg}(\mathcal{D} - \mathcal{N}) - \text{deg}^-(\mathcal{D} - \mathcal{N}) = \text{deg}(\mathcal{D}) - g + 1 + \text{deg}^+(\mathcal{N} - \mathcal{D})$$

and observe that $\mathcal{N} - \mathcal{D} = \mathcal{K} - \mathcal{D} - (\mathcal{K} - \mathcal{N})$. If \mathcal{N} is linearly equivalent to the moderator \mathcal{M} , c.f. Corollary 3.12, then $\mathcal{K} - \mathcal{N}$ is linearly equivalent to the dual moderator $\bar{\mathcal{M}}$, and so belongs to \mathfrak{N} . Thus, by Proposition 3.13, we have

$$\begin{aligned} r_{\mathfrak{C}}(\mathcal{D}) &= \text{deg}(\mathcal{D}) - g + 1 + \min_{\mathcal{N} \in \mathfrak{N}} \text{deg}^+(\mathcal{K} - \mathcal{D} - (\mathcal{K} - \mathcal{N})) - 1 \\ &= \text{deg}(\mathcal{D}) - g + 1 + \min_{\mathcal{N}' \in \mathfrak{N}} \text{deg}^+(\mathcal{K} - \mathcal{D} - \mathcal{N}') - 1 = \text{deg}(\mathcal{D}) - g + 1 + r_{\mathfrak{C}}(\mathcal{K} - \mathcal{D}). \end{aligned}$$

4. THE SPECIALIZATION MAP AND SPECIALIZATION INEQUALITY

Let \mathbb{K} be a complete and algebraically closed non-Archimedean field with non-trivial absolute value $|\cdot|$, R the valuation ring of \mathbb{K} , and κ its (algebraically closed) residue field. Let $\mathcal{G} = \text{val}(\mathbb{K}^\times)$ be the value group of \mathbb{K} . Let X be a smooth, proper, connected curve over \mathbb{K} and let X^{an} be the Berkovich analytic space associated to X . (We assume the reader is familiar with the theory of Berkovich analytic curves, see e.g. [BPR, Section 5] which contains everything we need.)

Our first goal will be to define a metrized complex $\mathfrak{C}\mathfrak{X}$ associated to a strongly semistable R -model \mathfrak{X} for X . We then define a specialization homomorphism from $\text{Div}(X)$ to $\text{Div}(\mathfrak{C}\mathfrak{X})$ and prove a specialization inequality which refines Lemma 2.8 and Corollary 2.11 from [B]. Finally, we give some applications to specialization of canonical divisors, Brill-Noether theory, and Weierstrass points as in *loc. cit.*

4.1. The metrized complex associated to a semistable model. Recall that a connected reduced algebraic curve over κ is called *semistable* if all of its singularities are ordinary double points, and is called *strongly semistable* if in addition its irreducible components are all smooth. A (*strongly*) *semistable model* for X is a flat and integral proper relative curve \mathfrak{X} over R whose generic fiber is isomorphic to X and whose special fiber $\bar{\mathfrak{X}}$ is a (strongly) semistable curve.

Given a semistable model \mathfrak{X} for X , there is a canonical associated reduction map $\text{red} : X(\mathbb{K}) \rightarrow \bar{\mathfrak{X}}(\kappa)$ which is defined using the natural bijection between $X(\mathbb{K})$ and $\mathfrak{X}(R)$. This extends naturally to a map $\text{red} : X^{\text{an}} \rightarrow \bar{\mathfrak{X}}$.

We define a metrized complex associated to a strongly semistable model \mathfrak{X} for X as follows. Let G be the dual graph of $\bar{\mathfrak{X}}$, so that vertices of G correspond to irreducible components of $\bar{\mathfrak{X}}$ and edges of G correspond to intersections between irreducible components. If x^e is the ordinary double point of $\bar{\mathfrak{X}}$ corresponding to an edge e of G , the *formal fiber* $\text{red}^{-1}(x^e)$ is isomorphic to an open annulus \mathbf{A} . We define the *length of the edge* e to be the length of the skeleton of \mathbf{A} , i.e., the *modulus* $\log(b) - \log(a)$ of $\mathbf{A} \cong \{x \in (\mathbf{A}^1)^{\text{an}} \mid a < |T|_x < b\}$. (The modulus of an open annulus is well-defined independent of the choice of such an analytic isomorphism.) In this way, we have defined a metric graph $\Gamma = \Gamma_{\bar{\mathfrak{X}}}$ associated to $\bar{\mathfrak{X}}$ together with a model G . The irreducible components C_v of $\bar{\mathfrak{X}}$ correspond bijectively to the vertices v of G , and we let $\mathcal{A}_v \subset C_v$ be the finite set of double points of $\bar{\mathfrak{X}}$ contained in C_v , so that there is a natural bijection between \mathcal{A}_v and the edges of G incident to v . In this way we have defined a metrized complex $\mathfrak{C}\mathfrak{X}$ canonically associated to $\bar{\mathfrak{X}}$.

One can show that essentially every metrized complex of curves comes from this construction. The following result is proved in [ABBR1, Theorem 3.24].

Theorem 4.1. *Let \mathfrak{C} be a metrized complex of κ -curves whose edge lengths are contained in the value group of \mathbb{K} . There exists a smooth, proper, connected curve X over \mathbb{K} and a semistable model \mathfrak{X} for X such that $\mathfrak{C} \cong \mathfrak{C}\mathfrak{X}$.*

4.2. The metrized complex associated to a semistable vertex set. It is sometimes useful to define skeleta and metrized complexes in terms of semistable vertex sets rather than semistable models. (By [BPR, Theorem 5.38], there is a bijective correspondence between the latter two objects.) A *semistable vertex set* for X^{an} is a finite set V of type 2 points of X^{an} such that the complement of V in X^{an} is isomorphic (as a \mathbb{K} -analytic space) to the disjoint union of a finite number of open annuli and an infinite number of open

balls. (Such a disjoint union is called the *semistable decomposition* of X^{an} associated to V .) Any finite set of type 2 points of X^{an} is contained in a semistable vertex set ([BPR, Proposition 5.27]). The *skeleton* $\Gamma = \Sigma(X^{\text{an}}, V)$ of X^{an} with respect to a semistable vertex set V is the union (inside X^{an}) of V and the skeletons of each of the open annuli in the semistable decomposition associated to V . Using the canonical metric on the skeletons of these open annuli, Γ can be naturally viewed as a (finite) metric graph contained in X^{an} (see [BPR, Definition 5.17] for more details). The skeleton Γ comes equipped with a natural model G whose vertices are the points of V and whose edges correspond bijectively to the open annuli in the semistable decomposition associated to V . A semistable vertex set V is called *strongly semistable* if the graph G has no loop edges. Every semistable vertex set is contained in a strongly semistable vertex set.

Given a strongly semistable vertex set V for X^{an} , there is a canonical corresponding metrized complex $\mathfrak{C}V$ of κ -curves. Indeed, we have already defined a metric graph $\Gamma = \Sigma(X^{\text{an}}, V)$ corresponding to V , together with a corresponding model G . For $v \in V$, let C_v be the unique smooth projective curve over κ with function field $\widetilde{\mathcal{H}(v)}$. (For $x \in X^{\text{an}}$ of type 2, recall that the residue field $\widetilde{\mathcal{H}(x)}$ of the completed residue field $\mathcal{H}(x)$ of x has transcendence degree one over κ .) It remains to specify, for each $v \in V$, a bijection ψ_v from the edges of G incident to v to a subset \mathcal{A}_v of $C_v(\kappa)$. Given such an edge e , we define $\psi(e)$ to be the point of $C_v(\kappa)$ corresponding to the tangent direction at v defined by e . (Recall from [BPR, Paragraph 5.67] that if x is of type 2, there is a canonical bijection between T_x and the set of discrete valuations on $\widetilde{\mathcal{H}(x)} = \kappa(C_x)$ which are trivial on κ .)

Remark 4.2. Passing to a larger semistable vertex set is compatible with linear equivalence of divisors and does not change the rank of divisors. One can thus associate a canonical group $\text{Pic } \mathfrak{C}X^{\text{an}}$ to X^{an} , defined as $\text{Pic}(\mathfrak{C}V)$ for any semistable vertex set V , together with a canonical rank function $r : \text{Pic } \mathfrak{C}X^{\text{an}} \rightarrow \mathbb{Z}$.

One can define in a similar way semistable vertex sets and skeleta for an affine curve X' (see [BPR, Definition 5.19]). In this case, one must also allow a finite number of punctured open balls in the semistable decomposition and the skeleton is a topologically finite but not necessarily finite length metric graph – it will contain a finite number of infinite rays corresponding to the points of $X \setminus X'$, where X is the projective completion of X' .

4.3. Specialization of divisors from curves to metrized complexes. There is a canonical embedding of $\Gamma_{\mathfrak{X}}$ in the Berkovich analytification X^{an} of X , as well as a canonical retraction map $\tau : X^{\text{an}} \rightarrow \Gamma_{\mathfrak{X}}$. In addition, there is a canonical reduction map $\text{red} : X^{\text{an}} \rightarrow \tilde{\mathfrak{X}}$ sending $X(\mathbb{K})$ surjectively onto the closed points of $\tilde{\mathfrak{X}}$. The retraction map τ induces by linearity a specialization map $\tau_* : \text{Div}(X) \rightarrow \text{Div}(\Gamma_{\mathfrak{X}})$ which is studied in [B]. We can promote this to a map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ whose target is the larger group $\text{Div}(\mathfrak{C}\mathfrak{X})$ as follows. If $P \in X^{\text{an}}$ satisfies $\tau(P) = v \in V$ then either P is the unique point of X^{an} with $\text{red}(P)$ equal to the generic point of C_v , or else $\text{red}(P)$ is a nonsingular closed point of C_v . The map $\tau_*^{\mathfrak{C}\mathfrak{X}} : \text{Div}(X) \rightarrow \text{Div}(\mathfrak{C}\mathfrak{X})$ is obtained by linearly extending the map $\tau^{\mathfrak{C}\mathfrak{X}} : X(\mathbb{K}) \rightarrow \text{Div}(\mathfrak{C}\mathfrak{X})$ defined by

$$\tau^{\mathfrak{C}\mathfrak{X}}(P) = \begin{cases} \tau(P) & \tau(P) \notin V \\ \text{red}(P) & \tau(P) \in V. \end{cases}$$

4.4. Reduction of rational functions and specialization of principal divisors. In this section, we show that if \mathcal{D} is a principal divisor on X , then $\tau_*^{\mathfrak{C}\mathfrak{X}}(\mathcal{D})$ is a principal divisor on $\mathfrak{C}\mathfrak{X}$. In fact, we show a more precise result (Theorem 4.5 below).

Let $x \in X^{\text{an}}$ be a point of type 2. Given a nonzero rational function f on X , choose $c \in \mathbb{K}^\times$ such that $|f(x)| = |c|$. Define $f_x \in \kappa(C_x)^\times$ to be the image of $c^{-1}f$ in $\widetilde{\mathcal{H}(x)} \cong \kappa(C_x)$. Although f_x is only well-defined up to multiplication by an element of κ^\times , its divisor $\text{div}(f_x)$ is canonical and the resulting map $\text{Prin}(X) \rightarrow \text{Prin}(C_x)$ is a homomorphism. We call f_x the *normalized reduction* of f .

If H is a \mathbb{K} -linear subspace of $\mathbb{K}(X)$, the collection of all possible reductions of nonzero elements of H , together with $\{0\}$, forms a κ -vector space H_x . For later use we note the following elementary lemma, which says that $\dim H = \dim H_x$:

Lemma 4.3. *Let X be a smooth proper curve over \mathbb{K} , and $x \in X^{\text{an}}$ a point of type 2. The κ -vector space H_x defined by the reduction to $\widetilde{\mathcal{H}(x)}$ of an $(r+1)$ -dimensional \mathbb{K} -subspace $H \subset \mathbb{K}(X)$ has dimension $r+1$.*

Proof. The inequality $\dim_\kappa(H_x) \leq r+1$ follows from the observation that if f_0, \dots, f_n are linearly dependent in H , with $|f_i| = 1$ for all i , then the reductions $f_{0,x}, \dots, f_{n,x}$ are linearly dependent in H_x .

To prove $\dim_\kappa(H_x) \geq r+1$, choose a \mathbb{K} -basis f_0, \dots, f_r for H consisting of elements of norm 1 and let $f_{0,x}, \dots, f_{r,x}$ in H_x be the corresponding reductions. It is not true that $f_{0,x}, \dots, f_{r,x}$ necessarily form a basis of H_x : for example, if f and g are linearly independent functions of norm 1 and c is any scalar with $|c| < 1$, then f and $f+cg$ are linearly independent but they have the same reduction. However, by the following procedure one can construct another basis g_0, \dots, g_r of H with $|g_i| = 1$ for all i such that the reductions $g_{i,x}$ are linearly independent over κ . Define $g_0 = f_0$, and recursively define g_{i+1} as follows: among all the rational functions of the form $f_{i+1} - \sum_{j \leq i} a_j g_j$ with $|a_j| \leq 1$ for all j , choose one, call it h_{i+1} , whose norm is minimal. Note that $|f_{i+1} - \sum_{j \leq i} a_j g_j| \leq 1$ so such an element exists by compactness. Let c_{i+1} be a scalar with $|c_{i+1}| = |h_{i+1}|$ and define $g_{i+1} = c_{i+1}^{-1} h_{i+1}$. It is clear that the $\{g_i\}$ form a basis of H . We claim that the reductions $g_{i,x}$ form a basis of H_x . Indeed, if not then there exists a minimal index i and scalars b_j with $|b_j| \leq 1$ for $j = 0, \dots, i$ such that $|g_{i+1} - \sum_{j \leq i} b_j g_j| < 1$. But this implies that $|f_{i+1} - \sum_{j \leq i} (a_j + c_{i+1} b_j) g_j| < c_{i+1}$, contradicting the choice of h_{i+1} . \square

We recall the following useful result from [BPR, Theorem 5.69].

Theorem 4.4 (Slope Formula). *Let f be a nonzero rational function on X , let X' be an open affine subset of X on which f has no zeros or poles, and let $F = -\log |f| : (X')^{\text{an}} \rightarrow \mathbb{R}$. Let V be a semistable vertex set for X' , and let $\Sigma = \Sigma(X', V)$ be the corresponding skeleton. Then:*

- (1) $F = F \circ \tau_\Sigma$ where $\tau_\Sigma : (X')^{\text{an}} \rightarrow \Sigma$ is the retraction.
- (2) F is piecewise linear with integer slopes, and F is linear on each edge of Σ .
- (3) If x is a type-2 point of X^{an} and $\vec{v} \in T_x$ corresponds to the discrete valuation ord_v on $\kappa(C_x)$, then $d_{\vec{v}} F(x) = \text{ord}_v(f_x)$.
- (4) F is harmonic at all $x \in \mathbf{H}(X^{\text{an}})$.
- (5) Let $x \in X \setminus X'$, let e be the ray in Σ whose closure in X^{an} contains x , let $y \in V$ be the other endpoint of e , and let $\vec{v} \in T_y$ be the tangent direction represented by e . Then $d_{\vec{v}} F(y) = \text{ord}_x(f)$.

To each nonzero rational function f on X and each strongly semistable model \mathfrak{X} for X , one can associate a rational function f on $\mathfrak{C}\mathfrak{X}$ whose Γ -part f_Γ is the restriction to $\Gamma = \Gamma_{\mathfrak{X}}$

of the piecewise linear function $F = \log |f|$ on X^{an} and whose C_v -part is the normalized reduction f_v (which is well-defined up to multiplication by a non-zero scalar).²

As an application of Theorem 4.4, we obtain the following important formula:

Theorem 4.5. *For every nonzero rational function f on X ,*

$$\tau_*^{\mathfrak{C}\mathfrak{X}}(\text{div}(f)) = \text{div}(f).$$

Proof. Let D be the support of $\text{div}(f)$ in X . We apply Theorem 4.4 to $X' = X \setminus D$, the invertible rational function f on X' , and a semistable vertex set V' for X' which contains V . Let $\Sigma = \Sigma(X'^{\text{an}}, V')$ and $\Gamma = \Sigma(X, V)$. Since V is a semistable vertex set for X , the closure of $\Sigma \setminus \Gamma$ in Σ is a disjoint union of metric trees $\mathcal{T}_1, \dots, \mathcal{T}_s$. For each tree \mathcal{T}_i , denote by $x_1^i, \dots, x_{n_i}^i$ all the points in $D \subset X(\mathbb{K})$ which are in the closure of rays in \mathcal{T}_i , and let $\overline{\mathcal{T}}_i = \mathcal{T}_i \cup \{x_1^i, \dots, x_{n_i}^i\}$. Denote by y_i the unique point of $\Gamma \cap \overline{\mathcal{T}}_i$. In addition, if $y_i \in V$, let $z_j^i = \tau_*^{\mathfrak{C}\mathfrak{X}}(x_j^i) \in C_{y_i}(\kappa)$.

