TECHNICAL REPORTS

Effective divisor classes on metric graphs

Andreas Gross, Farbod Shokrieh, and

Lilla Tóthmérész

July 2018

Effective divisor classes on metric graphs

Andreas Gross[*], Farbod Shokrieh[**], and Lilla Tóthmérész[***]

Abstract

We introduce the notion of semibreak divisors on metric graphs (tropical curves) and prove that every effective divisor class (of degree at most the genus) has a semibreak divisor representative. This appropriately generalizes the notion of break divisors (in degree equal to genus). Our method of proof is new, even for the special case of break divisors. We provide an algorithm to efficiently compute such semibreak representatives. Semibreak divisors provide the tool to establish some basic properties of effective loci inside Picard groups of metric graphs. We prove that effective loci are pure-dimensional polyhedral sets. We also prove that a 'generic' divisor class (in degree at most the genus) has rank zero, and that the Abel-Jacobi map is 'birational' onto its image. These are analogues of classical results for Riemann surfaces.

Contents

[*]Imperial College London, Department of Mathematics, South Kensington Campus, London, SW7 2AZ, UK, Email: `a.gross@imperial.ac.uk`

[**]Cornell University, Ithaca, New York 14853-4201, USA Email: `farbod@math.cornell.edu`

[***]Cornell University, Ithaca, New York 14853-4201, USA, MTA-ELTE Egerváry Research Group, Pázmány Péter sétány 1/C, Budapest, Hungary Email: `tmlilla@math.cornell.edu`

1 Introduction

Metric graphs, in many respects, are tropical (or non-Archimedean) analogues of Riemann surfaces. For example, there is a well-behaved theory of divisors and Jacobians for metric graphs (see e.g. [BN07, MZ08, GK08, ABKS14]). There is also an interesting interaction between the theories of divisors on metric graphs and on algebraic curves, with numerous applications in algebraic geometry (see e.g. [Bak08, CDPR12, JP14, JP16, CJP15, BR15, BPR16, BJ16]). The purpose of this work is to study *tropical effective loci* and establish some of their basic properties.

Let Γ be a compact *metric graph* (*abstract tropical curve*) of genus g. Fix an integer $0 \leq d \leq g$. There is a canonical (Abel-Jacobi) map $S^{(d)} \colon \operatorname{Div}^d_+(\Gamma) \to \operatorname{Pic}^d(\Gamma)$ taking an effective divisor D of degree d on Γ to its linear equivalence class $[D]$. The image of this map, denoted by W_d, is the *locus of effective divisor classes*. In the language of chip-firing games on metric graphs, this is the collection of chip configuration classes (up to chip-firing moves) which are 'winnable'.

We provide 'nice' representatives for equivalence classes $[D] \in W_d$. In the case $d = g$, this is done by Mikhalkin and Zharkov in [MZ08] using the theory of tropical theta functions. They introduce the notion of 'break divisors', and prove that every $[D] \in W_g = \operatorname{Pic}^g(\Gamma)$ has a unique break divisor representative. The notion of break divisors is further studied in [ABKS14] from a more combinatorial point of view related to orientations on graphs.

A break divisor can be described as follows: pick g disjoint open edge segments in Γ so that, if we remove them from Γ, the remaining space becomes contractible (see the gray edges in Figure 1). A break divisor is a divisor obtained by picking one point from *the closure* of each of these open edge segments (see Figure 1 (a), (b)). So a break divisor has degree equal to g by construction.
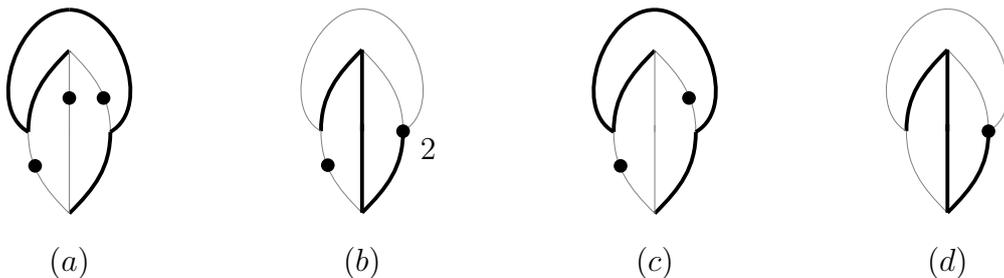
Figure 1: Semibreak divisors on a metric graph of genus 3.

  (a) a 'generic' break divisors.
  (b) a break divisor with some endpoints of open edge segments.
  (c) a semibreak divisor in degree 2 dominated by the break divisor in (a).
  (d) a semibreak divisor in degree 1 dominated by the break divisor in (b).

We define a *semibreak divisor* to be a divisor obtained from a break divisor after removing some points in its support. More precisely, a semibreak divisor is an effective

divisor 'dominated' by a break divisor (see Figure 1 (c), (d)). In particular, a break divisor is a semibreak divisor in degree g.

Most of our work is devoted to proving the following result.

Theorem A. There exists a semibreak divisor in each $[D] \in W_d$.

See Theorem 6.1.

Our techniques are new, even in the special case that $d = g$. We rely mostly on the geometry (e.g. a notion of 'convexity') and algebraic topology (e.g. various Mayer-Vietoris sequences) of Γ and its subspaces. We avoid the use of tropical theta functions, orientations, semimodels, reduced divisors, etc.

We start by a key result which states, roughly, that a given divisor is a break divisor if and only if a certain inequality holds for every 'admissible' subset of Γ (see Proposition 3.3 for a precise statement). By a Mayer-Vietoris argument, this characterization relates break divisors to submodular functions (see §4). So, in its core, our approach resembles classical combinatorial proofs using submodularity, however, we exploit these ideas directly in a 'continuous setting'. We find it remarkable that the 'discrete theory' of submodular functions fits so naturally into our tropical setting.

Having this characterization, the naive strategy to prove Theorem A is straightforward: given an effective divisor D, we should first find a linearly equivalent divisor D' that satisfies all the desired inequalities. We then would like to add points to D' carefully in a way that all the desired inequalities are preserved. This process should eventually stop and output a break divisor that dominates a semibreak divisor linearly equivalent to D. It turns out that the construction of a suitable break divisor is more subtle than by simply adding points to D'. The process will rely on understanding certain canonical subsets $\mathfrak{I}(E) \subsetneq \Gamma$ attached to effective divisors E. In a key result (Proposition 5.8) we will describe exactly how various invariants change as points are moved in relation to $\mathfrak{I}(E)$.

We then turn our attention to the question of uniqueness.

Theorem B. There is a unique break divisor in each $[D] \in W_g$. If $d < g$ there can be distinct semibreak divisors in $[D] \in W_d$. If Γ is $(d+1)$-edge connected then there is a unique semibreak divisor in each $[D] \in W_d$.

See Proposition 6.4, Figure 2, Proposition 6.5 for precise statements.

As mentioned earlier, the first statement is already proved in [MZ08, ABKS14]. Our proof indicates that the result can be thought of as a consequence of a 'maximum principle' (Lemma 6.3).

We will then prove that an 'integral' version of Theorem A also holds. A finite unweighted graph G may be thought of as a metric graph whose edges have length 1. We define an integral divisor to be a divisor supported on the vertices of G.

Theorem C. Let $[D] \in W_d$ and assume D is integral. Then every semibreak divisor in $[D]$ is also integral.

See Proposition 6.8 for a precise statement.

The fact that there *exists* an integral semibreak divisor in $[D]$ is immediate from our method of proof of Theorem A. More work is needed to show that *all* semibreak divisors in the equivalence class are indeed integral. We note that the 'existence' part of Theorem C may be stated purely in terms of finite graphs (avoiding metric graphs). We also remark that one could modify our proof to give a purely combinatorial proof of the finite graph version of Theorem C. One could also give another combinatorial proof (of the 'existence' part) by using the theory of partial orientations in [Bac17] and results in [ABKS14] (see Remark 6.9).

We will show that everything is efficiently computable.

Theorem D. Given $[D] \in W_d$, there is an efficient algorithm that computes a semibreak divisor $D' \in [D]$.

See Theorem 7.2, and §7 for the description of the algorithm.

We will reduce our computational problem to the theory of *submodular optimization* as in [Sch00, IFF01]. For break divisors on finite graphs an algorithm is presented in [Bac17, §7] relating the computation to the max-flow min-cut problem in graph theory.

Finally we apply the theory of semibreak divisors to prove the following tropical (non-Archimedean) analogues of some classical results on Riemann surfaces.

Theorem E.

(a) W_d is a *purely* d-dimensional polyhedral subset of $\text{Pic}^d(\Gamma)$.

(b) The tropical (Abel-Jacobi) map $S^{(d)} \colon \text{Div}_+^d(\Gamma) \to \text{Pic}^d(\Gamma)$ is 'birational' onto its image.

(c) There exists an open dense subset $U_d \subseteq W_d$ such that $r(D) = 0$ whenever $[D] \in U_d$.

See Theorem 8.3 and Theorem 8.5 for precise statements.

The analogous statements for Riemann surfaces essentially follow from simple linear algebraic facts applied to 'Brill–Noether' matrices (see e.g. [GH94, p.245]). The situation in tropical geometry is different. While the fact that W_d is a d-dimensional polyhedral subset of $\text{Pic}^d(\Gamma)$ is elementary and well-known, the *pure*-dimensionality of W_d is much more subtle (see Remark 8.4 (i)). We remark that one could give a highly non-constructive proof of pure-dimensionality by appealing to Berkovich's theory of non-Archimedean analytic spaces ([Ber90]) and combining the results in [BR15] and [Gub07] (see Remark 8.4 (ii)). Finally, the statement analogous to part (c) for Riemann surfaces (see e.g. [GH94, p.245]) is usually stated as $r(D) = 0$ for a generic *effective divisor* D. In algebraic geometry, this is equivalent to saying $r(D) = 0$ for a generic *effective divisor class* $[D]$. In tropical geometry these two statements are *not* equivalent and, in fact, the former statement is not true (see Remark 8.6).

Structure of the paper. In §2 we will review some basic definitions and set our notations and terminology. In §3 the notion of semibreak divisors is introduced. We will also state and prove a key result (Proposition 3.3) characterizing break divisors in terms

of certain inequalities arising from the topology of 'admissible' subsets of the metric graph. In §4 and §5 we study the functions and subsets related to the inequalities in Proposition 3.3. In §6 the existence of semibreak divisors (Theorem A) is proved. We will also discuss the uniqueness issues (Theorem B) and consider the integral version of semibreak divisors (Theorem C). In §7 we show how one can efficiently compute a semibreak divisor linearly equivalent to a given effective divisor (Theorem D). In §8 we apply the theory of semibreak divisors to prove basic generic properties of tropical effective loci (Theorem E).

Acknowledgments

2 Definitions and background

2.1 Metric graphs and vertex sets

Definition 2.1. A *metric graph* (or an *abstract tropical curve*) is a compact connected metric space such that every point has a neighborhood isometric to a star-shaped set, endowed with the path metric.

