A RIEMANN-ROCH THEOREM IN TROPICAL GEOMETRY

ANDREAS GATHMANN AND MICHAEL KERBER

ABSTRACT. Recently, Baker and Norine have proven a Riemann-Roch theorem for finite graphs. We extend their results to metric graphs and thus establish a Riemann-Roch theorem for divisors on (abstract) tropical curves.

Tropical algebraic geometry is a recent branch of mathematics that establishes deep relations between algebro-geometric and purely combinatorial objects. Ideally, every construction and theorem of algebraic geometry should have a tropical (i.e. combinatorial) counterpart that is then hopefully easier to understand — e.g. the tropical counterpart of n-dimensional varieties are certain n-dimensional polyhedral complexes. In this paper we will establish a tropical counterpart of the well-known Riemann-Roch theorem for divisors on curves.

Let us briefly describe the idea of our result. Following Mikhalkin, an (abstract) tropical curve is simply a connected metric graph Γ. A rational function on Γ is a continuous, piecewise linear real-valued function f with integer slopes. For such a function and any point $P \in \Gamma$ the order $\operatorname{ord}_P f$ of f in P is the sum of the slopes of f for all edges emanating from P. For example, the following picture shows a rational function f on a tropical curve Γ with simple zeroes at P_2 and P_5 (i.e. $\operatorname{ord}_{P_2} f = \operatorname{ord}_{P_5} f = 1$), and simple poles at P_3 and P_4 (i.e. $\operatorname{ord}_{P_3} f = \operatorname{ord}_{P_4} f = -1$).

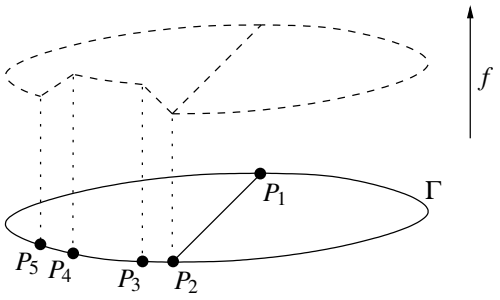

As expected from classical geometry, a divisor on Γ will simply be a formal $\mathbb{Z}$-linear combination of points of Γ. Any rational function f on Γ gives rise to a divisor $(f) := \sum_{P\in\Gamma} \operatorname{ord}_P f \cdot P$ (so that $(f) = P_2 - P_3 - P_4 + P_5$ in the above example).

For a given divisor D we denote by $R(D)$ the space of all rational functions f on Γ such that $(f)+D$ is effective, i.e. contains only non-negative coefficients (e.g. $f \in R(P_3+P_4)$ in the example above). A Riemann-Roch theorem should make a statement about the dimension of these spaces. However, we will see that in general $R(D)$ is a polyhedral complex which is not of pure dimension. As a replacement for the dimension of $R(D)$ we define $r(D)$ to be the biggest integer n such that $R(D-P_1-\cdots-P_n)$ is non-empty for all choices of

$P_1,\dots,P_n \in \Gamma$ (a number that is closely related to the dimension of the cells of $R(D)$ as we will see).

With these notations our Riemann-Roch theorem now simply and expectedly states that

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

where g is the first Betti number of Γ, $\deg D$ is the degree of D, and K is the canonical divisor of Γ following Zhang [Z], i.e. the sum of all vertices of Γ counted with multiplicity equal to their respective valence minus 2 (so that $K = P_1 + P_2$ in our example above).

Our proof relies heavily on a recent result of Baker and Norine that establishes an analogous result for integer-valued functions on the vertices of a (non-metric) graph [BN]. Basically, we will interpret this result as a statement about tropical curves whose edge lengths are integers (so-called $\mathbb{Z}$-graphs) and rational functions on them whose divisors consist of points with integer coordinates (so-called $\mathbb{Z}$-divisors). We then pass from integer to rational and finally real coordinates, as well as to possibly infinite edge lengths, to establish our Riemann-Roch theorem for tropical curves.

More precisely, we will first introduce our basic objects of study, namely divisors and rational functions (and their moduli spaces) on tropical curves in section 1. We then use the result of Baker and Norine to prove a Riemann-Roch theorem for $\mathbb{Z}$- and $\mathbb{Q}$-divisors in section 2 and extend this result in section 3 to arbitrary divisors and graphs (with possibly unbounded edges), with the main result being corollary 3.8.

Shortly after this manuscript had appeared on the e-print archive, Mikhalkin and Zharkov published a preprint that also includes a proof of the Riemann-Roch theorem for tropical curves [MZ]. Their results have been obtained independently and without our knowledge, and in fact their method of proof is entirely different from ours, using Jacobians of tropical curves.

1. Tropical rational functions and divisors

We start by introducing the basic notations used in this paper, in particular the notions of (abstract) tropical curves as well as rational functions and divisors on them.

Definition 1.1 (Graphs)**.** A *graph* Γ will always mean a finite and connected multigraph, not necessarily loop-free (i.e. there may be edges that connect a vertex to itself). The sets of vertices and edges of Γ are denoted $V(\Gamma)$ and $E(\Gamma)$, respectively. The valence of a vertex $P \in V(\Gamma)$ will be denoted $\mathrm{val}(P)$.

(a) A *metric graph* is a pair (Γ, l) consisting of a graph Γ together with a length function $l : E(\Gamma) \to \mathbb{R}_{>0}$. We identify an edge e with the real interval $[0, l(e)]$, leading to a "geometric representation" of the graph by gluing these intervals together at their boundary points according to the combinatorics of Γ. By abuse of notation we will usually denote this geometric representation also by Γ. In this metric space the distance between points as well as the distance from a point to a subset will be written as $\mathrm{dist}(\,\cdot\,,\,\cdot\,)$. The first Betti number of Γ will be called the *genus* of Γ.

(b) If all edge lengths of a metric graph Γ are integers (resp. rational numbers) we call Γ a $\mathbb{Z}$-*graph* (resp. $\mathbb{Q}$-*graph*). In this case the points of (the geometric representation of) Γ with integer (resp. rational) distance to the vertices are called $\mathbb{Z}$-*points* (resp. $\mathbb{Q}$-*points*) of Γ. We denote the set of these points by $\Gamma_{\mathbb{Z}}$ and $\Gamma_{\mathbb{Q}}$, respectively.

(c) A *tropical curve* is a "metric graph with possibly unbounded ends", i.e. a pair (Γ, l) as in (a) where the length function takes values in $\mathbb{R}_{>0} \cup \{\infty\}$, and where each edge of length ∞ is identified with the real interval $[0,\infty] = \mathbb{R}_{\geq 0} \cup \{\infty\}$ in such a way that the ∞ end of the edge has valence 1. These infinity points of will be called the *(unbounded) ends* of Γ.