For each $i \in \{1, \dots, s\}$, the restriction of $F = \log |f|$ to \mathcal{T}_i is harmonic at every point of $\overline{\mathcal{T}}_i \setminus \{y_i, x_1^i, \dots, x_{n_i}^i\}$ by (4). By (5), $\log |f|$ is linear in a sufficiently small neighborhood of x_j^i in \mathcal{T}_i , and has (constant) slope $\text{ord}_{x_j^i}(f)$ along the unique tangent direction in this interval toward the point y_i . We infer that the order of $F|_{\mathcal{T}_i}$ at y_i is equal to $-\sum_{1 \leq j \leq n_i} \text{ord}_{x_j^i}(f)$. By (1), F is constant on any connected component of $X^{\text{an}} \setminus \Sigma$. By (4), the restriction f_Γ of F to Γ satisfies $\text{ord}_{y_i}(f_\Gamma) = -\text{ord}_{y_i}(F|_{\mathcal{T}_i}) = \sum_{1 \leq j \leq n_i} \text{ord}_{x_j^i}(f)$. Since f_Γ is linear on each edge of Γ by (2), it follows that $\tau_*(\text{div}(f)) = \text{div}(f_\Gamma)$.

Fix $i \in \{1, \dots, s\}$ and suppose that $y_i \in V$. For any tangent direction $\vec{\nu} \in T_{y_i}$, by (3) we have $d_{\vec{\nu}}F = -\text{ord}_{\vec{\nu}}(f_{y_i})$. Since F is locally constant on $X^{\text{an}} \setminus \Sigma$, we infer that $d_{\vec{\nu}}F = 0$ for any $\vec{\nu} \in T_x$ not corresponding to a direction in Σ . This shows that $\text{div}(f_{y_i})$ is supported on $\{z_1^i, \dots, z_{n_i}^i\} \cup \mathcal{A}_{y_i}$. By (3), for any point ν in \mathcal{A}_{y_i} , corresponding to a tangent direction $\vec{\nu}$, we have $d_{\vec{\nu}}f_\Gamma + \text{ord}_{\vec{\nu}}(f_{y_i}) = 0$. This shows that the coefficient of ν in $\text{div}(f)$ is zero. For a point $\nu \in \{z_1^i, \dots, z_{n_i}^i\}$, let $\overline{\mathcal{T}}_{i,\nu}$ be the closure in Σ of the connected component of $\Sigma \setminus y_i$ which contains all the points x_j^i with $z_j^i = \nu$. By (1) and (4), F is harmonic at any point of $\overline{\mathcal{T}}_{i,\nu} \setminus (D \cup \{y_i\})$, and is linear of slope $\text{ord}_{x_j^i}(f)$ in a sufficiently small neighborhood of x_j^i in $\overline{\mathcal{T}}_{i,\nu}$ along the unique tangent direction in this interval toward the point y_i . By (3), we infer that $\text{ord}_{\nu}(f_{y_i}) = -d_{\vec{\nu}}(F) = \sum_{x_j^i: z_j^i = \nu} \text{ord}_{x_j^i}(f)$. We conclude that $\tau_*^{\mathfrak{C}\mathfrak{X}}\text{div}(f) = \text{div}(f)$. \square

In particular, it follows that $\tau_*^{\mathfrak{C}\mathfrak{X}}(\text{Prin}(X)) \subseteq \text{Prin}(\mathfrak{C}\mathfrak{X})$.

4.5. The specialization inequality. Fix a strongly semistable R -model \mathfrak{X} for X and let $\mathfrak{C}\mathfrak{X}$ be the metrized complex associated to \mathfrak{X} . The following is a variant of Lemma 2.8 and Corollary 2.11 from [B]:

Theorem 4.6 (Specialization Inequality). *For every divisor $D \in \text{Div}(X)$,*

$$r_X(D) \leq r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)).$$

²We take $F = \log |f|$ instead of $-\log |f|$ because our convention is that $\text{ord}_u(F)$ is the sum of the outgoing slopes of F at u . One could equally well take the opposite convention, defining $\text{ord}_u(F)$ to be *minus* the sum of the outgoing slopes of F at u , and then defining F to be $-\log |f|$. Such a modification would also necessitate a change of sign in our definition of $\text{div}_v(F)$.

Proof. Let $\mathcal{D} := \tau_*^{\mathfrak{C}\mathfrak{X}}(D)$. It suffices to prove that if $r_X(D) \geq k$, then $r_{\mathfrak{C}\mathfrak{X}}(\mathcal{D}) \geq k$ as well. The base case $k = -1$ is obvious. Now suppose $k = 0$, so that $r_X(D) \geq 0$. Then there exists an effective divisor $E \in \text{Div}(X)$ with $D - E \in \text{Prin}(X)$. Since $\tau_*^{\mathfrak{C}\mathfrak{X}}$ is a homomorphism and takes principal (resp. effective) divisors on X to principal (resp. effective) divisors on $\mathfrak{C}\mathfrak{X}$, we have $\mathcal{D} = \tau_*^{\mathfrak{C}\mathfrak{X}}(D) \sim \tau_*^{\mathfrak{C}\mathfrak{X}}(E) \geq 0$, so that $r_{\mathfrak{C}\mathfrak{X}}(\mathcal{D}) \geq 0$ as well.

We may therefore assume that $k \geq 1$. For each $v \in V$, let $\mathcal{R}_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ be a subset of size $g_v + 1$ and denote by \mathcal{R} the union of the \mathcal{R}_v . By Theorem A.1, \mathcal{R} is a rank-determining set in \mathfrak{C} . Let \mathcal{E} be an effective divisor of degree k with support in \mathcal{R} . For each $x \in \text{supp}(\mathcal{E})$, there exists a point $P \in X(\mathbb{K})$ whose reduction is x , so there is an effective divisor E of degree k on X with $\tau_*^{\mathfrak{C}\mathfrak{X}}(E) = \mathcal{E}$. By our assumption on the rank of D , $r_X(D - E) \geq 0$, and so $r_{\mathfrak{C}\mathfrak{X}}(\mathcal{D} - \mathcal{E}) \geq 0$. Since this is true for any effective divisor of degree k with support in \mathcal{R} , and since \mathcal{R} is rank-determining, we infer that $r_{\mathfrak{C}\mathfrak{X}}(\mathcal{D}) \geq k$ as desired. \square

Remark 4.7. There are many examples where the inequality in Theorem 4.6 can be strict. For example, if $g_v = 0$ for all $v \in V$ then $r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) = r_{\Gamma}(\tau_*(D))$ and thus the examples of strict inequality from [B] apply. At the other extreme, if $C = \bar{X}$ is smooth (so that G is a point) then the specialization inequality becomes the well-known semicontinuity statement $h^0(D) \leq h^0(\bar{D})$, where $\bar{D} \in \text{Div}(C)$ is the reduction of D , and it is clear that such an inequality can be strict; for example, take $D = (P) - (Q)$ where X has genus at least 1 and $P, Q \in X(\mathbb{K})$ are distinct points with the same reduction in $C(\kappa)$.

4.6. Refined versions of weighted specialization lemma. In this section we use the specialization lemma for divisors on metrized complexes to provide a framework for obtaining stronger versions of the weighted specialization lemma from [AC]. Our key technical tool for this will be the “ η -function” on a metrized complex. To help motivate the definition, we first study an analogous η -function which controls the rank of divisors on connected sums.

4.6.1. The η -function and the rank of divisors on connected sums. Suppose the underlying graph of a metrized complex \mathfrak{C} has a bridge edge $e = \{v_1, v_2\}$, and denote by \mathfrak{C}_1 and \mathfrak{C}_2 the two metrized complexes obtained by removing the edge e from \mathfrak{C} (v_i is a vertex of \mathfrak{C}_i). Denote by x_1 and x_2 the points of \mathfrak{C}_1 and \mathfrak{C}_2 , respectively, corresponding to the edge e . We say that \mathfrak{C} is a *connected sum* of \mathfrak{C}_1 and \mathfrak{C}_2 .

There is an addition map $\text{Div}(\mathfrak{C}_1) \oplus \text{Div}(\mathfrak{C}_2) \rightarrow \text{Div}(\mathfrak{C})$ which associates to any pair of divisors $(\mathcal{D}_1, \mathcal{D}_2) \in \text{Div}(\mathfrak{C}_1) \oplus \text{Div}(\mathfrak{C}_2)$ the divisor $\mathcal{D}_1 + \mathcal{D}_2$ on \mathfrak{C} whose restriction to each \mathfrak{C}_i is \mathcal{D}_i .

The following proposition provides a precise description of the \mathfrak{C} -rank of $\mathcal{D}_1 + \mathcal{D}_2$ in terms of the two rank functions $r_{\mathfrak{C}_1}(\cdot)$ and $r_{\mathfrak{C}_2}(\cdot)$. In order to state the result, we introduce a function $\eta = \eta_{x_2, \mathcal{D}_2} : \mathbb{N} \cup \{0\} \rightarrow \mathbb{Z}$ defined by the condition that $\eta(k)$ is the smallest integer n such that $r_{\mathfrak{C}_2}(\mathcal{D}_2 + n(x_2)) = k$. (By the Riemann-Roch theorem for \mathfrak{C}_2 , η is well-defined.)

Proposition 4.8. *With notation as above, we have*

$$(2) \quad r_{\mathfrak{C}}(\mathcal{D}) = \min_{k \in \mathbb{N} \cup \{0\}} \left\{ k + r_{\mathfrak{C}_1}(\mathcal{D}_1 - \eta(k)(x_1)) \right\}.$$

Proof. We prove the equality of the two sides of Equation (2) by showing that the inequalities \leq and \geq both hold. By Theorem A.1, we can assume that all the effective divisors we consider below are supported on geometric points (i.e., on points of the curves C_v).

(\geq) Let r be the right-hand term in (2) and let \mathcal{E} be an effective divisor of degree r on \mathfrak{C} supported on geometric points. We need to show the existence of a rational function f on

\mathfrak{C} such that $\operatorname{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} \geq 0$. There is a (unique) decomposition of \mathcal{E} as $\mathcal{E} = \mathcal{E}_1 + \mathcal{E}_2$ with \mathcal{E}_1 and \mathcal{E}_2 effective divisors of degree $r - k$ and k , respectively, having support on geometric points. By the definition of η , there exists a rational function \mathfrak{f}_2 on \mathfrak{C}_2 such that $\operatorname{div}(\mathfrak{f}_2) + \eta(k)(x_2) + \mathcal{D}_2 - \mathcal{E}_2 \geq 0$. By the choice of r as the minimum in Equation (2), we have $r_{\mathfrak{C}_1}(\mathcal{D}_1 - \eta(k)(x_1)) \geq r - k$, so there exists a rational function \mathfrak{f}_1 on \mathfrak{C}_1 such that $\operatorname{div}(\mathfrak{f}_1) + \mathcal{D}_1 - \eta(k)(x_1) - \mathcal{E}_1 \geq 0$. We can in addition suppose that the Γ_1 -part of \mathfrak{f}_1 takes value zero at v_1 . Up to adding a constant to the Γ_1 -part of \mathfrak{f}_1 , we can define a rational function \mathfrak{f} which restricts to \mathfrak{f}_1 and \mathfrak{f}_2 on \mathfrak{C}_1 and \mathfrak{C}_2 , respectively, and which is linear of slope $-\eta(k)$ on the oriented edge (v_1, v_2) in Γ . For this rational function, we clearly have

$$\begin{aligned} \operatorname{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} &= \operatorname{div}(\mathfrak{f}_1) + \mathcal{D}_1 - \eta(k)(x_1) - \mathcal{E}_1 \\ &\quad + \operatorname{div}(\mathfrak{f}_2) + \mathcal{D}_2 + \eta(k)(x_2) - \mathcal{E}_2 \geq 0. \end{aligned}$$

(\leq) It will be enough to show that for each $k \in \mathbb{N} \cup \{0\}$, the inequality

$$r_{\mathfrak{C}}(\mathcal{D}) \leq k + r_{\mathfrak{C}_1}(\mathcal{D}_1 - \eta(k)(x_1))$$

holds. Let $\tilde{r} = r_{\mathfrak{C}}(\mathcal{D})$. We can obviously restrict to the case $k \leq \tilde{r}$. For each effective divisor \mathcal{E}_2 of degree k supported on geometric points of \mathfrak{C}_2 , let $\mathfrak{f}_{\mathcal{E}_2}$ be a rational function on \mathfrak{C}_2 with the property that the coefficients of $\operatorname{div}(\mathfrak{f}_{\mathcal{E}_2}) + \mathcal{D}_2 - \mathcal{E}_2$ outside x_2 are all non-negative, and such that in addition, the coefficient of x_2 in $\operatorname{div}(\mathfrak{f}_{\mathcal{E}_2}) + \mathcal{D}_2 - \mathcal{E}_2$ is maximized among all the rational functions with this property. Denote by $n_{\mathcal{E}_2}$ the coefficient of x_2 in $\operatorname{div}(\mathfrak{f}_{\mathcal{E}_2}) + \mathcal{D}_2 - \mathcal{E}_2$ and let $\bar{\mathcal{E}}_2$ be an effective divisor of degree k on \mathfrak{C}_2 , supported on geometric points of \mathfrak{C}_2 , for which $n := n_{\bar{\mathcal{E}}_2}$ is minimal, i.e.,

$$n_{\bar{\mathcal{E}}_2} = \min\{n_{\mathcal{E}_2} \mid \mathcal{E}_2 \geq 0 : \deg(\mathcal{E}_2) = k\}.$$

Note that by the choice of n , for an effective divisor \mathcal{E}_2 of degree k supported on geometric points of \mathfrak{C}_2 and the rational function $\mathfrak{f}_{\mathcal{E}_2}$, we have $\operatorname{div}(\mathfrak{f}_{\mathcal{E}_2}) + \mathcal{D}_2 - \mathcal{E}_2 - n(x_2) \geq 0$, which shows that $r_{\mathfrak{C}_2}(\mathcal{D}_2 - n(x_2)) \geq k$. By the definition of η , it follows that $\eta(k) \leq -n$.

By the definition of $\mathfrak{f}_{\bar{\mathcal{E}}_2}$, for any rational function \mathfrak{f}_2 on \mathfrak{C}_2 with the property that $\operatorname{div}(\mathfrak{f}_2) + \mathcal{D}_2 - \bar{\mathcal{E}}_2$ has non-negative coefficients outside x_2 , the coefficient of x_2 in $\operatorname{div}(\mathfrak{f}_2) + \mathcal{D}_2 - \bar{\mathcal{E}}_2$ is at most $n \leq -\eta(k)$.