By a *star-shaped* set of finite *valency* n and *radius* r we mean a set of the form

$$S(n,r) = \left\{ z \in \mathbb{C} \colon z = te^{k\frac{2\pi i}{n}} \text{ for some } 0 \leq t < r \text{ and some } k \in \mathbb{Z} \right\}.$$

In other words, a metric graph is a pair (Γ, d) consisting of a compact connected topological graph Γ, together with an *inner metric* d.

As all of our results are easy to show for metric circles, we will always assume that Γ is neither a point, nor a circle.

Definition 2.2.

(i) The points of Γ that have valency different from 2 are called *branch points* of Γ.

(ii) A *vertex set* for Γ is a finite set of points of Γ containing all the branch points.

We denote the minimal vertex set (i.e. the set of branch points) of Γ by V_Γ. Note that V_Γ is finite because Γ is assumed to be compact. Also, V_Γ is nonempty, because Γ is not a circle.

We denote the set of components of $\Gamma \backslash V_\Gamma$ by E_Γ and call its elements the *open edges* of Γ. By a *closed edge* we mean the closure $\bar{e}$ of an open edge $e \in E_\Gamma$. An open connected subset of an open edge of Γ is called an *open edge segment*, and the

closure of such a segment is called a *closed edge segment.* If e is a (closed or open) edge segment, the points in the topological boundary ∂e of e are called its *endpoints.* Every edge has either 1 or 2 endpoints.

Every finite combinatorial graph G whose edges are weighted with positive real numbers naturally determines a metric graph Γ_G. A *model* of a metric graph Γ is a finite combinatorial weighted graph G, together with an isometry $\Gamma_G \xrightarrow{\phi} \Gamma$. Up to isomorphisms, the model G is completely determined by the set $\phi(V(G))$, which is a vertex set for Γ, and conversely every vertex set determines a unique (up to isomorphism) model. If G is the model of Γ corresponding to the vertex set $V = V(G)$ (the isometry ϕ being implicit), then we call the elements of the set $E(G)$ of components of $\Gamma\backslash V$ the open edges of G. Of course, every open edge of G is an open edge segment of Γ.

Remark 2.3. Open edge segments of Γ may be thought of as open edges of *some* model of Γ.

Definition 2.4. Let S be a subset of a metric graph Γ. The *convex hull* of S, denoted by $\mathrm{conv}(S)$, is defined to be the union of S and all closed edge segments whose endpoints are contained in S.

It is not difficult to see that one can obtain $\Gamma\backslash\mathrm{conv}(S)$ by removing all connected components from $\Gamma\backslash S$ that are contained in some open edge of Γ. A set $S \subseteq \Gamma$ is called *convex*, if $\mathrm{conv}(S) = S$.

Definition 2.5. We call a subset $S \subseteq \Gamma$ *admissible* if it has only finitely many path-connected components.

Remark 2.6.

(i) A subset $S \subseteq \Gamma$ is admissible if and only if there is a model G of Γ such that S is a finite union of vertices and open edges of G. We say such a model G is *compatible* with the admissible subset S.

(ii) The collection of admissible sets is closed under finite Boolean combinations.

Let S be a closed admissible subset of Γ and let $p \in S$. For any sufficiently small star-shaped open neighborhood B of p, the set $B\backslash S$ is a disjoint union of finitely many open edge segments, the number of which only depends on p. We denote this number by $\mathrm{val}_S(p)$. Informally, this is the number of edges emanating from the admissible set S at a point p. Clearly $\mathrm{val}_S(p) = 0$ if $p \notin \partial S$, where ∂S denotes the topological boundary of S.

Definition 2.7. Let S be an admissible subset of a metric graph Γ.

(i) The *arithmetic genus* of S is defined as $p_a(S) = 1 - \chi(S)$, where $\chi(S) = \dim H^0(S;\mathbb{R}) - \dim H^1(S;\mathbb{R})$ is the usual topological Euler characteristic of S. Here, $H^i(S;\mathbb{R})$ is the i-th singular cohomology group of S with real coefficients.

(ii) The *genus contribution* of S is defined as $\psi(S) = p_a(\Gamma) - p_a(\Gamma\backslash S)$.

Remark 2.8.

(i) A graph theorist might want to think of $H^1(S; \mathbb{R})$ as the vector space of $\mathbb{R}$-valued *flows* on S. Moreover, $\dim H^0(S; \mathbb{R})$ is the number of connected components of S. For a fine enough model G, the arithmetic genus $p_a(S)$ is equal to the number of closed edges corresponding to G lying entirely inside S, minus the number of vertices of G lying inside S, plus 1.

(ii) Recall that the geometric (or topological) genus of S is defined as $p_g(S) = \dim H^1(S; \mathbb{R})$. Clearly, $p_g(S) = p_a(S)$ if and only if S is connected. If S is connected, we will refer to $p_g(S) = p_a(S)$ as the *genus* of S.

(iii) If S_1 and S_2 are two open subsets of Γ, then the Euler characteristic obeys a version of the inclusion–exclusion principle:

$$\chi(S_1 \cup S_2) = \chi(S_1) + \chi(S_2) - \chi(S_1 \cap S_2).$$

This follows from the Mayer-Vietoris sequence:

$$0 \to H^0(S_1 \cup S_2, \mathbb{R}) \to H^0(S_1, \mathbb{R}) \oplus H^0(S_2, \mathbb{R}) \to H^0(S_1 \cap S_2, \mathbb{R}) \to$$
$$\to H^1(S_1 \cup S_2, \mathbb{R}) \to H^1(S_1, \mathbb{R}) \oplus H^1(S_2, \mathbb{R}) \to H^1(S_1 \cap S_2, \mathbb{R}) \to 0.$$

2.2   Divisor theory on metric graphs

Let $\mathrm{Div}(\Gamma)$ denote the free abelian group generated by the points of Γ. Denoting the generator corresponding to $p \in \Gamma$ by (p), an element of $\mathrm{Div}(\Gamma)$, called a *divisor* on Γ, can be uniquely represented as

$$D = \sum_{p \in \Gamma} a_p(p),$$

where $a_p \in \mathbb{Z}$ and all but finitely many of the a_p are zero. It is convenient to denote the coefficient a_p in D by $D(p)$. The *support* of D is $\mathrm{supp}(D) = \{p \in \Gamma \colon D(p) \neq 0\}$. A divisor $D \in \mathrm{Div}(\Gamma)$ is called *effective* if $D(p) \geq 0$ for all $p \in \Gamma$. For $D, E \in \mathrm{Div}(\Gamma)$, we write $E \leq D$ if $D - E$ is effective. The *degree* of a divisor D on Γ is defined as $\deg(D) = \sum_{p \in \Gamma} D(p)$. More generally, the degree of a divisor D on an admissible subset S is

$$\deg(D|_S) = \sum_{p \in S} D(p).$$

The set of divisors of a given degree d is denoted by $\mathrm{Div}^d(\Gamma)$. The set of effective divisors of a given degree d is denoted by $\mathrm{Div}^d_+(\Gamma)$.

Let $R(\Gamma)$ be the group of *continuous piecewise affine functions with integer slopes*. These are continuous functions $\phi \colon \Gamma \to \mathbb{R}$ such that for every isometric map $\gamma \colon [0, \epsilon] \to \Gamma$, the pullback $\phi \circ \gamma$ is piecewise-linear with integral slopes in the usual sense. They are the tropical analogues of meromorphic functions on Riemann surfaces ([MZ08]). Note that such a function ϕ can only change its slope finitely many times on each closed edge.

Let
$$\mathrm{div}\colon R(\Gamma) \to \mathrm{Div}(\Gamma)$$
denote the Laplacian operator in the sense of distributions; for $\phi \in R(\Gamma)$, we have
$$\mathrm{div}(\phi) = \sum_{p \in \Gamma} \sigma_p(\phi)(p)\,,$$
where $\sigma_p(\phi)$ is the sum of incoming slopes of ϕ at p. It is easy to check that the group of *principal divisors* $\mathrm{Prin}(\Gamma) = \mathrm{div}(R(\Gamma))$ is contained in $\mathrm{Div}^0(\Gamma)$.

Two divisors D_1 and D_2 are called *linearly equivalent*, written $D_1 \sim D_2$, if there exists $\phi \in R(\Gamma)$ such that $D_1 - D_2 = \mathrm{div}(\phi)$. It is immediate that $\sim$ defines an equivalence relation. We denote the equivalence class of a divisor D by $[D]$. The *complete linear system* $|D|$ is the set of all effective divisors linearly equivalent to D.

Remark 2.9. Given an effective divisor D, it is useful to think of $D(p)$ as the number of *chips* placed at the point $p \in \Gamma$. For an admissible subset S of Γ and (sufficiently small) $\epsilon > 0$, the rational function
$$\phi_{S,\epsilon}\colon \Gamma \to \mathbb{R}, \quad x \mapsto \min\{\epsilon, d(x, S)\}\,,$$
where d is the metric on Γ, has value 0 on S and ϵ outside an ϵ-neighborhood of S, with slope 1 in each outgoing direction from S. Replacing D with $D + \mathrm{div}(\phi_{S,\epsilon})$ has the effect of moving a chip to distance ϵ along each outgoing direction from S. This is often called 'firing' the subset S to distance ϵ. One can check that every element of $R(\Gamma)$ can be written as a finite integer linear combination of functions of the form $\phi_{S,\epsilon}$. Therefore, one can describe linear equivalence of divisors on Γ in terms of 'chip-firing games'.

3   Break and semibreak divisors

The notion of break divisors was introduced by Mikhalkin and Zharkov in [MZ08], and further studied in [ABKS14]. Here, we introduce a natural generalization of this concept.

Definition 3.1. Let Γ be a metric graph of genus g.

(i) A divisor D on Γ is called a *break divisor* if there exist g disjoint open edge segments $e_1, \ldots, e_g$ and points $p_i \in \overline{e}_i$ such that $D = (p_1) + \cdots + (p_g)$ and $\Gamma \backslash \bigcup_{i=1}^g e_i$ is contractible.

(ii) A *semibreak divisor* is an effective divisor E such that $E \leq D$ for some break divisor D.

Note that a break divisor is also a semibreak divisor.

Remark 3.2.

(i) One can alternatively think of $\Gamma \backslash \bigcup_{i=1}^g e_i$ in Definition 3.1(i) as a spanning tree of some model G of the metric graph Γ.

(ii) The term 'break' comes from the fact that a break divisor gives a recipe for fixing a fundamental domain in the maximal abelian covering of the metric graph Γ.

(iii) It is proved in [MZ08] and [ABKS14] that any equivalence class in degree g has a unique break divisor representative. The uniqueness in this statement is relatively straightforward, and much of the work goes into proving the existence. In both [MZ08] and [ABKS14] the statement is proved by a careful study of the notion of 'orientations'. One of our main results is that every effective divisor class of degree at most g contains a semibreak divisor. Our approach is more topological in nature and, in particular, we completely avoid the use of orientations.

We have the following useful characterization of break divisors. In the context of finite graphs, related statements can be found in the literature (see for example [ABKS14, Proposition 4.11] and [Ќ13, Theorem 3.4]). Some of the ideas in our proof below are adapted from the proof of [Ќ13, Theorem 3.4].

