

A Tropical Point of View on Weil Reciprocity

Pearson, Anthony

Citation

Pearson, A. (2025). A Tropical Point of View on Weil Reciprocity.

Version: Not Applicable (or Unknown)

License: License to inclusion and publication of a Bachelor or Master Thesis,

2023

Downloaded from: https://hdl.handle.net/1887/4262232

Note: To cite this publication please use the final published version (if applicable).

Master's Thesis

A Tropical Point of View on Weil Reciprocity

by

Anthony Pearson

27 August 2025

Supervisor

Dr. R.S. de Jong



Leiden University Mathematical Institute

Contents

1	Introduction	2
2	Polyhedral Geometry 2.1 Polyhedra and their properties	4
3	Tropical Spaces 3.1 Building Blocks	26 28
4	Metric Graphs	33
5	Weil-Reciprocity	38
6	Conclusion and Discussion	44
\mathbf{R}	eferences	45

1 Introduction

Tropical geometry has emerged over the past two decades as a vibrant area bridging combinatorial methods and classical algebraic geometry. At its core, it studies geometric objects over the tropical semiring $\mathbb{T} := \mathbb{R} \cup \{-\infty\}$ equipped with the binary operations

$$x \oplus y := \max\{x, y\}, \text{ and } x \odot y := x + y,$$

where \oplus plays the role of addition and \odot plays the role of multiplication. As the name suggests, these operations make \mathbb{T} into a semiring since the only element with an additive inverse is $-\infty$ since addition is idempotent.

One of the main applications is to study varieties in a valued field by tropicalizing said variety, which retains some of its original structure and is easier to study. We will not go into this aspect; however, a thorough discussion of this can be found in [9].

As a combinatorial analogue, one can expect to be able to take concepts and theorems from algebraic geometry to the world of tropical geometry. Among the rich landscape of tropical analogues of classical theorems, one can find: the Riemann-Roch theorem as discussed in [5], the tropical analogue to vector bundles as seen in [1] and a tropical analogue for ψ classes as found in [4]. In this paper, we focus on tropicalizing the classical Weil-reciprocity law, a foundational reciprocity statement in algebraic geometry concerning divisors and evaluations of rational functions on curves. Although previous work (notably [8]) has established a tropical statement and proof for tropical curves (metric graphs), the existing argument is somewhat hands-on and does not fully illuminate the analogy with the classical push-pull approach on algebraic curves. Our aim is to develop a proof of the tropical Weil-reciprocity law that closely parallels the classical strategy, thus making explicit the interplay between algebraic and tropical viewpoints and reinforcing the conceptual unity between these worlds.

Concretely, we shall prove:

Theorem 1.0.1 (Tropical Weil-reciprocity). Let Γ be a compact metric graph and f, g be rational functions on Γ . Then

$$f(\operatorname{div}(g)) = g(\operatorname{div}(f))$$

where if $D = \sum_{p \in \Gamma} n_p p$ we define

$$f(D) := \sum_{p \in \Gamma} f(p) \cdot n_p.$$

To achieve a proof that closely parallels the classical strategy, we first recall necessary foundations in polyhedral geometry and define tropical spaces using an atlas of balanced polyhedral charts. Specializing to dimension one, we identify a special class of tropical curves with metric graphs and develop divisors, rational functions, and (pseudo-)harmonic morphisms with push forward and pull back operations analogous to the algebraic case. The core of the proof constructs, for a given rational function, a suitable morphism to a line segment and reduces reciprocity to an explicit verification of reciprocity on said line segment; assembling these pieces yields the general statement on compact metric graphs.

2 Polyhedral Geometry

Unless otherwise stated, all the definitions and theorems throughout this section are taken from [10], except lemma 2.1.6, corollary 2.1.8, corollary 2.1.9, proposition 2.1.11, lemma 2.1.18, proposition 2.1.22. However, all proofs are done by myself with a few exceptions that are referenced when the said exceptions are mentioned.

2.1 Polyhedra and their properties

A function $\kappa : \mathbb{R}^n \to \mathbb{R}$ is called integer affine if $\kappa(x) = \langle m, x \rangle + c$ where $c \in \mathbb{R}$, $m \in \mathbb{Z}^n$ and $\langle \cdot, \cdot \rangle$ is the standard inner product on \mathbb{R}^n . We define the rational half-space corresponding to an integer affine function κ by

$$H_{\kappa} := \{ x \in \mathbb{R}^n : \kappa(x) \ge 0 \}.$$

Then a polyhedron σ in \mathbb{R}^n is the intersection of finitely many rational halfspaces and the faces of σ are given by intersections $\sigma \cap H_{-\kappa}$ where κ is an integer affine function with $\sigma \subset H_{\kappa}$. We call the maximal faces of σ with respect to inclusion facets. It is then clear from the definition that intersections of polyhedra is a polyhedron and the faces of any polyhedron is a polyhedron. We can also define a polyhedron as a set of the form

$$\sigma = \{ x \in \mathbb{R}^n : Ax < b \}$$

for some integer matrix A and a vector $b \in \mathbb{R}^m$ where m is the number of integer affine functions. This notation means that we are taking those $x \in \mathbb{R}^n$ such that the first entry of Ax is at most the first entry of b and so on. The rows of A correspond to the negative of the linear part of each integer affine function making up σ and the entries of b correspond to the constant part. We shall denote a polyhedron given in this form by P(A, b).