Let \mathcal{E}_1 be an effective divisor of degree $\tilde{r} - k$ on \mathfrak{C}_1 supported on geometric points. There exists a rational function \mathfrak{f} on \mathfrak{C} such that $\operatorname{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E}_1 - \bar{\mathcal{E}}_2 \geq 0$. Let \mathfrak{f}_1 and \mathfrak{f}_2 be the restrictions of \mathfrak{f} to \mathfrak{C}_1 and \mathfrak{C}_2 , respectively. By the preceding paragraph, the coefficient of x_2 in $\operatorname{div}(\mathfrak{f}_2) + \mathcal{D}_2 - \bar{\mathcal{E}}_2$ is at most $-\eta(k)$. Since the coefficient of x_2 in $\operatorname{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E}_1 - \bar{\mathcal{E}}_2$ is non-negative, the Γ -part of \mathfrak{f} has slope at least $\eta(k)$ along the tangent direction emanating from v_2 which corresponds to x_2 . Since the coefficients of $\operatorname{div}(\mathfrak{f})$ at the interior points of the edge $\{v_1, v_2\}$ are all non-negative, it follows that the slope of the Γ -part of \mathfrak{f} along the tangent direction emanating from v_1 on the edge $\{v_1, v_2\}$ is at most $-\eta(k)$, which by the assumption $\operatorname{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E}_1 - \bar{\mathcal{E}}_2 \geq 0$ implies that the coefficient of x_1 in $\operatorname{div}(\mathfrak{f}_1) + \mathcal{D}_1 - \mathcal{E}_1$ must be at least $\eta(k)$. It follows that $\operatorname{div}(\mathfrak{f}_1) + \mathcal{D}_1 - \eta(k)(x_1) - \mathcal{E}_1 \geq 0$. Since \mathcal{E}_1 was an arbitrary effective divisor of degree $\tilde{r} - k$ supported on geometric points of \mathfrak{C}_1 , it follows that $r_{\mathfrak{C}_1}(\mathcal{D}_1 - \eta(k)(x_1)) \geq \tilde{r} - k$, and the claim follows. \square

Remark 4.9. In the same spirit, consider two metric graphs Γ_1 and Γ_2 , and suppose that two distinguished points $v_1 \in \Gamma_1$ and $v_2 \in \Gamma_2$ are given. The *wedge* or *direct sum* of (Γ_1, v_1) and (Γ_2, v_2) , denoted $\tilde{\Gamma} = \Gamma_1 \vee \Gamma_2$, is the metric graph obtained by identifying the points v_1 and v_2 in the disjoint union of Γ_1 and Γ_2 . Denote by $v \in \tilde{\Gamma}$ the image of v_1 and v_2 in $\tilde{\Gamma}$. (By

abuse of notation, we will use v to denote both v_1 in Γ_1 and v_2 in Γ_2 .) We refer to $v \in \tilde{\Gamma}$ as a *cut-vertex* and to $\tilde{\Gamma} = \Gamma_1 \vee \Gamma_2$ as the decomposition corresponding to the cut-vertex v . Note that there is an addition map $\text{Div}(\Gamma_1) \oplus \text{Div}(\Gamma_2) \rightarrow \text{Div}(\tilde{\Gamma})$ which associates to any pair of divisors D_1 and D_2 in $\text{Div}(\Gamma_1)$ and $\text{Div}(\Gamma_2)$ the divisor, $D_1 + D_2$ on $\tilde{\Gamma}$ defined by pointwise addition of coefficients in D_1 and D_2 . Let $r_1(\cdot)$, $r_2(\cdot)$, and $\tilde{r}(\cdot) = r_{\Gamma_1 \vee \Gamma_2}(\cdot)$ be the rank functions in Γ_1, Γ_2 , and $\tilde{\Gamma}$, respectively. For any non-negative integer k , denote by $\eta_{v, D_2}(k)$, or simply $\eta(k)$, the smallest integer n such that $r_2(D_2 + n(v)) = k$. Then for any divisor D_1 in $\text{Div}(\Gamma_1)$, we have (by an entirely similar argument)

$$(3) \quad \tilde{r}(D_1 + D_2) = \min_{k \in \mathbb{N} \cup \{0\}} \left\{ k + r_1(D_1 - \eta(k)(v)) \right\}.$$

Remark 4.10. The above formalism is very handy for studying metric graphs such as the chain of cycles used in [CDPR]. In particular, the lingering lattice paths studied in [CDPR] in connection with a tropical proof of the Brill-Noether theorem can be understood in a natural way as values of η -functions. The η -function formalism is also useful for studying limit linear series, see for example the proof of Theorem 5.4 below.

We now define a partial analogue of the η -function introduced above for a general metrized complex \mathfrak{C} and provide an analogue of Proposition 4.8.

Consider a divisor \mathcal{D} on \mathfrak{C} with C_v -part D_v for $v \in V$. For every vertex v of G , define the function $\eta_v : \mathbb{N} \cup \{0\} \rightarrow \mathbb{N} \cup \{0\}$ as follows. For any $k \geq 0$, $\eta_v(k)$ is the smallest integer $n \geq 0$ such that there exists a divisor \tilde{D}_v of degree $n - \deg(D_v)$ supported on the points of $\mathcal{A}_v \subset C_v(\kappa)$ such that $r_{C_v}(D_v + \tilde{D}_v) = k$. For any effective divisor $E = \sum_{v \in V(G)} E(v)(v)$ on G , define

$$\eta(E) := \sum_{v \in V(G)} \eta_v(E(v))(v).$$

Note that the definition of η depends on the data of the divisors $D_v \in \text{Div}(C_v)$.

Proposition 4.11. *Let \mathfrak{C} be a metrized complex of algebraic curves over κ . For any divisor \mathcal{D} on \mathfrak{C} with Γ -part D_Γ , we have*

$$r_{\mathfrak{C}}(\mathcal{D}) \leq \min_{E \geq 0} \left(\deg(E) + r_\Gamma(D_\Gamma - \eta(E)) \right),$$

where the minimum is over all effective divisors E on G .

Proof. Fix an effective divisor E supported on the vertices of G . We need to prove that

$$r_{\mathfrak{C}}(\mathcal{D}) \leq \deg(E) + r_\Gamma(D_\Gamma - \eta(E)).$$

For the sake of contradiction, suppose this is not the case, i.e., that $\deg(E) + r_\Gamma(D_\Gamma - \eta(E)) < r_{\mathfrak{C}}(\mathcal{D})$. Since $r_{\mathfrak{C}}(\mathcal{D}) - \deg(E) > r_\Gamma(D_\Gamma - \eta(E))$, Luo's Theorem (which is a special case of Theorem A.1) implies that there exists a divisor E^0 of degree $r_{\mathfrak{C}}(\mathcal{D}) - \deg(E)$ supported on $\Gamma \setminus V(G)$ such that $|D_\Gamma - E_\Gamma^0 - \eta(E)| = \emptyset$.

Consider all the effective divisors on \mathfrak{C} of the form $E^0 + \sum_{v \in V} E_v$, where each E_v is an effective divisor of degree $E(v)$ on C_v . Note that $E^0 + \sum_{v \in V} E_v$ has degree $r_{\mathfrak{C}}(\mathcal{D})$, so that by the definition of the rank function, for any choice of effective divisors E_v on C_v there exists a rational function f on \mathfrak{C} such that $\text{div}(f) + \mathcal{D} - E^0 - \sum_v E_v \geq 0$. Denote by f_Γ and f_v the different parts of f , so in particular we have $\text{div}(f_\Gamma) + D_\Gamma - E^0 - E \geq 0$. One can assume as in the proof of Proposition 2.1 that the slope of f_Γ along the edges incident to each vertex

of G is fixed. This shows the existence of the divisor $\tilde{D}_v = \text{div}_v(f_\Gamma)$ supported on the points $\mathcal{A}_v \subset C_v(\kappa)$ such that $D_v + \tilde{D}_v$ has rank at least $E(v) = \deg(E_v)$ on C_v . By the definition of η_v , \tilde{D}_v has degree at least $\eta_v(E(v)) - \deg(D_v)$. Since the degree of \tilde{D}_v is the sum of the slopes of f_Γ along the edges adjacent to v , which is by definition $\text{ord}_v(f_\Gamma)$, for each $v \in V(G)$ we obtain $\text{ord}_v(f_\Gamma) \geq \eta_v(E(v)) - \deg(D_v)$, i.e., $\text{ord}_v(f_\Gamma) + D_\Gamma(v) - \eta_v(E(v)) \geq 0$. In addition, for each non-vertex point $u \in \Gamma$ we must have $\text{ord}_u(f_\Gamma) + D_\Gamma(u) - E^0(u) \geq 0$. This implies that $|D_\Gamma - E^0 - \eta(E)| \neq \emptyset$, a contradiction. \square

4.6.2. Weighted specialization lemma. Let \mathbb{K} be a complete and algebraically closed non-Archimedean field with non-trivial absolute value $|\cdot|$, R the valuation ring of \mathbb{K} , and κ its residue field. Let $\mathcal{G} = \text{val}(\mathbb{K}^\times)$ be the value group of \mathbb{K} . Let X be a smooth, proper, and connected \mathbb{K} -curve and fix a semistable R -model \mathfrak{X} . Let (Γ, ω) denote the skeleton of \mathfrak{X} together with the weight function $\omega : V \rightarrow \mathbb{N}$ defined by $\omega(v) = g_v$. Let $\mathcal{W} = \sum_{v \in V} g_v(v)$. Following [AC], we define a metric graph $\Gamma^\#$ by attaching g_v loops of arbitrary positive length at each point $v \in \Gamma$ (see [AC] for a discussion of the intuition behind this construction.)

Any divisor D in Γ defines a divisor on $\Gamma^\#$ with the same support and coefficients (by viewing Γ as a subgraph of $\Gamma^\#$). By an abuse of notation, we will also denote this divisor by D . However, we distinguish the ranks of D in Γ and $\Gamma^\#$ by using the notation $r_\Gamma(D)$ for the former and $r_{\Gamma^\#}(D)$ for the latter. Following [AC], the *weighted rank* $r^\#$ of D is defined by $r^\#(D) := r_{\Gamma^\#}(D)$. Applying the formalism of the η -function from the previous subsection, we obtain the following useful formula for $r^\#$:

Corollary 4.12. *For every $D \in \text{Div}(\Gamma)$, we have*

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

Proof. Proceeding by induction, we are reduced to the case $\deg(\mathcal{W}) = 1$. This case follows from (3) applied to $\Gamma_1 = \Gamma$, Γ_2 a circle, $\tilde{\Gamma} = \Gamma_1 \vee \Gamma_2$, $D_1 = D$, and $D_2 = 0$, since in this case one has $\eta(0) = 0$, $\eta(1) = 2$, and $\eta(k) = k + 1$ for $k > 1$. \square

The following result is a generalization to metric graphs Γ , and to smooth curves X over not necessarily discrete non-Archimedean fields, of the corresponding statement for graphs which was obtained by different methods in [AC].

Theorem 4.13 (Weighted specialization inequality). *For every divisor $D \in \text{Div}(X)$, we have $r_X(D) \leq r^\#(\tau_*(D))$.*

Proof. Let $\mathcal{D} = \tau_*^{\mathfrak{C}\mathfrak{X}}(D)$ and note that the Γ -part D_Γ of \mathcal{D} is equal to $\tau_*(D)$. Applying Proposition 4.11, we obtain

$$(4) \quad r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) \leq \min_{E \geq 0} \left(\deg(E) + r_\Gamma(\tau_*(D) - \eta(E)) \right).$$

For each vertex v of G , define the function $\bar{\eta}_v : \mathbb{N} \cup \{0\} \rightarrow \mathbb{N} \cup \{0\}$ by $\bar{\eta}_v(k) = 2k$ for $k \leq g_v$ and $\bar{\eta}_v(k) = k + g_v$ for $k \geq g_v$. If E is an effective divisor supported on vertices of G , define $\bar{\eta}(E) = \sum_v \bar{\eta}(E(v))(v)$.

By the definition of η_v and Clifford's theorem for the curve C_v , for any integer $k \geq 0$ we have $\eta_v(k) \geq \bar{\eta}_v(k)$. Combining this with (4), we get the inequality

$$(5) \quad r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) \leq \min_{E \geq 0} \left(\deg(E) + r_\Gamma(\tau_*(D) - \bar{\eta}(E)) \right).$$

Note that for $E(v) > g_v$, we have $\bar{\eta}_v(E(v)) = g_v + E(v)$ and so, whatever the other coefficients of E are, we have

$$\begin{aligned} \deg(E) + r_\Gamma(\tau_*(D) - \bar{\eta}(E)) &= \deg(E - (v)) + 1 + r_\Gamma(\tau_*(D) - \bar{\eta}(E - (v)) - (v)) \\ &\geq \deg(E - (v)) + r_\Gamma(\tau_*(D) - \bar{\eta}(E - (v))). \end{aligned}$$

In other words, in taking the minimum in (5), one can restrict to effective divisors E with $E(v) \leq g_v$ for all v , i.e., $0 \leq E \leq \mathcal{W}$. Finally, applying (5), Theorem 4.6, and Corollary 4.12, we obtain the inequality $r_X(D) \leq r^\#(\tau_*(D))$. \square

4.7. Some applications. We give some direct applications of the results of this section.

4.7.1. Specialization of canonical divisors. Let K_X be a canonical divisor on the curve X/\mathbb{K} . By Theorem 4.6, we have $r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X)) \geq r(K_X) = g(X) - 1 = g(\mathfrak{C}\mathfrak{X}) - 1$. By the Riemann-Roch theorem for metrized complexes (Theorem 3.2), the divisor $\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X) - \mathcal{K}$ has degree zero and non-negative rank, and thus $\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X) \sim \mathcal{K}$. In particular, we have $\tau_*(K_X) \sim K^\#$.

Remark 4.14. As mentioned in the Introduction, for discretely valued R the fact that $\tau_*(K_X) \sim K^\#$ follows from the adjunction formula for arithmetic surfaces. But it does not seem straightforward to generalize that argument to the not necessarily discretely valued case.

4.7.2. Brill-Noether theory. The following result is an immediate consequence of Theorem 4.1 combined with Theorem 4.6.

Corollary 4.15 (Brill-Noether existence theorem). *Let κ be a field and let g, r, d be nonnegative integers. If the Brill-Noether number $\rho_d^r(g) := g - (r+1)(g-d+r)$ is nonnegative, then for every metrized complex \mathfrak{C} of κ -curves with $g(\mathfrak{C}) = g$, there exists a divisor $\mathcal{D} \in \text{Div}_+^d(\mathfrak{C})$ such that $r_{\mathfrak{C}}(\mathcal{D}) \geq r$.*

Remark 4.16. In particular, combining the original Specialization Lemma from [B] with Theorem 4.1 yields a new proof of [B, Theorem 3.20], which says that if $\rho_d^r(g) \geq 0$ then for every metrized graph of genus g there exists $D \in \text{Div}_+^d(\Gamma)$ such that $r(D) \geq r$. The present proof is more direct in that one does not need the approximation arguments given in Lemma 3.17 and Corollary 3.18 of *loc. cit.*

Remark 4.17. The notion of Brill-Noether rank on metric graphs introduced in [LPP] and the corresponding specialization lemma extend to the context of metrized complexes as well. Let \mathfrak{C} be a metrized complex of κ -curves, and let $W_d^r(\mathfrak{C}) \subset \text{Pic}^d(\mathfrak{C})$ be the subset of all divisors of degree d and rank at least r . Define $w_d^r(\mathfrak{C})$ to be the largest integer n such that for every effective divisor \mathcal{E} of degree $r+n$ on \mathfrak{C} , there exists $\mathcal{D} \in W_d^r(\mathfrak{C})$ such that $\mathcal{D} - \mathcal{E} \geq 0$. If $W_d^r(\mathfrak{C}) = \emptyset$, define $w_d^r(\mathfrak{C}) = -1$. If \mathfrak{X} is a strongly semistable R -model of a proper smooth curve X/\mathbb{K} such that $\mathfrak{C}\mathfrak{X} = \mathfrak{C}$, and $W_d^r(X) \subset \text{Pic}^d(X)$ denotes the Brill-Noether locus of X , then we have the inequality $w_d^r(\mathfrak{C}) \geq \dim W_d^r(X)$.

4.7.3. Weierstrass points on metrized complexes of κ -curves. Let \mathfrak{C} be a metrized complex of κ -curves. We say that a point x in \mathfrak{C} is a *Weierstrass point* if $r_{\mathfrak{C}}(g(\mathfrak{C})(x)) \geq 1$. By Theorem 4.6, the specialization of a Weierstrass point on X is a Weierstrass point on \mathfrak{C} . And by Theorem 4.1 combined with Theorem 4.6, any metrized complex of κ -curves of genus at least two has at least one Weierstrass point.

5. LIMIT LINEAR SERIES

Our aim in this section is to compare our divisor theory on metrized complexes with the Eisenbud-Harris theory of limit linear series for curves of compact type [EH]. As we will see, our comparison results lead naturally to a generalization of the notion of limit linear series to nodal curves which are not necessarily of compact type. Note that since Eisenbud and Harris work over a discrete valuation ring, in order to compare our theory to theirs we will sometimes restrict to this setting; we note, however, that most of our results hold over the valuation ring of an arbitrary non-trivially valued non-Archimedean field.

5.1. Degeneration of linear series. We start by reviewing some basic facts and definitions concerning the degeneration of linear series in families. Let Y be a smooth projective curve over κ . Recall that a linear series L of (projective) dimension r and degree d , or simply a \mathfrak{g}_d^r , over Y consists of a pair (\mathcal{L}, H) with \mathcal{L} an invertible sheaf of degree d on Y and H a subspace of $H^0(Y, \mathcal{L})$ of κ -dimension $r + 1$.

Let $\phi : \mathcal{X} \rightarrow B = \text{Spec } R$ be a regular smoothing of a strongly semistable curve X_0 over a discrete valuation ring R with fraction field K . This means that ϕ is proper, \mathcal{X} is regular, the generic fiber \mathcal{X}_η of ϕ is smooth, and the special fiber of ϕ is X_0 . Denote by $G = (V, E)$ the dual graph of X_0 and let $\{X_v\}_{v \in V(G)}$ be the set of irreducible components of X_0 .