Remark 1.2. Note that (in contrast to some other conventions on abstract tropical curves found in the literature) our definition allows vertices of valence 1 and 2, and adds "points at infinity" at each unbounded edge. Note also that every metric graph is a tropical curve.

Definition 1.3 (Divisors)**.** A *divisor* on a tropical curve Γ is an element of the free abelian group generated by the points of (the geometric representation of) Γ. The group of all divisors on Γ is denoted $\mathrm{Div}(\Gamma)$. The *degree* $\deg D$ of a divisor $D = \sum_i a_i P_i$ (with $a_i \in \mathbb{Z}$ and $P_i \in \Gamma$) is defined to be the integer $\sum_i a_i$ and obviously gives rise to a morphism $\deg : \mathrm{Div}(\Gamma) \to \mathbb{Z}$. The *support* $\mathrm{supp} D$ of D is defined to be the set of all points of Γ occurring in D with a non-zero coefficient. A divisor is called *effective* if all its coefficients a_i are non-negative. On a $\mathbb{Z}$-graph (resp. $\mathbb{Q}$-graph) a divisor D will be called a *$\mathbb{Z}$-divisor* (resp. *$\mathbb{Q}$-divisor*) if $\mathrm{supp} D \subset \Gamma_{\mathbb{Z}}$ (resp. $\mathrm{supp} D \subset \Gamma_{\mathbb{Q}}$). Following Zhang [Z] we define the canonical divisor of Γ to be

$$K_\Gamma := \sum_{P \in V(\Gamma)} (\mathrm{val}(P) - 2) \cdot P;$$

on a $\mathbb{Z}$-graph (resp. $\mathbb{Q}$-graph) it is obviously a $\mathbb{Z}$-divisor (resp. $\mathbb{Q}$-divisor).

Definition 1.4 (Rational functions)**.** A *rational function* on a tropical curve Γ is a continuous function $f : \Gamma \to \mathbb{R} \cup \{\pm\infty\}$ such that the restriction of f to any edge of Γ is a piecewise linear integral affine function with a finite number of pieces. In particular, f can take on the values $\pm\infty$ only at the unbounded ends of Γ.

For a rational function f as above and a point $P \in \Gamma$ the *order* $\mathrm{ord}_P f \in \mathbb{Z}$ of f at P will be the sum of the outgoing slopes of all segments of Γ emanating from P (of which there are $\mathrm{val}(P)$ if $P \in V(\Gamma)$ and 2 otherwise). In particular, if P is an unbounded end of Γ lying on an unbounded edge e then the order of f at P equals the negative of the slope of f at a point on e sufficiently close to P.

Note that $\mathrm{ord}_P f = 0$ for all points $P \in \Gamma \backslash V(\Gamma)$ at which f is locally linear and thus for all but finitely many points. We can therefore define the *divisor associated to f*

$$(f) := \sum_{P \in \Gamma} \mathrm{ord}_P f \cdot P \in \mathrm{Div}(\Gamma)$$

as in classical geometry.

Remark 1.5. If f is a rational function on a tropical curve Γ then the degree of its associated divisor (f) is $\deg(f) = \sum_{P \in \Gamma} \mathrm{ord}_P f$. By definition of the order this expression can be written as a sum over all segments of Γ on which f is linear, where each such segment counts with the sum of the outgoing slopes of f on it at the two end points of the segment. But as these two slopes are obviously just opposite numbers on each such edge we can conclude that $\deg(f) = 0$ — again analogous to the case of compact curves in classical geometry.

Definition 1.6 (Spaces of functions associated to a divisor)**.** Let D be a divisor of degree n on a tropical curve Γ.

(a) We denote by $R(D)$ the set of all rational functions f on Γ such that the divisor $(f)+D$ is effective. Note that for any such $f \in R(D)$ the divisor $(f)+D$ is a sum of exactly $\deg((f)+D) = \deg D = n$ points by remark 1.5. So if we define

$$S(D) := \{(f, P_1, \dots, P_n);\ f \text{ a rational function on } \Gamma, \\ P_1, \dots, P_n \in \Gamma \text{ such that } (f)+D = P_1 + \dots + P_n\}$$

then we obviously have $R(D) = S(D)/S_n$, where the symmetric group S_n acts on $S(D)$ by permutation of the points P_i.

(b) If Γ is a $\mathbb{Z}$-graph and D a $\mathbb{Z}$-divisor we define a "discrete version" of (a) as follows: let $\tilde{R}(D)$ be the set of all rational functions f on Γ such that $(f)+D$ is an effective $\mathbb{Z}$-divisor, and set

$$\tilde{S}(D) := \{(f, P_1, \dots, P_n);\ f \text{ a rational function on } \Gamma, \\ P_1, \dots, P_n \in \Gamma_{\mathbb{Z}} \text{ such that } (f)+D = P_1 + \dots + P_n\},$$

so that again $\tilde{R}(D) = \tilde{S}(D)/S_n$.

If we want to specify the curve Γ in the notation of these spaces we will also write them as $R_\Gamma(D)$, $S_\Gamma(D)$, $\tilde{R}_\Gamma(D)$, and $\tilde{S}_\Gamma(D)$, respectively.

Remark 1.7. The spaces $R(D)$, $S(D)$, $\tilde{R}(D)$, $\tilde{S}(D)$ of definition 1.6 have the following obvious properties:

(a) all of them are empty if $\deg D < 0$;

(b) $R(D-P) \subset R(D)$ and $\tilde{R}(D-P) \subset \tilde{R}(D)$ for all $P \in \Gamma$;

(c) $\tilde{R}(D) \subset R(D)$ and $\tilde{S}(D) \subset S(D)$ if D is a $\mathbb{Z}$-divisor on a $\mathbb{Z}$-graph Γ.

We want to see now that $R(D)$ and $S(D)$ are *polyhedral complexes* in the sense of [GM], i.e. spaces that can be obtained by gluing finitely many polyhedra along their boundaries, where a polyhedron is defined to be a subset of a real vector space given by finitely many linear equalities and strict inequalities. To do this we first need a lemma that limits the combinatorial possibilities for the elements of $R(D)$ and $S(D)$. For simplicity we will only consider the case of metric graphs here (but it is in fact easy to see with the same arguments that lemmas 1.8 and 1.9 hold as well for tropical curves, i.e. in the presence of unbounded ends).

Lemma 1.8. *Let $p > 0$ be an integer, and let f be a rational function on a metric graph Γ that has at most p poles (counted with multiplicities). Then the absolute value of the slope of f at any point of Γ (which is not a vertex and where f is differentiable) is bounded by a number that depends only on p and the non-metric graph Γ (i.e. the combinatorics of Γ).*

Proof. To simplify the notation of this proof we will consider all zeroes and poles of f to be vertices of Γ (by making them into 2-valent vertices in case they happen to lie in the interior of an edge).