Proposition 3.3. *Let Γ be a metric graph of genus g, and let $D \in \mathrm{Div}^g(\Gamma)$. The following are equivalent:*

(i) D is a break divisor.

(ii) $\deg(D|_S) \geq p_a(S)$ for all open admissible subsets $\emptyset \neq S \subseteq \Gamma$.

(iii) $\deg(D|_S) \leq \psi(S)$ for all closed admissible subsets $S \subsetneq \Gamma$.

Proof. (ii)$\Longleftrightarrow$(iii) follows immediately from exchanging S with $\Gamma \backslash S$ and the definition of the genus contribution function ψ (see Definition 2.7 (ii)).

(i)$\Rightarrow$(ii): Assume D is a break divisor and let $e_1, \ldots, e_g$ be open edge segments in Γ with $p_i \in \overline{e}_i$ such that $T = \Gamma \backslash \bigcup e_i$ is contractible and $D = \sum (p_i)$. Let S be a nonempty open admissible subset of Γ. If $S \cap T = \emptyset$, then S is a union of open edge segments and we have $p_a(S) \leq 0$ and there is nothing to prove. We may thus assume that $S \cap T \neq \emptyset$.

Any sufficiently small open neighborhood of T will intersect each $S \cap e_i$ in zero, one, or two connected components. Let U be an open neighborhood of T such that $S \cap U \cap e_i$ has as few connected components as possible, for all $1 \leq i \leq g$.

- If $S \cap U \cap e_i$ has two connected components, then $\chi(S \cap U \cap e_i) = 2$ and $\chi(S \cap e_i) \geq 1$.

- If $S \cap U \cap e_i$ has one connected component, then $\chi(S \cap U \cap e_i) = 1$ and $\chi(S \cap e_i) \geq 1$.

- If $S \cap U \cap e_i = \emptyset$, then $\chi(S \cap U \cap e_i) = 0$ and $\chi(S \cap e_i) \geq 0$.

In any case, we have

$$\chi(S \cap e_i) - \chi(S \cap U \cap e_i) + 1 \geq 0, \tag{1}$$

Equality in (1) occurs if and only if $e_i \subseteq S$.

- If $e_i \subseteq S$ and $p_i \notin S$, then p_i is a boundary point of e_i and at least one component of $S \cap U$ is contained entirely in e_i.

- As U is a metric tree, every component of $S \cap U$ is a metric tree as well so that $\chi(S \cap U)$ equals the number of components of $S \cap U$.

Together with the fact that $S \cap U$ has a component intersecting T (recall that we assumed $S \cap T \neq \emptyset$), it follows that

$$\chi(S \cap U) + \sum_{i \,:\, p_i \notin S} \big(\chi(S \cap e_i) - \chi(S \cap U \cap e_i)\big) \geq 1 \tag{2}$$

By the Mayer-Vietoris sequence (see Remark 2.8 (ii)), applied to the pair $(S \cap U, S \cap (\bigcup_i e_i))$, together with the fact that the Euler characteristic is additive with respect to disjoint unions of open sets, we have:

$$\chi(S) = \chi(S \cap U) + \chi\bigg(S \cap \bigg(\bigcup_i e_i\bigg)\bigg) - \chi\bigg(S \cap U \cap \bigg(\bigcup_i e_i\bigg)\bigg)$$

$$= \chi(S \cap U) + \sum_i \chi(S \cap e_i) - \sum_i \chi(S \cap U \cap e_i) \,.$$

Therefore,

$$\chi(S) + \deg(D|_S) = \chi(S \cap U) + \sum_{i \,:\, p_i \notin S} \big(\chi(S \cap e_i) - \chi(S \cap U \cap e_i)\big)$$
$$+ \sum_{i \,:\, p_i \in S} \big(\chi(S \cap e_i) - \chi(S \cap U \cap e_i) + 1\big) \,,$$

which is at least 1 by (1) and (2). Therefore, we have $\deg(D|_S) \geq p_a(S)$ as required.

(ii)$\Rightarrow$(i): Each point $p \in \Gamma$ has a neighborhood B homeomorphic to a star-shaped set which is open and admissible. Since $\deg(D|_B) \geq p_a(B) = 0$, by taking B to be sufficiently small so $\deg(D|_B) = D(p)$ we conclude that $D(p) \geq 0$. Therefore, D must be effective. Let us denote $D = (p_1) + \cdots + (p_g)$.

Let G be the model of Γ corresponding to the vertex set $V_\Gamma \cup \operatorname{supp}(D)$. For each p_i in the support of D, choose an open edge $e_i \in E(G)$ such that $p_i \in \bar{e}_i$. Note that for $i \neq j$ we might choose $e_i = e_j$. We define the *multiplicity* of the open edge $e \in E(G)$, denoted by $\operatorname{mult}(e)$, to be the number of indices $1 \leq i \leq g$ with $e_i = e$.

Let $T = \Gamma \backslash \bigcup_{i=1}^g e_i$. Because we remove at most g edges from the graph Γ of genus g, we have $p_a(T) \geq 0$. Therefore, if the connected components of T all have genus 0, then T must be connected. In this case, it is contractible and D is a break divisor. We may thus assume that T has a connected component of positive genus. We will show how to modify the choice of $\{e_1, \ldots, e_g\}$ such that, eventually, T becomes a spanning tree of G.

Let us denote by $m(T)$ the minimal number of edges of a component of T of positive genus, and let C be a connected component of T of positive genus with precisely $m(T)$ edges. After relabeling, we may assume that there exists $0 \leq k \leq g$ such that for $1 \leq i \leq g$ we have $p_i \in C$ if and only if $i \leq k$ ($k = 0$ thus means that $\deg(D|_C) = 0$).

Claim. There exists $1 \leq j \leq k$ such that either $\mathrm{mult}(e_j) \geq 2$, or e_j connects C to a different connected component of T.

Proof of the Claim. Suppose, to the contrary, that we have $\mathrm{mult}(e_i) = 1$ and $\partial e_i \subseteq C$ for all $1 \leq i \leq k$. Let $C' = C \cup \bigcup_{i=1}^{k} e_i$. Then C' is closed and we have

$$p_a(C') = k + p_a(C) \geq k + 1 \,.$$

For any sufficiently small connected open neighborhood V of C' we then have

$$\deg(D|_V) = \deg(D|_{C'}) = \deg(D|_C) = k < k + 1 \leq p_a(C') = p_a(V) \,,$$

contradicting the assumption (ii).

Let $1 \leq j \leq k$ be as in the claim. We will substitute an edge e'_j for e_j as follows:

(1) If $\mathrm{mult}(e_j) \geq 2$, then choose an arbitrary edge e of C incident to p_j and set $e'_j = e$.

(2) If $\mathrm{mult}(e_j) = 1$ and there exists an edge e of C incident to p_j such that $C \backslash e$ is connected, we set $e'_j = e$.

(3) If $\mathrm{mult}(e_j) = 1$, and every edge incident to p_j disconnects C, we proceed as follows: for an edge e of C incident to p_j we denote by C_e the connected component of $C \backslash e$ that does not contain p_j. Then the sum $\sum_e p_g(C_g)$, where e runs over all edges of C incident to p_j, is equal to $p_g(C)$. In particular, this sum is positive, and hence there exists an edge e of C incident to p_j such that $p_g(C_e) > 0$. We set $e'_j = e$.

Setting $e'_i = e_i$ for $i \neq j$, we have defined a new choice of edges. Denote $T' = \Gamma \backslash \bigcup e'_j$. If we are in case (1), then clearly $p_a(T') = p_a(T) - 1$, whereas $p_a(T') = p_a(T)$ in cases (2) and (3). If we are in case (2), then the number of connected components of T' is one less than the number of components of T. Therefore, we have $p_g(T') = p_g(T) - 1$. In case (3), the number of components stays the same, so $p_g(T') = p_g(T)$. Finally, in case (3) the component $C_{e'_j}$ of $C \backslash e'_j$ is a component of T' of positive genus. It clearly has less edges than C, and hence $m(T') < m(T)$. Summarizing this, we see that the triple $(p_a(T'), p_g(T'), m(T'))$ is strictly smaller than $(p_a(T), p_g(T), m(T))$ in the lexicographic order on $\mathbb{N}^3$. As this is a well-order, we see that after modifying the choice (e_i) of edges finitely many times as described above, we obtain a choice of edges whose complement has geometric genus 0, finishing the proof. $\square$

4   (Sub)modularity of ψ

We will need the characterization of break divisors described in Proposition 3.3 (iii) in terms of the genus contribution function ψ. Here, we record two important (sub)modularity properties of ψ.

Lemma 4.1. *Let Γ be a metric graph. For any two closed admissible subsets S_1 and S_2 of Γ we have*

$$\psi(S_1) + \psi(S_2) = \psi(S_1 \cap S_2) + \psi(S_1 \cup S_2).$$

Proof. By Definition 2.7, it suffices to show that

$$\chi(\Gamma \backslash S_1) + \chi(\Gamma \backslash S_2) = \chi(\Gamma \backslash (S_1 \cap S_2)) + \chi(\Gamma \backslash (S_1 \cup S_2)).$$

This follows from the Mayer-Vietoris sequence, applied to the couple of open subsets $(\Gamma \backslash S_1, \Gamma \backslash S_2)$ of $\Gamma \backslash (S_1 \cap S_2)$ (see Remark 2.8 (ii)). □

We will frequently need to pass to the convex hull. To express $\psi(\mathrm{conv}(S))$ in terms of $\psi(S)$, we need the following definition.

Definition 4.2.

(i) Let S be an admissible subset of Γ. We define $e(S)$ to be the number of open edge segments in $\Gamma \backslash S$ whose endpoints are contained in S.

(ii) For two admissible subsets S_1 and S_2 of Γ we define $e(S_1, S_2) = e(S_1 \cup S_2)$.

Remark 4.3. If S_1 and S_2 are admissible and *convex*, then $e(S_1, S_2)$ is precisely the number of open edge segments contained in $\Gamma \backslash (S_1 \cup S_2)$ that have one endpoint in $S_1 \backslash S_2$ and one endpoint in $S_2 \backslash S_1$.

Lemma 4.4. *Let Γ be a metric graph.*

(a) Let S be a closed admissible subset of Γ. Then we have

$$\psi(\mathrm{conv}(S)) = \psi(S) - e(S).$$

(b) For any two closed admissible subsets S_1 and S_2 of Γ we have

$$\psi(S_1) + \psi(S_2) = \psi(S_1 \cap S_2) + \psi(\mathrm{conv}(S_1 \cup S_2)) + e(S_1, S_2)$$

Proof. As noted after Definition 2.4, one obtains $\Gamma \backslash \mathrm{conv}(S)$ from $\Gamma \backslash S$ by removing from it all connected components that are contained in an open edge. These connected components are precisely the open edge segments in $\Gamma \backslash S$ whose endpoints are contained in S, so there are $e(S)$ of them. The topological Euler characteristic of each open edge segment is 1. So by the additivity of the topological Euler characteristic on disjoint unions, we have $\chi(\Gamma \backslash \mathrm{conv}(S)) = \chi(\Gamma \backslash S) - e(S)$. Part (a) now follows from the definition of ψ (Definition 2.7 (ii)). Part (b) follows from part (a) and Lemma 4.1. □

5 Error functions and error sets

Definition 5.1. Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}_+^d(\Gamma)$, with $0 \le d \le g$.