Proposition 5.1. *Any line bundle \mathcal{L}_η on \mathcal{X}_η extends to \mathcal{X} . The restriction of any two extensions \mathcal{L}_1 and \mathcal{L}_2 of \mathcal{L}_η to the special fiber X_0 are combinatorially equivalent (c.f. Section 2.2).*

Proof. The first part is a well-known consequence of the regularity of \mathcal{X} (see for example [H77, Chapter II, 6.5(a) and 6.11]). For the second part, note that any two extensions \mathcal{L}_1 and \mathcal{L}_2 of \mathcal{L}_η differ by a divisor supported on the special fiber X_0 . In other words, there exist integers n_v for each vertex v of the dual graph of X_0 such that $\mathcal{L}_1 \simeq \mathcal{L}_2(\sum_v n_v X_v)$. It is now easy to see that, in the notation of Section 2.2, $\pi(\mathcal{L}_1|_{X_0})$ and $\pi(\mathcal{L}_2|_{X_0})$ differ by the divisor of the function $f_G : V(G) \rightarrow \mathbb{Z}$ which sends v to n_v . \square

Consider now a $\mathfrak{g}_d^r L_\eta = (\mathcal{L}_\eta, H_\eta)$ on \mathcal{X}_η and choose an extension \mathcal{L} of \mathcal{L}_η to \mathcal{X} . Since cohomology commutes with flat base change, there is a natural isomorphism $H^0(X_\eta, \mathcal{L}_\eta) \cong H^0(X, \mathcal{L}) \otimes_R K$ and H_η corresponds to a free R -submodule H of $H^0(X, \mathcal{L})$ such that $H \otimes_R K \cong H_\eta$. The restriction H_0 of H to the special fiber X_0 is an $(r + 1)$ -dimensional subspace of $H^0(X_0, \mathcal{L}_0)$. This shows the existence of the limit of a \mathfrak{g}_d^r once the extension \mathcal{L} of \mathcal{L}_η is fixed. However, since the extension is not unique, the limit pair (\mathcal{L}_0, H_0) is not unique. If one were to restrict to a single such extension, there would inevitably be a loss of important information; this can be seen, for example, in the study of limits of ramification points and in the study of limits of smooth hyperelliptic, and more generally, smooth d -gonal curves, see for example [HM82, EH, EM, R]. Some of the difficulties in formulating a good notion of limits for linear series are already manifest in the atypical behavior of linear series on reducible curves, since (as mentioned earlier) they do not satisfy many of the classical theorems which govern the behavior of (global sections of) linear series on irreducible curves.

The Eisenbud-Harris theory of limit linear series on curves of compact type provides a way to keep track of the geometry in the limit by choosing an *aspect* of the limit for each irreducible component, with some extra conditions relating these different aspects at nodes (see below for a precise definition). We refer to [O1, O2] for recent refinements of the

Eisenbud-Harris theory, and to [E, EM] for generalizations of this theory to certain curves not of compact type.

5.2. Limit linear series for curves of compact type. We recall the definition of limit linear series for curves of compact type, following [EH].

We first need to recall the sequence of orders of vanishing of a $\mathfrak{g}_d^r L = (\mathcal{L}, H)$ at a point on a smooth projective curve Y over κ . For any closed point p on Y and any section $f \in H$, denote by $\text{ord}_p(f)$ the order of vanishing of f at p . The orders of vanishing at p of all the sections of \mathcal{L} in H define a sequence of integers $0 \leq a_0^L(p) < a_1^L(p) < \cdots < a_{r-1}^L(p) < a_r^L(p)$. (This sequence is obtained by induction: let f_r be the section in H with the highest order of vanishing, and define f_{r-i} inductively as a section in H with the highest possible order of vanishing at p among all sections which are linearly independent of the first i sections f_r, \dots, f_{r-i+1} . Then $a_i(p) = \text{ord}_p(f_i)$.)

Let X_0 be a strongly semistable curve over κ , let $G = (V, E)$ be the dual graph of X_0 , and for any $v \in V$, let X_v be the corresponding irreducible component of X_0 . When G is a tree, the curve X_0 is called of *compact type*: in this case, $\text{Pic}^0(X_0)$ is compact. (Recall that in general, we have an exact sequence of algebraic groups

$$0 \rightarrow H^1(G, \mathbb{Z}) \otimes \mathbb{G}_m \rightarrow \text{Pic}^0(X_0) \rightarrow \prod_{v \in V_0} \text{Pic}^0(X_v) \rightarrow 0.)$$

Let X_0 be a curve of compact type. A *crude limit* $\mathfrak{g}_d^r L$ over X_0 is by definition the data of a $\mathfrak{g}_d^r L_v = (\mathcal{L}_v, H_v)$ over X_v for each vertex $v \in V$, called the X_v -*aspect* of L , such that the following property holds: If the two components X_u and X_v of X_0 meet at a node p corresponding to an edge $\{u, v\}$ in E , then

$$(6) \quad a_i^{L_v}(p) + a_{r-i}^{L_u}(p) \geq d$$

for all $0 \leq i \leq r$. A crude limit linear series is a *refined limit linear series* if all the inequalities in (6) are equalities.

We are now going to associate to L a degree d divisor on the regularization of X_0 (c.f. Section 2.2 for the definition). Roughly speaking, this amounts to choosing a divisor of degree d on each X_v in the divisor class of L_v . However, the global divisor on $\mathfrak{C}X_0$ defined by such a collection of divisors does not have degree d . To fix this problem, we need to choose a root \mathfrak{r} for the tree G and modify the local divisors using the root. This is done as follows.

Let Γ be the metric graph associated to G and let $\mathfrak{C}X_0$ be the regularization of X_0 . For each vertex v of G , let \overline{D}_v be a divisor of degree d on X_v in the divisor class defining \mathcal{L}_v (i.e., $\mathcal{L}(\overline{D}_v) \simeq \mathcal{L}_v$). The subspace $H_v \subseteq H^0(X_v, \mathcal{L}(\overline{D}_v))$ can be naturally regarded as a subspace of $\kappa(X_v)$. Let \mathfrak{r} be a fixed vertex of G and let $G_{\mathfrak{r}}$ denote G considered as a rooted tree with root \mathfrak{r} . For any vertex $u \neq \mathfrak{r}$ in G , consider the unique path from u to \mathfrak{r} in G and let e_u be the edge adjacent to u along this path. Let x_u be the node of X_u corresponding to the edge e_u . Consider the divisor \mathcal{D} of degree d on $\mathfrak{C}X_0$ defined by

$$(7) \quad D_{\Gamma} := \overline{D}_{\mathfrak{r}} + \sum_{u \neq \mathfrak{r}} D_u, \quad \text{where } D_u := \overline{D}_u - d(x_u).$$

We will see below that \mathcal{D} has rank at least r in $\mathfrak{C}X_0$. In order to derive a more precise characterization of crude limits in terms of a suitable rank function, we need to introduce a refined notion of rank on a metrized complex which takes into account the spaces H_v .

5.2.1. *A refined notion of rank for divisors on a metrized complex.* Let \mathfrak{C} be a metrized complex of algebraic curves, Γ the underlying metric graph with model G , and $\{C_v\}$ the corresponding family of smooth projective curves over κ . Suppose we are given, for each $v \in V$, a non-empty κ -linear subspace F_v of $\kappa(C_v)$. We denote by \mathcal{F} the collection of all F_v . Define the \mathcal{F} -rank of a divisor \mathcal{D} on \mathfrak{C} , denoted $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D})$, to be the maximum integer r such that for any effective divisor \mathcal{E} of degree r , there exists a nonzero rational function f on \mathfrak{C} , with C_v -parts $f_v \in F_v$ for all $v \in V$, such that $\mathcal{D} + \text{div}(f) - \mathcal{E} \geq 0$.

The following proposition provides an upper bound for $r_{\mathfrak{C},\mathcal{F}}$ in terms of the maximal dimension of F_v for $v \in V(G)$.

Proposition 5.2. *Let $s = \max(0, \min_{v \in V(G)}(\dim_{\kappa}(F_v) - 1))$. Then for any divisor \mathcal{D} on \mathfrak{C} , we have $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) \leq s$.*

Proof. If $\dim_{\kappa}(F_v) = 0$ for some v , then either $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) = 0$ (if $\mathcal{D} \geq 0$) or $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) = -1$ (otherwise), and the proposition holds. For the sake of contradiction, suppose now that $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) \geq s + 1 \geq 1$. Let v be a vertex with $s = \dim_{\kappa}(F_v) - 1$. For any effective divisor E_v of degree $s + 1$ on C_v , there exists a rational function f on \mathfrak{C} with $f_v \in F_v$ such that $\mathcal{D} - E_v + \text{div}(f) \geq 0$. Taking Γ -parts shows that $D_{\Gamma} - (s + 1)(v) + \text{div}(f_{\Gamma}) \geq 0$. For a generic choice of E_v , we can assume as in the proof of Proposition 2.1 that the slopes of f_{Γ} on the edges adjacent to v are fixed. For the divisor $D'_v = D_v + \text{div}_v(f_{\Gamma})$, this shows that for any choice of E_v on C_v , there exists an element $f_v \in F_v$ with $D'_v - E_v + \text{div}(f_v) \geq 0$. This contradicts the assumption that $\dim_{\kappa}(F_v) = s + 1$, since F_v can define a linear system of projective dimension at most s on C_v . \square

In Section 4.6.1, we provided a formula for the rank of a divisor on a metrized complex \mathfrak{C} which is a connected sum of two metrized complexes \mathfrak{C}_1 and \mathfrak{C}_2 . We now state a straightforward generalization of Proposition 4.8 to restricted ranks, which will be used in the next section.

Suppose that the underlying graph of \mathfrak{C} has a bridge edge $e = \{v_1, v_2\}$ and let \mathfrak{C}_1 and \mathfrak{C}_2 be the two metrized complexes obtained by removing the edge e . Let Γ_1, Γ_2 , and Γ be the underlying metric graphs of $\mathfrak{C}_1, \mathfrak{C}_2$, and \mathfrak{C} respectively (with v_i a vertex of Γ_i). Denote by x_1 and x_2 the points of \mathfrak{C}_1 and \mathfrak{C}_2 , respectively, corresponding to the edge e . Let \mathcal{F} be a family of spaces of rational functions F_v for vertices v of Γ , and denote by \mathcal{F}_1 (resp. \mathcal{F}_2) the collection of those F_v for v a vertex of Γ_1 (resp. Γ_2).

Proposition 5.3. *Let $\mathcal{D} = \mathcal{D}_1 + \mathcal{D}_2$ be the sum of divisors \mathcal{D}_1 and \mathcal{D}_2 on \mathfrak{C}_1 and \mathfrak{C}_2 , respectively. For a non-negative integer k , define $\eta(k)$ to be the smallest integer n such that $r_{\mathfrak{C}_1,\mathcal{F}_1}(\mathcal{D}_1 + n(x_1)) = k$. Then*

$$(8) \quad r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) = \min_{k \geq 0} \{ k + r_{\mathfrak{C}_2,\mathcal{F}_2}(\mathcal{D}_2 - \eta(k)(x_2)) \}.$$

5.2.2. *Characterization of limit linear series in terms of the refined rank function.* We retain the terminology from the previous sections. Let X_0 be a curve of compact type, $G = (V, E)$ the dual graph, $\mathfrak{C}X_0$ the regularization of X_0 , and let L be a collection of \mathfrak{g}_d^r 's $L_v = (\mathcal{L}_v, H_v)$ on X_v , one for each $v \in V$. Fix a root vertex $\mathfrak{r} \in V$ and divisors $\overline{D}_v \in \text{Div}(X_v)$ such that

$\mathcal{L}(\overline{D}_v) \sim \mathcal{L}_v$, and let \mathcal{D} be the divisor defined by (7). Denote by \mathcal{H} the family of all $H_v \subseteq H^0(X_v, \mathcal{L}(D_v)) \subset \kappa(X_v)$.

Theorem 5.4. *The following two assertions are equivalent:*

- (i) L is a crude limit \mathfrak{g}_d^r on X_0 .
- (ii) $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$.

Proof. Denote by $\mathfrak{r}_1, \dots, \mathfrak{r}_h$ all the children of \mathfrak{r} in the rooted tree $G_{\mathfrak{r}}$. For each $1 \leq i \leq h$, denote by $x_i = x_{\mathfrak{r}_i}$ and y_i the κ -points of $X_{\mathfrak{r}_i}$ and $X_{\mathfrak{r}}$, respectively, which correspond to the edge $\{\mathfrak{r}, \mathfrak{r}_i\}$.

Proof of (i) \Rightarrow (ii): Let L be a crude limit \mathfrak{g}_d^r and let $s \leq r$ be a non-negative integer. Let $F_{\mathfrak{r}} \subset H_{\mathfrak{r}}$ be a subspace of dimension $s + 1$, and denote by \mathcal{F} the collection of $F_{\mathfrak{r}}$ and H_v for $v \neq \mathfrak{r}$ in V . We will prove that $r_{\mathfrak{C}X_0, \mathcal{F}}(\mathcal{D}) = s$. The result then follows by taking $\mathcal{F} = \mathcal{H}$.

By Proposition 5.2, we have $r_{\mathfrak{C}X_0, \mathcal{F}}(\mathcal{D}) \leq s$, so it will be enough to prove the reverse inequality $r_{\mathfrak{C}X_0, \mathcal{F}}(\mathcal{D}) \geq s$. The proof proceeds by induction on $|V|$, and by applying Proposition 5.3. For the base case of our induction, we have $|V| = 1$, and thus $\mathfrak{C}X_0 = X_{\mathfrak{r}}$ and $r_{X_{\mathfrak{r}}, \{F_{\mathfrak{r}}\}}(D_{\mathfrak{r}}) = s$.

Now suppose that the claim holds for all metrized complexes on at most $|V| - 1$ vertices. For each $1 \leq i \leq h$, denote by G_i the rooted tree at \mathfrak{r}_i obtained by removing the edge $\{\mathfrak{r}, \mathfrak{r}_i\}$ from G , and let Γ_i be the metric graph obtained from G_i by taking all edge lengths to be 1. Let \mathfrak{C}_i be the sub-metrized complex of $\mathfrak{C}X_0$ with underlying metric graph Γ_i and curves X_v for v a vertex of G_i , and denote by \mathcal{D}_i the restriction of \mathcal{D} to \mathfrak{C}_i , i.e., $D_{i,v} = D_v$ for all vertices v of G_i . Denote by \mathcal{H}_i the family of all H_v for v a vertex of G_i . Let η_i be the η -function associated to the bridge edge $\{\mathfrak{r}, \mathfrak{r}_i\}$, so that $\eta_i(k)$ is the smallest integer n such that $r_{\mathfrak{C}_i, \mathcal{H}_i}(\mathcal{D}_i + n(x_i)) = k$. Repeatedly applying Proposition 5.3 reduces the calculation of $r_{\mathfrak{C}X_0, \mathcal{F}}(\mathcal{D})$ to the calculation of the restricted ranks of certain divisors on the metrized complex $X_{\mathfrak{r}}$ (with a single vertex \mathfrak{r} and a unique curve $X_{\mathfrak{r}}$) with respect to the subspace $F_{\mathfrak{r}} \subset \kappa(X_{\mathfrak{r}})$; more precisely,

$$r_{\mathfrak{C}X_0, \mathcal{F}}(\mathcal{D}) = \min_{s_1, \dots, s_h \geq 0} \left\{ s_1 + \dots + s_h + r_{X_{\mathfrak{r}}, \{F_{\mathfrak{r}}\}} \left(D_{\mathfrak{r}} - \sum_{i=1}^h \eta_i(s_i)(y_i) \right) \right\}.$$

We are thus led to prove that each term $s_1 + \dots + s_h + r_{X_{\mathfrak{r}}, \{F_{\mathfrak{r}}\}} \left(D_{\mathfrak{r}} - \sum_{i=1}^h \eta_i(s_i)(y_i) \right)$ in the above formula, for non-negative integers s_i , is bounded below by s . We may clearly assume that $s_1 + \dots + s_h \leq s + 1$.