Let e be any edge of Γ on which f is not constant. Construct a path γ along Γ starting with e in the direction in which f is increasing, and then successively following the edges of Γ, at each vertex continuing along an edge on which the outgoing slope of f is maximal.

By our convention on 2-valent vertices above the function f is affine linear on each edge of Γ. Let us now study how the slope of f changes along γ when we pass a vertex $P \in \Gamma$. By definition we have $\lambda_1 + \dots + \lambda_n = \operatorname{ord}_P f$, where $\lambda_1, \dots, \lambda_n$ are the outgoing slopes of f on

the edges $e_1,\dots,e_n$ adjacent to P. Now let N be the maximal valence of a vertex occurring in Γ, and assume that our path γ approaches P along the edge e_1 on which f has incoming slope $-\lambda_1$ greater or equal to $(N+p)^\alpha$ for some $\alpha \geq 1$. It then follows that

$$\begin{aligned}\lambda_2+\dots+\lambda_n &= -\lambda_1 + \operatorname{ord}_P f \\ &\geq (N+p)^\alpha - p \\ &= N(N+p)^{\alpha-1} + p((N+p)^{\alpha-1}-1) \\ &\geq N(N+p)^{\alpha-1},\end{aligned}$$

which means that the biggest of the numbers $\lambda_2,\dots,\lambda_n$, i.e. the outgoing slope of f along γ when leaving P, is at least $(N+p)^{\alpha-1}$ (recall that $n \leq N$ and that λ_1 can never be the biggest of the $\lambda_1,\dots,\lambda_n$ since it is negative by assumption whereas at least one of the $\lambda_2,\dots,\lambda_n$ is positive).

So if we assume that the slope of f is at least $(N+p)^\alpha$ on the edge e this means by induction that the slope of f on γ is at least $(N+p)^{\alpha-i}$ after crossing i vertices, i.e. in particular that f is strictly increasing on the first $\alpha+1$ edges of γ. But this is only possible if α is less than the number of edges of Γ: otherwise at least one edge must occur twice among the first $\alpha+1$ edges of γ, in contradiction to f being strictly increasing on γ in this range. As the initial edge e was arbitrary this means that the slope of f on any edge is bounded by $(N+p)^\alpha$, with α being the number of edges of Γ. □

Lemma 1.9. *For any divisor D on a metric graph Γ the spaces $R(D)$ and $S(D)$ are polyhedral complexes.*

Proof. We will start with $S(D)$. For each edge e of Γ we choose an adjacent vertex that we will call the starting point of e. To each element $(f,P_1,\dots,P_n)$ of $S(D)$ we associate the following discrete data:

(a) the information on which edge or vertex P_i lies for all $i=1,\dots,n$;

(b) the (integer) slope of f on each edge at its starting point;

and the following continuous data:

(c) the distance of each P_i that lies on an edge from the starting point of this edge;

(d) the value of f at a chosen vertex.

These data obviously determine f uniquely: on each edge we know the starting slope of f as well as the position and orders of all zeroes and poles, so f can be reconstructed on each edge if its starting value on the edge is given. As Γ is connected by assumption we can thus reconstruct the whole function from the starting value (d).

Since there are only finitely many choices for (a) and (b) (use lemma 1.8 for (b)), we get a stratification of $S(D)$ with finitely many strata. The data (c) and (d) are given by finitely many real variables in each stratum, so each stratum is a subset of a real vector space. Finally, the condition on the given data to be compatible is given by several linear equalities and inequalities (the distances (c) must be positive and less than the length of the corresponding edges, and the values of f at the boundary points of the edges must be so that we get a well-defined continuous function on Γ), so that $S(D)$ is indeed a polyhedral complex.

The space $R(D)$ is then simply the quotient of $S(D)$ by the affine linear action of the permutation group of the $P_1,\dots,P_n$, and hence is a polyhedral complex as well. □

Remark 1.10. For the spaces $\tilde{R}(D)$ and $\tilde{S}(D)$ the same argument as in the proof of lemma 1.9 holds, with the only exception that the data (c) becomes discrete since the points in (f) are required to be $\mathbb{Z}$-points. Hence the only continuous parameter left is the additive constant (d), i.e. both $\tilde{R}(D)$ and $\tilde{S}(D)$ are finite unions of real lines. We can thus regard $\tilde{R}(D)$ and $\tilde{S}(D)$ as "discrete versions" of the spaces $R(D)$ and $S(D)$.

The following example shows that the polyhedral complexes $R(D)$ and $S(D)$ are in general not pure-dimensional, i.e. there may exist inclusion-maximal cells of different dimensions:

Example 1.11. Consider the canonical divisor $K_\Gamma = P+Q$ of the metric graph Γ obtained by connecting two cycles C_1 and C_2 of length 1 by an edge e of length $l(e) \in \mathbb{Z}_{>0}$ (see the picture below). Furthermore, let f be a rational function on Γ such that $(f)+K_\Gamma = P_1+P_2$.

Assume first that both P_1 and P_2 lie in the interior of the edge e. Note that for all such choices of the points P_i there exists (up to an additive constant) exactly one rational function with zeros at P_1 and P_2 and poles at the prescribed points P and Q. It follows that the corresponding cell in $S(K_\Gamma)$ can be identified with $[0,l(e)]\times[0,l(e)]\times\mathbb{R}$, where the first two factors represent the position of the points P_1 and P_2, and the last factor parametrizes the additive constant. Hence the dimension of this cell in $S(K_\Gamma)$ is 3.

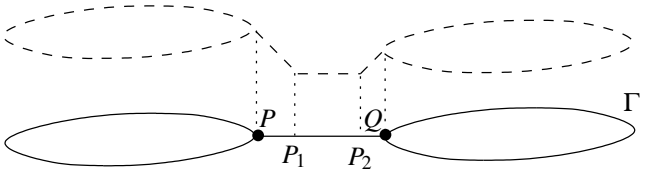

Next, assume that $(f,P_1',P_2') \in S(K_\Gamma)$ such that P_1' is not on the closure of e but rather in the interior of a cycle C_i. We will see in lemma 2.2 that P_2' must then lie on the same cycle. Moreover, it is easy to check that this requires P_2' to be the point on C_i "opposite" to P_1' as in the following picture:

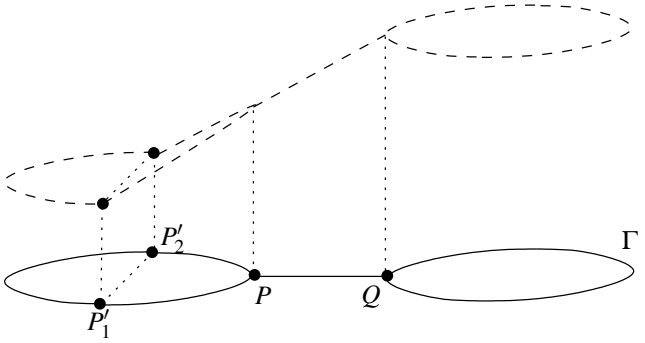

Hence for each choice of P_1' on one of the cycles there exists exactly one point P_2' such that $(f,P_1',P_2') \in S(K_\Gamma)$. It follows that this cell of $S(K_\Gamma)$ can be identified with $C_i\times\mathbb{R}$, where the second factor parametrizes the additive constant as above. In particular, the dimension of this cell is 2.