   (i) For a closed admissible subset $S \subseteq \Gamma$ we define the *D-error of S* as the integer

$$\mathrm{Err}(D, S) = \deg(D|_S) - \psi(S).$$

   (ii) The *D-max error* is defined to be the integer

$$\mathrm{ME}(D) = \max\{\mathrm{Err}(D, S) \colon S \subsetneq \Gamma \text{ closed and admissible}\}.$$

   (iii) A *D-max error set* is a closed and admissible (not necessarily proper) subset $S \subseteq \Gamma$ with $\mathrm{Err}(D, S) = \mathrm{ME}(D)$.

Remark 5.2. In the context of finite graphs, the function χ studied in [ABKS14, §4.1] is closely related to a combinatorial analogue of our function Err.

Lemma 5.3. *Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}_+^d(\Gamma)$, with $0 \le d \le g$.*

   (a) For any two closed admissible subsets S_1 and S_2 of Γ we have:

$$\mathrm{Err}(D, S_1) + \mathrm{Err}(D, S_2) = \mathrm{Err}(D, S_1 \cap S_2) + \mathrm{Err}(D, S_1 \cup S_2)$$

   and

$$\mathrm{Err}(D, S_1) + \mathrm{Err}(D, S_2) + e(S_1, S_2) \le \mathrm{Err}(D, S_1 \cap S_2) + \mathrm{Err}(D, \mathrm{conv}(S_1 \cup S_2)).$$

   (b) We have $\mathrm{ME}(D) \ge 0$.

   (c) D is a break divisor if and only if $\deg(D) = g$ and $\mathrm{ME}(D) = 0$.

   (d) $\mathrm{Err}(D, \Gamma) > \mathrm{ME}(D)$ if and only if D is a break divisor.

Proof. The equality in part (a) is a combination of Lemma 4.1 and the obvious fact that

$$\deg(D|_{S_1}) + \deg(D|_{S_2}) = \deg(D|_{S_1 \cap S_2}) + \deg(D|_{S_1 \cup S_2}).$$

To prove the inequality we use the same fact about degrees and combine it with Lemma 4.4 (b) and the fact that $\deg(D|_{\mathrm{conv}(S_1 \cup S_2)}) \ge \deg(D_{S_1 \cup S_2})$ because D is effective.

   Part (b) follows from $\mathrm{Err}(D, \emptyset) = 0$.

   Part (c) follows Proposition 3.3 and part (b).

   For part (d), we first note that $\mathrm{Err}(D, \Gamma) = d - g + 1$. If D is a break divisor, we have $\mathrm{Err}(D, \Gamma) = 1$ but $\mathrm{ME}(D) = 0$ by part (c). Conversely, if $\mathrm{Err}(D, \Gamma) > \mathrm{ME}(D)$, then the nonnegativity of $\mathrm{ME}(D)$ implies that $d - g + 1 \ge 1$. This implies that $d = g$, and hence that $\mathrm{ME}(D) = 0$. By part (c), we conclude that D is a break divisor. $\square$

Lemma 5.4. *Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}_+^d(\Gamma)$, with $0 \le d \le g$. Assume $\mathrm{ME}(D) > 0$.*

(a) *If S is a D-max error set, then S is convex.*

(b) *If S_1 and S_2 are two D-max error sets, then $S_1 \cup S_2$ and $S_1 \cap S_2$ are also D-max error sets. Moreover $e(S_1, S_2) = 0$.*

Proof. (a) If S is not convex, then $\psi(\mathrm{conv}(S)) < \psi(S)$ by Lemma 4.4 (a). Since $S \subseteq \mathrm{conv}(S)$ we always have $\deg(D|_{\mathrm{conv}(S)}) \geq \deg(D|_S)$. It follows that

$$\mathrm{Err}(D, \mathrm{conv}(S)) > \mathrm{Err}(D, S) = \mathrm{ME}(D).$$

If $\mathrm{conv}(S) \neq \Gamma$ this is a contradiction (see Definition 5.1 (ii)). If $\mathrm{conv}(S) = \Gamma$ then

$$1 \geq \mathrm{Err}(D, \Gamma) > \mathrm{ME}(D) > 0,$$

which, again, is a contradiction.

(b) First, we observe that $\mathrm{Err}(D, S_1 \cup S_2) \leq \mathrm{ME}(D)$. If $S_1 \cup S_2 \neq \Gamma$ this follows directly from Definition 5.1 (ii). If $S_1 \cup S_2 = \Gamma$, then it follows from Lemma 5.3(d) and the fact that D is not a break divisor by assumption.

Together with Definition 5.1 (ii) and Lemma 5.3 (a), it follows that

$$\begin{aligned} 2\,\mathrm{ME}(D) &\geq \mathrm{Err}(D, S_1 \cap S_2) + \mathrm{Err}(D, S_1 \cup S_2) \\ &= \mathrm{Err}(D, S_1) + \mathrm{Err}(D, S_2) \\ &= 2\,\mathrm{ME}(D). \end{aligned}$$

Therefore, $\mathrm{Err}(D, S_1 \cap S_2) = \mathrm{Err}(D, S_1 \cup S_2) = \mathrm{ME}(D)$, that is $S_1 \cap S_2$ and $S_1 \cup S_2$ are D-max error sets. By part (a), it follows that $S_1 \cup S_2$ is convex and hence that $e(S_1, S_2) = e(S_1 \cup S_2) = 0$. □

Proposition 5.5. *Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}^d_+(\Gamma)$, with $0 \leq d \leq g$. Then there exists a unique smallest (with respect to inclusion) D-max error set in Γ. If G is the model corresponding to the vertex set $V(G) = V_\Gamma \cup \mathrm{supp}(D)$, then this smallest D-max error set is of the form*

$$I \cup \bigcup_{e \in J} \bar{e} \tag{3}$$

for some $I \subseteq V(G)$ and $J \subseteq E(G)$.

Proof. Let S be a D-max error set. Suppose that there exists a point in the boundary ∂S of S that is contained in an open edge $e \in E(G)$. Since S is convex, $e \backslash S$ is either an open edge segment or a disjoint union of two open edge segments. So, by the Mayer-Vietoris sequence (see Remark 2.8 (ii)), we have

$$\chi(\Gamma \backslash (S \backslash e)) = \chi(\Gamma \backslash S) + \chi(e) - \chi(e \backslash S) \leq \chi(\Gamma \backslash S)$$

and hence $\psi(S \backslash e) \leq \psi(S)$. As we also have $\deg(D|_S) = \deg(D|_{S \backslash e})$ by definition of G, we see that $\mathrm{Err}(D, S \backslash e) \geq \mathrm{Err}(D, S)$ and hence that $S \backslash e$ is a D-max error set. This shows that every D-max error set contains a D-max error set of the form (3). As there are only finitely many sets of this form, we see that every D-max error set contains an inclusion minimal D-max error set, and that all of these are of the form (3). By Lemma 5.4 (b), the intersection of all minimal D-max error sets is also a D-max error set, hence this is the unique smallest D-max error set. □

Definition 5.6. We will denote the unique minimal D-max error subset of Γ (as in Proposition 5.5) by $\mathfrak{I}(D)$.

Remark 5.7. $\mathfrak{I}(D)$ is always a proper subset of Γ, as the maximum error is taken for proper (admissible) subsets. We have $\mathrm{ME}(D) = 0$ if and only if $\mathfrak{I}(D) = \emptyset$.

Our next goal is to prove a key result about $\mathfrak{I}(D)$. We recall two standard notations. Let (X, d) be a metric space.

- For two nonempty subsets $A, B \subseteq X$, one defines

$$\mathrm{dist}(A, B) = \inf\{d(x, y) \colon x \in A, y \in B\}.$$

- For a nonempty subset $A \subseteq X$, one defines its ϵ-fattening by

$$A_\epsilon = \bigcup_{x \in A} \{z \in X \colon d(z, x) \leq \epsilon\}.$$

Proposition 5.8. *Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}_+^d(\Gamma)$, with $0 \leq d \leq g$. Let $S = \mathfrak{I}(D)$, and assume that $\mathrm{ME}(D) > 0$.*

(a) Let $\epsilon = \mathrm{dist}\left(S, V_\Gamma \backslash S\right)$. Let D_1 be the divisor obtained from D by 'firing' S to distance ϵ. In other words, $D_1 = D + \mathrm{div}(\phi_{S,\epsilon})$, where $\phi_{S,\epsilon}$ is as in Remark 2.9. Then

- *(i) D_1 is effective.*
- *(ii) $\mathrm{ME}(D_1) \leq \mathrm{ME}(D)$.*
- *(iii) If $\mathrm{ME}(D_1) = \mathrm{ME}(D)$, then $S_\epsilon \subseteq \mathfrak{I}(D_1)$.*

(b) Let $e^- \in \partial S$ and $e^+ \in V_\Gamma \backslash S$ be endpoints of an open edge segment $e \subseteq \Gamma \backslash S$. Let D_2 be the divisor obtained from D by moving a chip from e^- to e^+. In other words, $D_2 = D + (e^+) - (e^-)$. Then

- *(i) D_2 is effective.*
- *(ii) $\mathrm{ME}(D_2) \leq \mathrm{ME}(D)$.*
- *(iii) If $\mathrm{ME}(D_2) = \mathrm{ME}(D)$, then $S \cup \{e^+\} \subseteq \mathfrak{I}(D_2)$.*

Proof. We observe that neither can S be empty, since $\mathrm{ME}(D) > 0$, nor can it contain V_Γ, as this would imply $S = \mathrm{conv}(S) = \Gamma$, a contradiction (see Remark 5.7).

Claim 1. For every $p \in S$ we have $\mathrm{val}_S(p) \leq D(p)$.

Proof of Claim 1. If $p \notin \partial S$ we have $\mathrm{val}_S(p) = 0$ and there is nothing to prove. Let $p \in \partial S$. Let B be a sufficiently small open neighborhood of p, isometric to a star-shaped set and not containing any point in $V_\Gamma \cup \mathrm{supp}(D)$ aside from p.

By the minimality of S, we know $S \backslash B$ is not a D-max error set, so

$$\mathrm{Err}(D, S) \geq \mathrm{Err}(D, S \backslash B) + 1. \tag{4}$$

By the choice of B, we have

$$\deg(D|_{S\backslash B}) = \deg(D|_S) - D(p) \, . \tag{5}$$

Furthermore, we have

$$\chi(\Gamma\backslash(S\backslash B)) = \chi(\Gamma\backslash S) + \chi(B) - \chi(B\backslash S)$$

by Mayer-Vietoris (Remark 2.8 (ii)), which equals $\chi(\Gamma\backslash S) + 1 - \operatorname{val}_S(p)$ by the choice of B. Therefore,

$$\psi(S\backslash B) = \psi(S) + 1 - \operatorname{val}_S(p) \, . \tag{6}$$

Combining (4), (5), and (6), we obtain

$$\begin{aligned}
\operatorname{Err}(D, S) &\geq \operatorname{Err}(D, S\backslash B) + 1 \\
&= \deg(D|_{S\backslash B}) - \psi(S\backslash B) + 1 \\
&= \big(\deg(D|_S) - D(p)\big) - \big(\psi(S) + 1 - \operatorname{val}_S(p)\big) + 1 \\
&= \operatorname{Err}(D, S) - D(p) + \operatorname{val}_S(p) \, ,
\end{aligned}$$

from which we deduce that $\operatorname{val}_S(p) \leq D(p)$.