Claim 5.5. *For each $1 \leq i \leq h$, $\eta_i(s_i) \leq d - a_{r-s_i}^{L_{\mathfrak{r}_i}}(x_i)$.*

Proof of the claim. By definition, the subspace $F_{\mathfrak{r}_i}$ of $H_{\mathfrak{r}_i}$ consisting of all global sections of $\mathcal{L}_{\mathfrak{r}_i}$ having an order of vanishing at least $a_{r-s_i}^{L_{\mathfrak{r}_i}}(x_i)$ at x_i has dimension $s_i + 1$. Let \mathcal{D}'_i be the divisor defined on \mathfrak{C}_i by the data of the linear series $L_v = (\mathcal{L}(\overline{D}_v), H_v)$, for v a vertex of G_i , and the root \mathfrak{r}_i , as in (7). Note that $\mathcal{D}'_i = \mathcal{D}_i + d(x_i)$. Let \mathcal{F}_i be the collection of $F_{\mathfrak{r}_i}$ and H_v , for $v \neq \mathfrak{r}_i$ a vertex of G_i . By the induction hypothesis applied to the subcurve of X_0 consisting of all curves X_v for v a vertex of G_i and the crude limit linear series $\{L_v\}_{v \in V(G_i)}$, we have $r_{\mathfrak{C}_i, \mathcal{F}_i}(\mathcal{D}'_i) = s_i$. On the other hand, by the definition of $F_{\mathfrak{r}_i}$, we have

$$r_{\mathfrak{C}_i, \mathcal{F}_i}(\mathcal{D}'_i) = r_{\mathfrak{C}_i, \mathcal{F}_i}(\mathcal{D}'_i - a_{r-s_i}^{L_{\mathfrak{r}_i}}(x_i)(x_i))$$

and thus

$$r_{\mathfrak{C}_i, \mathcal{F}_i}(\mathcal{D}_i + (d - a_{r-s_i}^{L_{\tau_i}}(x_i))(x_i)) = r_{\mathfrak{C}_i, \mathcal{F}_i}(\mathcal{D}'_i - a_{r-s_i}^{L_{\tau_i}}(x_i)(x_i)) = s_i.$$

By the definition of $\eta_i(s_i)$, we infer that $\eta_i(s_i) \leq d - a_{r-s_i}^{L_{\tau_i}}(x_i)$, which proves the claim.

To finish the proof of (i) \Rightarrow (ii), note that since L is a crude limit \mathfrak{g}_d^r , we have

$$a_{s_i}^{L_{\tau}}(y_i) \geq d - a_{r-s_i}^{L_{\tau_i}}(x_i).$$

This shows that the space of all global sections of \mathcal{L}_{τ} with an order of vanishing at least $\eta(s_i)$ at y_i has codimension at most s_i . The intersection of all these spaces with F_{τ} has dimension at least $s + 1 - \sum_{i=1}^h s_i$, which shows that $r_{X_{\tau}, \{F_{\tau}\}}(\mathcal{D}_{\tau} - \sum_{i=1}^h \eta_i(s_i)(y_i)) \geq s - \sum_{i=1}^h s_i$.

Proof of (ii) \Rightarrow (i): By Corollary A.6, in calculating $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D})$ we can restrict to effective divisors in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$. In this case, since G is a tree, we can suppose that all the rational functions on Γ which appear in the calculation of $r_{\mathfrak{C}X_0, \mathcal{H}}$ arise via linear interpolation from integral functions $f_G : V \rightarrow \mathbb{Z}$. Since G is a tree, the data of an integral function f_G on V is equivalent, up to an additive constant, to an assignment of labels $a_{uv} \in \mathbb{Z}$ to each oriented edge uv of G such that $a_{uv} = -a_{vu}$ for any $\{u, v\} \in E$. (Set $a_{uv} = f_G(u) - f_G(v)$.) In what follows, we denote by p_v the unique parent of a vertex $v \neq \tau$ in the rooted tree G_{τ} .

Suppose now that $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$. We first observe that, for any other choice of a root vertex τ' , if \mathcal{D}' denotes the corresponding divisor on $\mathfrak{C}X_0$ defined as in (7) then the two divisors \mathcal{D} and \mathcal{D}' differ by the principal divisor $\text{div}(f)$, where f is a rational function with Γ -part f_G and constant X_v -part for all $v \in V$. Thus $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}') = r$, so it will be enough to prove the validity of Condition (6) for the root τ and a child u of τ among τ_1, \dots, τ_h .

Denote by x_u and y_u the κ -points of X_u and X_{τ} , respectively, corresponding to the edge $\{\tau, u\}$, and let $0 \leq i \leq d$. For any effective divisor $\mathcal{E} = E_{\tau} + E_u$, where E_{τ} and E_u are effective divisors of degree $r - i$ and i on X_{τ} and X_u , respectively, there must exist a rational function f with $f_w \in H_w$ for all w such that $\mathcal{D} - \mathcal{E} + \text{div}(f) \geq 0$. For any two adjacent vertices w and z in G , let $a_{wz} = f(w) - f(z)$ as above. Since for any vertex $w \neq \tau$, D_w has degree zero, a simple induction starting from the leaves and going towards the root shows that $a_{wpw} \geq 0$. This in particular implies that $D_{\tau} + \text{div}(f_{\tau}) - E_{\tau} - a_{\tau u}(y_u) \geq 0$ and $D_u + \text{div}(f_u) - E_u - (d - a_{\tau u})(x_u) \geq 0$. For a generic choice of E_{τ} and E_u , we can assume as in the proof of Proposition 2.1 that $a_{\tau u} = a$ is a constant. In other words, there exists an integer $0 \leq a \leq d$ such that the sublinear system of H_{τ} (resp., H_u) which consists of those sections with an order of vanishing at least a (resp. $d - a$) at y_u (resp. x_u) has projective dimension at least $r - i$ (resp. i). By the definition of the sequences $a^{L_{\tau}}(\cdot)$ and $a^{L_u}(\cdot)$, this simply means that $a_i^{L_{\tau}}(y_u) \geq a$ and $a_{r-i}^{L_u}(x_u) \geq d - a$. It follows that $a_i^{L_{\tau}}(y_u) + a_{r-i}^{L_u}(x_u) \geq d$, which is Condition (6) for the two irreducible components X_{τ} and X_u of X_0 . \square

Corollary 5.6. *Let κ be of characteristic zero. Then L is a refined limit series iff $r_{\mathfrak{C}, \mathcal{H}}(\overline{D}) = r$ and all the ramification points of L_v are smooth points of X_v for all $v \in V(G)$.*

Proof. This is a direct consequence of Theorem 5.4 above and the Plücker formula, c.f. [EH, Proposition 1.1]. \square

We note that the *a priori* dependence of the family \mathcal{H} on the X_v -aspects L_v can be removed, and recovered from the condition on the rank, in the following sense. For any crude limit $\mathfrak{g}_d^r L$ on X_0 with X_v -aspect $L_v = (\mathcal{L}_v, H_v)$, choose a divisor $D_v \in \text{Div}^d(X_v)$ with $\mathcal{L}_v \simeq \mathcal{L}(D_v)$. Two divisors \mathcal{D} and \mathcal{D}' in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ are called *combinatorially equivalent* if they differ by the divisor of a rational function f on $\mathfrak{C}X_0$ with all f_v constant.

The following result is essentially a reformulation of Theorem 1.6 from the Introduction and easily implies the result stated there.

Theorem 5.7. *Let X_0 be a curve of compact type and $\mathfrak{C}X_0$ the regularization of X_0 . Then there is a bijective correspondence between the following:*

- Pairs $(L, \{D_v\})$ consisting of a crude limit $\mathfrak{g}_d^r L$ on X_0 and a collection of divisors D_v on X_v with $L_v \simeq \mathcal{L}(D_v)$ for any irreducible component X_v of X_0 ; and
- Pairs $(\mathcal{H}, [\mathcal{D}])$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r+1)$ -dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and $[\mathcal{D}]$ is the combinatorial linear equivalence class of a divisor $\mathcal{D} \in \text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ of degree d on $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$.

In particular, the data of a crude limit \mathfrak{g}_d^r on X_0 is equivalent to the data of a pair $(\mathcal{H}, [\mathcal{D}])$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$.

Proof. We already know by Theorem 5.4 that any crude limit \mathfrak{g}_d^r gives a pair $(\mathcal{H}, [\mathcal{D}])$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$.

Let \mathcal{H} be a collection of $(r+1)$ -dimensional subspaces $H_v \subset \kappa(X_v)$ for $v \in V$, and let \mathcal{D} be a divisor of degree d in $\text{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ with $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$. We show the existence of a $\mathfrak{g}_d^r \mathcal{L}_v = (L_v, H_v)$ on X_v for each $v \in V$ such that the collection $L = \{L_v\}$ defines a crude limit \mathfrak{g}_d^r on X_0 , and such that \mathcal{D} is (combinatorially) linearly equivalent to the divisor \mathcal{D}^L associated to L by (7).

Let D_Γ (resp. D_v) denote the Γ -part (resp. X_v -part) of \mathcal{D} . Since Γ has genus zero, for any vertex v in $V(G)$ there exists a rational function f_Γ on Γ such that $D_\Gamma + \text{div}(f_\Gamma) = d(v)$. Define $\tilde{D}_v := D_v + \text{div}_v(f_\Gamma)$, where $\text{div}_v(f_\Gamma)$ is defined as in (1). We claim that $H_v \subset H^0(X_v, \mathcal{L}_v)$; this amounts to showing that $\text{div}(f) + \tilde{D}_v \geq 0$ for any $f \in H_v$. If we denote by H'_v the subspace of H_v consisting of all rational functions in H_v with this property, it will be enough to show that $\dim_\kappa(H'_v) = r+1$.

First, we note that for any rational function f'_Γ on Γ with $D_\Gamma + \text{div}(f'_\Gamma) \geq 0$, we have $\text{div}_v(f_\Gamma) - \text{div}_v(f'_\Gamma) \geq 0$. Indeed, $D_\Gamma + \text{div}(f'_\Gamma) \geq 0$ implies that for any edge e adjacent to v , the slope of f'_Γ in the tangent direction corresponding to e is bounded above by the sum of the coefficients $D_\Gamma(w)$ of the points w lying on the connected component of $\Gamma \setminus \{v\}$ which contains the edge e . This sum is precisely the slope of f_Γ in the tangent direction corresponding to e .

For any effective divisor E of degree r on X_v , let \mathcal{E} denote the corresponding divisor on $\mathfrak{C}X_0$. Since $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$, there exists a rational function \mathfrak{f}^E such that $\text{div}(\mathfrak{f}^E) + \mathcal{D} - \mathcal{E} \geq 0$. Since $\text{div}(f_\Gamma^E) + D_\Gamma \geq 0$, the preceding remark implies that $\tilde{D}_v = D_v + \text{div}_v(f_\Gamma) \geq D_v + \text{div}_v(f_\Gamma^E)$. By the definition of $\text{div}(\mathfrak{f})$, we conclude that $\text{div}(f_v^E) + \tilde{D}_v - E \geq 0$, and in particular f_v^E belongs to H'_v . Since this holds for any E , we infer that the linear series on X_v defined by H'_v has (projective) dimension at least r and so H'_v has dimension at least $r+1$, thus $H'_v = H_v$. We conclude that $L_v = (\mathcal{L}_v, H_v)$ is a \mathfrak{g}_d^r on X_v , and thus the collection of L_v for $v \in V$ defines a crude limit $\mathfrak{g}_d^r L$ on X_0 . If \mathcal{D}^L is the divisor associated to L by (7), it is straightforward to check that \mathcal{D}^L and \mathcal{D} differ by the divisor of a rational function \mathfrak{f} on $\mathfrak{C}X_0$ whose X_v -parts are all constant and whose Γ -part is the linear interpolation of an integer-valued function $f : V \rightarrow \mathbb{Z}$. \square

5.3. Limit linear series for metrized complexes. Let \mathfrak{C} be a metrized complex of algebraic curves over κ , Γ the underlying metric graph with model $G = (V, E)$, and $\{C_v\}$ the corresponding collection of smooth projective curves over κ .

We define a (*crude*) limit \mathfrak{g}_d^r on \mathfrak{C} to be an equivalence class of pairs $(\mathcal{H}, \mathcal{D})$ consisting of a divisor \mathcal{D} of degree d on \mathfrak{C} and a collection \mathcal{H} of $(r+1)$ -dimensional subspaces $H_v \subset \kappa(C_v)$, for $v \in V(G)$, such that $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) = r$. Two pairs $(\mathcal{H}, \mathcal{D})$ and $(\mathcal{H}', \mathcal{D}')$ are considered equivalent if there is a rational function f on \mathfrak{C} such that $\mathcal{D}' = \mathcal{D} + \text{div}(f)$ and $H_v = H'_v \cdot f_v$ for all $v \in V$, where f_v denotes the C_v -part of f .

Remark 5.8. Since $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$ for all \mathcal{H} as above and all $\mathcal{D} \in \text{Div}(\mathfrak{C})$, it follows immediately from Theorem 3.2 that for any limit \mathfrak{g}_d^r on \mathfrak{C} we have if $d \geq 2g$ or $r \geq g$, then $r + g \leq d$ (Riemann's inequality), and if \mathcal{D} is *special* then $r \leq d/2$ (Clifford's inequality). For curves of compact type, these inequalities are established (by a completely different proof) in [EH].

Theorem 5.9. *Let X be a smooth proper curve over \mathbb{K} , \mathfrak{X} a strongly semistable model for X , and $\mathfrak{C}\mathfrak{X}$ the metrized complex associated to \mathfrak{X} . Let D be a divisor on X and let $L_\eta = (\mathcal{L}(D), H_\eta)$, for $H_\eta \subseteq H^0(X, \mathcal{L}(D)) \subset \mathbb{K}(X)$, be a \mathfrak{g}_d^r on X . For any vertex $v \in V$, define H_v as the κ -vector space defined by the reduction to $\kappa(C_v)$ of all the rational functions in H_η (c.f. Section 4.4), and let $\mathcal{H} = \{H_v\}_{v \in V}$. Then the pair $(\tau_*^{\mathfrak{C}\mathfrak{X}}(D), \mathcal{H})$ defines a limit \mathfrak{g}_d^r on $\mathfrak{C}\mathfrak{X}$.*

Proof. By Lemma 4.3, all the subspaces H_v have dimension $r+1$ over κ . By Proposition 5.2, we only have to show that $r_{\mathfrak{C}\mathfrak{X}, \mathcal{H}} \geq r$. By Theorem A.1, it suffices to show that for any effective divisor $\mathcal{E} \in \text{Div}(\mathfrak{C}\mathfrak{X})$ of degree r supported on a subset $\mathcal{R} = \bigcup_v \mathcal{R}_v$ of $\mathfrak{C}\mathfrak{X}$ with $\mathcal{R}_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ of size $g_v + 1$, there exists a rational function f on $\mathfrak{C}\mathfrak{X}$ such that $f_v \in H_v$ for all $v \in V$ and $\text{div}(f) + \tau_*^{\mathfrak{C}\mathfrak{X}}(D) - \mathcal{E} \geq 0$. For any such \mathcal{E} , there exists an effective divisor of degree r on X such that $\tau_*^{\mathfrak{C}\mathfrak{X}}(E) = \mathcal{E}$. Since L_η is a \mathfrak{g}_d^r on X , there exists a rational function $f \in H_\eta$ such that $D - E + \text{div}(f) \geq 0$. Let \mathfrak{f} be the corresponding rational function on $\mathfrak{C}\mathfrak{X}$ (as in the paragraph preceding Theorem 4.5). We conclude that $\tau_*^{\mathfrak{C}\mathfrak{X}}(D - E + \text{div}(\mathfrak{f})) = \tau_*^{\mathfrak{C}\mathfrak{X}}(D) - \mathcal{E} + \tau_*^{\mathfrak{C}\mathfrak{X}}(\text{div}(\mathfrak{f})) = \tau_*^{\mathfrak{C}\mathfrak{X}}(D) - \mathcal{E} + \text{div}(\mathfrak{f}) \geq 0$. Since $f_v \in H_v$, this shows that $r_{\mathfrak{C}\mathfrak{X}, \mathcal{H}} \geq r$ as desired. \square

For the regularization $\mathfrak{C}X_0$ of a strongly semistable curve X_0 over κ , Theorem 5.9 can be reformulated as follows. Let $\phi : \mathcal{X} \rightarrow \text{Spec } R$ be a regular smoothing of X_0 over a discrete valuation ring R and let $L_\eta = (\mathcal{L}_\eta, H_\eta)$ be a \mathfrak{g}_d^r over the generic fiber \mathcal{X}_η of ϕ . Consider an extension \mathcal{L} of L_η to \mathcal{X} , and let \mathcal{L}_0 be the restriction of \mathcal{L} to X_0 . Let \mathcal{D} be a divisor on $\mathfrak{C}X_0$ in the linear equivalence class of $\pi(\mathcal{L}_0)$, where $\pi : \text{Pic}(X_0) \rightarrow \bigoplus_{v \in V(G)} \text{Pic}(X_v)$ is as in Section 2. There exists a family $\mathcal{H} = \{H_v\}_{v \in V(G)}$ of $(r+1)$ -dimensional subspaces $H_v \subset \kappa(X_v)$ such that $r_{\mathfrak{C}X_0, \mathcal{H}}(\mathcal{D}) = r$. In addition, H_v can be explicitly constructed as follows. For any vertex v , choose a rational function \mathfrak{f}_v on $\mathfrak{C}X_0$ such that the divisor $\mathcal{D} + \text{div}(f_{G,v})$ is v -reduced and let $f_{G,v} : V(G) \rightarrow \mathbb{Z}$ be the Γ -part of \mathfrak{f}_v . Consider the extension $\mathcal{L}(\sum_v f_{G,v}(u)X_u)$ of \mathcal{L}_η to \mathcal{X} and let $H_{f_{G,v}}$ be the closure of H_η in $H^0(\mathcal{X}, \mathcal{L}(\sum_v f_{G,v}(u)X_u))$. Define $H_v := H_{f_{G,v}}|_{X_v}$ and $\tau_*(L_\eta) := (\mathcal{D}, \mathcal{H})$. Finally, for any $v \in V$ let \mathcal{D}^v be the v -reduced divisor linearly equivalent to \mathcal{D} , and define $d_v := \mathcal{D}_\Gamma(v) = \text{deg}(\mathcal{D}_v^v)$.

