

CANONICAL REPRESENTATIVES FOR DIVISOR CLASSES ON TROPICAL CURVES AND THE MATRIX–TREE THEOREM

YANG AN¹, MATTHEW BAKER²,
GREG KUPERBERG³ and FARBOD SHOKRIEH⁴

¹ Columbia University, USA;

email: yangan@math.columbia.edu

² Georgia Institute of Technology, USA;

email: mbaker@math.gatech.edu

³ University of California, Davis, USA;

email: greg@math.ucdavis.edu

⁴ Cornell University, USA;

email: farbod@math.cornell.edu

Received 2 May 2013; accepted 9 September 2014

Abstract

Let Γ be a compact tropical curve (or *metric graph*) of genus g . Using the theory of tropical theta functions, Mikhalkin and Zharkov proved that there is a canonical effective representative (called a *break divisor*) for each linear equivalence class of divisors of degree g on Γ . We present a new combinatorial proof of the fact that there is a unique break divisor in each equivalence class, establishing in the process an ‘integral’ version of this result which is of independent interest. As an application, we provide a ‘geometric proof’ of (a dual version of) Kirchoff’s celebrated matrix–tree theorem. Indeed, we show that each weighted graph model G for Γ gives rise to a canonical polyhedral decomposition of the g -dimensional real torus $\text{Pic}^g(\Gamma)$ into parallelotopes C_T , one for each spanning tree T of G , and the dual Kirchoff theorem becomes the statement that the volume of $\text{Pic}^g(\Gamma)$ is the sum of the volumes of the cells in the decomposition.

2010 Mathematics Subject Classification: 05A19, 05C25, 05E45, 14T05

1. Introduction

Let Γ be a compact tropical curve (or *metric graph*) of genus g . There is a canonical continuous map $\pi : \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ taking an effective divisor

© The Author(s) 2014. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

of degree g on Γ to its linear equivalence class. Using tropical theta functions, Mikhalkin and Zharkov [MZ08] showed that there is a canonical continuous section σ to the map π . In particular, every divisor class of degree g has a canonical effective representative. (This is in sharp contrast to the situation for compact Riemann surfaces, where the analogous map π does not admit such a section.) Although not stated explicitly in the paper, one can deduce easily from the results in [MZ08] that the image of σ is the set of *break divisors* in $\text{Div}_+^g(\Gamma)$ (combine Theorem 6.5, Corollary 6.6, and Lemma 8.3 from [MZ08]).

In this paper, we study break divisors in detail, and give some applications. One application is a new combinatorial proof of the existence of the section σ which does not make use of tropical theta functions; this proof has the advantage that it yields an integral version of the Mikhalkin–Zharkov theorem. Another application is a ‘geometrization’ of Kirchhoff’s celebrated matrix–tree theorem: we show that, for each weighted graph model G for Γ , there is a canonical polyhedral decomposition of the g -dimensional real torus $\text{Pic}^g(\Gamma)$ into parallelotopes C_T , one for each spanning tree T of G ; from this point of view, Kirchhoff’s theorem (or rather its *matroid dual*) becomes the statement that the volume of $\text{Pic}^g(\Gamma)$ is the sum of the volumes of the cells in the decomposition.

In order to define break divisors, it is convenient to fix a model G for Γ . For each spanning tree T of G , let Σ_T (respectively Σ_T°) be the image of the canonical map $\prod_{e \notin T} \bar{e} \rightarrow \text{Div}_+^g(\Gamma)$ (respectively $\prod_{e \notin T} e^\circ \rightarrow \text{Div}_+^g(\Gamma)$) sending (p_1, \dots, p_g) to $(p_1) + \dots + (p_g)$. (Here, \bar{e} denotes a closed edge and e° denotes the corresponding open edge with the endpoints removed.) We call $\Sigma := \bigcup_T \Sigma_T$ the set of *break divisors* on Γ . It does not depend on the choice of the model G . Our first main result is a new combinatorial proof of the following theorem of Mikhalkin and Zharkov.

THEOREM 1.1. *Every degree g divisor on Γ is linearly equivalent to a unique break divisor.*

Since Σ is compact, and a continuous bijection between compact Hausdorff spaces is a homeomorphism, the theorem implies that π induces a homeomorphism from Σ onto its image. The canonical section σ is the inverse of this homeomorphism.

Our combinatorial proof of Theorem 1.1 is based on a study of *orientable divisors*. If \mathcal{O} is a (not necessarily acyclic) orientation of Γ , we define a corresponding divisor $D_{\mathcal{O}}$ of degree $g - 1$ by the formula

$$D_{\mathcal{O}} = \sum_{p \in \Gamma} (\text{indeg}_{\mathcal{O}}(p) - 1)(p).$$

We call divisors of this form *orientable*. We will show by a constructive procedure that every divisor of degree $g - 1$ is linearly equivalent to an orientable divisor. More precisely, fix $q \in \Gamma$. We say that an orientation is *q-connected* if there is a directed path from q to p for every $p \in \Gamma$. A divisor of the form $D_{\mathcal{O}}$, with \mathcal{O} a *q-connected* orientation, is called *q-orientable*. We will prove the following.

THEOREM 1.2. *Fix $q \in \Gamma$. Every divisor of degree $g - 1$ on Γ is linearly equivalent to a unique q -orientable divisor.*

There is a close connection between break divisors and *q-orientable* divisors. Indeed, we will see that the map sending a degree g divisor D to the degree $g - 1$ divisor $D - (q)$ induces a bijection between break divisors and *q-orientable* divisors. Using this observation, one deduces in a completely formal way that Theorems 1.1 and 1.2 are in fact equivalent.

One advantage of our approach to Theorems 1.1 and 1.2 is that it enables us to prove an *integral version* of Theorem 1.1 for finite unweighted graphs G (or equivalently, finite weighted graphs in which all edges have length 1). Indeed, if we define an *integral break divisor* to be a break divisor supported on the vertices of G , then the constructive procedure furnished by our proof of Theorem 1.2 shows the following.

THEOREM 1.3. *Every degree g divisor supported on the vertices of G is linearly equivalent to a unique integral break divisor.*

As a consequence of Theorem 1.3, the set $\text{Pic}^g(G)$ of divisors of degree g on G modulo linear equivalence is canonically in bijection with the set of integral break divisors. In particular, the number of integral break divisors is equal to the number of spanning trees of G .

We now turn to a more detailed discussion of the connection between break divisors, polyhedral decompositions of $\text{Pic}^g(\Gamma)$, and Kirchhoff's matrix–tree theorem. Let G be a weighted graph with associated metric graph Γ . Recall that $\Sigma := \bigcup_T \Sigma_T$, where the union is over all spanning trees T of G . If we define $\mathcal{C}_T = \pi(\Sigma_T)$, then clearly $\text{Pic}^g(\Gamma) = \bigcup_T \mathcal{C}_T$ by Theorem 1.1. It turns out that each cell \mathcal{C}_T is a parallelotope and that, if $T \neq T'$, then the relative interiors of \mathcal{C}_T and $\mathcal{C}_{T'}$ are disjoint. Thus $\text{Pic}^g(\Gamma)$ has a polyhedral decomposition depending only on the choice of the model G . It is not hard to check that, if G' is a refinement of G , then the cell decomposition coming from G' is a refinement of the cell decomposition coming from G .

EXAMPLE 1.4. Let Γ be the metric graph consisting of two vertices joined by three edges of lengths 2, 1, and 2, respectively. Fix a model G for Γ in which all

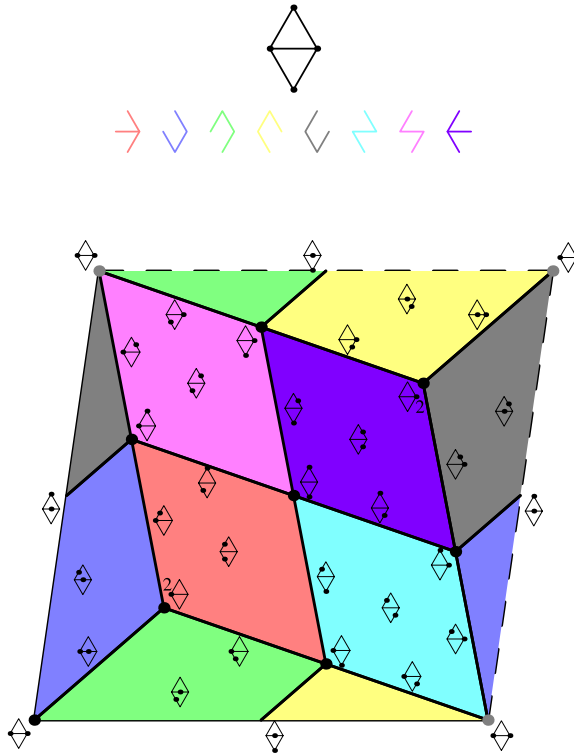


Figure 1. A fixed model for the metric graph Γ and the corresponding decomposition of $\text{Pic}^2(\Gamma)$.

edge lengths have length 1. In Figure 1, we have listed all spanning trees of G and the corresponding cell decomposition of $\text{Pic}^2(\Gamma)$. We have labeled the center of each cell with its corresponding break divisor.