(a) Note that $\epsilon = \operatorname{dist}(S, V_\Gamma\backslash S)$ is well-defined because S and $V_\Gamma\backslash S$ are nonempty. For any $p \in S$, we have $D_1(p) = D(p) - \operatorname{val}_S(p)$. So the effectiveness of D_1 follows directly from Claim 1.

Let $\mathbf{U}$ denote the the set of all closed edge segments e of length ϵ with endpoints $\partial e = \{e^-, e^+\}$ such that $e^- \in S$ and $\operatorname{dist}(e^+, S) = \epsilon$. 'Firing' S to distance ϵ has the effect of sending one chip from e^- to e^+ for each $e \in \mathbf{U}$. So, for any admissible subset $R \subseteq \Gamma$ we have

$$\deg(D_1|_R) = \deg(D|_R) - |\{e \in \mathbf{U} \colon e^- \in R, e^+ \notin R\}| + |\{e \in \mathbf{U} \colon e^- \notin R, e^+ \in R\}| \, . \tag{7}$$

Let $R = \mathfrak{I}(D_1)$. By Lemma 5.3 (a) and (7) we have:

$$\begin{aligned}
\operatorname{Err}(D, S \cap R) + \operatorname{Err}(D, \operatorname{conv}(S \cup R)) &\geq \operatorname{Err}(D, S) + \operatorname{Err}(D, R) + e(S, R) \\
&\geq \operatorname{Err}(D, S) + \operatorname{Err}(D_1, R) + e(S, R) - |\{e \in \mathbf{U} \colon e^- \notin R, e^+ \in R\}| \, ,
\end{aligned} \tag{8}$$

with equality only if $\{e \in \mathbf{U} \colon e^- \in R, e^+ \notin R\} = \emptyset$.

Claim 2. $|\{e \in \mathbf{U} \colon e^- \notin R, e^+ \in R\}| \leq e(S, R)$.

Proof of Claim 2. For any $e \in \mathbf{U}$ with $e^- \notin R$ and $e^+ \in R$ we have $e\backslash\{e^-\} \nsubseteq R$ because R is closed. Since R is also convex (Lemma 5.4 (a)), e contains a unique connected component of the complement of $S \cup R$ in its interior. This component is an open edge segment which has one endpoint in $S\backslash R$ and one endpoint in $R\backslash S$. It therefore contributes with 1 to $e(S, R)$, proving the claim.

By (8), Claim 2, and the fact that $\operatorname{Err}(D, \operatorname{conv}(S \cup R)) \leq \operatorname{ME}(D)$ (see proof of Lemma 5.4 (b)), we obtain:

$$\begin{aligned}
2\operatorname{ME}(D) &\geq \operatorname{Err}(D, S \cap R) + \operatorname{Err}(D, \operatorname{conv}(S \cup R)) \\
&\geq \operatorname{Err}(D, S) + \operatorname{Err}(D_1, R) \\
&= \operatorname{ME}(D) + \operatorname{ME}(D_1) \, .
\end{aligned}$$

It follows that $\mathrm{ME}(D_1) \leq \mathrm{ME}(D)$. In case of equality, R is also a D-max error set and thus, by Lemma 5.4 (b), so is $S \cap R$. Because $S = \mathfrak{I}(D)$, we must have $S \subseteq R$. And since (8) is an equality, we have $\{e \in \mathbf{U} \colon e^- \in R, e^+ \notin R\} = \emptyset$. Therefore, R must contain all points of Γ that have distance ϵ to S. By the convexity of R (Lemma 5.4 (a)), it follows that R does in fact contain all points of distance at most ϵ to S.

(b) It follows directly from Claim 1 that D_2 is also effective. Let $Q = \mathfrak{I}(D_2)$. We have four cases:

Case 1: $e^-, e^+ \in Q$. We have

$$\mathrm{ME}(D_2) = \mathrm{Err}(D_2, Q) = \mathrm{Err}(D, Q) \leq \mathrm{ME}(D).$$

In case of equality Q is D-max error set and thus contains both e^+ and S.

Case 2: $e^-, e^+ \notin Q$. As Q does not contain S, it is not a D-max error set. Therefore,

$$\mathrm{ME}(D_2) = \mathrm{Err}(D_2, Q) = \mathrm{Err}(D, Q) < \mathrm{ME}(D).$$

Case 3: $e^- \in Q$, $e^+ \notin Q$. We have

$$\mathrm{ME}(D_2) = \mathrm{Err}(D_2, Q) = \mathrm{Err}(D, Q) - 1 < \mathrm{ME}(D).$$

Case 4: $e^- \notin Q$, $e^+ \in Q$. We have $\mathrm{ME}(D_2) = \mathrm{Err}(D_2, Q) = \mathrm{Err}(D, Q) + 1$. Since Q is closed, the open edge segment e is not contained in $S \cup Q$. It follows that e contains a connected component of $\Gamma \backslash (S \cup Q)$. As such an edge segment is automatically an open edge segment with endpoints in $S \cup Q$, we have $e(S, Q) > 0$ and hence

$$\mathrm{Err}(D, S \cap Q) + \mathrm{Err}(D, \mathrm{conv}(S \cup Q)) \geq \mathrm{Err}(D, S) + \mathrm{Err}(D, Q) + e(S, Q)$$
$$\geq \mathrm{ME}(D) + \mathrm{ME}(D_2) - 1 + 1.$$

From this, and the fact $\mathrm{Err}(D, \mathrm{conv}(S \cup Q)) \leq \mathrm{ME}(D)$ (see proof of Lemma 5.4 (b)), we conclude

$$2\,\mathrm{ME}(D) \geq \mathrm{Err}(D, S \cap Q) + \mathrm{Err}(D, \mathrm{conv}(S \cup Q)) \geq \mathrm{ME}(D) + \mathrm{ME}(D_2).$$

It follows that $\mathrm{ME}(D_2) \leq \mathrm{ME}(D)$. Equality is not possible in this case because Q does not contain S.

<div align="right">□</div>

6   Semibreak divisors in effective divisor classes

6.1   Existence of semibreak divisors

We are now ready to prove our main theorem about the existence of semibreak divisors in effective classes.

Theorem 6.1. *Let Γ be a metric graph of genus g. Let $D \in \mathrm{Div}^d(\Gamma)$, with $0 \leq d \leq g$, and assume $|D| \neq \emptyset$. Then $|D|$ contains a semibreak divisor.*

Proof. The result is straightforward for metric circles so, as before, we will assume Γ is not homeomorphic to a circle. Since $|D| \neq \emptyset$, we may also assume that D is effective.

Let E be any effective divisor of degree $g - d$. If $\mathrm{ME}(D + E) = 0$, we are done by Lemma 5.3 (c). If $\mathrm{ME}(D + E) > 0$, let $S = \mathfrak{I}(D + E)$. We will show how to construct a pair (D', E') of effective divisors such that

- $D' \sim D$ and

- either $\mathrm{ME}(D' + E') < \mathrm{ME}(D + E)$, or $\mathrm{ME}(D' + E') = \mathrm{ME}(D + E)$ and $\mathfrak{I}(D' + E')$ contains more branch points of Γ than S.

Since Γ has only finitely many branch points, and a convex subset of Γ containing all its branch points must be equal to Γ, this will prove the theorem. We consider two cases:

(1) $\mathrm{supp}(E) \cap \partial S \neq \emptyset$. Then there exists an open edge segment $e \subseteq \Gamma \backslash S$ with endpoints $e^- \in \mathrm{supp}(E) \cap \partial S$ and $e^+ \in V_\Gamma \backslash S$. Set $D' = D$ and $E' = E - (e^-) + (e^+)$. Both D' and E' are effective by construction, and $D' \sim D$. By Proposition 5.8 (b), we have $\mathrm{ME}(D' + E') \leq \mathrm{ME}(D + E)$, and if there is equality, then $\mathfrak{I}(D' + E')$ contains more branch points of Γ than S.

(2) $\mathrm{supp}(E) \cap \partial S = \emptyset$. Let $\epsilon = \mathrm{dist}(S, V_\Gamma \backslash S)$, and consider the divisor obtained from $D + E$ by 'firing' S to distance ϵ, i.e. $D + E + \mathrm{div}(\phi_{S,\epsilon})$ (see Remark 2.9). By Proposition 5.8 (a), $D + E + \mathrm{div}(\phi_{S,\epsilon})$ is effective. Since $\partial S \cap \mathrm{supp}(E) = \emptyset$, this implies that $D' = D + \mathrm{div}(\phi_{S,\epsilon})$ is also effective. Let $E' = E$. Then, again by Proposition 5.8 (a), we have $\mathrm{ME}(D' + E') \leq \mathrm{ME}(D + E)$, and if there is equality, then $\mathfrak{I}(D' + E')$ contains more branch points of Γ than S.

$\square$

6.2   Uniqueness issues

The existence of semibreak divisors (Theorem 6.1) is sufficient for the applications considered in §8. However, it is natural to wonder about uniqueness of such representatives.

By a *cut C* in a metric graph Γ we mean a disjoint union of open edge segments that disconnects Γ. The *size* of a cut C, denote by $\mathrm{size}(C)$, is the number of connected components (maximal open edge segments) of C. If $S \subseteq \Gamma$ is a closed admissible subset then, for sufficiently small $\epsilon > 0$, the set $\{x \in \Gamma \backslash S : \mathrm{dist}(x, S) < \epsilon\}$ forms a cut. We say that such a cut is determined by S. The size of a cut determined by S does not depend on any choices, and will be denoted by $c(S)$.

We start with two useful lemmas.

Lemma 6.2. *Let Γ be a metric graph, and let $S \subseteq \Gamma$ be a closed admissible set. Then*

$$c(S) = \psi(S) - p_a(S) + 1.$$

Proof. Let C be a cut determined by S. Because all components of C are open edge segments we have

$$\chi(C) = c(S). \tag{9}$$

Let $S' = S \cup C$. Then S is a deformation retract of the open and admissible set S'. Therefore

$$\chi(S) = \chi(S'). \tag{10}$$

Applying the Mayer-Vietoris sequence (see Remark 2.8 (ii)) to the pair $(S', \Gamma \backslash S)$ yields

$$\chi(C) = \chi(S') + \chi(\Gamma \backslash S) - \chi(\Gamma). \tag{11}$$

The result follows from (9), (10), and (11). $\qquad\qquad\qquad\qquad\qquad\qquad\qquad \square$

The following result is a generalized version of the 'maximum principle' (see e.g. [BS13, Lemma 3.7])

Lemma 6.3. *Let Γ be a metric graph and $\phi \in R(\Gamma)$. Let S be the subset of Γ where ϕ attains its minimum. Then*

(a) S is closed and admissible.