Theorem 5.10. *Let X_0 be a strongly semistable curve over κ and $\mathfrak{C}X_0$ the regularization of X_0 . For any regular smoothing $\pi : X \rightarrow \text{Spec } R$ of X_0 over a discrete valuation ring R , $\tau_*(L_\eta)$ as defined above is a limit \mathfrak{g}_d^r on X_0 . In addition, for each vertex $v \in V(G)$ we have*

$H_v \subseteq H^0(X_v, \mathcal{L}(D_v^\circ))$, so that $(\mathcal{L}(D_v^\circ), H_v)$ defines a $\mathfrak{g}_{d_v}^r$ on the irreducible component X_v of X_0 .

We conclude this section with the following strengthening of Theorem 5.9 which will be used in §5.5.

Theorem 5.11. *Let X be a smooth proper curve over \mathbb{K} , \mathfrak{X} a strongly semistable model for X , and let $\mathfrak{C}\mathfrak{X}$ be the metrized complex associated to \mathfrak{X} . Let D be a divisor on X and set $\mathcal{D} = \tau_*^{\mathfrak{C}\mathfrak{X}}(D)$. Let $L_\eta = (\mathcal{L}(D), H_\eta)$ be a \mathfrak{g}_d^r on X and let $\mathcal{H} = \{H_v\}_{v \in V}$ be as in the statement of Theorem 5.9. Let E° be an effective divisor of degree e supported on the smooth locus of the special fiber of \mathfrak{X} , and define $\mathcal{E}^\circ = \sum_{v \in V} E_v^\circ \in \text{Div}(\mathfrak{C}\mathfrak{X})$, where E_v° is the restriction of E° to C_v . Suppose that $H_v \subseteq L(D_v - E_v^\circ)$ for all v . Then the pair $(\mathcal{D} - \mathcal{E}^\circ, \mathcal{H})$ is a limit \mathfrak{g}_{d-e}^r on $\mathfrak{C}\mathfrak{X}$.*

Proof. In the view of Theorem A.1, the proof is similar to the proof of Theorem 5.9, by requiring each subset $\mathcal{R}_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ to be disjoint from the support of E_v° . For any effective divisor \mathcal{E} of degree r supported on $\mathcal{R} = \bigcup \mathcal{R}_v$, the proof of Theorem 5.9 shows that there exists a rational function f such that $\mathcal{D} - \mathcal{E} + \text{div}(f) \geq 0$. We claim that in fact $\mathcal{D} - \mathcal{E} + \text{div}(f) - \mathcal{E}^\circ \geq 0$, which implies that $(\mathcal{D} - \mathcal{E}^\circ, \mathcal{H})$ is a limit \mathfrak{g}_{d-e}^r on $\mathfrak{C}\mathfrak{X}$ as desired. To see this, let E_Γ° be the Γ -part of \mathcal{E}° and note that since $E_\Gamma^\circ = \sum_{v \in V} \deg(E_v^\circ)(v)$ and $\text{div}(f_\Gamma) \geq E_\Gamma - D_\Gamma$ by hypothesis, it suffices to prove that the C_v -part of $\text{div}(f)$ is at least the C_v -part of $\mathcal{E} + \mathcal{E}^\circ - \mathcal{D}$ for all $v \in V$, i.e.,

$$(9) \quad \text{div}(f)_v := \text{div}(f_v) + \text{div}_v(f_\Gamma) \geq E_v + E_v^\circ - D_v.$$

If $z \in \text{supp}(E_v^\circ)$, then since $\text{supp}(E_v^\circ)$ is disjoint from $\mathcal{A}_v \supset \text{supp}(\text{div}_v(f_\Gamma))$ we have $\text{div}(f)_v(z) = \text{div}(f_v)(z)$, and since $\text{supp}(E_v^\circ)$ is disjoint from $\mathcal{R}_v \supset \text{supp}(E_v)$ we have $E_v(z) = 0$. Thus (9) holds in this case since $H_v \subset L(D_v - E_v^\circ)$ implies that $\text{div}(f_v)(z) \geq E_v^\circ(z) - D_v(z)$. On the other hand, if $z \notin \text{supp}(E_v^\circ)$, then $E_v^\circ(z) = 0$ and thus (9) holds because by assumption we have $\text{div}(f) \geq \mathcal{E} - \mathcal{D}$ and in particular $\text{div}(f)_v(z) \geq E_v(z) - D_v(z)$.

Thus (9) is valid for all $z \in C_v(\kappa)$ as desired. \square

5.4. Remarks on completion and smoothing. Given a divisor \mathcal{D} of degree d and rank r on \mathfrak{C} , one may ask if there exists a collection $\mathcal{H} = \{H_i\}$, with H_i an $(r+1)$ -dimensional subspace of $\kappa(C_i)$ for all i , such that $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) = r$. (Recall that we always have $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$.) If so, we say that \mathcal{D} can be *completed to a limit \mathfrak{g}_d^r* on \mathfrak{C} . Moreover, Theorem 5.9 shows that if X is a smooth proper curve over \mathbb{K} , (D, L_η) is a \mathfrak{g}_d^r on X , and $\mathfrak{C}\mathfrak{X}$ is the metrized complex associated to a strongly semistable model \mathfrak{X} for X , then the associated pair $(\tau_*^{\mathfrak{C}\mathfrak{X}}(D), \mathcal{H})$ (which we call the *specialization* of $(\mathcal{D}, \mathcal{H})$) is a limit \mathfrak{g}_d^r on $\mathfrak{C}\mathfrak{X}$. We say that a limit $\mathfrak{g}_d^r(\mathcal{D}, \mathcal{H})$ on a metrized complex \mathfrak{C} over κ is *smoothable* if there exists a smooth proper curve X over \mathbb{K} such that $\mathfrak{C} = \mathfrak{C}\mathfrak{X}$ for some strongly semistable model \mathfrak{X} of X and $(\mathcal{D}, \mathcal{H})$ arises the specialization of a limit \mathfrak{g}_d^r on X . Note that by [ABBR1, Theorem 3.24], every $\text{val}(\mathbb{K}^*)$ -rational metrized complex \mathfrak{C} over κ is of the form $\mathfrak{C} = \mathfrak{C}\mathfrak{X}$ for some smooth proper curve X over \mathbb{K} and some strongly semistable model \mathfrak{X} of X , where $\text{val}(\mathbb{K}^*)$ -rational means that the edge lengths in G all belong to the value group of \mathbb{K} .

In this section, we show by example that not every divisor \mathcal{D} of degree d and rank r on a metrized complex \mathfrak{C} can be completed to a limit \mathfrak{g}_d^r , and not every limit \mathfrak{g}_d^r on a $\text{val}(\mathbb{K}^*)$ -rational metrized complex \mathfrak{C} is smoothable.




FIGURE 2. A weighted metric graph underlying a metrized complex possessing a degree 2 rank 1 divisor which cannot be completed to a limit \mathfrak{g}_2^1 .

Example 5.12 (A degree d rank r divisor which cannot be completed to a limit \mathfrak{g}_d^r). Consider the weighted graph Γ depicted in Figure 2 (with arbitrary edge lengths), and let \mathfrak{C} be a metrized complex whose underlying weighted metric graph is Γ (so each C_{v_i} has genus 1 and C_v has genus 0). For any point p in $C_v \simeq \mathbb{P}^1$, consider the degree 2 divisor $\mathcal{D} = 2(p)$. We claim that \mathcal{D} has rank 1 on \mathfrak{C} but cannot be completed to a limit \mathfrak{g}_2^1 . The fact that \mathcal{D} has rank 1 is straightforward and left to the reader. Assume for the sake of contradiction that there exists a limit $\mathfrak{g}_2^1(\mathcal{H}, \mathcal{D})$ on \mathfrak{C} . Let x_1, x_2, x_3 (resp. y_1, y_2, y_3) be the points of $C_v = \mathbb{P}^1$ (resp. $C_i := C_{v_i}$) corresponding to the three edges of Γ . Fix $i \in \{1, 2, 3\}$, choose a point $q_i \neq y_i$ in C_i and let $\mathcal{E} = (q_i)$. For any rational function f on \mathfrak{C} such that $\mathcal{D} - \mathcal{E} + \text{div}(f) \geq 0$, the fact that C_i has genus one implies that $\text{div}_{v_i}(f_\Gamma) \geq 2(y_i)$, since there is no rational function on C_i whose divisor is $(q_i) - (y_i)$. Therefore $\text{div}(f_\Gamma) = 2(v_i) - 2(v)$. Restricting f to C_v gives a rational function f_i on C_v with $\text{div}(f_i) = 2(x_i) - 2(p)$. Since $(\mathcal{H}, \mathcal{D})$ is by assumption a limit \mathfrak{g}_2^1 , we must have $f_1, f_2, f_3 \in H_v$. But a simple calculation shows that the f_i are linearly independent, implying that $\dim_\kappa(H_v) \geq 3$, a contradiction.

Remark 5.13. Suppose the residue field κ of \mathbb{K} has characteristic different from 2. Then there exists a metrized complex \mathfrak{C} over κ (with all curves C_v isomorphic to \mathbb{P}^1) and a divisor \mathcal{D} of degree 2 and rank 1 on \mathfrak{C} which cannot be completed to a limit \mathfrak{g}_2^1 . Indeed, any hyperelliptic metric graph which cannot be realized as the skeleton of a hyperelliptic curve over \mathbb{K} provides an example of this phenomenon (and one can construct many examples of such metric graphs, using the characterization of those hyperelliptic metric graphs which can be lifted to a hyperelliptic curve over \mathbb{K} given in [C2] and [ABBR2].) Here is a sketch of the argument. It is not difficult to show that a limit \mathfrak{g}_2^1 on a metrized complex \mathfrak{C} gives rise, in a natural way, to a finite harmonic morphism of degree two (in the terminology of [ABBR1, Definition 2.19]) from \mathfrak{C} to a metrized complex of genus zero. By the lifting result [ABBR1, Theorem 7.7], such a finite harmonic morphism can be lifted to a degree two map from a smooth proper curve X/\mathbb{K} with associated metrized complex \mathfrak{C} to \mathbb{P}^1 , which implies that the underlying metric graph of \mathfrak{C} is the skeleton of a hyperelliptic curve over \mathbb{K} .

Example 5.14 (A limit \mathfrak{g}_d^r which is not smoothable). This phenomenon occurs already in the compact type case: the original paper of Eisenbud-Harris on limit linear series [EH] provides an example (Example 3.2 of *loc. cit.*) of a limit \mathfrak{g}_4^2 which cannot be smoothed; we recall their example here.

The metrized complex \mathfrak{C} in question is associated to a semistable curve X_0 with two components meeting transversely at a point p ; one of the components is a hyperelliptic curve Y of genus at least 4 and the other component Z has genus zero. The point p is a ramification point of the \mathfrak{g}_2^1 on Y , and without loss of generality we may assume that p corresponds to the point 0 on $\mathbb{P}^1 \cong Z$. The dual graph G of X_0 has two vertices v_Y and v_Z

connected by a single edge, which may take to be of length one in Γ . Now define a divisor \mathcal{D} on \mathfrak{C} by setting $\mathcal{D} = 4(\infty)$ for the point $\infty \in \mathbb{P}^1 \simeq Z$. If we define $\mathcal{H} = \{H_Y, H_Z\}$, where $H_Y = H^0(Y, 4(p))$ and H_Z is the span of $1, t^2 + t^3, t^4$ (with t a local parameter at 0 on \mathbb{P}^1), then by [EH, Example 3.2] and Theorem 5.4 above, the pair $(\mathcal{D}, \mathcal{H})$ defines a limit \mathfrak{g}_4^2 on \mathfrak{C} . (It is also not hard to show this directly using our definitions.) This limit \mathfrak{g}_4^2 cannot be smoothed.

Remark 5.15. Using Mnev universality theorem, Dustin Cartwright has recently proved that the question of smoothing g_d^2 can be very dependent on the characteristic of the field [Car14].

5.5. Application to bounding the number of rational points on curves. In this section, we explain how limit linear series on metrized complexes of curves can be used to illuminate the proof of a recent theorem due to E. Katz and D. Zureick-Brown [KZB] and put it into a broader context.

Let K be a number field and suppose X is a smooth, proper, geometrically integral curve over K of genus $g \geq 2$. Let J be the Jacobian of X , which is an abelian variety of dimension g defined over K . If the Mordell-Weil rank r of $J(K)$ is less than g , Coleman [Co] adapted an old method of Chabauty to prove that if $p > 2g$ is a prime which is unramified in K and \mathfrak{p} is a prime of good reduction for X lying over p , then $\#X(K) \leq \#\bar{X}(\mathbf{F}_{\mathfrak{p}}) + 2g - 2$. Here \bar{X} denotes the special fiber of a smooth proper model for X over the completion $\mathcal{O}_{\mathfrak{p}}$ of \mathcal{O}_K at \mathfrak{p} and $\mathbf{F}_{\mathfrak{p}} = \mathcal{O}_K/\mathfrak{p}$. Stoll [St] improved this bound by replacing $2g - 2$ with $2r$. Lorenzini and Tucker [LT] (see also [MP]) proved the same bound as Coleman without assuming that X has good reduction at \mathfrak{p} ; in their bound, $\bar{X}(\mathbf{F}_{\mathfrak{p}})$ is replaced by $\tilde{\mathfrak{X}}^{\text{sm}}(\mathbf{F}_{\mathfrak{p}})$ where \mathfrak{X} is a proper regular model for X over $\mathcal{O}_{\mathfrak{p}}$ and $\tilde{\mathfrak{X}}^{\text{sm}}$ is the smooth locus of the special fiber of \mathfrak{X} . Katz and Zureick-Brown combine the improvements of Stoll and Lorenzini-Tucker by proving:

Theorem 5.16. *Let K be a number field and suppose X is a smooth, proper, geometrically integral curve over K of genus $g \geq 2$. Suppose the Mordell-Weil rank r of $J(K)$ is less than g , and that $p > 2g$ is a prime which is unramified in K . Let \mathfrak{p} be a prime of \mathcal{O}_K lying over p and let \mathfrak{X} be a proper regular model for X over $\mathcal{O}_{\mathfrak{p}}$. Then*

$$\#X(K) \leq \#\tilde{\mathfrak{X}}^{\text{sm}}(\mathbf{F}_{\mathfrak{p}}) + 2r.$$

In order to explain the main new idea in the paper of Katz and Zureick-Brown, we first quickly recall the basic arguments used by Coleman, Stoll, and Lorenzini-Tucker. (See [MP] for a highly readable and more detailed overview.) Assume first that we are in the setting of Coleman's paper, so that $r < g$, $p > 2g$ is a prime which is unramified in K , and X has good reduction at the prime \mathfrak{p} lying over p . Fix a rational point $P \in X(K)$ (if there is no such point, we are already done!). Coleman associates to each regular differential ω on X over $K_{\mathfrak{p}}$ (the \mathfrak{p} -adic completion of K) a "definite p -adic integral" $\int_P^Q \omega \in K_{\mathfrak{p}}$. If V_{chab} denotes the vector space of all ω such that $\int_P^Q \omega = 0$ for all $Q \in X(K)$, Coleman shows that $\dim V_{\text{chab}} \geq g - r > 0$. Locally, p -adic integrals are obtained by formally integrating a power series expansion for ω with respect to a local parameter. Using this observation and an elementary Newton polygon argument, Coleman proves that

$$\#X(K) \leq \sum_{\tilde{Q} \in \bar{X}(\mathbf{F}_{\mathfrak{p}})} (1 + n_{\tilde{Q}}),$$

where $n_{\tilde{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\text{ord}_{\tilde{Q}} \tilde{\omega}$; here $\tilde{\omega}$ denotes the reduction of a suitable rescaling $c\omega$ of ω to \bar{X} , where the scaling factor is chosen so that $c\omega$ is regular

and non-vanishing along the special fiber \bar{X} . If we choose any nonzero $\omega \in V_{\text{chab}}$, then the fact that the canonical divisor class on \bar{X} has degree $2g - 2$ gives

$$\sum_{\tilde{Q} \in \bar{X}(\mathbf{F}_p)} n_{\tilde{Q}} \leq \sum_{\tilde{Q} \in \bar{X}(\mathbf{F}_p)} \text{ord}_{\tilde{Q}} \tilde{\omega} \leq 2g - 2,$$

which yields Coleman's bound.