We define the dimension of a polyhedron σ to be the dimension of the linear subspace of \mathbb{R}^n spanned by vectors of the form x-y where $x,y \in \sigma$. We denote the linear subspace by $L(\sigma)$ and the dimension by $\dim(\sigma)$, and also we denote $L_{\mathbb{Z}}(\sigma) := L(\sigma) \cap \mathbb{Z}^n$.

According to [12, p.60], we can equivalently define the dimension of a polyhedron σ by considering the dimension of its affine hull:

$$\operatorname{aff}(\sigma) := \left\{ \sum_{i=1}^{k} \alpha_i x_i : 0 < k < \infty, \ x_i \in \sigma, \ \alpha_i \in \mathbb{R}, \ \sum_{i=1}^{k} \alpha_i = 1 \right\}$$

An affine set is a translation of a vector subspace and the dimension of such an affine set is the dimension of the vector subspace. In this case $\operatorname{aff}(\sigma) = s + L(\sigma)$

for some $s \in \mathbb{R}^n$, this is because

$$\operatorname{aff}(\sigma) = \left\{ s + \sum_{i=1}^{k} \alpha_i(x_i - s) : 0 < k < \infty, \ x_i \in \sigma, \ \alpha_i \in \mathbb{R}, \ \sum_{i=1}^{k} \alpha_i = 1 \right\}$$

and each $x_i - s$ is an element of $L(\sigma)$ for any choice of $s \in \sigma$ so $aff(\sigma) \subset s + L(\sigma)$. Conversely, since any difference x - y for $x, y \in \sigma$ can be written as x - y = (x - s) - (y - s) we may write any $v \in L(\sigma)$ as

$$v = \sum_{i=1}^{k} \alpha_i (x_i - s)$$

for some coefficients $\alpha_i \in \mathbb{R}$. Therefore, we may write

$$s + v = s + \sum_{i=1}^{k} \alpha_i (x_i - s) = \sum_{i=1}^{k} \alpha_i x_i + s \left(1 - \sum_{i=1}^{k} \alpha_i \right)$$

which lies inside $\operatorname{aff}(\sigma)$ by definition since the sum of the coefficients is one; so we conclude that $\operatorname{aff}(\sigma) = s + L(\sigma)$. The advantage of defining the dimension like this is that we always have $\sigma \subset \operatorname{aff}(\sigma)$ and so by taking the image of σ under an affine transformation $\operatorname{aff}(\sigma) \cong \mathbb{R}^{\dim(\sigma)}$ we may always assume that σ is full-dimensional in its ambient space.

Remark 2.1.1. If σ is a polyhedron in \mathbb{R}^n then the empty set and σ itself are faces of σ since we have

$$\sigma = \sigma \cap \mathbb{R}^n = \sigma \cap \{x \in \mathbb{R}^n : \langle 0, x \rangle \ge 0\},$$

$$\emptyset = \sigma \cap \emptyset = \sigma \cap \{x \in \mathbb{R}^n : \langle 0, x \rangle \ge 1\}.$$

Then we see that σ lies inside $H_{-\kappa}$ for both, because in either case $H_{-\kappa} = \mathbb{R}^n$ as $H_{-\kappa} = \{x \in \mathbb{R}^n : \langle 0, x \rangle \leq c\}$ where c = 0 or 1 depending on the case.

Example 2.1.2. A square in \mathbb{R}^2 with vertices $(\pm 1, \pm 1)$ is a polyhedron since it is the intersection of rational half-spaces corresponding to the functions $\kappa_1(x,y) = x+1$, $\kappa_2(x,y) = -x+1$, $\kappa_3(x,y) = y+1$ and $\kappa_4(x,y) = -y+1$. Then, as one might expect, the faces of the square include the edges, as we can see by looking at the rational half-spaces corresponding to $-\kappa_i$ for each i. However, we note that the vertices are also faces, for example, by looking at the rational half-space corresponding to -y-x-2 which gives the vertex (-1,-1).

The following definition is taken from [6, p.26].

Definition 2.1.3. Let $\sigma = H_{\kappa_1} \cap \cdots \cap H_{\kappa_t}$ then we say that this is an irredundant representation of σ , provided that for each $i = 1, \ldots, t$ we have $\sigma \subsetneq \sigma_i := \cap_{i=1, j \neq i}^t H_{\kappa_j}$ is a strict inclusion of sets.