(b) For any $p \in S$ we have $-\operatorname{div}(\phi)(p) \geq \operatorname{val}_S(p)$.

Proof. Part (a) follows from the fact that ϕ is continuous, and only changes its slope finitely many times on each closed edge. For part (b), note that if $p, q \in S$ then $\phi(p) = \phi(q)$ whereas if $p \in S$ but $q \in S_\epsilon \backslash S$ (for a sufficiently small ϵ), then the outgoing slope of ϕ from p to q is at least 1. Therefore $-\operatorname{div}(\phi)(p) \geq \operatorname{val}_S(p)$. $\quad \square$

It is known that there is a unique break divisor representative in any equivalence class of divisors in degree g ([MZ08, ABKS14]). Here we give a new proof of this fact, which is better suited for the study of semibreak divisors.

Proposition 6.4. *Let Γ be a metric graph of genus g. If $D, D' \in \operatorname{Div}^g(\Gamma)$ are two distinct break divisors then $D \not\sim D'$.*

Proof. Suppose, for a contradiction, that there exist two distinct linearly equivalent break divisors D and D'. Then $D' = D + \operatorname{div}(\phi)$ for some $\phi \in R(\Gamma)$. Let S be the closed admissible subset of Γ where ϕ attains its minimum, and let C be a sufficiently small cut determined by S such that

$$C \cap \operatorname{supp}(D) = \emptyset \quad , \quad C \cap \operatorname{supp}(\operatorname{div}(\phi)) = \emptyset. \tag{12}$$

Let $S' = S \cup C$. Then S is a deformation retract of the admissible open set S', and in particular $p_a(S') = p_a(S)$. By (12) and Proposition 3.3 we obtain

$$\deg(D'|_S) = \deg(D'|_{S'}) \geq p_a(S') = p_a(S).$$

By the definition of S and Lemma 6.3 we have

$$-\deg(\operatorname{div}(\phi)|_S) \geq c(S).$$

Together with Lemma 6.2, we conclude:

$$\deg(D|_S) = \deg(D'|_S) - \deg(\operatorname{div}(\phi)|_S) \geq p_a(S) + \psi(S) - p_a(S) + 1 \geq \psi(S) + 1 \,,$$

This implies, by Proposition 3.3 (iii), that D cannot be a break divisor, which is a contradiction. □

The above argument does not guarantee the uniqueness of semibreak representatives even for degree $g - 1$. Notice that indeed, an effective divisor class can have more than one semibreak divisor (see Figure 2).
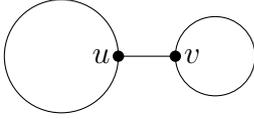

Figure 2: Linearly equivalent semibreak divisors: $(u) \sim (v)$.

By a slight modification of our proof of Proposition 6.4, we obtain the following sufficient condition for the uniqueness of semibreak divisors in their equivalence classes.

Proposition 6.5. *Let Γ be a metric graph of genus g. Fix $0 \leq d \leq g$. Assume for each cut C of Γ we have $\operatorname{size}(C) \geq d+1$. If $D, D' \in \operatorname{Div}^d(\Gamma)$ are two distinct semibreak divisors then $D \not\sim D'$.*

Proof. Suppose, for the sake of contradiction, that there exist two linearly equivalent semibreak divisors D and D' of degree d. Then $D' = D + \operatorname{div}(\phi)$ for some $\phi \in R(\Gamma)$. Let S be the closed admissible subset of Γ where ϕ attains its minimum. By the definition of S and Lemma 6.3 we have

$$- \deg(\operatorname{div}(\phi)|_S) \geq c(S) \,.$$

As D' is a semibreak divisor, it is effective, thus

$$\deg(D|_S) + \deg(\operatorname{div}(\phi)|_S) = \deg(D'|_S) \geq 0 \,.$$

Consequently,

$$d \geq \deg(D|_S) \geq - \deg(\operatorname{div}(\phi)|_S) \geq c(S) \geq d + 1 \,,$$

which is a contradiction. □

Remark 6.6. Proposition 6.5 should be compared with [BN07, Theorem 1.8]. In fact, in the proof, we only use the fact that D and D' are effective. In other words, we are precisely proving the metric graph analogue of [BN07, Theorem 1.8].

The following result will be useful later, and its proof is a slight modification of the previous proof.

Lemma 6.7. *Let Γ be a metric graph, and let D and D' be distinct effective divisors on Γ with $D \sim D'$. Then there exists a closed admissible subset S of Γ such that for every $p \in S$ we have $D(p) \geq \operatorname{val}_S(p)$.*

Proof. By assumption, we have $D' = D + \operatorname{div}(\phi)$ for some $\phi \in R(\Gamma)$. Let S be the closed admissible subset of Γ where ϕ attains its minimum. By the definition of S and Lemma 6.3, for any $p \in S$ we have

$$- \operatorname{div}(\phi)(p) \geq \operatorname{val}_S(p) \,.$$

As D' is effective, it follows that

$$D(p) + \operatorname{div}(\phi)(p) = D'(p) \geq 0 \,.$$

Consequently,

$$D(p) \geq - \operatorname{div}(\phi)(p) \geq \operatorname{val}_S(p) \,.$$

$\square$

6.3 Integral semibreak divisors

Let $G = (V(G), E(G))$ be a finite (unweighted) graph. Let Γ be a metric graph of genus g, obtained from G by turning each edge to an edge segment of length 1. Let $\operatorname{Div}^d(G)$ denote those elements of $\operatorname{Div}^d(\Gamma)$ that are supported on $V(G)$. We will refer to such divisors as *integral*. Let $|D|_G$ denote the set of all effective integral divisors linearly equivalent to D. It is known that, for an integral divisor D, we have $|D|_G \neq \emptyset$ if and only if $|D| \neq \emptyset$ [HKN13, Luo11].

The following result implies that there is an entirely integral version of Theorem 6.1.

Proposition 6.8. *Let $D \in \operatorname{Div}^d(G)$, with $0 \leq d \leq g$, and assume $|D|_G \neq \emptyset$. Then there exists an integral semibreak divisor in $|D|$. Moreover, each semibreak divisor in $|D|$ is integral.*

Proof. The existence part follows from the proof of Theorem 6.1: we can choose E such that $\operatorname{supp}(E) \subseteq V(G)$. Then $D + E$ is still integral. To obtain our semibreak divisor, we successively apply steps (1) or (2). It is enough to show that in the above case these both give integral break divisors. For step (1), this is trivial. For step (2), notice that for an integral divisor $D + E$, the set $\partial \mathfrak{I}(D + E)$ is contained in $\operatorname{supp}(D + E)$, and hence is contained in $V(G)$. It follows that $\epsilon = \operatorname{dist}(S, V_\Gamma \backslash S)$ is an integer and hence that D' and E' will still be integral divisors, proving the existence part of the statement.

Let us suppose for a contradiction that there also exists a non-integral break divisor D'' in $|D|$. This means that there exists $p \notin V(G)$ such that $D''(p) > 0$, which implies that $D''(p) = 1$, as D'' is a semibreak divisor. Let us suppose that the two vertices of the edge of p are u and v. Take some $\phi \in R(\Gamma)$ such that $D'' = D' + \operatorname{div}(\phi)$. Without loss of generality we can assume that $\phi(u) = k$ is an integer. Then $D''(p) = 1$ implies that the slope of the segment between u and p and the slope of the segment between p and v differ by one. Let us suppose that the slope of the segment between u and p is $t \in \mathbb{Z}$, and let $\operatorname{dist}(u, p) = a$ (which is not an integer). Then $\phi(p) = k + at$ and $\phi(v) = k + at + (1-a)(t-1) = k + t - 1 + a$. Hence $\phi(v)$ is not an integer. Let $S = \{w \in V(G) : \phi(w) \in \mathbb{Z}\}$. Then $S \subsetneq V(G)$. Let C be the cut determined by S, i.e.,

the set of edges of G where exactly one endpoint is from S. As S is a proper subset of $V(G)$ and G is connected, C is nonempty. If $\operatorname{div}(\phi)$ is constant zero on an edge of G, then the value of ϕ on the two endpoints differ by an integer, since the slopes are integer and the length of each edge is one. Hence each edge of the cut C needs to have an interior point q where $\operatorname{div}(\phi)(q) \neq 0$. As D' is integral, this means that $\operatorname{div}(\phi)(q) > 0$ on each of these points, and thus D'' has positive number of chips in the interior of each edge of C, contradicting the fact that D'' is a semibreak divisor. $\qquad\square$

Remark 6.9. It follows from Proposition 6.8 that there is an entirely finite graph version of the theory of semibreak divisors. For example, any effective divisor of degree at most the genus on a finite graph G is linearly equivalent to some semibreak divisor. One could directly use a combinatorial analogue of our constructions to prove this (slightly weaker) result. Alternatively, one can give a completely combinatorial proof using the theory of orientations on graphs as described in [ABKS14] and [Bac17] (see loc. cit. for the notation and terminology): let $D \in \operatorname{Div}_+^d(G)$ with $0 \leq d \leq g$. Pick a point $q \in \Gamma$. By [Bac17, Theorem 5.7] we have $D - (q) \sim D_{\mathcal{O}}$ for some 'q-connected partial orientation' $\mathcal{O}$. For any unoriented edge in $\mathcal{O}$, pick an arbitrary orientation to obtain a (full) q-connected orientation $\mathcal{O}'$. By construction $D_{\mathcal{O}} \leq D_{\mathcal{O}'}$. Let $E = D_{\mathcal{O}'} + (q)$. This is a break divisor by [ABKS14, Lemma 3.3]. Let $D' = D_{\mathcal{O}} + (q)$ which is dominated by E.

7 Computational aspects

In this section we show that for an effective divisor, we can find a linearly equivalent semibreak divisor in polynomial time if the input data can be given by rational numbers.

Suppose that the metric graph Γ has rational edge lengths. We call a divisor D rational, if for all p with $D(p) \neq 0$, the distance of p from each branch point is rational.

We will encode a rational number a/b by the pair (a, b), where a and b are integers, but they need not be relatively prime. We encode a metric graph in the following way: For each branch point we write down the list of edges incident to it, along with the edge lengths. We encode a point of an edge by writing down which edge it is on, and what is its distance from one of the endpoints. We encode a divisor D by writing down $D(p)$ along with the encoding of p for each point $p \in \Gamma$ such that $D(p) \neq 0$. We only work with effective divisors of degree at most g, hence we can suppose that $D(p) \leq g$ for each p. We might need more space for encoding a divisor than for encoding the graph if the distance of some p with $D(p) \neq 0$ has a large denominator. However, as the numbers $D(p)$ are at most g, this is the only factor that can make the code of a divisor large. We will need to encode closed convex sets. Let the spanning set of a closed convex set S be the following:

$$\operatorname{spset}(S) = \bigcup_{e \in E_\Gamma} \{\partial(e \cap S)\},$$

where $\partial\emptyset = \emptyset$. Then $\operatorname{spset}(S)$ is a finite set of points, and $S = \operatorname{conv}(\operatorname{spset}(S))$. We encode S by giving the points of $\operatorname{spset}(S)$.