Stoll observed that one could do better than this by adapting the differential ω to the point \tilde{Q} rather than using the same differential ω on all residue classes. Define the *Chabauty divisor*

$$D_{\text{chab}} = \sum_{\tilde{Q} \in \bar{X}(\mathbf{F}_p)} n_{\tilde{Q}}(\tilde{Q}).$$

Then D_{chab} and $K_{\bar{X}} - D_{\text{chab}}$ are both equivalent to effective divisors, so by Clifford's inequality (applied to the smooth proper curve \bar{X}) we have $r(D_{\text{chab}}) := h^0(D_{\text{chab}}) - 1 \leq \frac{1}{2} \deg(D_{\text{chab}})$. On the other hand, by the semicontinuity of h^0 under specialization we have $h^0(D_{\text{chab}}) \geq \dim V_{\text{chab}} \geq g - r$. Combining these inequalities gives

$$\sum_{\tilde{Q} \in \bar{X}(\mathbf{F}_p)} n_{\tilde{Q}} \leq 2r$$

which leads to Stoll's refinement of Coleman's bound.

Lorenzini and Tucker observed that one can generalize Coleman's bound to the case of bad reduction as follows. Let \mathfrak{X} be a proper regular model for X over \mathcal{O}_p and note that points of $X(K)$ specialize to $\mathfrak{X}^{\text{sm}}(\mathbf{F}_p)$. One obtains by a similar argument the bound

$$(10) \quad \#X(K) \leq \sum_{\tilde{Q} \in \mathfrak{X}^{\text{sm}}(\mathbf{F}_p)} (1 + n_{\tilde{Q}}),$$

where $n_{\tilde{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\text{ord}_{\tilde{Q}} \tilde{\omega}$; here $\tilde{\omega}$ denotes the reduction of (a suitable rescaling of) ω to the unique irreducible component of the special fiber of \mathfrak{X} containing \tilde{Q} and $\dim V_{\text{chab}} \geq g - r > 0$ as before. Choosing a nonzero $\omega \in V_{\text{chab}}$ as in Coleman's bound, the fact that the relative dualizing sheaf for \mathfrak{X} has degree $2g - 2$ gives the Lorenzini-Tucker bound.

In order to combine the bounds of Stoll and Lorenzini-Tucker, we see that it is natural to form the Chabauty divisor

$$D_{\text{chab}} = \sum_{\tilde{Q} \in \mathfrak{X}^{\text{sm}}(\mathbf{F}_p)} n_{\tilde{Q}}(\tilde{Q})$$

and try to prove, using some version of semicontinuity of h^0 and Clifford's inequality, that its degree is at most $2r$. This is the main technical innovation of Katz and Zureick-Brown, so we state it as a theorem:

Theorem 5.17 (Katz–Zureick-Brown). *The degree of D_{chab} is at most $2r$.*

Combining Theorem 5.17 with (10) yields Theorem 5.16. As noted by Katz and Zureick-Brown, if one makes a base change from K_p to an extension field K' over which there is a regular semistable model \mathfrak{X}' for X dominating the base change of \mathfrak{X} , then the corresponding Chabauty divisors satisfy $D'_{\text{chab}} \geq D_{\text{chab}}$. (Here D'_{chab} is defined relative to the K' -vector space $V'_{\text{chab}} = V_{\text{chab}} \otimes_K K'$; one does not want to look at the Mordell-Weil group of J over

extensions of K .) In order to prove Theorem 5.17, we may therefore assume that \mathfrak{X} is a regular semistable model for X (and also that the residue field of K' is algebraically closed).

Let $d = \deg(D_{\text{chab}})$. We now explain how to prove that $d \leq 2r$ when \mathfrak{X} is a semistable regular model using limit linear series on metrized complexes of curves. Our proof is different from (and arguably streamlined than) the one given in [KZB].

Proof of Theorem 5.17. Let $s = \dim_{K'} V'_{\text{chab}} - 1 \geq g - r - 1 \geq 0$. We can identify V'_{chab} with an $(s + 1)$ -dimensional space W of rational functions on X in the usual way by identifying $H^0(X, \Omega_X^1)$ with $L(K_X) = \{f : \text{div}(f) + K_X \geq 0\}$ for a canonical divisor K_X on X . The divisor D_{chab} on $\tilde{\mathfrak{X}}^{\text{sm}}$ defines in a natural way a divisor $\mathcal{D} = \sum_v D_v$ of degree d on the metrized complex $\mathfrak{C}\mathfrak{X}$ associated to \mathfrak{X} , where D_v is the restriction of D_{chab} to C_v . We promote the divisor $K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}$ to a limit linear series $\mathcal{L} = (K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}, \{H_v\})$ on $\mathfrak{C}\mathfrak{X}$ by defining H_v to be the reduction of W to C_v for each $v \in V(G)$. By the definition of D_{chab} , each element of H_v vanishes to order at least $n_{\tilde{Q}}$ at each point \tilde{Q} in $\text{supp}(D_v)$. By Corollary 1.5 and Theorem 5.11, \mathcal{L} is a limit \mathfrak{g}_{2g-2-d}^s on $\mathfrak{C}\mathfrak{X}$. In particular, we must have $r_{\mathfrak{C}\mathfrak{X}}(K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}) \geq g - r - 1$. On the other hand, Clifford's inequality for metrized complexes (Theorem 3.4) implies that $r_{\mathfrak{C}\mathfrak{X}}(K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}) \leq \frac{1}{2}(2g - 2 - d)$. Combining these inequalities gives $d \leq 2r$ as desired. \square

APPENDIX A. RANK-DETERMINING SETS FOR METRIZED COMPLEXES

We retain the terminology from Section 2. Let \mathfrak{C} be a metrized complex of algebraic curves, Γ the underlying metric graph, $G = (V, E)$ a model of Γ and $\{C_v\}_{v \in V}$ the collection of smooth projective curves over κ corresponding to \mathfrak{C} . In this section, we generalize some basic results concerning rank-determining sets [L, HKN] from metric graphs to metrized complexes by following and providing complements to the arguments of [A] (to which we refer for a more detailed exposition).

Let \mathcal{R} be a set of geometric points of \mathfrak{C} (i.e., a subset of $\bigcup_{v \in V} C_v(\kappa)$). The set \mathcal{R} is called *rank-determining* if for any divisor \mathcal{D} on \mathfrak{C} , $r_{\mathfrak{C}}(\mathcal{D})$ coincides with $r_{\mathfrak{C}}^{\mathcal{R}}(\mathcal{D})$, defined as the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all degree k effective divisors \mathcal{E} on \mathfrak{C} with support in \mathcal{R} . In other words, \mathcal{R} is rank-determining if in the definition of rank given in Section 2, one can restrict to effective divisors \mathcal{E} with support in \mathcal{R} .

The following theorem is a common generalization of (a) Luo's theorem [L] (see also [HKN]) that V is a rank-determining set for any loopless model $G = (V, E)$ of a metric graph Γ and (b) the classical fact (also reproved in [L]) that for any smooth projective curve C of genus g over κ , every subset of $C(\kappa)$ of size $g + 1$ is rank-determining.

Theorem A.1. *Let \mathfrak{C} be a metrized complex of algebraic curves, and suppose that the given model G of Γ is loopless. Let $\mathcal{R}_v \subset C_v(\kappa)$ be a subset of size $g_v + 1$ and let $\mathcal{R} = \cup_{v \in V} \mathcal{R}_v$. Then \mathcal{R} is a rank-determining subset of \mathfrak{C} .*

Let \mathcal{D} be a divisor on \mathfrak{C} . For any point $P \in \Gamma$, let \mathcal{D}^P be the quasi-unique P -reduced divisor on \mathfrak{C} linearly equivalent to \mathcal{D} , and denote by D_{Γ}^P (resp. D_v^P) the Γ -part (resp. C_v -part) of \mathcal{D}^P .

Lemma A.2. *A divisor \mathcal{D} on \mathfrak{C} has rank at least one if and only if*

- (1) *For any point P of Γ , $D_{\Gamma}^P(P) \geq 1$, and*

(2) For any vertex $v \in V(G)$, the divisor D_v^ν has rank at least one on C_v .

Proof. The condition $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$ is equivalent to requiring that $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$ for every effective divisor \mathcal{E} of degree 1 on \mathfrak{C} . For $P \in \Gamma \setminus V$, the divisor $\mathcal{D} - (P)$ has non-negative rank if and only if $D_\Gamma^P(P) \geq 1$ (by Lemma 3.11). Similarly, for $v \in V$ and $x \in C_v(\kappa)$, the divisor $\mathcal{D} - (x)$ has non-negative rank in \mathfrak{C} if and only if $D_\Gamma^v(v) \geq 1$ and $D_v^\nu - (x)$ has non-negative rank on C_v (by Lemma 3.11). These are clearly equivalent to (1) and (2). \square

Lemma A.3. A subset $\mathcal{R} \subseteq \bigcup_{v \in V} C_v(\kappa)$ which has non-empty intersection with each $C_v(\kappa)$ is rank-determining if and only if for every divisor \mathcal{D} of non-negative rank on \mathfrak{C} , the following two assertions are equivalent:

- (i) $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$.
- (ii) For any vertex $u \in V$, and for any point $z \in \mathcal{R} \cap C_u(\kappa)$, $D_u^u - (z)$ has non-negative rank on C_u .

Proof. In view of Lemma A.2, for a rank-determining set the two conditions (i) and (ii) are equivalent. Suppose now that (i) and (ii) are equivalent for any divisor \mathcal{D} on \mathfrak{C} . By induction on r , we prove that $r_{\mathfrak{C}}(\mathcal{D}) \geq r$ if and only if for every effective divisor \mathcal{E} of degree r with support in \mathcal{R} , $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$. This will prove that \mathcal{R} is rank-determining.

The case $r = 1$ follows by the hypothesis and Lemma A.2. Supposing now that the statement holds for some integer $r \geq 1$, we prove that it also holds for $r + 1$.

Let \mathcal{D} be a divisor with the property that $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$ for every effective divisor \mathcal{E} of degree $r + 1$ with support in \mathcal{R} . Fix an effective divisor \mathcal{E} of degree r with support in \mathcal{R} . By the base case $r = 1$, the divisor $\mathcal{D} - \mathcal{E}$ has rank at least 1 on \mathfrak{C} because $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E} - (x)) \geq 0$ for any $x \in \mathcal{R}$. Thus $r_{\mathfrak{C}}(\mathcal{D} - (x) - \mathcal{E}) \geq 0$ for any point of $|\mathfrak{C}|$. This holds for any effective divisor \mathcal{E} of degree r with support in \mathcal{R} , and so from the inductive hypothesis we infer that $\mathcal{D} - (x)$ has rank at least r on \mathfrak{C} . Since this holds for any $x \in |\mathfrak{C}|$, we conclude that \mathcal{D} has rank at least $r + 1$. \square

Let \mathcal{D} be a divisor of degree d and non-negative rank on \mathfrak{C} , and let D_Γ and D_v be the Γ and C_v -parts of \mathcal{D} , respectively. Define

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

Note that $|D_\Gamma|$ is a non-empty subset of the symmetric product $\Gamma^{(d)}$ of d copies of Γ .

Consider the *reduced divisor map* $\text{Red}_{\mathcal{D}}^{\mathfrak{C}} : \Gamma \rightarrow \Gamma^{(d)}$ which sends a point $P \in \Gamma$ to D_Γ^P , the Γ -part of the P -reduced divisor \mathcal{D}^P . The following theorem extends [A, Theorem 3] to divisors on metrized complexes.

Theorem A.4. For any divisor \mathcal{D} of degree d and non-negative rank on \mathfrak{C} , the reduced divisor map $\text{Red}_{\mathcal{D}}^{\mathfrak{C}} : \Gamma \rightarrow \Gamma^{(d)}$ is continuous.

Proof. This is based on an explicit description of the reduced divisor map $\text{Red}_{\mathcal{D}}^{\mathfrak{C}}$ in a small neighborhood around any point of Γ , similar to the description provided in [A, Theorem 3] in the context of metric graphs. We merely give the description by providing appropriate modifications to [A, Theorem 3], referring to *loc. cit.* for more details.

Let P be a point of Γ and let $\vec{\mu}$ be a (unit) tangent direction in Γ emanating from P . For $\epsilon > 0$ sufficiently small, we denote by $P + \epsilon\vec{\mu}$ the point of Γ at distance ϵ from P in the direction of $\vec{\mu}$. We will describe the restriction of $\text{Red}_{\mathcal{D}}^{\mathfrak{C}}$ to the segment $[P, P + \epsilon\vec{\mu}]$ for sufficiently small $\epsilon > 0$. One of the two following cases can happen:

(1) For all sufficiently small $\epsilon > 0$, the P -reduced divisor \mathcal{D}^P is also $(P + \epsilon\vec{\mu})$ -reduced. In this case, the map $\text{Red}_{\mathcal{D}}^{\mathcal{C}}$ is constant (and so obviously continuous) on a small segment $[P, P + \epsilon_0\vec{\mu}]$ with $\epsilon_0 > 0$.

(2) There exists a cut S in Γ which is saturated with respect to \mathcal{D}^P such that $P \in \partial S$ and $P + \epsilon\vec{\mu} \notin S$ for all sufficiently small $\epsilon > 0$.

We note that there is a maximum saturated cut S (i.e., containing any other saturated cut) with the property described in (2) (see the proof of [A, Theorem 3] for details). In the following S denotes the maximum saturated cut with property (2). In this case, there exists an $\epsilon_0 > 0$ such that for any $0 < \epsilon < \epsilon_0$, the reduced divisor $\mathcal{D}^{P+\epsilon\vec{\mu}}$ has the following description (the proof mimics that of [A, Theorem 3] and is omitted).

Let $\mu, \vec{\mu}_1, \dots, \vec{\mu}_s$ be all the different tangent vectors in Γ (based at the boundary points $P, x_1, \dots, x_s \in \partial S$, respectively) which are outgoing from S . (It might be the case that $x_i = x_j$ for two different indices i and j). Let $\gamma_0 > 0$ be small enough so that for any point $x \in \partial S$ and any tangent vector \vec{v} to Γ at x which is outgoing from S , the entire segment $(x, x + \gamma_0\vec{v}]$ lies outside S and does not contain any point of the support of D_Γ .

For any $0 < \gamma < \gamma_0$ and any positive integer α , we will define below a rational function $f_\Gamma^{(\gamma, \alpha)}$ on Γ . Appropriate choices of $\gamma = \gamma(\epsilon)$ and α will then give the $(P + \epsilon\vec{\mu})$ -reduced divisor $\mathcal{D}^{P+\epsilon\vec{\mu}} = \mathcal{D}^P + \text{div}(f^{\gamma, \alpha})$, for any $\epsilon < \epsilon_0 := \frac{\gamma_0}{\alpha}$, where $f^{\gamma, \alpha}$ is the rational function on \mathcal{C} given by $f_\Gamma^{(\gamma, \alpha)}$ on Γ and $f_v = 1$ on each C_v .