Since the canonical map from $\Sigma = \cup \Sigma_T$ to $\text{Pic}^g(\Gamma)$ is a homeomorphism and \mathcal{C}_T is the projection to $\text{Pic}^g(\Gamma)$ of the cube Σ_T , it follows easily that, if T, T' are spanning trees of G , then the dimension of $\mathcal{C}_T \cap \mathcal{C}_{T'} \subset \text{Pic}^g(\Gamma)$ is equal to the dimension of $\Sigma_T \cap \Sigma_{T'}$ inside $\prod_{e \in E} \bar{e}$. In particular, $\mathcal{C}_T \cap \mathcal{C}_{T'}$ is nonempty if and only if $\prod_{e \notin T} \bar{e} \cap \prod_{e \notin T'} \bar{e}$ is nonempty.

Moreover, recall from [MZ08] that $\text{Pic}^g(\Gamma)$ is canonically a principal homogeneous space for the Picard group $\text{Pic}^0(\Gamma)$ and that there is a canonical isomorphism (the ‘tropical Abel–Jacobi map’) between $\text{Pic}^0(\Gamma)$ and $\text{Jac}(\Gamma) = H_1(\Gamma, \mathbb{R})/H_1(\Gamma, \mathbb{Z})$, which is a real torus of dimension g . The intersection form

on $H_1(\Gamma, \mathbb{Z})$ gives rise to a canonical translation-invariant Riemannian metric on $\text{Jac}(\Gamma)$, and via translation invariance $\text{Pic}^g(\Gamma)$ is equipped with a canonical metric as well. In particular, the volume of $\text{Pic}^g(\Gamma)$ and of each of the cells C_T is well defined.

THEOREM 1.5. *For each spanning tree T of G , the volume of the parallelotope C_T is $w(T)/\text{Vol}(\text{Jac}(\Gamma))$, where $w(T) := \prod_{e \notin E(T)} \ell(e)$ is the product of the lengths all edges of G not in T . Moreover, the volume of $\text{Jac}(\Gamma)$ is $\sqrt{\det(\mathfrak{C}\mathfrak{C}^t)}$, where \mathfrak{C} is any matrix whose rows form a \mathbb{Z} -basis for $H_1(G, \mathbb{Z})$.*

Since distinct cells C_T intersect in positive codimension, we have $\text{Vol}(\text{Jac}(\Gamma)) = \sum_T \text{Vol}(C_T)$, and thus Theorem 1.5 implies the following dual version of Kirchhoff's matrix-tree theorem.

COROLLARY 1.6. *For any weighted graph G ,*

$$\det(\mathfrak{C}\mathfrak{C}^t) = \sum_T w(T).$$

The usual version of Kirchhoff's matrix-tree theorem is (a special case of) the dual statement that, for any weighted graph G , we have

$$\det(\mathfrak{B}\mathfrak{B}^t) = \sum_T w'(T),$$

where $w'(T) := \prod_{e \in E(T)} \ell(e)$ is the product of the lengths all edges of G in T and \mathfrak{B} is any matrix whose rows form a \mathbb{Z} -basis for the cocycle lattice of G (which is the intersection of $C_1(G, \mathbb{Z})$ with the orthogonal complement of $H_1(G, \mathbb{R})$ in $C_1(G, \mathbb{R})$).

Note that the dual version of Kirchhoff's theorem, like the cycle lattice $H_1(G, \mathbb{Z})$, is unchanged if we replace G by a different model G' for the same metric graph Γ . This is not true of the usual version of Kirchhoff's theorem, or of the cocycle lattice.

2. Definitions and background

2.1. Metric graphs.

DEFINITION 2.1. A metric graph (or abstract tropical curve) Γ is a compact connected metric space such that every point $p \in \Gamma$ has a neighborhood isometric to a star-shaped set.

For the purposes of this paper, we define a *weighted graph* to be a finite connected graph G with no loop edges, together with a collection of positive weights L_e (which we also call *lengths*), one for each edge.

A weighted graph G gives rise to a metric space Γ in the following way. To each edge e , associate a line segment of length L_e , and identify the ends of distinct line segments if they correspond to the same vertex of G . The points of these line segments are the points of Γ . We call G a *model* for Γ . The distance between two points x and y in Γ is defined to be the length of the shortest path between them.

It is easy to check that every metric graph arises from this construction, though the weighted graph model is not unique. Write $G \sim G'$ if the two weighted graphs G, G' admit a *common refinement*, where we *refine* a weighted graph by subdividing its edges in a manner that preserves the total length. This yields an equivalence relation on the collection of weighted graphs, and two weighted graphs are equivalent if and only if their associated metric graphs are isometric. There is thus a bijective correspondence between metric graphs and equivalence classes of weighted graphs (see [BF11] for details).

2.2. Reduced divisors. Let $\text{Div}(\Gamma)$ (the group of *divisors* on Γ) be the free abelian group on Γ . An element of $\text{Div}(\Gamma)$ is of the form $D = \sum_{p \in \Gamma} a_p(p)$, where $a_p \in \mathbb{Z}$ and all but finitely many of the a_p are zero. The degree of D is by definition $\deg(D) = \sum_{p \in \Gamma} a_p$. Let $\text{Div}^0(\Gamma)$ be the subgroup of divisors of degree zero on Γ . A divisor $D = \sum a_p(p)$ is called *effective* if $a_p \geq 0$ for all p , and is called *effective outside q* if $a_p \geq 0$ for all $p \neq q$. The support of D is by definition $\text{Supp}(D) = \{p \in \Gamma \mid D(p) \neq 0\}$.

Let $R(\Gamma)$ be the group consisting of continuous piecewise affine functions with integer slopes. This can be viewed as the space of *tropical rational functions* on Γ ; see [GK08, MZ08]. Let Δ be the Laplacian operator on Γ ; for $f \in R(\Gamma)$, we have

$$\Delta(f) = \sum_{p \in \Gamma} \sigma_p(f)(p),$$

where $-\sigma_p(f)$ is the sum of the slopes of f in all tangent directions emanating from p . (A *tangent direction* at a point $p \in \Gamma$ is an equivalence class of paths emanating from p , where two paths are equivalent if they share a common initial segment.) Let $\text{Prin}(\Gamma)$ be the subgroup $\{\Delta(f) \mid f \in R(\Gamma)\}$ of $\text{Div}^0(\Gamma)$ consisting of *principal divisors*. We write $D \sim D'$ if $D - D'$ belongs to $\text{Prin}(\Gamma)$, and say that D and D' are *linearly equivalent*. For $D \in \text{Div}(\Gamma)$, we define the complete linear system $|D|$ to be the set of all effective divisors E equivalent to D ; that is, $|D| = \{E \in \text{Div}(\Gamma) \mid E \geq 0, E \sim D\}$. Similarly, we define $|D|_q$ to be the set of divisors equivalent to D which are effective outside q :

$$|D|_q = \{E \in \text{Div}(\Gamma) \mid E(p) \geq 0, \forall p \neq q, E \sim D\}.$$

REMARK 2.2. Given an effective divisor D , it is customary to think of $D(v)$ as the number of ‘chips’ placed at the point $v \in \Gamma$. For a subset X of Γ consisting of a finite union of closed intervals, one can construct a rational function $f \in R(\Gamma)$ which is 0 on X and ϵ outside an ϵ -neighborhood of X (for some sufficiently small ϵ), with slope 1 in each outgoing direction from X . Replacing D with $D + \Delta(f)$ has the effect of moving a chip a distance of ϵ along each outgoing direction from X , and is often called ‘firing’ the subset X . Every element of $R(\Gamma)$ can be written as a finite integer-affine combination of functions of this form, and therefore one can describe linear equivalence of divisors on Γ in terms of chip firing.

DEFINITION 2.3. Fix $q \in \Gamma$. A divisor D on Γ is called q -reduced if it satisfies the following two conditions.

(R1) D is effective outside q .

(R2) For every closed connected set $X \subseteq \Gamma$ not containing q , there exists a point $p \in \partial X$ such that $D(p) < \text{outdeg}_X(p)$.