Remark 7.1. The encoding of a metric graph outlined above has polynomial size in the genus plus the number of bits needed for writing down the edge lengths. A more succinct encoding of metric graphs would be to only write down how many edges of each length connect two given vertices. However, to write down divisors we need to be able to distinguish between edges connecting the same pair of branch points, hence we need the first type of encoding.

In the case of discrete graphs (i.e., in the model where edge lengths are one and we are not allowed to place chips on the interior points of edges), the model of only encoding edge multiplicities is more natural. Using this more succinct encoding, computational problems are potentially more difficult. In particular, in the integral setting, deciding whether $|D| \neq \emptyset$ can be done in polynomial time in the less succinct model ([BS13]), but it is open whether it can be done in polynomial time in the succinct model.

Theorem 7.2. *If Γ is a metric graph with rational edge lengths, and D is an effective rational divisor, then a semibreak divisor linearly equivalent to D can be found in polynomial time.*

Proof. The trivial case of Γ being a circle can once again be excluded.

We need to be able to do the procedure in Theorem 6.1 algorithmically: at the first step, we can choose $E = (g - d) \cdot p$ where $d = \deg(D)$ and p is an arbitrary branch point of Γ. Then $D + E$ is still rational, and the largest denominator in the encoding did not grow.

The next step is to find $\mathrm{ME}(D + E)$ and $S = \mathfrak{I}(D + E)$. We address this issue later. If we have $\mathfrak{I}(D + E)$, we have to decide whether $\mathrm{supp}(E) \cap \partial S \neq \emptyset$. This can be done since ∂S has at most $|E_\Gamma|$ points. If $\mathrm{supp}(E) \cap \partial S \neq \emptyset$, the operations of case (1) can be trivially done in polynomial time and the resulting D and E are rational. Moreover, the largest denominator in the encoding does not grow. If $\mathrm{supp}(E) \cap \partial S = \emptyset$, then we need to find $\epsilon = \mathrm{dist}(S, V_\Gamma \setminus S)$. For this, we need to check distances along polynomially many edges. Note that ∂S is contained in $\mathrm{supp}(D + E)$, hence all the distances between ∂S and $V_\Gamma \setminus S$ are rational. Thus, ϵ is also rational, and so are the updates D' and E'. The effect of the firing can also be computed in polynomial time, and as we only add or subtract distances, the largest denominator in the encoding of $D + E$ does not grow.

We need to update the divisors D and E polynomially many times (i.e., the loop in the proof of Theorem 6.1 is executed polynomially many times): after any update, $\mathrm{ME}(D + E)$ does not increase, and if it does not decrease, then the number of branch points in $\mathfrak{I}(D + E)$ increases. At the beginning, $\mathrm{ME}(D + E)$ is at most g, as $\deg(D|_S) \leq \deg(D) \leq g$, and $\psi(S) \geq 0$. Hence there are at most $g \cdot |V_\Gamma|$ updates.

It is left to show that one can find $\mathrm{ME}(D+E)$ and $\mathfrak{I}(D+E)$ in polynomial time. For any divisor $D+E$ and convex set S, $\mathrm{Err}(D+E, S)$ can be computed in polynomial time, hence it is enough to find $\mathfrak{I}(D + E)$, and then $\mathrm{ME}(D + E) = \mathrm{Err}(D + E, \mathfrak{I}(D + E))$.

For finding $\mathfrak{I}(D + E)$, we will use submodular minimization. For a finite set A, a set function $f : 2^A \to \mathbb{R}$ is called submodular if

$$f(X \cap Y) + f(X \cup Y) \leq f(X) + f(Y),$$

for all subsets X and Y of A. It is known [Sch00, IFF01] (see also [Sch03, Chapter 45]) that, if f takes rational values and for any set $X \subseteq A$ the value of $f(X)$ can be computed in polynomial time, then a set minimizing f can be found in polynomial time. Schrijver and Iwata–Fleischer–Fujishige give combinatorial strongly polynomial algorithms [Sch00, IFF01] that achieve this goal.

By Lemma 5.3 (a), for two closed convex sets S_1 and S_2,

$$\mathrm{Err}(D + E, S_1 \cap S_2) + \mathrm{Err}(D + E, \mathrm{conv}(S_1 \cup S_2)) \geq \mathrm{Err}(D + E, S_1) + \mathrm{Err}(D + E, S_2).$$

This essentially means that $(-1) \cdot \mathrm{Err}(D + E, .)$ is a submodular set function, and we need to find a smallest minimizing set for it. Though Err is defined for infinitely many sets, we show how to turn the problem into a finite setting, and then we can apply a submodular minimization algorithm.

Let us take a more refined model G of Γ, where $V(G) = V_\Gamma \cup \mathrm{supp}(D + E)$. Then $|V(G)|$ is still polynomial in the input size. By Proposition 5.5, $\mathfrak{I}(D + E)$ is a convex set which is the union of vertices and closed edges of G, hence it is enough to look for $\mathfrak{I}(D + E)$ among these sets. We can introduce conv_G, as the convex hull with respect to the model G, i.e., $\mathrm{conv}_G(S)$ is the union of S and all closed edge segments of the model G where both endpoints are from S.

For a set S which is the union of vertices and closed edges of G, let $\mathrm{spset}_G(S) = S \cap V(G)$. Notice that if S is convex in G, and it is the union of vertices and closed edges of G, then $S = \mathrm{conv}_G(\mathrm{spset}_G(S))$. Also, if $S_1 \subseteq S_2$, then $\mathrm{conv}_G(S_1) \subseteq \mathrm{conv}_G(S_2)$.

Let us define the set function

$$f \colon 2^{V(G)} \to \mathbb{R}, \quad X \mapsto (-1) \cdot \mathrm{Err}(D + E, \mathrm{conv}_G(X)).$$

By the argument above, $\mathfrak{I}(D + E) = \mathrm{conv}(X)$ for the smallest f-minimizing set X. We claim that f is submodular. This follows from the variant of Lemma 4.4, where we replace conv with conv_G (the proof is completely analogous) and the elementary fact that

$$\mathrm{conv}_G(X \cup Y) = \mathrm{conv}_G(\mathrm{conv}_G(X) \cup \mathrm{conv}_G(Y)) \quad, \text{ and}$$
$$\mathrm{conv}_G(X \cap Y) = \mathrm{conv}_G(X) \cap \mathrm{conv}_G(Y)$$

for all subsets X and Y of $V(G)$.

A submodular minimization algorithm only gives us a minimizing set, and not necessarily a smallest one. To solve this problem, take $f'(X) = f(X) + \frac{|X|}{2|V(G)|}$. Then only the fractional part of $f'(X)$ depends on the additional term, hence a set minimizing f' is a minimizing set for f, and an inclusion-minimal (and hence smallest by Proposition 5.5) among those. It is easy to check that f' is also submodular, and the values of f' are still computable in polynomial time. $\qquad\square$

8 Generic effective divisor classes

8.1 Effective loci

Let Γ be a metric graph of genus g. Recall, for $D \in \mathrm{Div}(\Gamma)$, its *rank* $r(D)$ is defined by the properties that $r(D) = -1$ if $|D| = \emptyset$, and $r(D) \geq s \geq 0$ if for all $E \in \mathrm{Div}_+^s(\Gamma)$

we have $|D - E| \neq \emptyset$. Clearly, the rank of a divisor D only depends on its linear equivalence class $[D]$.

The *tropical Riemann-Roch theorem* of [GK08, MZ08] (see also [BN07]) states that

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

where $K = \sum_{p \in \Gamma} (\mathrm{val}(p) - 2)(p)$.

The *tropical Jacobian* of Γ can be defined as the g-dimensional real torus

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

For each choice of a base point $q \in \Gamma$, there is a natural, continuous map $\Phi_q \colon \Gamma \to \mathrm{Jac}(\Gamma)$ sending q to 0, the *Abel-Jacobi map* (cf. [MZ08, BF11]). This map is piecewise linear, which means that locally on Γ it factors through a piecewise linear map to the vector space $H_1(\Gamma, \mathbb{R})$. As $\mathrm{Jac}(\Gamma)$ is a topological group, we may use its addition and the map Φ_q to define maps

$$\Phi_q^{(d)} \colon \mathrm{Div}^d(\Gamma) \to \mathrm{Jac}(\Gamma)$$

for $d \geq 0$. Of course, the composite $\Gamma^d \to \mathrm{Div}^d(\Gamma) \to \mathrm{Jac}(\Gamma)$ will still be piecewise linear. As Γ^d is compact and $\mathrm{Jac}(\Gamma)$ is Hausdorff, it follows from *closed map lemma* (see e.g. [Lee13, Lemma A.52]) that $\Gamma^d \to \mathrm{Jac}(\Gamma)$ is a closed map. In particular, the *effective locus* $\widetilde{W}_d = \Phi_q^{(d)}(\mathrm{Div}_+^d(\Gamma))$ is a closed polyhedral subset of $\mathrm{Jac}(\Gamma)$. It follows from the tropical Riemann-Roch theorem that, for $d \geq g$, we have $\widetilde{W}_d = \mathrm{Jac}(\Gamma)$.

If we denote

$$\mathrm{Pic}^d(\Gamma) = \mathrm{Div}^d(\Gamma)/\mathrm{Prin}(\Gamma) \, ,$$

then it is the content of the tropical Abel-Jacobi theorem that $\Phi_q^{(d)}$ factors through the natural map $S^{(d)} \colon \mathrm{Div}^d(\Gamma) \to \mathrm{Pic}^d(\Gamma)$, and that the induced morphism $\mathrm{Pic}^d(\Gamma) \to \mathrm{Jac}(\Gamma)$ is a bijection ([MZ08]). We endow $\mathrm{Pic}^d(\Gamma)$ the topology inherited from this bijection. Under this bijection, the effective locus $\widetilde{W}_d$ corresponds to the locus of *effective divisors classes* W_d, i.e. those divisors classes $[D] \in \mathrm{Pic}^d(\Gamma)$ in degree d such that $|D| \neq \emptyset$.

8.2   Generic semibreak divisors in effective loci

We are interested in *generic* properties of W_d, i.e. properties that hold on a dense open subset of W_d.

Let $0 \leq d \leq g$. The set $\mathrm{Div}_+^d(\Gamma)$ is endowed with the quotient topology coming from its identification with Γ^d modulo the action of the symmetric group $\mathfrak{S}_d$.

Let $\mathbb{SB}_d \subseteq \mathrm{Div}_+^d(\Gamma)$ denote the set of all semibreak divisors of degree d. Its preimage in Γ^d is the union of all sets of the form $\overline{e}_1 \times \cdots \times \overline{e}_d$, where $e_1, \ldots, e_g$ are distinct open edges of Γ such that $\Gamma \backslash \bigcup_{i=1}^g e_i$ is connected. Therefore, the set $\mathbb{SB}_d$ is *closed* in $\mathrm{Div}_+^d(\Gamma)$. Let Σ_d denote the interior of $\mathbb{SB}_d$. The preimage of Σ_d in Γ^d is the union of all sets of the form $e_1 \times \cdots \times e_d$, with $e_1, \ldots, e_g$ as above. In particular, Σ_d is *open* in $\mathbb{SB}_d$.