Putting all this we obtain the following schematic picture of the polyhedral complex $S(K_\Gamma)$, where for simplicity we have omitted the factor $\mathbb{R}$ corresponding to the additive constant in all cells:

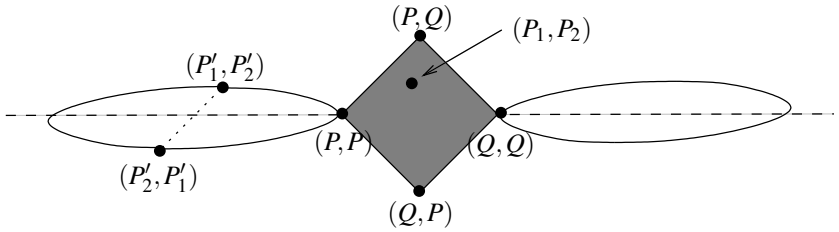

The space $R(K_\Gamma)$ is then obtained from this by dividing out the action of the symmetric group on two elements, which can be realized geometrically by "folding $S(K_\Gamma)$ along the dashed line above". In particular, both $R(K_\Gamma)$ and $S(K_\Gamma)$ are not pure-dimensional, but rather have components of dimensions 2 and 3.

The above example shows that when formulating a Riemann-Roch type statement about the dimensions of the spaces $R(D)$ we have to be careful since these dimensions are ill-defined in general. The following definition will serve as a replacement:

Definition 1.12. Let D be a divisor of degree n on a tropical curve Γ.

(a) We define $r(D)$ to be the biggest integer k such that for all choices of (not necessarily distinct) points $P_1,\dots,P_k \in \Gamma$ we have $R(D-P_1-\cdots-P_k) \neq \emptyset$ (or equivalently $S(D-P_1-\cdots-P_k) \neq \emptyset$), where $r(D)$ is understood to be -1 if $R(D)$ (or equivalently $S(D)$) itself is empty.

(b) If D is a $\mathbb{Z}$-divisor on a $\mathbb{Z}$-graph Γ there is also a corresponding "discrete version": we let $\tilde{r}(D)$ be the biggest integer such that $\tilde{R}(D-P_1-\cdots-P_k) \neq \emptyset$ for all choices of k points $P_1,\dots,P_k \in \Gamma_{\mathbb{Z}}$.

If we want to specify the curve Γ in the notation of these numbers we will also write them as $r_\Gamma(D)$ and $\tilde{r}_\Gamma(D)$, respectively.

Example 1.13.

(a) By remark 1.7 (a) it is clear that $r(D) = -1$ if $\deg D < 0$, and $r(D) \leq \deg D$ otherwise. The same statement holds for $\tilde{r}(D)$ for $\mathbb{Z}$-divisors on $\mathbb{Z}$-graphs.

(b) For the canonical divisor of the metric graph in example 1.11 we have $r(K_\Gamma) = 1$ since we have seen that

- for all points $P_1 \in \Gamma$ there is a rational function f with $(f) + K_\Gamma = P_1 + P_2$ (i.e. $f \in R(K_\Gamma - P_1)$);
- for some choice of $P_1, P_2 \in \Gamma$ (e.g. P_1 and P_2 in the interior of the circles C_1 and C_2, respectively) there is no rational function f with $(f) + K_\Gamma = P_1 + P_2$.

(c) Let Γ be a metric graph, and let $\lambda \in \mathbb{R}_{>0}$. By a *rescaling* of Γ by λ we mean the metric graph of the same combinatorics as Γ where we replace each edge e of length $l(e)$ by an edge of length $\lambda \cdot l(e)$. Note that any divisor (resp. rational function) on Γ gives rise to an induced divisor (resp. rational function) on the rescaling by also rescaling the positions of the points (resp. the values of the function). In particular, the numbers $r(D)$ for a divisor D on Γ remain constant under rescalings.

Note that rescalings by positive integers take $\mathbb{Z}$-graphs and $\mathbb{Z}$-divisors again to $\mathbb{Z}$-graphs and $\mathbb{Z}$-divisors, but that $\tilde{r}(D)$ may change in this case since the rescaling introduces new $\mathbb{Z}$-points.

Remark 1.14. By the proof of lemma 1.9 the continuous parameters for the elements $(f, P_1, \ldots, P_n)$ of $S(D)$ are the positions of the points P_i and the value of f at a chosen vertex. In particular, when passing from $S(D)$ to $S(D-P)$ for a generic choice of P this fixes one of the P_i and thus makes each cell of $S(D)$ (disappear or) one dimension smaller. It follows that the maximal dimension of the cells of $S(D)$ (and $R(D)$) is always at least $r(D)+1$ (with the $+1$ coming from the additive constant, i.e. the value of the functions at the chosen vertex).

Remark 1.15. There is another interpretation of the numbers $r(D)$ that we will need later: let D be a divisor of degree n on a tropical curve Γ, let $i \in \{0, \ldots, n\}$, and assume that $S(D) \neq \emptyset$. Consider the forgetful maps

$$\pi_i : S(D) \to \Gamma^i, \quad (f, P_1, \ldots, P_n) \mapsto (P_1, \ldots, P_i).$$

Note that these maps are morphisms of polyhedral complexes in the sense of [GM] (i.e. they map each cell of the source to a single cell in the target by an affine linear map). It is clear by definition that the number $r(D)$ can be interpreted using these maps as the biggest integer i such that π_i is surjective.

Example 1.16. Consider again the metric graph Γ of example 1.11, but now the spaces $R(D)$ and $S(D)$ for the divisor $D = P' + Q'$, where P' are Q' are interior points of the cycles C_1 and C_2, respectively. In this case lemma 2.2 will tell us that (f, P_1, P_2) can only be in $S(D)$ if each cycle C_i contains one of the points P_1, P_2, which is then easily seen to require that in fact $\{P, Q\} = \{P_1, P_2\}$, i.e. that f is a constant function. It follows that $R(D)$ is simply the real line, whereas $S(D)$ is two disjoint copies of $\mathbb{R}$ (i.e. both spaces have pure dimension 1). It also follows in the same way that $r(D) = 0$.

In particular, when comparing this to the result of examples 1.11 and 1.13 (b) (which can be regarded as the limit case when $P' \to P$ and $Q' \to Q$) we see that $r(D)$ can jump, and that the spaces $R(D)$ and $S(D)$ can change quite drastically under "continuous deformations of D". So as in the classical case it is really only the number $r(D) - r(K_\Gamma - D)$, and not $r(D)$ alone, that will turn out to depend on the degree of D and the genus of Γ only.