The following is the metric graph analog of [BS13, Lemma 4.11].

LEMMA 2.4 (Principle of least action). *If D is q -reduced and $f \in R(\Gamma)$ is a tropical rational function with $D + \Delta(f) \in |D|_q$, then f has a global minimum at q .*

Proof. Consider the set of points where f achieves its (global) minimum value. If this set is not $\{q\}$, then we may find a closed connected component X not containing q . By Definition 2.3, there exists $p \in \partial X$ such that $D(p) < \text{outdeg}_X(p)$. On the other hand, we have $\Delta(f)(p) < -\text{outdeg}_X(p)$, and thus $(D + \Delta(f))(p) < 0$, contradicting the assumption that $D + \Delta(f) \in |D|_q$. \square

The importance of reduced divisors is given by the following theorem (see, for example, [HKN13, Luo11]).

THEOREM 2.5. *Fix $q \in \Gamma$. There is a unique q -reduced divisor in each linear equivalence class of divisors on Γ .*

2.3. Tropical Picard group, Jacobian, and Abel–Jacobi map. The (degree 0) Picard group of Γ is by definition the quotient

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

More generally, for $d \geq 0$, let $\text{Div}^d(\Gamma)$ be the subset of divisors of degree d on Γ , and define

$$\text{Pic}^d(\Gamma) := \text{Div}^d(\Gamma) / \text{Prin}(\Gamma).$$

The set $\text{Pic}^d(\Gamma)$ is canonically a $\text{Pic}^0(\Gamma)$ -torsor.

The tropical Abel–Jacobi theory identifies (as topological groups) $\text{Pic}^0(\Gamma)$ with the **Jacobian torus**

$$\text{Jac}(\Gamma) = \Omega^*(\Gamma) / H_1(\Gamma, \mathbb{Z}).$$

Here, $\Omega(\Gamma)$ denotes the space of *harmonic 1-forms* on Γ . A harmonic 1-form on Γ is obtained by assigning a real-valued slope to each edge in Γ in such a way that the sum of the incoming slopes is zero at every vertex. The homology group $H_1(\Gamma, \mathbb{Z})$ embeds as a lattice in $\Omega^*(\Gamma)$ (the dual vector space of $\Omega(\Gamma)$) by integration of 1-forms along 1-cycles. There is a canonical identification of $\Omega(\Gamma)^*$ with the singular cohomology space $H_1(\Gamma, \mathbb{R})$, so we will often write

$$\text{Jac}(\Gamma) = H_1(\Gamma, \mathbb{R}) / H_1(\Gamma, \mathbb{Z}).$$

The vector space $H_1(\Gamma, \mathbb{R})$ is equipped with a natural translation-invariant Riemannian metric which induces a canonical metric, and in particular a canonical volume form, on the quotient torus $\text{Jac}(\Gamma)$.

For each positive integer d , the corresponding **Abel–Jacobi map**

$$S^{(d)} : \text{Div}_+^d(\Gamma) \rightarrow \text{Pic}^d(\Gamma)$$

sends a divisor $D \in \text{Div}_+^d(\Gamma)$ to the divisor class $[D] \in \text{Pic}^d(\Gamma)$.

We will denote $S^{(g)}$ by π .

Fixing a basepoint $q \in \Gamma$, the torus $\text{Pic}^0(\Gamma)$ is identified with the group $\text{Pic}^0(\Gamma)$ and $S^{(d)}$ is identified with the map

$$S_q^{(d)} : \text{Div}_+^d(\Gamma) \rightarrow \text{Pic}^0(\Gamma)$$

taking a divisor $D \in \text{Div}_+^d(\Gamma)$ to the divisor class $[D - d(q)] \in \text{Pic}^0(\Gamma)$.

All maps $S^{(d)}$ are *piecewise linear* in the appropriate sense; fixing the basepoint q , the real vector space $\Omega^*(\Gamma)$ is identified with the universal cover of $\text{Pic}^d(\Gamma)$, and the restriction of the Abel–Jacobi map to any contractible subset factors through a piecewise-linear map to $\Omega^*(\Gamma)$. Note that $\text{Div}_+^d(\Gamma)$ is endowed with a natural integral affine structure induced from Γ (see, for example, [Ami13, Section 2.1]).

3. Spanning trees and divisors

3.1. Break divisors, break pairs, fundamental domains, and orientations.

Recall from Section 1 that a *break divisor* on Γ is a divisor of the form




Figure 2. Examples of break divisors.




Figure 3. Examples of break sets.

$(p_1) + \dots + (p_g)$, where $p_i \in \bar{e}_i$ and $\Gamma \setminus T = \bigcup_{i=1}^g e_i^\circ$ for some spanning tree T of G . (Here, G is any fixed model of Γ). See Figure 2.

Here is another way to characterize break divisors, following [MZ08]. If p is a vertex of some model G for Γ , there is a natural bijection between tangent directions at p and edges of G incident to p . For $\varepsilon > 0$ sufficiently small, there is a well-defined point $p + \varepsilon\eta$ at distance ε from p in the direction η . Let p_1, \dots, p_g be (not necessarily distinct) points on Γ , and for each i let η_i be a tangent direction at p_i . If $\Gamma \setminus \{q_1, \dots, q_g\}$ is a tree (that is, is connected and simply connected) for $\varepsilon > 0$ sufficiently small, where $q_i = p_i + \varepsilon\eta_i$, we call the collection $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$ a *break set* for Γ . See Figure 3.

LEMMA 3.1. *If $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$ is a break set, then $(p_1) + \dots + (p_g)$ is a break divisor. Conversely, if $(p_1) + \dots + (p_g)$ is a break divisor, then there exist (not necessarily unique) tangent directions η_i at p_i such that $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$ is a break set.*

Proof. Let $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$ be a break set, and let $\varepsilon > 0$ be sufficiently small that $\Gamma \setminus \{q_1, \dots, q_g\}$, where $q_i = p_i + \varepsilon\eta_i$, is a tree and all the q_i have valence 2. Fix a model G for Γ such that $p_i \in V(G)$ and $q_i \notin V(G)$ for all i . It follows that $\Gamma \setminus \{q_1, \dots, q_g\}$ contains a spanning tree T of G , and that $p_i \in \bar{e}_i$, where $\Gamma \setminus T = \bigcup_{i=1}^g e_i^\circ$.




Figure 4. Examples of fundamental domains.

Conversely, assume that G is a model of Γ and that, for some spanning tree T of G , we have $\Gamma \setminus T = \bigcup_{i=1}^g e_i^\circ$ and that we are given a break divisor $(p_1) + \dots + (p_g)$, where $p_i \in \bar{e}_i$. If $p_i \in e_i^\circ$, we may choose η_i to be either of the two tangent directions at p_i . If $p_i \in \partial e_i = \bar{e}_i \setminus e_i^\circ$, then we choose the tangent direction directed toward e_i° . □

Because of the lemma, we will sometimes denote a break set by (D, η) , where $D = (p_1) + \dots + (p_g)$ is a break divisor and $\eta = \{\eta_1, \dots, \eta_g\}$ is the corresponding set of *break directions*.

REMARK 3.2. If we think of a metric graph as made of wires, we can think of a break set as a rule for snipping the wires so that the resulting network is connected and simply connected. More formally, define a *fundamental domain* for Γ to be a pair (F, ψ) consisting of a (not necessarily compact) connected and simply connected topological space F and a continuous bijection $\psi : F \rightarrow \Gamma$. We identify two fundamental domains (F, ψ) and (F', ψ') if there is a homeomorphism $\phi : F \rightarrow F'$ such that $\psi = \psi' \circ \phi$. Then one can show that there is a natural bijection between fundamental domains for Γ and break sets. (This is the point of view taken in [MZ08, Section 4.5].) See Figure 4.

An *orientation* \mathcal{O} of a metric graph Γ is an equivalence class of pairs (G, O) , where G is a model for Γ and O is an orientation of the edges of G , where the equivalence relation is generated by the operation of replacing G by a refinement G' and letting O' be the orientation induced by O .

For $q \in \Gamma$, we say that an orientation \mathcal{O} is *q-connected* if there is a directed path from q to p for every $p \in \Gamma$.

Given $q \in \Gamma$, there is a canonical way to associate a *q-connected* orientation \mathcal{O} to any break set $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$, as follows. Choose $\varepsilon > 0$ small enough so that, setting $q_i = p_i + \varepsilon \eta_i$, we have (i) $\Gamma \setminus \{q_1, \dots, q_g\}$ is a tree, (ii) all the q_i have valence 2, and (iii) $q \neq q_i$ for any i . We get an associated *q-connected* orientation \mathcal{O}_ε by orienting all edges away from q in this tree, and letting $\varepsilon \rightarrow 0$

gives the desired q -connected orientation \mathcal{O} . (In terms of the corresponding fundamental domain (F, ψ) described in Remark 3.2, \mathcal{O} corresponds to orienting all paths away from $\psi^{-1}(q)$ in the connected and simply connected space F .)