Lemma 8.1. *For any $D \in \Sigma_d$, we have $|D| = \{D\}$ and $r(D) = 0$.*

Proof. Let S be a closed admissible subset of Γ, and let C be a cut determined by S. Because $\Gamma \backslash C$ is disconnected, there exists a component e of C such that $D(p) = 0$ for all $p \in \overline{e}$. In particular, for the unique point $p \in \overline{e} \cap S$ we have $D(p) < \mathrm{val}_S(p)$. The statement $|D| = \{D\}$ now follows from Lemma 6.7. If $q \notin \mathrm{supp}(D)$, it follows from $|D| = \{D\}$ that we must have $|D - (q)| = \emptyset$ and therefore $r(D) = 0$. $\qquad\square$

Remark 8.2. One can alternatively show, using the burning algorithm (see e.g. [Luo11]), that any $D \in \Sigma_d$ is D is *universally reduced* (i.e. q-reduced for all $q \in \Gamma$) which is equivalent to having $|D| = \{D\}$ (see [ABKS14, Lemma 4.19]).

Theorem 8.3. *Let Γ be a metric graph of genus g, and fix $0 \leq d \leq g$.*

(a) *The tropical Abel-Jacobi map $S^{(d)} \colon \mathrm{Div}_+^d(\Gamma) \to \mathrm{Pic}^d(\Gamma)$ is 'birational' onto its image. More precisely, there exists an open dense subset $U_d \subseteq W_d$ such that the induced map $(S^{(d)})^{-1}(U_d) \to U_d$ is a homeomorphism.*

(b) *$W_d \subseteq \mathrm{Pic}^d(\Gamma)$ is of pure dimension d.*

Proof. (a) Let $U_d = S^{(d)}(\Sigma_d)$. It follows from Lemma 8.1 that $(S^{(d)})^{-1}(U_d) = \Sigma_d$. By Theorem 6.1 the induced map

$$\mathbb{SB}_d \to W_d$$

is surjective. It is also a closed map (by the *closed map lemma*), because $\mathbb{SB}_d$ is compact and W_d is Hausdorff. In particular, the topology on W_d coincides with the quotient topology.

As Σ_d is dense in $\mathbb{SB}_d$, it follows from the closedness of the map that $U_d = S^{(d)}(\Sigma_d)$ is dense in W_d. It is a direct consequence of Lemma 8.1 that $\Sigma_d = (S^{(d)})^{-1}(U_d)$ and $\Sigma_d \to U_d$ is a bijection. Since Σ_d is open in $\mathbb{SB}_d$, this implies that that U_d is open as well. As closedness is local on the target, the induced map $\Sigma_d \to U_d$ is closed again. And as a continuous and closed bijection it must be a homeomorphism.

For part (b), note that Σ_d, and hence U_d, are purely d-dimensional; every component of Σ_d can be identified with a d-dimensional open polyhedron in Γ^d under the quotient map $\Gamma^d \to \mathrm{Div}_+^d(\Gamma)$. It follows immediately that the closure W_d of U_d is also purely d-dimensional. $\qquad\square$

Remark 8.4.

(i) As W_d is the d-fold sum of W_1, which is easily seen to be purely 1-dimensional, it follows directly from the subadditivity of dimensions of sums that W_d is at most d-dimensional. With the additional ingredient that $W_g = \mathrm{Pic}^g(\Gamma)$ is g-dimensional it even follows that the dimension of W_d is equal to d (cf. [LPP12, Proposition 3.6]). Note that this argument does *not* immediately imply that W_d is purely d-dimensional, as sums of pure-dimensional polyhedral sets are not pure-dimensional in general. However, with some extra care this approach will yield a different proof of the pure-dimensionality, and with some more work even of Theorem 6.1. But unlike our proof, this approach does not yield an algorithm that calculates a semibreak divisor that is linear equivalent to a given effective divisor.

(ii) Another way to prove the pure-dimensionality of W_d is by tropicalization. It is well-known that there exists a Mumford curve C whose Berkovich analytification has Γ as its skeleton. Combining the results [BR15, Theorem 1.3] and [Gub07, Theorem 6.9] then yields the statement. Of course, this approach is highly non-constructive.

Theorem 8.5. *Let Γ be a metric graph of genus g and let d be a nonnegative integer. Then there exists an open dense subset $U_d \subseteq W_d$ of the effective locus such that, for $[D] \in U_d$, we have*

$$r(D) = \begin{cases} d - g & \text{if } d > g, \\ 0 & \text{if } 0 \leq d \leq g. \end{cases}$$

Proof. The case $d > g$ is an elementary consequence of the tropical Riemann-Roch Theorem:

- If $d \geq 2g - 1$ we can take

$$U_d = W_d = \mathrm{Pic}^d(\Gamma).$$

  Since $\deg(K) = 2g - 2$, for every degree d divisor D we have $r(K - D) = -1$, hence $r(D) = d - g$ by the tropical Riemann-Roch.

- If $g < d \leq 2g - 2$ we can take

$$U_d = \mathrm{Pic}^d(\Gamma) - \left(S^{(2g-2)}(K) - W_{2g-2-d} \right).$$

  We claim U_d is a dense open subset of $W_d = \mathrm{Pic}^d(\Gamma)$. This follows from the fact that W_{2g-2-d} is a closed polyhedral subset of $\mathrm{Pic}^d(\Gamma)$, of dimension *at most* $2g - 2 - d < g$. This certainly follows from Theorem 8.3 (b), but it is more elementary and follows directly from definitions.

  If $[D] \in U_d$ then, by definition of U_d, $K - D$ is not equivalent to an effective divisor. Therefore, $r(K - D) = -1$ and $r(D) = d - g$ by the tropical Riemann-Roch Theorem.

Assume $d \leq g$. In this case, we may take $U_d = S^{(d)}(\Sigma_d)$ which is open and dense in W_d by Theorem 8.3. If $[D] \in U_d$, then D is linear equivalent to some $D' \in \Sigma_d$. It follows from Lemma 8.1 that $r(D') = 0$, finishing the proof. $\square$

Remark 8.6. Unlike in algebraic geometry, a property that holds generically on W_d does not automatically hold generically for $\mathrm{Div}^d_+(\Gamma)$. This is because the tropical Abel-Jacobi map $S^{(d)}$ may contract facets of $\mathrm{Div}^d_+(\Gamma)$. For example, if Γ is a chain of two loops, as depicted in Figure 2, and e is the *bridge* (i.e the edge connecting the two circles) then all divisors of the form $(p) + (q)$ with $p, q \in e$ are linear equivalent and of rank 1. Of course, the set of these divisors has nonempty interior and hence there does not exist a dense open subset of $\mathrm{Div}^2_+(\Gamma)$ where the rank is 0. On the other hand, by Theorem 8.5, there exists a dense open subset $U_2 \subseteq W_2$ such that $r(D) = 0$ whenever $[D] \in U_2$.

References

[ABKS14]   Yang An, Matthew Baker, Greg Kuperberg, and Farbod Shokrieh. Canonical representatives for divisor classes on tropical curves and the matrix-tree theorem. *Forum Math. Sigma*, 2:e24, 25, 2014.

[Bac17]   Spencer Backman. Riemann-Roch theory for graph orientations. *Adv. Math.*, 309:655–691, 2017.

[Bak08]   Matthew Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008. With an appendix by Brian Conrad.

[Ber90]   Vladimir G. Berkovich. *Spectral theory and analytic geometry over non-Archimedean fields*, volume 33 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1990.

[BF11]   Matthew Baker and Xander Faber. Metric properties of the tropical Abel-Jacobi map. *J. Algebraic Combin.*, 33(3):349–381, 2011.

[BJ16]   Matthew Baker and David Jensen. Degeneration of linear series from the tropical point of view and applications. In *Nonarchimedean and tropical geometry*, Simons Symp., pages 365–433. Springer, [Cham], 2016.

[BN07]   Matthew Baker and Serguei Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215(2):766–788, 2007.

[BPR16]   Matthew Baker, Sam Payne, and Joseph Rabinoff. Nonarchimedean geometry, tropicalization, and metrics on curves. *Algebr. Geom.*, 3(1):63–105, 2016.

[BR15]   Matthew Baker and Joseph Rabinoff. The skeleton of the Jacobian, the Jacobian of the skeleton, and lifting meromorphic functions from tropical to algebraic curves. *Int. Math. Res. Not. IMRN*, (16):7436–7472, 2015.

[BS13]   Matthew Baker and Farbod Shokrieh. Chip-firing games, potential theory on graphs, and spanning trees. *J. Combin. Theory Ser. A*, 120(1):164–182, 2013.

[CDPR12]   Filip Cools, Jan Draisma, Sam Payne, and Elina Robeva. A tropical proof of the Brill-Noether theorem. *Adv. Math.*, 230(2):759–776, 2012.

[CJP15]   Dustin Cartwright, David Jensen, and Sam Payne. Lifting divisors on a generic chain of loops. *Canad. Math. Bull.*, 58(2):250–262, 2015.

[GH94]   Phillip Griffiths and Joseph Harris. *Principles of algebraic geometry*. Wiley Classics Library. John Wiley & Sons, Inc., New York, 1994. Reprint of the 1978 original.

[GK08]   Andreas Gathmann and Michael Kerber. A Riemann-Roch theorem in tropical geometry. *Math. Z.*, 259(1):217–230, 2008.

[Gub07]   Walter Gubler. Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.*, 169(2):321–376, 2007.

[HKN13]   Jan Hladký, Daniel Král, and Serguei Norine. Rank of divisors on tropical curves. *J. Combin. Theory Ser. A*, 120(7):1521–1538, 2013.

[IFF01]   Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777, 2001.

[JP14]   David Jensen and Sam Payne. Tropical independence I: Shapes of divisors and a proof of the Gieseker-Petri theorem. *Algebra Number Theory*, 8(9):2043–2066, 2014.

[JP16]   David Jensen and Sam Payne. Tropical independence II: The maximal rank conjecture for quadrics. *Algebra Number Theory*, 10(8):1601–1640, 2016.

[Kĺ3]   Tamás Kálmán. A version of Tutte's polynomial for hypergraphs. *Adv. Math.*, 244:823–873, 2013.

[Lee13]   John M. Lee. *Introduction to smooth manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2013.

[LPP12]   Chang Mou Lim, Sam Payne, and Natasha Potashnik. A note on Brill-Noether theory and rank-determining sets for metric graphs. *Int. Math. Res. Not. IMRN*, (23):5484–5504, 2012.

[Luo11]   Ye Luo. Rank-determining sets of metric graphs. *J. Combin. Theory Ser. A*, 118(6):1775–1793, 2011.

[MZ08]   Grigory Mikhalkin and Ilia Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and abelian varieties*, volume 465 of *Contemp. Math.*, pages 203–230. Amer. Math. Soc., Providence, RI, 2008.

[Sch00]   Alexander Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *J. Combin. Theory Ser. B*, 80(2):346–355, 2000.

[Sch03]   Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.