2. Riemann-Roch for $\mathbb{Q}$-divisors

We will now start with the study of Riemann-Roch theorems. Our basic ingredient is the Riemann-Roch theorem for finite (non-metric) graphs of Baker and Norine ([BN] theorem 1.11) that is easily translated into our set-up:

Theorem 2.1 (Baker and Norine)**.** *Let Γ be a $\mathbb{Z}$-graph of genus g all of whose edge lengths are bigger than 1. Then for every $\mathbb{Z}$-divisor D on Γ we have $\tilde{r}(D) - \tilde{r}(K_\Gamma - D) = \deg D + 1 - g$.*

Sketch of proof. We start by replacing each edge e of Γ by a chain of $l(e)$ edges of length 1, arriving at a graph whose geometric representation is the same as before, and where all $\mathbb{Z}$-points that were in the interior of an edge have been turned into 2-valent vertices. Note that by the condition that all edge lengths of the original graph are bigger than 1 this implies that the new graph has no loops, i.e. no edges whose two boundary points coincide (an assumption made throughout in [BN]). As it is clear by definition that none of the

terms in the Riemann-Roch equation changes under this transformation it suffices to prove the theorem for the new graph. By abuse of notation we will also denote it by Γ.

Note that every rational function f on Γ whose divisor is a $\mathbb{Z}$-divisor is uniquely determined by its values on the vertices (since it is just given by linear interpolation on the edges). Moreover, up to a possibly non-integer global additive constant all these values of f on the vertices are integers. Conversely, every integer-valued function on the vertices of Γ gives rise to a rational function on Γ (by linear interpolation) whose divisor is a $\mathbb{Z}$-divisor. As all edge lengths in Γ are 1 the divisor (f) can then be rewritten using this correspondence as

$$(f) = \sum_{\overline{PQ}} (f(Q) - f(P)) \cdot (P - Q) \qquad (*)$$

where the sum is taken over all edges of Γ (and P and Q denote the boundary vertices of these edges in any order). In particular, for a $\mathbb{Z}$-divisor D the number $\tilde{r}(D)$ can also be defined as the maximum number k such that for each choice of vertices $P_1, \dots, P_k$ of Γ there is an integer-valued function f on the vertices of Γ such that $(f) + D$ is effective, where (f) is defined by $(*)$.

This is the approach that Baker and Norine take in [BN]. They establish the Riemann-Roch theorem in this set-up, thus proving the theorem as stated above. To prove their theorem their first step is to show its equivalence to the following two statements:

- $\tilde{r}(K_\Gamma) \geq g - 1$; and
- for any $\mathbb{Z}$-divisor $D \in \mathrm{Div}(\Gamma)$ there exists a $\mathbb{Z}$-divisor $E \in \mathrm{Div}(\Gamma)$ with $\deg(E) = g - 1$ and $\tilde{r}(E) = -1$ such that exactly one of the sets $\tilde{R}(D)$ and $\tilde{R}(E - D)$ is empty.

The central idea in the proof of these two statements is then to consider total orderings on the vertices of Γ. For each such ordering there is an associated divisor

$$E = \sum_{e \in E(\Gamma)} m(e) - \sum_{P \in V(\Gamma)} P$$

where $m(e)$ denotes the boundary point of e that is the bigger one in the given ordering — the divisor E in the second statement above can for example be taken to be of this form for a suitable ordering (that depends on D). For details of the proof see [BN]. □

In order to pass from the "discrete case" (the spaces $\tilde{R}(D)$) to the "continuous case" (the spaces $R(D)$) we need a few lemmas first.

Lemma 2.2. *Let D be a $\mathbb{Z}$-divisor on a $\mathbb{Z}$-graph Γ, and let $(f, P_1, \dots, P_n) \in S(D)$. Assume moreover that some P_i is not a $\mathbb{Z}$-point. Then on every cycle of Γ containing P_i there is another point P_j (with $i \neq j$) that is also not a $\mathbb{Z}$-point.*

Proof. Assume that C is a cycle containing exactly one simple zero $P = P_i \in \Gamma \setminus \Gamma_\mathbb{Z}$ (note that if P is a multiple zero then we are done). Consider the cycle C to be the interval $[0, l(C)]$ with the endpoints identified such that the zero point lies on a vertex, and let $x \in \{1, \dots, l(C)\}$ be the integer such that $P \in (x-1, x)$ with this identification. By adding a suitable constant to f we may assume that $f(x-1) \in \mathbb{Z}$. Since $P \in (x-1, x)$ and the slope of f on the interval $[x-1, P]$ differs from that on the interval (P, x) by 1 we conclude that $f(x) \notin \mathbb{Z}$. As all other points of non-differentiability of f on $[0, l(C)]$ are $\mathbb{Z}$-points by assumption it follows that $f(Q) \in \mathbb{Z}$ for all $Q = 0, \dots, x-1$ and $f(Q) \notin \mathbb{Z}$ for all $Q = x, \dots, l(C)$. In particular, we see that $f(0) \neq f(l(C))$, in contradiction to the continuity of f. □

Lemma 2.3. *For every $\mathbb{Z}$-divisor D on a $\mathbb{Z}$-graph Γ with $R(D) \neq \emptyset$ we have $\tilde{R}(D) \neq \emptyset$.*

Proof. We will prove the statement by induction on $n := \deg D$.

Let $f \in R(D)$, so that $(f) + D = P_1 + \cdots + P_n$ for some (not necessarily distinct) points $P_i \in \Gamma$. In particular, this requires of course that $n \geq 0$. Moreover, if $n = 0$ then $(f) = -D$ is a $\mathbb{Z}$-divisor and hence $f \in \tilde{R}(D)$. As this finishes the proof in the case $n \leq 0$ we can assume from now on that $n > 0$, and that the statement of the lemma is true for all divisors of degree less than n.

If $P_i \in \Gamma_{\mathbb{Z}}$ for some i then $f \in R(D - P_i)$ and hence $\tilde{R}(D - P_i) \neq \emptyset$ by the induction assumption. As this implies $\tilde{R}(D) \neq \emptyset$ we have proven the lemma in this case and may thus assume from now on that none of the P_i is a $\mathbb{Z}$-point of the curve.