Conversely, there is a canonical way to associate a break divisor $D_{(\mathcal{O}, q)}$ to a q -connected orientation \mathcal{O} :

$$D_{(\mathcal{O}, q)} := (q) + D_{\mathcal{O}} = (q) + \sum_{p \in \Gamma} (\text{indeg}_{\mathcal{O}}(p) - 1)(p).$$

LEMMA 3.3. *Fix $q \in \Gamma$. The map $\phi_q : \text{Div}^g(\Gamma) \rightarrow \text{Div}^{g-1}(\Gamma)$ sending D to $D - (q)$ induces a bijection between break divisors and q -orientable divisors.*

Proof. Let $D = (p_1) + \dots + (p_g)$ be a break divisor, and consider an associated break set $\{(p_1, \eta_1), \dots, (p_g, \eta_g)\}$ as in Lemma 3.1. Let \mathcal{O} be the associated q -connected orientation as above.

Let $p \in \Gamma$. If $p \neq q$ then $\text{indeg}_{\mathcal{O}}(p) = 1 + s(p)$, where $s(p)$ is the number of tangent directions η_i with $p_i = p$, and if $p = q$ then $\text{indeg}_{\mathcal{O}}(p) = s(p)$. Thus, for every $p \in \Gamma$, the coefficient of (p) in $D_{\mathcal{O}} + (q)$ is equal to $s(p)$, which is equal to $D(p)$ by construction. This proves that ϕ_q induces a map from break divisors to q -orientable divisors.

The map ϕ_q is clearly injective. To see that it is surjective, suppose that $D_{\mathcal{O}}$ is a q -orientable divisor corresponding to a q -connected orientation \mathcal{O} . We will equip the effective divisor $D_{\mathcal{O}} + (q)$ with a set of tangent directions η so that $(D_{\mathcal{O}} + (q), \eta)$ is a break set. By *breaking* an edge e adjacent to p , we will mean adding the tangent direction at p which corresponds to the edge e to the set η . If $\text{indeg}_{\mathcal{O}}(q) \geq 1$, break all the incoming edges at q . For each $p \neq q$ with $\text{indeg}_{\mathcal{O}}(p) \geq 2$, break all but one of the incoming edges at p . (The unbroken edge can be chosen arbitrarily.) Then one easily checks that $|\eta| = g$ and $(D_{\mathcal{O}} + (q), \eta)$ is a break set. \square

Note that the set of break directions is not uniquely determined by \mathcal{O} , so we do not get a bijection between q -connected orientations and break sets (see Figure 5).

As a formal consequence of Lemma 3.3, we have the following.

COROLLARY 3.4. *Theorems 1.1 and 1.2 are equivalent.*

3.2. The canonical cell decomposition of $\text{Pic}^g(\Gamma)$. Let G be a model for Γ . Given any spanning tree T of G , we let $\Sigma_T^o \subset \Sigma$ be the product of the interiors of all edges of Γ not in T :

$$\Sigma_T^o := \prod_{e \notin E(T)} e^o.$$




Figure 5. Two different break sets, their associated fundamental domains, and their identical associated q -connected orientation.

An element of Σ_T^o determines the spanning tree T of G uniquely. Therefore for distinct spanning trees T and T' of G we have $\Sigma_T^o \cap \Sigma_{T'}^o = \emptyset$. Since $\{e^o\}_{e \notin E(T)}$ consists of mutually disjoint segments, no two elements of Σ_T^o are identified under the action of the symmetric group, and we may consider Σ_T^o as an (open) *parallelootope* inside $\text{Div}_+^g(\Gamma) = \Gamma^{(g)}$.

We let $\Sigma^o \subset \text{Div}_+^g(\Gamma)$ be the (disjoint) union of the sets Σ_T^o for all spanning trees T of G :

$$\Sigma^o := \bigcup_T \Sigma_T^o.$$

Note that Σ^o , unlike Σ , depends on the choice of the model G .

Divisors in Σ^o have the following useful property.

LEMMA 3.5. *For any $D \in \Sigma^o$, we have $|D| = \{D\}$.*

Proof. The proof is an application of the least action principle (Lemma 2.4). Let $D \in \Sigma_T^o$ for some spanning tree T of the model G . Pick $p, q \in \Gamma$. Then one sees easily that D is both p -reduced and q -reduced. If $E = D + \Delta(f) \in |D|$ then by the least action principle we get $f(p) \leq f(q)$ and $f(q) \leq f(p)$. Thus f is constant, and $E = D$. \square

Recall that the map $\pi : \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ is piecewise linear. It follows that the open cell $\mathcal{C}_T^o := \pi(\Sigma_T^o)$ is also an (open) parallelotope inside the torus $\text{Pic}^g(\Gamma)$. By Lemma 3.5, the restriction of π to each Σ_T^o is injective with image \mathcal{C}_T^o . Moreover, for distinct spanning trees T and T' of G , we have $\mathcal{C}_T^o \cap \mathcal{C}_{T'}^o = \emptyset$ because $\Sigma_T^o \cap \Sigma_{T'}^o = \emptyset$. It follows that the restriction of the map π to $\Sigma^o \subset \text{Div}_+^g(\Gamma)$ is injective.

Let $\mathcal{C}^o \subset \text{Pic}^g(\Gamma)$ be the (disjoint) union of the sets \mathcal{C}_T^o for all spanning trees T of G . Let $\mathcal{C} := \overline{\mathcal{C}^o} = \bigcup_T \overline{\mathcal{C}_T^o}$ and $\overline{\mathcal{C}_T} := \overline{\mathcal{C}_T^o}$ denote the topological closures of \mathcal{C}^o and \mathcal{C}_T^o , respectively, inside $\text{Pic}^g(\Gamma)$. By Theorem 1.1 (stated in Section 1 and proved in Section 4 below), we have $\mathcal{C} = \text{Pic}^g(\Gamma)$. Since $\Sigma = \bigcup \Sigma_T$ and $\text{Pic}^g(\Gamma)$ are compact Hausdorff spaces, it also follows from Theorem 1.1 that the canonical map from Σ to $\text{Pic}^g(\Gamma)$ is a homeomorphism.

4. Orientations and divisors

4.1. Finite graphs. Assume that G is a finite connected multigraph. As usual, we denote the set of vertices by $V(G)$ and the set of edges by $E(G)$. In what follows, S will always denote a subset of $V(G)$. We denote by $G[S]$ the induced subgraph of G with vertex set S . Let $\chi(S)$ denote the topological Euler characteristic of $G[S]$, which is equal to the number of vertices of $G[S]$ minus the number of edges of $G[S]$. (If $G[S]$ is connected, then $\chi(S) = 1 - g(S)$, where $g(S)$ is the *genus*, or first Betti number, of $G[S]$.)

Given a divisor $D \in \text{Div}(G)$, we define

$$\chi(S, D) = \deg(D|_S) + \chi(S),$$

where $D|_S$ denotes the restriction of D to $G[S]$; that is, if $D = \sum_{p \in V(G)} a_p(p)$, then $D|_S = \sum_{p \in S} a_p(p)$. For $S = V(G)$, we write $\chi(G, D)$ instead of $\chi(V(G), D)$.

If S and T are disjoint subsets of $V(G)$, we define $e(S, T)$ to be the number of edges of G with one end in S and the other end in T . We define $e(S)$ to be the number of edges both of whose endpoints belong to S .

A *submodular function* is function f from the collection of subsets of $V(G)$ to \mathbb{R} such that, for any subsets S, T of $V(G)$, we have

$$f(S) + f(T) \geq f(S \cup T) + f(S \cap T). \quad (4.1)$$

If equality holds, then f is called *modular*. The following lemma is an immediate consequence of Equation (4.1).

LEMMA 4.2. *Subsets where a submodular function takes its minimum value are closed under intersection and union.*

LEMMA 4.3. *For any divisor D , the function $\chi(\cdot, D)$ is submodular.*

Proof. By definition, we have $\chi(S, D) = \deg(D|_S) + |S| - e(S)$. One easily checks that $f_1(S) = \deg(D|_S)$ and $f_2(S) = |S|$ are modular and that $f_3(S) = -e(S)$ is submodular, and the result follows. \square

More precisely, we have the following quantitative refinement of submodularity.