For $0 < \gamma < \gamma_0$ and integer $\alpha \geq 1$, define $f_\Gamma^{(\gamma, \alpha)}$ as follows:

- $f_\Gamma^{(\gamma, \alpha)}$ takes value zero at any point of S ;
- On any outgoing interval $[x_i, x_i + \gamma\vec{\mu}_i]$ from S , $f_\Gamma^{(\gamma, \alpha)}$ is linear of slope -1 ;
- The restriction of $f_\Gamma^{(\gamma, \alpha)}$ to the interval $[P, P + (\frac{\gamma}{\alpha})\vec{\mu}]$ is linear of slope $-\alpha$;
- $f_\Gamma^{(\gamma, \alpha)}$ takes value $-\gamma$ at any other point of Γ .

Note that the values of $f_\Gamma^{(\gamma, \alpha)}$ at the points $(x_i + \gamma\vec{\mu}_i)$ and $P + (\frac{\gamma}{\alpha})\vec{\mu}$ are all equal to $-\gamma$, so $f_\Gamma^{(\gamma, \alpha)}$ is well-defined.

It remains to determine the values of α and γ . Once the value of α is determined, γ will be defined as $\alpha\epsilon$ so that the point $P + (\frac{\gamma}{\alpha})\vec{\mu}$ coincides with the point $P + \epsilon\vec{\mu}$. We consider the following two cases, depending on whether or not P is a vertex of G :

- If $P \in \Gamma \setminus V$, then $\alpha = D_\Gamma^P(P) - \text{outdeg}_S(P) + 1$. (Note that since S is saturated with respect to D_Γ^P , we have $D_\Gamma^P(P) \geq \text{outdeg}_S(P)$ and thus $\alpha \geq 1$.)
- If $P = v$ for a vertex $v \in V(G)$, let e_0, e_1, \dots, e_l be the outgoing edges at v with respect to S , and consider the points $x_v^{e_0}, x_v^{e_1}, \dots, x_v^{e_l}$ in $C_v(\kappa)$ indexed by these edges. Suppose in addition that e_0 is the edge which corresponds to the tangent direction $\vec{\mu}$. Since S is a saturated cut with respect to D_Γ^P , the divisor $D_v - \text{div}_v(\partial S) = D_v - \sum_{i=0}^l (x_v^{e_i})$ has non-negative rank in C_v . Define α to be the largest integer $n \geq 1$ such that $D_v - n(x_v^{e_0}) - \sum_{i=1}^l (x_v^{e_i})$ has non-negative rank.

Now for any $0 \leq \epsilon < \epsilon_0 = \frac{\gamma_0}{\alpha}$, the divisor $\mathcal{D}^{P+\epsilon\vec{\mu}}$ is $(P + \epsilon\vec{\mu})$ -reduced. (The argument is similar to [A, Proof of Theorem 3].) It follows immediately that the reduced divisor map is continuous on the interval $[P, P + \epsilon_0\vec{\mu}]$, and the result follows. \square

We are now ready to give the proof of Theorem A.1.

Proof of Theorem A.1. By Lemma A.3, it is enough to check the equivalence of the following two properties for any divisor \mathcal{D} on \mathfrak{C} :

- (i) $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$.
- (ii) For any $u \in V$ and any point $z \in \mathcal{R}_u = \mathcal{R} \cap C_u(\kappa)$, the divisor $D_u^u - (z)$ has non-negative rank on C_u .

It is clear that (i) implies (ii). So we only need to prove that (ii) implies (i). In addition, by Lemma A.2, Property (i) is equivalent to:

- (1) for any point P of Γ , $D_{\Gamma}^P(P) \geq 1$; and
- (2) for any vertex $v \in V(G)$, the divisor D_v^v has rank at least one on C_v .

So it suffices to prove that (ii) \Rightarrow (1) and (2). Since cardinality of \mathcal{R}_v is $g_v + 1$, \mathcal{R}_v is rank-determining in C_v . Therefore, (ii) implies (2). We now show that (2) implies (1). Let Γ_0 be the set of all $P \in \Gamma$ such that $D_{\Gamma}^P(P) \geq 1$. By the continuity of the map $\text{Red}_{\mathcal{D}}^{\mathfrak{C}}$, Γ_0 is a closed subset of Γ . In addition, since D_v^v has rank at least one on C_v and $D_{\Gamma}^v(v) = \deg(D_v^v)$ for every vertex $v \in C$, we have $V \subset \Gamma_0$. This shows that $\Gamma \setminus \Gamma_0$ is a disjoint union of open segments contained in edges of G . Suppose for the sake of contradiction that $\Gamma_0 \subsetneq \Gamma$, and let $I = (P, Q)$ be a non-empty segment contained in the edge $\{u, v\}$ of G such that $I \cap \Gamma_0 = \emptyset$.

Claim: $\text{Red}_{\mathcal{D}}^{\mathfrak{C}}$ is constant on the closed interval $[P, Q]$.

To see this, note that for any point $Z \in [P, Q]$ and any tangent direction $\vec{\mu}$ for which $Z + \epsilon\vec{\nu} \in [P, Q]$ for all sufficiently small $\epsilon > 0$, we are always in case (1) in the description of $\text{Red}_{\mathcal{D}}^{\mathfrak{C}}$. Otherwise, there would be an integer $\alpha > 0$ such that $\mathcal{D}^{Z+\epsilon\vec{\mu}} = \mathcal{D}^Z + \text{div}(f^{(\epsilon\alpha, \alpha)})$ for all sufficiently small $\epsilon > 0$. In particular, this would imply (by the definition of $f^{(\eta, a)}$) that $D_{\Gamma}^{Z+\epsilon\vec{\mu}}(Z + \epsilon\vec{\mu}) = \alpha \geq 1$, which implies that $Z + \epsilon\vec{\nu} \in \Gamma_0$, a contradiction. This proves the claim.

A case analysis (depending on whether P and Q are vertices or not) shows that for a point $Z \in (P, Q)$, the cut $S = \Gamma \setminus (P, Q)$ is saturated for $\mathcal{D}^P = \mathcal{D}^Q$. Since $\mathcal{D}^Z = \mathcal{D}^P = \mathcal{D}^Q$, and S does not contain Z , this contradicts the assumption that \mathcal{D}^Z is Z -reduced. \square

Theorem A.1 has the following direct corollaries.

Corollary A.5. *Let \mathcal{G} be a subgroup of \mathbb{R} which contains all the edge lengths in G . For any divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})_{\mathcal{G}}$, we have*

$$r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D}) = r_{\mathfrak{C}}(\mathcal{D}).$$

Proof. Fix a rank-determining set $\mathcal{R} \subset \cup_{v \in V} C_v(\kappa)$ as in Theorem A.1. Since \mathcal{R} is rank-determining and any effective divisor \mathcal{E} with support in \mathcal{R} obviously belongs to $\text{Div}(\mathfrak{C})_{\mathcal{G}}$, to prove the equality of $r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D})$ and $r_{\mathfrak{C}}(\mathcal{D})$ it will be enough to show that the two statements $r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D}) \geq 0$ and $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ are equivalent. Obviously, the former implies the latter, so we only need to show that if $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ then $r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D}) \geq 0$. Let v be a vertex of G and \mathcal{D}^v the v -reduced divisor linearly equivalent to \mathcal{D} . By Lemma 3.11, $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ is equivalent to $r_{C_v}(D_v^v) \geq 0$. Now let \mathcal{D} be an element of $\text{Div}(\mathfrak{C})_{\mathcal{G}}$ with $r_{C_v}(D_v^v) \geq 0$. Since $v \in V$ and \mathcal{G} contains all the edge-lengths in G , it is easy to see that \mathcal{D} and \mathcal{D}^v differ by the divisor of a rational function f with support in $\text{Div}(\mathfrak{C})_{\mathcal{G}}$. In other words, $\mathcal{D} \sim \mathcal{D}^v$ in $\text{Div}(\mathfrak{C})_{\mathcal{G}}$. Since \mathcal{D}^v is linearly equivalent to an effective divisor in $\text{Div}(\mathfrak{C})_{\mathcal{G}}$ (with constant rational function on Γ), we conclude that $r_{\mathfrak{C}, \mathcal{G}}(\mathcal{D}) \geq 0$. \square

Corollary A.6. *Let $\mathfrak{C}X_0$ be the regularization of a strongly semistable curve X_0 over κ . Let \mathcal{L} be a line bundle on X_0 corresponding to a divisor $\mathcal{D} \in \text{Div}(\mathfrak{C})$. Then $r_{\mathfrak{C}}(\mathcal{L}) = r_{\mathfrak{C}X_0}(\mathcal{D})$.*

Proof. This follows from the previous corollary with $\mathcal{G} = \mathbb{Z}$. □

REFERENCES

- [A14] O. Amini, *Equidistribution of Weierstrass points on curves over non-Archimedean fields*, preprint.
- [AB] O. Amini and M. Baker, *Limit linear series for a generic chain of genus one curves*, in preparation.
- [A] O. Amini, *Reduced divisors and embeddings of tropical curves*. Trans. Amer. Math. Soc. 365 (2013), no. 9, 4851–4880.
- [ABBR1] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff, *Lifting harmonic morphisms I: metrized complexes and Berkovich skeleta*. Preprint available at <http://arxiv.org/abs/1303.4812>.
- [ABBR2] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff, *Lifting harmonic morphisms II: tropical curves and metrized complexes*. Preprint available at <http://arxiv.org/abs/1303.4812>.
- [AC] O. Amini and L. Caporaso, *Riemann-Roch theory for weighted graphs and tropical curves*. Adv. Math. 240 (2013), 1–23.
- [B] M. Baker, *Specialization of linear systems from curves to graphs*. Algebra & Number Theory 2 (2008), no. 6, 613–653.
- [BN] M. Baker and S. Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph*. Adv. Math. 215 (2007), no. 2, 766–788.
- [BPR] M. Baker, S. Payne, and J. Rabinoff, *Non-Archimedean geometry, tropicalization, and metrics on curves*. Preprint available at [arXiv:1104.0320v1](https://arxiv.org/abs/1104.0320v1).
- [BS] M. Baker and F. Shokrieh, *Chip-firing games, potential theory on graphs, and spanning trees*. Journal of Combinatorial Theory, Series A 120 (2013), no. 1, 164–182.
- [Ber] V. G. Berkovich, *Spectral theory and analytic geometry over non-Archimedean fields*. Mathematical Surveys and Monographs, vol. 33, American Mathematical Society, Providence, RI, 1990.
- [Ber99] V. G. Berkovich, *Smooth p -adic analytic spaces are locally contractible*. Invent. Math. 137 (1999), no. 1, 1–84.
- [Ber04] V. G. Berkovich, *Smooth p -adic analytic spaces are locally contractible II*. Geometric aspects of Dwork theory, Walter de Gruyter and Co. KG, Berlin, 2004, 293–370.
- [Bi] M. T. I. Bigas, *Brill-Noether theory for stable vector bundles*. Duke Math. J. 62 (1991), no. 2 385–400.
- [C1] L. Caporaso, *Linear series on semistable curves*. Int. Math. Res. Not. 13 (2011), 2921–2969.
- [C2] L. Caporaso, *Gonality of algebraic curves and graphs*. Preprint available at <http://arxiv.org/abs/1201.6246>.
- [Car14] D. Cartwright, *Lifting rank 2 tropical divisors*. Preprint available at <http://users.math.yale.edu/~dc597/lifting.pdf>
- [CR] T. Chinburg and R. Rumely, *The capacity pairing*. J. für die reine und angewandte Mathematik 434 (1993), 1–44.
- [CDPR] F. Cools, J. Draisma, S. Payne, and E. Robeva, *A tropical proof of the Brill-Noether Theorem*. Adv. Math. 230 (2012), no. 2, 759–776.
- [Co] R. F. Coleman, *Effective Chabauty*, Duke Math. J. 52 (1985), no. 3, 765–770.
- [DM] P. Deligne and D. Mumford, *The irreducibility of the space of curves of given genus*. Publications Mathématiques de l’IHES 36 (1969), no. 1, 75–109.
- [EH] D. Eisenbud and J. Harris, *Limit linear series: Basic theory*. Invent. Math. 85 (1986), 337–371.
- [EH2] D. Eisenbud and J. Harris, *The Kodaira dimension of the moduli space of curves of genus ≥ 23* . Invent. Math. 90 (1987), no. 2, 359–387.
- [EH3] D. Eisenbud and J. Harris, *Existence, decomposition, and limits of certain Weierstrass points*. Invent. Math. 87 (1987), no. 3, 495–515.
- [EH4] D. Eisenbud and J. Harris, *The monodromy of Weierstrass points*. Invent. Math. 90 (1987), no.2, 333–341.
- [E] E. Esteves, *Linear systems and ramification points on reducible nodal curves*. Mathematica Contemporanea 14 (1998), 21–35.

- [EM] E. Esteves and N. Medeiros, *Limit canonical systems on curves with two components*. Invent. Math. 149 (2002), no. 2, 267–338.
- [GK] A. Gathmann and M. Kerber, *A Riemann-Roch theorem in tropical geometry*. Math. Z. 259 (2008), no. 1, 217–230.
- [HM98] J. Harris and I. Morrison, *Moduli of Curves*. Graduate Texts in Mathematics 187, Springer Verlag, 1998.
- [HM82] J. Harris and D. Mumford, *On the Kodaira dimension of the moduli space of curves*. Invent. Math. 67 (1982), 23–88.
- [H77] R. Hartshorne, *Algebraic Geometry*. Springer Graduate Texts in Math **52**, 1977.
- [H] R. Hartshorne, *Residues and Duality*. Springer Lecture Notes **20**, 1966.
- [HKN] J. Hladký, D. Král', and S. Norine, *Rank of divisors on tropical curves*. J. Comb. Theory. Series B. 120 (2013), no. 7, 1521–1538.
- [KZB] E. Katz and D. Zureick-Brown, *The Chabauty-Coleman bound at a prime of bad reduction and Clifford bounds for geometric rank functions*, Compositio Math. 149 (2013), no. 11, 1818–1838.
- [LT] D. J. Lorenzini and T. J. Tucker, *Thue equations and the method of Chabauty-Coleman*, Invent. Math. 148 (2002), no. 1, 47–77.
- [LPP] C. M. Lim, S. Payne, and N. Potashnik, *A note on Brill-Noether theory and rank determining sets for metric graphs*. Int. Math. Res. Not. 23 (2012), 5484–5504.
- [L] Y. Luo, *Rank-determining sets of metric graphs*. J. Comb. Theory. Series A. 118 (2011), no. 6, 1775–1793.
- [MP] W. McCallum and B. Poonen, *The method of Chabauty and Coleman*, June 14, 2010. Preprint, <http://www-math.mit.edu/~poonen/papers/chabauty.pdf>, to appear in Panoramas et Synthèses, Société Math. de France.
- [MZ] G. Mikhalkin and I. Zharkov, *Tropical curves, their Jacobians and Theta functions*. Contemporary Mathematics 465: Proceedings of the International Conference on Curves and Abelian Varieties in honor of Roy Smith's 65th birthday. 203–231.
- [N] A. Neeman, *The Distribution of Weierstrass points on a Compact Riemann Surface*, Ann. Math. 120 (1984), 317–328.
- [O1] B. Osserman, *A limit linear series moduli scheme (Un schéma de modules de séries linéaires limites)*. Ann. Inst. Fourier 56 (2006), no. 4, 1165–1205.
- [O2] B. Osserman, *Linked Grassmannians and crude limit linear series*. Int. Math. Res. Not. 25 (2006), 1–27.
- [Par] B. Parker, *Exploded manifolds*. Adv. Math. 229 (2012), no. 6, 3256–3319.
- [Pay] S. Payne, *Fibers of tropicalization*. Math. Zeit. 262 (2009), 301–311.
- [R] Z. Ran, *Modifications of Hodge bundles and enumerative geometry I : the stable hyperelliptic locus*. Preprint available at arxiv.org/1011.0406.
- [St] M. Stoll, *Independence of rational points on twists of a given curve*, Compositio Math. 142 (2006), no. 5, 1201–1214.
- [Te] M. Temkin, *On local properties of non-Archimedean analytic spaces*, Math. Annalen 318 (2000), 585–607.
- [W] G. Welters, *A theorem of Gieseker-Petri type for Prym varieties*. Ann. Sci. Éc. Norm. Supér. 18 (1985), no. 4, 671–683.
- [Z] S-W. Zhang, *Admissible pairing on a curve*. Invent. math. 112 (1993), no. 1, 171–193

CNRS - DMA, ÉCOLE NORMALE SUPÉRIEURE, PARIS
E-mail address: oamini@math.ens.fr

SCHOOL OF MATHEMATICS, GEORGIA INSTITUTE OF TECHNOLOGY
E-mail address: mbaker@math.gatech.edu