After possibly relabeling the points P_i we may assume in addition that $0 < \operatorname{dist}(P_n, \Gamma_{\mathbb{Z}}) \leq \operatorname{dist}(P_i, \Gamma_{\mathbb{Z}})$ for all $i = 1, \ldots, n$, i.e. that P_n is a point among the P_i that minimizes the distance to the $\mathbb{Z}$-points of the curve. Let $P \in \Gamma_{\mathbb{Z}}$ be a point with $\operatorname{dist}(P_n, P) = \operatorname{dist}(P_n, \Gamma_{\mathbb{Z}}) =: d$, and let $\Gamma' \subset \Gamma$ be the connected component of $\Gamma \backslash \{P_1, \ldots, P_n\}$ that contains P. With this notation consider the rational function

$$h : \Gamma \to \mathbb{R}, \quad Q \mapsto \begin{cases} -\min(d, \operatorname{dist}(Q, \{P_1, \ldots, P_n\})) & \text{if } Q \in \Gamma', \\ 0 & \text{otherwise.} \end{cases}$$

The following picture shows an example of this construction. In this example we have assumed for simplicity that all edges of the graph have length 1 so that $\Gamma_{\mathbb{Z}}$ is just the set of vertices. The distance from P_5 to P is smallest among all distances from the P_i to a vertex, and the subset $\Gamma' \subset \Gamma$ is drawn in bold.

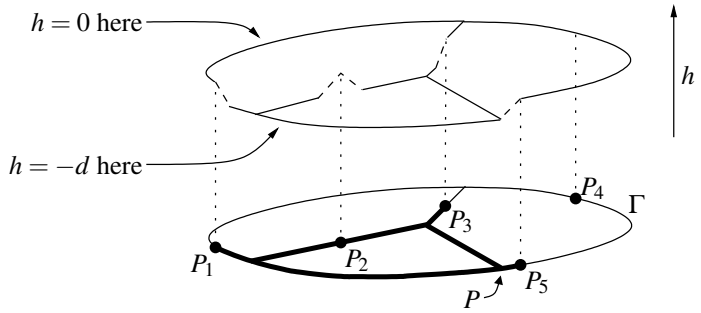

We claim that $f + h \in R(D - P)$. In fact, this will prove the lemma since $R(D - P) \neq \emptyset$ implies $\tilde{R}(D - P) \neq \emptyset$ and thus also $\tilde{R}(D) \neq \emptyset$ by the induction assumption.

To prove that $f + h \in R(D - P)$ we have to show that $(f + h) + D - P \geq 0$, or in other words that $(h) + P_1 + \cdots + P_n - P \geq 0$. Let us assume that this statement is false, i.e. that there is a point $Q \in \Gamma$ that is contained in the divisor $(h) + P_1 + \cdots + P_n - P$ with a negative coefficient. Note that Q cannot be the point P since $\operatorname{ord}_P h \geq 1$ by construction. So Q must be a pole of h. But again by construction h can only have poles at the points P_i, and the order of the poles can be at most 2 since the slope of h is 0 or ± 1 everywhere. So the only possibility is that Q is a point with $\operatorname{ord}_Q h = -2$ that occurs only once among the P_i (as it is the case for $Q = P_2$ in the example above). But this means that Γ' contains both sides of Q, and thus (since Γ' is connected) that $\Gamma' \cup \{Q\}$ contains a cycle on which Q is the only point

in (f) that is not a $\mathbb{Z}$-point. But this is a contradiction to lemma 2.2 and hence finishes the proof of the lemma. □

Proposition 2.4. *Let D be a $\mathbb{Z}$-divisor on a $\mathbb{Z}$-graph Γ. Then there is an integer $N \geq 1$ such that $r(D) = \tilde{r}(D)$ on every rescaling of Γ by an integer multiple of N (see example 1.13 (c)).*

Proof. Let $n := \deg D$ and $m := r(D)+1$, and assume first that $m \leq n$. Consider the map $\pi_m : S(D) \to \Gamma^m$ of remark 1.15. As π_m is a morphism of polyhedral complexes its image is closed in Γ^m. Since π_m is not surjective by remark 1.15 this means that $\Gamma^m \backslash \pi_m(S(D))$ is a non-empty open subset of Γ^m that consequently must contain an element $(P_1, \dots, P_m)$ with rational coordinates. For this element we have $S(D - P_1 - \cdots - P_m) = \emptyset$ by construction.

Now let N be the least common multiple of the denominators of these coordinates. Then $P_1, \dots, P_m$ become $\mathbb{Z}$-points on each rescaling of Γ by a multiple of N, and thus we also have $\tilde{S}(D - P_1 - \cdots - P_m) = \emptyset$ on each such rescaling by remark 1.7 (c). By definition this then means that $\tilde{r}(D) \leq m - 1 = r(D)$ on these rescalings. This proves the "$\tilde{r}(D) \leq r(D)$" part of the proposition in the case $m \leq n$. But note that this part is trivial if $m > n$, since then $\tilde{r}(D) \leq n \leq m - 1 = r(D)$ by example 1.13 (a) (on any rescaling). So we have in fact proven the "$\tilde{r}(D) \leq r(D)$" part of the proposition in any case.

To show the opposite inequality "$\tilde{r}(D) \geq r(D)$" (which in fact holds for any rescaling) we just have to show that $\tilde{R}(D - P_1 - \cdots - P_{r(D)}) \neq \emptyset$ for any choice of $\mathbb{Z}$-points $P_1, \dots, P_{r(D)}$. But this now follows immediately from lemma 2.3 since $R(D - P_1 - \cdots - P_{r(D)}) \neq \emptyset$ by definition. □

We are now ready to prove the Riemann-Roch theorem for $\mathbb{Q}$-divisors on $\mathbb{Q}$-graphs.

Corollary 2.5 (Riemann-Roch for $\mathbb{Q}$-graphs)**.** *Let D be a $\mathbb{Q}$-divisor on a $\mathbb{Q}$-graph Γ. Then $r(D) - r(K_\Gamma - D) = \deg D + 1 - g$.*

Proof. Note that it suffices by example 1.13 (c) to prove the statement after a rescaling of the curve.

As Γ has only finitely many edges and D contains only finitely many points we can assume after such a rescaling that D is in fact a $\mathbb{Z}$-divisor on a $\mathbb{Z}$-graph Γ, and that all edge lengths of Γ are bigger than 1. By proposition 2.4 we can then assume after possibly two more rescalings that both $r(D) = \tilde{r}(D)$ and $r(K_\Gamma - D) = \tilde{r}(K_\Gamma - D)$. The corollary now follows from theorem 2.1. □

3. Riemann-Roch for tropical curves

We will now extend our Riemann-Roch theorem for $\mathbb{Q}$-graphs (corollary 2.5) in two steps, first to metric graphs (i.e. graphs whose edge lengths need not be rational numbers) and then to tropical curves (i.e. graphs with possibly unbounded edges).

Proposition 3.1 (Riemann-Roch for metric graphs)**.** *For any divisor D on a metric graph Γ of genus g we have $r(D) - r(K_\Gamma - D) = \deg D + 1 - g$.*

Proof. Let $D = a_1 Q_1 + \cdots + a_m Q_m$, and let $n = \deg D$. The idea of the proof is to find a "nearby" $\mathbb{Q}$-graph Γ' with a $\mathbb{Q}$-divisor D' on it such that $r_{\Gamma'}(D') = r_\Gamma(D)$ and $r_{\Gamma'}(K_{\Gamma'} - D') = r(K_\Gamma - D)$, and then to apply the result of corollary 2.5 to this case.