LEMMA 4.4. *For any subsets S, T of $V(G)$, we have*

$$\chi(S, D) + \chi(T, D) = \chi(S \cup T, D) + \chi(S \cap T, D) + e(S \setminus T, T \setminus S). \quad (4.5)$$

In particular, if $S \cap T = \emptyset$, then

$$\chi(S \cup T, D) = \chi(S, D) + \chi(T, D) - e(S, T). \quad (4.6)$$

Proof. Since we have $\chi(S, D) = \deg(D|_S) + |S| - e(S)$, and both $\deg(D|_S)$ and $|S|$ are modular functions, it suffices to prove that $e(S) + e(T) + e(S \setminus T, T \setminus S) = e(S \cap T) + e(S \cup T)$. This is a well-known stronger version of the submodularity of $e(\cdot)$. \square

For a given divisor $D \in \text{Div}(G)$, we define

$$\begin{aligned} \chi_D &= \min\{\chi(S, D) : \emptyset \neq S \subsetneq V(G)\}, \\ \mathbf{S}(D) &= \{\emptyset \neq S \subsetneq V(G) : \chi(S, D) = \chi_D\}. \end{aligned}$$

COROLLARY 4.7. *Let $D \in \text{Div}(G)$ be a divisor such that $\chi(G, D) \geq 0$ and $\chi_D < 0$. Then $\mathbf{S}(D)$ has a unique minimal element (with respect to inclusion).*

Proof. Since $\chi(S, D)$ is a submodular function, by Lemma 4.2 there is a unique minimal subset S for which $\chi(\cdot, D)$ takes on its minimum value, namely the

intersection of all such subsets. Since $\chi(\emptyset, D) = 0$ and $\chi(G, D) \geq 0$, while $\chi_D < 0$, we conclude that the unique minimal element is a proper nonempty subset of $V(G)$, and hence lies in $\mathbf{S}(D)$. \square

Recall from Section 1 that a divisor $D \in \text{Div}(G)$ is called *orientable* if there exists an orientation \mathcal{O} on G such that at every point $p \in V(G)$ we have $D(p) = \text{indeg}_{\mathcal{O}}(p) - 1$. It is easy to see that if D is orientable then $\text{deg}(D) = g - 1$, which is equivalent to saying that $\chi(G, D) = 0$. The following result strengthens this observation. Although we discovered this result independently, Backman and the anonymous referee both informed us that Theorem 4.8 is in fact a reformulation of a classical result of Hakimi [Hak65] (see also [Sch03, Theorem 61.1]). Backman shows in [Bac13, Theorem 7.3] that Theorem 4.8 is equivalent to the well-known max-flow min-cut theorem.

THEOREM 4.8 (Hakimi). *A divisor $D \in \text{Div}(G)$ is orientable if and only if $\chi(G, D) = 0$ and $\chi(S, D) \geq 0$ for all nonempty subsets S of $V(G)$. (Equivalently, D is orientable if and only if $\text{deg}(D) = g - 1$ and $\chi_D \geq 0$.)*

REMARK 4.9. It follows from Lemma 4.4 that one only needs to check the condition $\chi(S, D) \geq 0$ for subsets S whose induced subgraph is *connected*: if the condition is satisfied for all connected components of a set, it is automatically satisfied for the whole set by Equation (4.6).

THEOREM 4.10. *Every divisor $D \in \text{Div}^{g-1}(G)$ is linearly equivalent to some orientable divisor.*

Proof. If $D_0 = D$ is not orientable, then Theorem 4.8 guarantees that there exists a subset $\emptyset \neq S \subsetneq V(G)$ with $\chi(S, D) < 0$. It follows from Corollary 4.7 that there is a unique minimal element (with respect to inclusion) in $\mathbf{S}(D_0) = \{\emptyset \neq S \subsetneq V(G) : \chi(S, D_0) \text{ is minimal}\}$. In other words, there is a unique nonempty set $S_0 \subsetneq V(G)$ such that $\chi(S_0, D_0)$ is minimal and which is contained in any other vertex set S with this property.

Let D_1 be the divisor obtained from D_0 by simultaneously firing all vertices in \bar{S}_0 .

CLAIM. *For all $\emptyset \neq S \subsetneq V(G)$, we have $\chi(S, D_1) \geq \chi(S_0, D_0)$. Moreover, if equality holds, then $S \supseteq S_0$.*

To prove this claim, we consider the following cases.

- (1) $S \subsetneq S_0$. In this case, $\chi(S, D_0) > \chi(S_0, D_0)$ (by the minimality of S_0) and $\chi(S, D_1) \geq \chi(S, D_0)$ (by the construction of D_1).

(2) $S = S_0$. In this case, it follows from definitions that $\chi(S, D_1) = \chi(S_0, D_0) + e(S_0, \bar{S}_0) > \chi(S_0, D_0)$.

(3) $S \subseteq \bar{S}_0$. In this case, it follows from definitions and Equation (4.6) that

$$\chi(S, D_1) = \chi(S, D_0) - e(S, S_0) = \chi(S \cup S_0, D_0) - \chi(S_0, D_0).$$

Therefore $\chi(S, D_1) \geq 0 > \chi(S_0, D_0)$.

(4) $S \cap S_0 \neq \emptyset$ and $S \cap \bar{S}_0 \neq \emptyset$. In this case, let $A = S \cap S_0$ and $B = S \cap \bar{S}_0$. Then

$$\chi(B, D_1) = \chi(B, D_0) - e(B, S_0) = \chi(B \cup S_0, D_0) - \chi(S_0, D_0) \geq 0,$$

and we have

$$\begin{aligned} \chi(S, D_1) &= \chi(A \cup B, D_1) \\ &= \chi(A, D_1) + \chi(B, D_1) - e(A, B) \\ &\geq \chi(A, D_1) - e(A, B) \\ &= \chi(A, D_0) + e(\bar{S}_0, A) - e(A, B) \\ &\geq \chi(A, D_0) \\ &\geq \chi(S_0, D_0). \end{aligned}$$

By the minimality of S_0 , equality can happen only if $A = S_0$, in which case $S = S_0 \cup B \supsetneq S_0$.

Let S_1 be the minimal element in $\mathbf{S}(D_1)$. It follows that either (i) $\chi(S_1, D_1) > \chi(S_0, D_0)$ or (ii) $\chi(S_1, D_1) = \chi(S_0, D_0)$ and $S_0 \subsetneq S_1$. If we repeat this procedure, we are therefore guaranteed to stop, at which point we will have an orientable divisor. \square

There can be many different orientations associated to a given orientable divisor. Also, in each equivalence class of divisors of degree $g - 1$ there can be many different orientable divisors. Our next goal is to show that one can obtain a uniqueness result by fixing a vertex q .

Recall from Section 1 that an orientation of G is called q -connected if there exists an oriented path from q to each vertex p of G . Also, a divisor $D \in \text{Div}(G)$ is called q -orientable if there exists a q -connected orientation \mathcal{O} on G such that at every point $p \in V(G)$ we have $D(p) = \text{indeg}_{\mathcal{O}}(p) - 1$.

PROPOSITION 4.11. *An orientable divisor $D \in \text{Div}(G)$ is q -orientable if and only if $\chi(S, D) > 0$ for all nonempty subsets $S \subseteq V(G) \setminus \{q\}$.*

Proof. If D is q -orientable, it is in particular orientable, and it follows from Theorem 4.8 that $\chi(S, D) \geq 0$ for all $\emptyset \neq S \subseteq V(G) \setminus \{q\}$. If $\chi(S, D) = 0$ for some subset S , then all edges connecting S to \bar{S} will be directed away from S in any associated orientation, and the orientation will not be q -connected.

Now suppose that D is an orientable divisor and that $\chi(S, D) > 0$ for all $\emptyset \neq S \subseteq V(G) \setminus \{q\}$. We will show that for any orientation associated to D there is a directed path from q to any other vertex. Let $p = p_1$ be an arbitrary vertex in $V(G) \setminus \{q\}$. Since $\chi(\{p_1\}, D) > 0$, there exists an edge oriented toward p_1 . Let p_2 be the other end of this oriented edge. If $p_2 = q$, we have found a directed path from q to p . Otherwise, $\{p_1, p_2\} \subseteq V(G) \setminus \{q\}$, and we have $\chi(\{p_1, p_2\}, D) > 0$. Continuing this procedure will yield a directed path from q to p . \square

Fix a vertex q . For a divisor $D \in \text{Div}(G)$, we define

$$\mathbf{S}_q(D) = \{\emptyset \neq S \subseteq V(G) \setminus \{q\} : \chi(S, D) = \chi_D\}.$$

LEMMA 4.12. *Fix a vertex q , and let $E \in \text{Div}(G)$ be an orientable divisor, but not q -orientable. If $S_1, S_2 \in \mathbf{S}_q(E)$, then $S_1 \cup S_2 \in \mathbf{S}_q(E)$. In particular, $\mathbf{S}_q(E)$ has a unique maximal element (with respect to inclusion).*

Proof. When restricted to subsets of $V(G) \setminus \{q\}$, $\chi(\cdot, D)$ is still a submodular function. By Lemma 4.2, there is a unique maximal subset of $V(G) \setminus \{q\}$ for which $\chi(S, D) = \chi_D$, namely the union of all such subsets. Since E is orientable but not q -orientable, it follows from Theorem 4.8 and Proposition 4.11 that the maximal set is nonempty, and hence lies in $\mathbf{S}_q(E)$. \square

THEOREM 4.13. *Fix a vertex q . Every divisor $D \in \text{Div}^{s-1}(G)$ is equivalent to a unique q -orientable divisor.*

Proof. Existence. By Theorem 4.10, we know that $D \sim D_1$ for some orientable divisor D_1 . If D_1 is not already q -orientable, let S_1 be the unique maximal element of $\mathbf{S}_q(D_1)$, which exists by Lemma 4.12. In any orientation associated to D_1 , all edges connecting S_1 to its complement are directed from S_1 to \bar{S}_1 . Also, it follows from the maximality of S_1 that there is a directed path from q to any vertex $p \in \bar{S}_1$. We now replace D_1 with the divisor D_2 obtained by firing all vertices in the subset \bar{S}_1 . This reverses the orientation of edges connecting \bar{S}_1 and S_1 , directing them toward S_1 , and leaves all other orientations unchanged. If D_2 is not already q -orientable, there exists a maximal element S_2 of $\mathbf{S}_q(D_2)$. Since there is a directed path from q to any vertex $p \in \bar{S}_1$, as well as at least one vertex in S_1 , it follows that S_2 is a proper subset of S_1 . We now fire the subset \bar{S}_2 and proceed as before. This procedure must eventually terminate in a q -orientable divisor.

Uniqueness. Let D_1 and D_2 be distinct orientable divisors, and write $D_1 = D_2 + \Delta(f)$. Consider the (nonempty) set $X \subsetneq V(G)$ where f achieves its (global) minimum value. We have

$$0 \leq \chi(X, D_1) \leq \chi(X, D_2) - e(X, \bar{X}) = -\chi(\bar{X}, D_2) \leq 0.$$

It follows that $\chi(X, D_1) = \chi(\bar{X}, D_2) = 0$. This means that in any orientation associated to D_1 all edges are directed away from X , and in any orientation associated to D_2 all edges are toward X . Thus there cannot be a vertex q for which D_1 and D_2 are both q -orientable. \square

REMARK 4.14. It follows from the proofs of Theorems 4.10 and 4.13 that we have the following algorithm for finding the unique q -orientable divisor equivalent to a given divisor D . Our algorithm for finding the unique q -orientable divisor equivalent to a given divisor D takes exponential time. Since this paper was first posted on the arXiv, Backman [Bac13] has given a polynomial-time algorithm for this problem, as well as for finding an associated orientation.

- (1) While there exists a subset $\emptyset \neq S \subsetneq V(G)$ with $\chi(S, D) < 0$, find the unique minimal element A of $\mathbf{S}(D)$, and replace D with the divisor obtained by firing all vertices of \bar{A} .
- (2) While there exists a subset $\emptyset \neq S \subseteq V(G) \setminus \{q\}$ with $\chi(S, D) = 0$, find the unique maximal element B of $\mathbf{S}_q(D)$, and replace D with the divisor obtained from firing all vertices of \bar{B} .

It seems difficult to deduce effective algorithms for these problems from the work of Mikhalkin and Zharkov.

4.2. Metric graphs. Fix a metric graph Γ and a divisor $D \in \text{Div}(\Gamma)$. A *model for* (Γ, D) is a (weighted graph) model for Γ whose vertex set contains the support of D . If a point $q \in \Gamma$ is also distinguished, we further assume that the vertex set of G contains q .

We call a (weighted graph) model G for Γ a *semimodel for* (Γ, D) if $\Gamma \setminus V(G)$ consists of a finite union of open intervals $\cup_{i=1}^r e_i^\circ$, and for each i the set $e_i^\circ \cap \text{Supp}(D)$ is either empty or consists of a single point p with $D(p) = 1$. In other words, a semimodel is allowed to ‘miss’ some points $p \in \text{Supp}(D)$ where $D(p) = 1$ and p is the only point of $\text{Supp}(D)$ lying in the corresponding open edge. Again, if a point $q \in \Gamma$ is also distinguished, we further assume that $q \in V(G)$.

It turns out that semimodels are more convenient to work with than models when we want to show that certain algorithms for metric graphs terminate. We will

obtain results similar to those in Section 4.1 by reducing the metric graph case to the case of finite graphs via semimodels.

Let $G = (V, E)$ be an arbitrary semimodel for (Γ, D) , and let G_D be the finite graph obtained from G by removing all open edges which contain a point of the support of D . Note that $V(G_D) = V(G)$ and $E(G_D) \subset E(G)$. The following lemma will help us compare the set of orientable divisors on Γ and on G_D .

LEMMA 4.15. *Let $D \in \text{Div}^{g-1}(\Gamma)$, and let G be a semimodel for (Γ, D) . Let D_G be the restriction of D to G ; that is, if $D = \sum_{p \in \Gamma} a_p(p)$, then $D_G = \sum_{p \in V(G)} a_p(p)$. Then, for any $q \in \Gamma$, the divisor D is (q) -orientable on Γ if and only if D_G is (q) -orientable on G_D .*

Proof. Suppose that D is given by the q -connected orientation \mathcal{O} on Γ . Then \mathcal{O} naturally induces a q -connected orientation on the semimodel G . Strictly speaking, the induced orientation is not an orientation for the semimodel G but for its refinement whose vertex set consists of $V(G)$ and the support of D . On each edge $e \in E(G) \setminus E(G_D)$, the orientation \mathcal{O} on Γ must look like two arrows pointing toward a single point in the support of D . Hence, after removing these edges, \mathcal{O} is still a q -connected orientation on G_D , and the resulting divisor is D_G .

Conversely, given a q -connected orientation on G_D for D_G , we obtain an orientation for D on Γ by directing every edge $e \in E(G) \setminus E(G_D)$ toward the corresponding point in the support of D . \square

By Lemma 4.15, in order to show that $D \in \text{Div}(\Gamma)$ is q -orientable, it suffices to show that there is a semimodel G for D such that D_G is q -orientable on G_D . This helps us reduce our problems to the case of finite graphs.

REMARK 4.16. In the following, since we are working with different finite graphs, we use the notation $\chi_G(S, D)$, $\mathbf{S}(G, D)$, and $\mathbf{S}_q(G, D)$ (instead of $\chi(S, D)$, $\mathbf{S}(D)$, and $\mathbf{S}_q(D)$) to identify the underlying graph we are working with at each step.

THEOREM 4.17. *Every divisor $D \in \text{Div}^{g-1}(\Gamma)$ is equivalent to an orientable divisor on Γ . More precisely, let G be a model for (Γ, D) . Then D is equivalent to a divisor D' on Γ such that G is a semimodel for D' and D'_G is orientable on $G_{D'}$.*

Proof. Fix a model G for (Γ, D) . Let $D^0 = D$, let $k \geq 0$, and assume that G is a semimodel for D^k . If D^k is not orientable, we inductively define a divisor $D^{k+1} \in \text{Div}^{g-1}(\Gamma)$ equivalent to D as follows. By Lemma 4.15, Corollary 4.7,

and Theorem 4.8, there is a unique minimal element $S_k \in \mathbf{S}(G_{D^k}, D^k)$. Define ℓ to be the minimal distance between $\Gamma[S_k]$ (the closed subset of Γ corresponding to the induced subgraph $G[S_k]$) and $T_k := (V(G) \cup \text{Supp}(D^k)) \setminus \Gamma[S_k]$. Let C be the cut consisting of all the closed intervals connecting S_k to points of T_k , and fire by moving each chip on an interval in C a distance ℓ toward S_k . We call the resulting divisor D^{k+1} .