To do so we will set up a relative version of the spaces $S(D)$ of definition 1.6 and the interpretation of $r(D)$ of remark 1.15 in terms of these spaces. We fix $\varepsilon \in \mathbb{Q}_{>0}$ smaller than all edge lengths of Γ and denote by $A(\Gamma)$ the set of all metric graphs that are of the same combinatorial type as Γ and all of whose edge lengths are greater or equal to ε. For such a metric graph $\Gamma' \in A(\Gamma)$ we denote by $B(\Gamma')$ the set of all divisors on Γ' that can be written as $a_1Q'_1 + \cdots + a_mQ'_m$ for some $Q'_1, \ldots, Q'_m \in \Gamma'$ and the same $a_1, \ldots, a_m$ as in D. With these notations we set

$$\begin{aligned}
S &:= \{(\Gamma', D', f, P_1, \ldots, P_n);\ \Gamma' \in A(\Gamma),\ D' \in B(\Gamma'),\ f \text{ a rational function on } \Gamma', \\
&\qquad\qquad P_1, \ldots, P_n \in \Gamma' \text{ such that } (f) + D' = P_1 + \cdots + P_n\}, \\
M_i &:= \{(\Gamma', D', P_1, \ldots, P_i);\ \Gamma' \in A(\Gamma),\ D' \in B(\Gamma'),\ P_1, \ldots, P_i \in \Gamma'\} \qquad \text{for } i = 0, \ldots, n, \\
M &:= \{(\Gamma', D');\ \Gamma' \in A(\Gamma),\ D' \in B(\Gamma')\}
\end{aligned}$$

In the same way as in lemma 1.9 we see that all these spaces are polyhedral complexes — the only difference is that there is some more discrete data (corresponding to fixing the edges or vertices on which the points in D' lie) and some more continuous data (corresponding to the edge lengths of Γ' and the positions of the points in D' on their respective edges). There are obvious forgetful morphisms of polyhedral complexes (i.e. continuous maps that send each cell of the source to a single cell of the target by an affine linear map)

$$\pi_i : S \to M_i, \quad (\Gamma', D', f, P_1, \ldots, P_n) \mapsto (\Gamma', D', P_1, \ldots, P_i)$$

and

$$p_i : M_i \to M, \quad (\Gamma', D', P_1, \ldots, P_i) \mapsto (\Gamma', D').$$

As in remark 1.15 we have $r_{\Gamma'}(D') \geq i$ for a divisor $D' \in B(\Gamma')$ on a metric graph $\Gamma' \in A(\Gamma)$ if and only if $\pi_i(S)$ contains $(\Gamma', D', P_1, \ldots, P_i)$ for all $P_1, \ldots, P_i \in \Gamma'$, or equivalently if and only if $(\Gamma', D') \in M \backslash p_i(M_i \backslash \pi_i(S))$.

Since S is a polyhedral complex and π_i a morphism of polyhedral complexes it follows that the image $\pi_i(S) \subset M_i$ is a union of closed polyhedra. Consequently, $M_i \backslash \pi_i(S)$ is a union of open polyhedra (i.e. an open subset of M_i whose intersection with each polyhedron of M_i can be written as a union of spaces given by finitely many strict linear inequalities).

Next, note that the map p_i is open as it is locally just a linear projection. It follows that $p_i(M_i \backslash \pi_i(S))$, i.e. the locus in M of all (Γ', D') such that $r_{\Gamma'}(D') < i$, is a union of open polyhedra as well. Consequently, its complement $M \backslash p_i(M_i \backslash \pi_i(S))$, i.e. the locus in M of all (Γ', D') such that $r_{\Gamma'}(D') \geq i$, is a union of closed polyhedra. Finally, note that all polyhedral complexes and morphisms involved in our construction are defined over $\mathbb{Q}$, so that the locus of all (Γ', D') with $r_{\Gamma'}(D') < i$ (resp. $r_{\Gamma'}(D') \geq i$) is in fact a union of *rational* open (resp. closed) polyhedra in M. Of course, the same arguments hold for $r_{\Gamma'}(K_{\Gamma'} - D')$ as well.

We are now ready to finish the proof of the proposition. By what we have said above the locus of all (Γ', D') in M such that $r_{\Gamma'}(D') < r_\Gamma(D) + 1$ and $r_{\Gamma'}(K_{\Gamma'} - D') < r_\Gamma(K_\Gamma - D) + 1$ is an open neighborhood U of (Γ, D). Conversely, the locus of all (Γ', D') in M such that $r_{\Gamma'}(D') \geq r_\Gamma(D)$ and $r_{\Gamma'}(K_{\Gamma'} - D') \geq r_\Gamma(K_\Gamma - D)$ is a union V of rational closed polyhedra. In particular, this means that the rational points of V are dense in V. As $U \cap V$ is non-empty (it contains the point (Γ, D)) it follows that there is a rational point in $U \cap V$, i.e. a $\mathbb{Q}$-graph Γ' with a $\mathbb{Q}$-divisor D' on it such that $r_{\Gamma'}(D') = r_\Gamma(D)$ and $r_{\Gamma'}(K_{\Gamma'} - D') = r_\Gamma(K_\Gamma - D)$. As Γ' and Γ have the same genus, and D' and D the same degree, the proposition now follows from corollary 2.5. □

So far we have only considered metric graphs, i.e. tropical curves in which every edge is of finite length. In our final step of the proof of the Riemann-Roch theorem we will now extend this result to arbitrary tropical curves (with possibly infinite edges). In order to do this we will first introduce the notion of equivalence of divisors.

Definition 3.2. Two divisors D and D' on a tropical curve Γ are called *equivalent* (written $D \sim D'$) if there exists a rational function f on Γ such that $D' = D + (f)$.

Remark 3.3. If $D \sim D'$, i.e. $D' = D + (f)$ for a rational function f, then it is obvious that the map $R(D') \to R(D)$, $g \mapsto g + f$ is a bijection. In particular, this means that $r(D) = r(D')$, i.e. that the function $r : \mathrm{Div}(\Gamma) \to \mathbb{Z}$ depends only on the equivalence class of D.