Clearly, G is again a semimodel for D^{k+1} . The claim in the proof of Theorem 4.10 also holds here (with a similar proof); that is, for each $k \geq 0$, either (i) $\chi_{G_{D^{k+1}}}(S_{k+1}, D^{k+1}) > \chi_{G_{D^k}}(S_k, D^k)$ or (ii) $\chi_{G_{D^{k+1}}}(S_{k+1}, D^{k+1}) = \chi_{G_{D^k}}(S_k, D^k)$ and $S_k \subsetneq S_{k+1}$. Therefore this procedure is guaranteed to stop, at which point we will have an orientable divisor. \square

THEOREM 4.18. *Every divisor $D \in \text{Div}^{g-1}(\Gamma)$ is equivalent to a unique q -orientable divisor.*

Proof. Existence. By Theorem 4.17, we may assume that $D = D^0$ is orientable. Fix a model G for (Γ, D) , and apply the following algorithm.

For $k \geq 0$, assume that G is a semimodel for D^k . If D^k is not already q -orientable, then by Proposition 4.11 and Lemma 4.12 there is a unique maximal element $S_k \in \mathbf{S}_q(G_{D^k}, D^k)$. Define ℓ and C exactly as in the proof of Theorem 4.17, and fire by moving each chip on an interval in C a distance ℓ toward S_k . Clearly, G is also a semimodel for the resulting divisor D^{k+1} .

In obtaining D^{k+1} from D^k , at least one chip must arrive at some vertex $v \in S_k$, so there is a directed path from q to v in the corresponding orientation \mathcal{O}_{k+1} . Thus S_{k+1} is a proper subset of S_k , and the algorithm will terminate to give a q -orientable divisor.

Uniqueness. This is identical to the proof of uniqueness in Theorem 4.13, starting with distinct orientable divisors D_1 and D_2 and letting G be a model for Γ such that $V(G)$ contains $\text{Supp}(D_1) \cup \text{Supp}(D_2)$. (Note that if $D_1 = D_2 + \Delta(f)$ then f is linear on every edge of G because $\Delta(f) = D_1 - D_2$ is supported on $V(G)$.) \square

4.3. Break divisors and universally reduced divisors. Break divisors, like q -reduced divisors, provide a way for us to pick out a distinguished representative from each linear equivalence class of divisors. In this section, we link the two notions by characterizing break divisors as limits of degree g effective divisors which are q -reduced for all $q \in \Gamma$.

LEMMA 4.19. *An effective divisor D on Γ is q -reduced for every $q \in \Gamma$ if and only if $|D| = \{D\}$.*

Proof. If D is q -reduced for every $q \in \Gamma$, it follows from the least action principle that $|D| = \{D\}$, exactly as in the proof of Lemma 3.5. Conversely, if D is effective, then for every q the unique q -reduced divisor D_q equivalent to D is also effective, so if $|D| = \{D\}$ we must have $D_q = D$. \square

We define a divisor on Γ to be *universally reduced* if it has degree g , is effective, and is q -reduced for every $q \in \Gamma$.

THEOREM 4.20. *The set Σ of break divisors on Γ is equal to the closure in $\text{Div}_+^g(\Gamma)$ of the set of universally reduced divisors.*

Proof. Let Ω be the set of universally reduced divisors. By Lemma 3.5, Σ contains a dense subset Σ° belonging to Ω . Since Σ is compact, we have $\Sigma \subseteq \bar{\Omega}$. To prove the reverse inclusion, note that, if $D \in \Omega$, then by Lemma 4.19 we have $|D| = \{D\}$. Since D is equivalent to a break divisor by Theorem 1.1, D must itself be a break divisor. Thus $\Omega \subseteq \Sigma$, and by taking closures we obtain $\bar{\Omega} \subseteq \Sigma$ as desired. \square

4.4. Integral break divisors on finite graphs. Suppose that G is a finite (unweighted) graph and that Γ is the associated metric graph in which all edges of G are assigned length 1. We let $\Sigma(G) := \Sigma \cap \text{Div}_+^g(G)$ denote the set of *integral break divisors*, that is, those break divisors which are supported on vertices of G . Inside $\text{Pic}^d(\Gamma)$, for each integer d , we have the finite subset $\text{Pic}^d(G)$ consisting of linear equivalence classes of divisors of degree d supported on the vertices of G . For $d = 0$, the set $\text{Pic}^0(G)$ is a group whose cardinality is equal to the number of spanning trees in G , and each $\text{Pic}^d(G)$ is a torsor for $\text{Pic}^0(G)$.

THEOREM 4.21. *The canonical map $\pi : \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ induces a bijection from $\Sigma(G)$ to $\text{Pic}^g(G)$. In particular, the number of integral break divisors is equal to the number of spanning trees of G .*

Proof. Choose $q \in V(G)$. Then the result follows from Theorem 4.13, which says that every element of $\text{Div}^{g-1}(G)$ is linearly equivalent to an *integral* divisor of the form $D_{\mathcal{O}}$ with \mathcal{O} a q -connected orientation, together with the equivalence of Theorems 1.1 and 1.2. \square

REMARK 4.22. One can interpret Theorem 4.21 as follows: although the set of spanning trees of G is not canonically a torsor for $\text{Pic}^0(G)$, the set $\Sigma(G)$ of integral break divisors is. Fixing a vertex q of G gives a bijection between $\text{Pic}^0(G)$ and $\text{Pic}^g(G)$, and for a generic choice of $\lambda \in \text{Jac}(\Gamma)$ there will be exactly one

element of $\text{Pic}^g(G) + \lambda$ in each open cell \mathcal{C}_T° . Thus the pair (q, λ) provides a bijection between elements of $\text{Pic}^0(G)$ and spanning trees.

5. The dual matrix–tree theorem

5.1. Weights of spanning trees and volumes. Let Γ be a metric graph, and let the weighted graph G be a model for Γ . Given any spanning tree T of G , we define the *weight* of T to be the product of the lengths of all edges of G **not in** T :

$$w(T) := \prod_{e \notin E(T)} \ell(e).$$

We also define

$$w(G) := \sum_T w(T), \tag{5.1}$$

the sum being over all spanning trees of G . It is easy to check that $w(G)$ is invariant under refinement, and therefore it depends only on the underlying metric graph Γ . We will therefore also denote $w(G)$ by $w(\Gamma)$.

Let Λ be a *lattice*, that is, a free \mathbb{Z} -module of some rank g equipped with a symmetric integer-valued bilinear form $\langle \cdot, \cdot \rangle$ whose corresponding quadratic form is positive definite on $\Lambda_{\mathbb{R}} := \Lambda \otimes \mathbb{R}$. We denote by $\text{Vol}(\Lambda)$ the volume of any fundamental domain for Λ , or, equivalently, the volume of the real torus $\Lambda_{\mathbb{R}}/\Lambda$. If M is any Gram matrix for Λ , that is, the matrix $(\langle \lambda_i, \lambda_j \rangle)$, where $\lambda_1, \dots, \lambda_g$ is a basis for Λ , then it is well known that $\text{Vol}(\Lambda) = \sqrt{\det(M)}$.

Our goal in this section is to give a geometric proof, via the decomposition of $\text{Pic}^g(\Gamma)$ into the cells \mathcal{C}_T , of the following ‘dual version’ of Kirchhoff’s matrix–tree theorem.

THEOREM 5.2 (Matrix–tree theorem, dual version). *The volume of the real torus $\text{Jac}(\Gamma)$ is $\text{Vol}(\text{Jac}(\Gamma)) = \sqrt{w(\Gamma)}$. Equivalently, if M is any Gram matrix for the cycle lattice $H_1(\Gamma, \mathbb{Z})$, then $\det(M) = w(\Gamma)$.*

One can explicitly calculate a Gram matrix for $H_1(\Gamma, \mathbb{Z})$ as follows. Fix an arbitrary orientation of the model G and a spanning tree T of G . For each $e \notin T$, the *fundamental cycle* $\mathbf{c}(T, e)$ is the unique element of $H_1(G, \mathbb{Z})$ contained in $T \cup e$. It is well known that the set

$$\mathbf{C}(T) := \{\mathbf{c}(T, e) : e \in E(G) \setminus E(T)\}$$

of fundamental cycles associated to T forms a basis for $H_1(\Gamma, \mathbb{Z})$.

Let $m = |E(G)|$, and let $g = |E(G)| - |V(G)| + 1$ be the rank of $H_1(\Gamma, \mathbb{Z})$. Let \mathcal{C}_T denote the $g \times m$ matrix whose rows correspond to basis

elements $\mathbf{c}(T, e) \in \mathbf{C}(T)$. For this, first fix a labeling $\{e_1, e_2, \dots, e_m\}$ of $E(G)$. The (i, j) -entry of \mathfrak{C}_T is

$$\begin{cases} +\sqrt{\ell(e_j)} & \text{if } +e_j \text{ appears in } \mathbf{c}(T, e_i); \\ -\sqrt{\ell(e_j)} & \text{if } -e_j \text{ appears in } \mathbf{c}(T, e_i); \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathfrak{C}_T \mathfrak{C}_T^t$ is a Gram matrix for $H_1(\Gamma, \mathbb{Z})$, and consequently for any spanning tree T we have

$$\text{Vol}(\text{Jac}(\Gamma)) = \sqrt{\det(\mathfrak{C}_T \mathfrak{C}_T^t)}. \quad (5.3)$$

Alternatively, let \mathfrak{D} denote the $m \times m$ diagonal matrix whose (i, i) -entry is $\sqrt{\ell(e_i)}$. Then $\mathfrak{C}_T = \mathfrak{C}'_T \mathfrak{D}$, where \mathfrak{C}'_T is the matrix whose (i, j) -entry is

$$\begin{cases} +1 & \text{if } +e_j \text{ appears in } \mathbf{c}(T, e_i); \\ -1 & \text{if } -e_j \text{ appears in } \mathbf{c}(T, e_i); \\ 0 & \text{otherwise.} \end{cases}$$

Fix an identification of $\text{Jac}(\Gamma)$ with $\text{Pic}^g(\Gamma)$, and let \mathcal{D}_T be the cell in $\text{Jac}(\Gamma)$ corresponding to the cell C_T in $\text{Pic}^g(\Gamma)$. In order to prove Theorem 5.2, it suffices to prove the following result.

PROPOSITION 5.4. $\text{Vol}(\mathcal{D}_T) = w(T)/\sqrt{\det(\mathfrak{C}_T \mathfrak{C}_T^t)} = w(T)/\text{Vol}(\text{Jac}(\Gamma))$.

Proof. Let \tilde{e} denote the orthogonal projection of an oriented edge e in $C_1(\Gamma, \mathbb{R})$ onto $H_1(\Gamma, \mathbb{R})$. Then the volume of \mathcal{D}_T is equal to the $\sqrt{\det(\mathfrak{L}_T \mathfrak{L}_T^t)}$, where \mathfrak{L}_T is the $g \times m$ matrix whose rows correspond to the basis elements \tilde{e} for $e \notin T$.

Let \mathfrak{D}_T denote the $g \times g$ diagonal matrix whose (i, i) -entry is $\sqrt{\ell(e_i)}$ for $e_i \notin T$. Then $\mathfrak{L}_T = \mathfrak{D}_T^2 (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} \mathfrak{C}_T$, since

$$\mathfrak{L}_T \mathfrak{C}_T^t = \mathfrak{D}_T^2 (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} (\mathfrak{C}_T \mathfrak{C}_T^t) = \mathfrak{D}_T^2.$$

Now we also have

$$\mathfrak{L}_T \mathfrak{L}_T^t = \mathfrak{D}_T^2 (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} (\mathfrak{C}_T \mathfrak{C}_T^t) (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} \mathfrak{D}_T^2 = \mathfrak{D}_T^2 (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} \mathfrak{D}_T^2,$$

and therefore

$$\begin{aligned} \text{Vol}(\mathcal{D}_T) &= \sqrt{\det(\mathfrak{D}_T^2 (\mathfrak{C}_T \mathfrak{C}_T^t)^{-1} \mathfrak{D}_T^2)} = \det(\mathfrak{D}_T)^2 / \sqrt{\det(\mathfrak{C}_T \mathfrak{C}_T^t)} \\ &= w(T) / \sqrt{\det(\mathfrak{C}_T \mathfrak{C}_T^t)}. \end{aligned} \quad \square$$

5.2. Matroid duality. We now explain the precise sense in which Theorem 5.2 is dual to the usual version of Kirchhoff’s matrix–tree theorem. In order to do this, we first give a slightly more canonical formulation of the latter.

Given any spanning tree T of G , we define the *coweight* of T to be the product of the lengths all edges of G in T :

$$w'(T) := \prod_{e \in E(T)} \ell^{-1}(e).$$

We also define

$$w'(G) := \sum_T w'(T)$$

to be the sum of $w'(T)$ over all spanning trees of G .

Note that, unlike $w(G)$, the quantity $w'(G)$ is not invariant under refinement and is therefore *not* an invariant of the metric graph Γ .

THEOREM 5.5 (Kirchhoff’s matrix–tree theorem, canonical version). *Let B be the cocycle lattice (or lattice of integer cuts) of G . Then $\text{Vol}(B)^2 = w'(G)$. Equivalently, if M' is any Gram matrix for B , then $\det(M') = w'(G)$.*

If we fix a vertex q of G , then the *reduced Laplacian matrix* $Q' = Q_q$ obtained by deleting the row and column corresponding to q in the usual weighted Laplacian matrix for G is the Gram matrix of the basis for B consisting of $\partial^*(p)$ for vertices $p \neq q$, where $\partial^* : C^0(G, \mathbb{Z}) \rightarrow C^1(G, \mathbb{Z})$ is adjoint to the usual boundary map $\partial : C_1(G, \mathbb{Z}) \rightarrow C_0(G, \mathbb{Z})$. We therefore obtain the following.

COROLLARY 5.6 (Kirchhoff’s matrix–tree theorem, usual version). *Fix $q \in V(G)$, and let Q' be the corresponding reduced Laplacian matrix. Then $\det(Q') = w'(G)$.*

Theorem 5.5 is dual to Theorem 5.2 in the precise sense that it is obtained by interchanging the cycle lattice with the cocycle lattice and weights with coweights. As is well known, interchanging the cycles and cocycles in a graph is a special case of matroid duality.

REMARK 5.7. The classical linear-algebraic proof of Kirchhoff’s matrix–tree theorem is an application of the Cauchy–Binet formula. One can also prove Theorem 5.2 via the Cauchy–Binet formula; we omit the details. We note in addition that a generalization of the computations in [KS00, Lemma 3.4] [CY09, Lemma 2] to the setting of weighted graphs can be used to prove that Kirchhoff’s matrix–tree theorem and its dual version are equivalent.

Acknowledgements

Matthew Baker and Farbod Shokrieh were partially supported by NSF grant DMS-1201473. Greg Kuperberg was partially supported by NSF grant CCF-1013079. We enthusiastically thank Spencer Backman and the anonymous referee for their helpful comments.

References

- [Ami13] O. Amini, ‘Reduced divisors and embeddings of tropical curves’, *Trans. Amer. Math. Soc.* **365** (9) (2013), 4851–4880.
- [Bac13] S. Backman, ‘Riemann–Roch theory for graph orientations’, available at [arXiv:1401.3309](https://arxiv.org/abs/1401.3309), 27 pages, 2014.
- [BF11] M. Baker and X. Faber, ‘Metric properties of the tropical Abel–Jacobi map’, *J. Algebraic Combin.* **33** (3) (2011), 349–381.
- [BS13] M. Baker and F. Shokrieh, ‘Chip-firing games, potential theory on graphs, and spanning trees’, *J. Combin. Theory Ser. A* **120** (1) (2013), 164–182.
- [CY09] S. Chen and S. K. Ye, ‘Critical groups for homeomorphism classes of graphs’, *Discrete Math.* **309** (1) (2009), 255–258.
- [GK08] A. Gathmann and M. Kerber, ‘A Riemann–Roch theorem in tropical geometry’, *Math. Z.* **259** (1) (2008), 217–230.
- [Hak65] S. L. Hakimi, ‘On the degrees of the vertices of a directed graph’, *J. Franklin Inst.* **279** (1965), 290–308.
- [HKN13] J. Hladký, D. Král and S. Norine, ‘Rank of divisors on tropical curves’, *J. Combin. Theory Ser. A* **120** (7) (2013), 1521–1538.
- [KS00] M. Kotani and T. Sunada, ‘Jacobian tori associated with a finite graph and its abelian covering graphs’, *Adv. Appl. Math.* **24** (2) (2000), 89–110.
- [Luo11] Y. Luo, ‘Rank-determining sets of metric graphs’, *J. Combin. Theory Ser. A* **118** (6) (2011), 1775–1793.
- [MZ08] G. Mikhalkin and I. Zharkov, ‘Tropical curves, their Jacobians and theta functions’, in *Curves and Abelian Varieties*, Contemporary Mathematics 465 (American Mathematical Society, Providence, RI, 2008), 203–230. Available at [arXiv:math/0612267](https://arxiv.org/abs/math/0612267).
- [Sch03] A. Schrijver, ‘Matroids, trees, stable sets’, in *Combinatorial Optimization. Polyhedra and Efficiency. Vol. B*, Algorithms and Combinatorics 24 (Springer, Berlin, 2003) (Chs 39–69).