Lemma 3.4. *Let $\bar{\Gamma}$ be a tropical curve, and let Γ be the metric graph obtained from $\bar{\Gamma}$ by removing all unbounded edges. Then every divisor $D \in \mathrm{Div}(\bar{\Gamma})$ is equivalent on $\bar{\Gamma}$ to a divisor D' with $\mathrm{supp}\, D' \subset \Gamma$. Moreover, if D is effective then D' can be chosen to be effective as well.*

Proof. To any $P \in \bar{\Gamma}$, we associate a rational function f_P as follows. If $P \in \Gamma$, we define f_P to be the zero function. Otherwise, if P lies on some unbounded edge E, we define

$$f_P : \bar{\Gamma} \to \mathbb{R} \cup \{\infty\}, \quad Q \mapsto \begin{cases} \min(\mathrm{dist}(P,\Gamma), \mathrm{dist}(Q,\Gamma)) & \text{if } Q \in E, \\ 0 & \text{if } Q \notin E. \end{cases}$$

If $P \notin \Gamma$, then the function f_P has a simple pole at P and no other zeros or poles away from Γ. The following picture shows an example of such a function, where the metric graph Γ is drawn in bold:

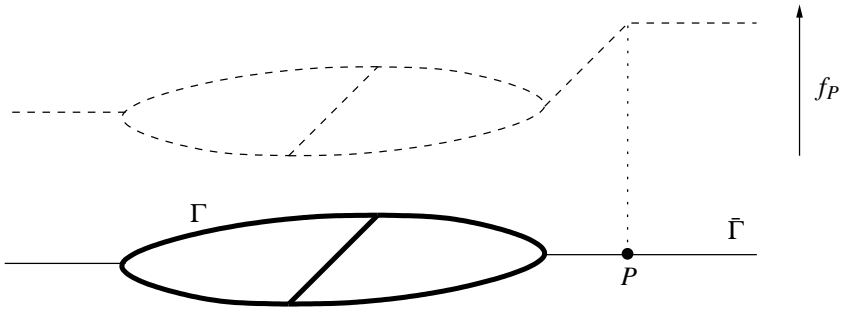

So if $D = a_1 P_1 + \cdots + a_n P_n$ and we set $f = \sum_i a_i f_{P_i}$ then $D + (f)$ is a divisor equivalent to D with no zeros or poles away from Γ. Moreover, if D is effective then $D + (f)$ is effective as well since all poles of f are cancelled by D by construction. □

Remark 3.5. With notations as above, let $P_1, \cdots, P_n$ denote the end points of the unbounded edges E_i of $\bar{\Gamma}$, and consider the function $f = \sum_i f_{P_i}$. Then f is zero on the graph Γ and has slope one on each unbounded edge. If we denote for all $i \in \{1, \ldots, n\}$ the point $E_i \cap \Gamma$ by Q_i, then $(f) = \sum Q_i - \sum P_i$. Hence $K_\Gamma + (f) = K_{\bar{\Gamma}}$, i.e. $K_\Gamma \sim K_{\bar{\Gamma}}$ on $\bar{\Gamma}$.

Lemma 3.6. *As in the previous lemma let $\bar{\Gamma}$ be a tropical curve, and let Γ be the metric graph obtained from $\bar{\Gamma}$ by removing all unbounded edges. Moreover, let D be a divisor on Γ (that can then also be thought of as a divisor on $\bar{\Gamma}$ with support on Γ). Then $R_\Gamma(D) \neq \emptyset$ if and only if $R_{\Gamma'}(D) \neq \emptyset$.*

Proof. "$\Rightarrow$": Let f be a rational function in $R_\Gamma(D)$. Extend f to a rational function $\bar{f}$ on $\bar{\Gamma}$ so that it is constant on each unbounded edge. Then $\bar{f} \in R_{\bar{\Gamma}}(D)$.

"$\Leftarrow$": Let $\bar{f} \in R_{\bar{\Gamma}}(D)$, and set $f = \bar{f}|_\Gamma$. Let e be an unbounded edge of $\bar{\Gamma}$, and let $P = \Gamma \cap e$ be the vertex where e is attached to Γ. Since $\bar{f}$ has no poles on e it follows that $\bar{f}|_e$ is (not necessarily strictly) decreasing if we identify e with the real interval $[0,\infty]$. Hence the order of f on Γ at P cannot be less than the order of $\bar{f}$ on $\bar{\Gamma}$ at P, and so it follows that $f \in R_\Gamma(D)$. □

Remark 3.7. Let $\bar{\Gamma}$, Γ, and D as in lemma 3.6. By lemma 3.4 any effective divisor $P_1 + \cdots + P_k$ on $\bar{\Gamma}$ is equivalent to an effective divisor $P'_1 + \cdots + P'_k$ with support on Γ. So by remark 3.3 the number $r_{\bar{\Gamma}}(D)$ can also be thought of as the biggest integer k such that $R_{\bar{\Gamma}}(D - P_1 - \cdots - P_k) \neq \emptyset$ for all $P_1, \ldots, P_k \in \Gamma$ (instead of for all $P_1, \ldots, P_k \in \bar{\Gamma}$). By lemma 3.6 we can therefore conclude that $r_{\bar{\Gamma}}(D) = r_\Gamma(D)$.

With these results we are now able to prove our main theorem:

Corollary 3.8 (Riemann-Roch for tropical curves)**.** *For any divisor D on a tropical curve $\bar{\Gamma}$ of genus g we have $r(D) - r(K_{\bar{\Gamma}} - D) = \deg D + 1 - g$.*

Proof. Let Γ be the metric graph obtained from $\bar{\Gamma}$ by removing all unbounded edges. By lemma 3.4 and remark 3.3 we may assume that $\operatorname{supp} D \subset \Gamma$. Moreover, by remark 3.5 we can replace $K_{\bar{\Gamma}}$ by K_Γ (which also has support in Γ) in the Riemann-Roch equation. Finally, remark 3.7 now tells us that we may replace $r_{\bar{\Gamma}}(D)$ and $r_{\bar{\Gamma}}(K_\Gamma - D)$ by $r_\Gamma(D)$ and $r_\Gamma(K_\Gamma - D)$ respectively, so that the statement follows from proposition 3.1. □

REFERENCES

[BN] M. Baker, S. Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph*, Adv. Math. (to appear), preprint math.CO/0608360.

[GM] A. Gathmann, H. Markwig, *Kontsevich's formula and the WDVV equations in tropical geometry*, preprint math.AG/0509628.

[MZ] G. Mikhalkin, I. Zharkov, *Tropical curves, their Jacobians and Theta functions*, preprint math.AG/0612267.

[Z] S. Zhang, *Admissible pairing on a curve*, Invent. Math. **112** (1993), 171–193.

ANDREAS GATHMANN, FACHBEREICH MATHEMATIK, TU KAISERSLAUTERN, POSTFACH 3049, 67653 KAISERSLAUTERN, GERMANY

E-mail address: andreas@mathematik.uni-kl.de

MICHAEL KERBER, FACHBEREICH MATHEMATIK, TU KAISERSLAUTERN, POSTFACH 3049, 67653 KAISERSLAUTERN, GERMANY

E-mail address: mkerber@mathematik.uni-kl.de