Proposition 2.1.4. If $\sigma = H_{\kappa_1} \cap \cdots \cap H_{\kappa_t}$ is an irredundant representation, then the collection $\{\sigma \cap H_{-\kappa_i}\}_{1 \le i \le t}$ consists of all facets of σ .

Proof. See [6, Theorem 2.6.3].
$$\Box$$

We call a collection $\mathcal{X} = \{\sigma_1, \ldots, \sigma_t\}$ of polyhedra in \mathbb{R}^n a polyhedral complex if every non-empty face of each σ_i is an element of \mathcal{X} and every non-empty intersection $\sigma_i \cap \sigma_j$ is a face of both σ_i and σ_j . The support of \mathcal{X} is given by $|\mathcal{X}| := \bigcup_i \sigma_i$ and if $|\mathcal{X}|$ is equal to a polyhedron σ then we call \mathcal{X} a polyhedral subdivision of σ . If all maximal polyhedra (with respect to inclusion) in \mathcal{X} have dimension n, then we say that \mathcal{X} is pure of dimension n. The elements of a polyhedral complex \mathcal{X} are called cells, and the maximal elements are called facets.

Example 2.1.5. If κ is an integer affine function, then the collection

$$\mathcal{H}_{\kappa} := \{H_{\kappa}, H_{-\kappa}, H_{\kappa} \cap H_{-\kappa}\}$$

is a polyhedral complex with support \mathbb{R}^n . This will be justified in the proof of corollary 2.1.9.

Lemma 2.1.6. Let σ, σ' be polyhedra in \mathbb{R}^n and F and F' be faces of σ and σ' respectively. Then $F \cap F'$ is a face of $\sigma \cap \sigma'$.

Proof. Let F and F' be associated with the half-spaces H_{κ} and $H_{\kappa'}$ respectively. We now note that the only way that $x \in F$ or F' is if $\kappa(x) = 0$ or $\kappa'(x) = 0$ since this is the only place where H_{κ} and $H_{-\kappa}$ coincide. Now we note that for $x \in \sigma$ that $\kappa(x) = 0$ and $\kappa'(x) = 0$ if and only if $\kappa(x) + \kappa'(x) = 0$ since $\kappa(x) \geq 0$ and $\kappa'(x) \geq 0$ by assumption. It follows that

$$F \cap F' = (\sigma \cap H_{\kappa}) \cap (\sigma \cap H_{\kappa'}) = \sigma \cap H_{\kappa + \kappa'}.$$

This implies $F \cap F'$ is a face of σ after noting $\kappa(x) + \kappa'(x) \ge 0$ for all $x \in \sigma$. \square

Lemma 2.1.7. If \mathcal{X} and \mathcal{Y} are two polyhedral complexes in \mathbb{R}^n then the collection

$$\mathcal{X}\cap\mathcal{Y}:=\{\sigma\cap\sigma':\sigma\in\mathcal{X},\sigma'\in\mathcal{Y},\sigma\cap\sigma'\neq\emptyset\}$$

is a polyhedral complex and $|\mathcal{X} \cap \mathcal{Y}| = |\mathcal{X}| \cap |\mathcal{Y}|$.

Proof. We first note that $\mathcal{X} \cap \mathcal{Y}$ is finite because \mathcal{X} and \mathcal{Y} are finite, and consist of polyhedra because intersections of polyhedra are polyhedra. Let $\sigma \cap \sigma' \in \mathcal{X} \cap \mathcal{Y}$ with $\sigma \in \mathcal{X}$ and $\sigma' \in \mathcal{Y}$ and F be a face of $\sigma \cap \sigma'$ then there

is a half-space H_{κ} such that $\sigma \cap \sigma' \subset H_{\kappa}$ and $\sigma \cap \sigma' \cap H_{-\kappa} = F$. Now we must have either $\sigma \subset H_{\kappa}$ or $\sigma' \subset H_{\kappa}$ because if H_{κ} cuts σ and σ' through their interiors then it will cut $\sigma \cap \sigma'$ through its interior contrary to the assumption that $\sigma \cap \sigma' \subset H_{\kappa}$. Now assume without loss of generality that $\sigma \subset H_{\kappa}$ then we have

$$F = \sigma \cap \sigma' \cap H_{-\kappa} = (\sigma \cap H_{-\kappa}) \cap \sigma'.$$

Now $\sigma \cap H_{-\kappa}$ is a face of σ hence in \mathcal{X} and $\sigma' \in \mathcal{Y}$ so this implies that $F \in \mathcal{X} \cap \mathcal{Y}$. Now let $\sigma_1 \cap \sigma'_1$ and $\sigma_2 \cap \sigma'_2$ be two elements of $\mathcal{X} \cap \mathcal{Y}$ with $\sigma_1, \sigma_2 \in \mathcal{X}$ and $\sigma'_1, \sigma'_2 \in \mathcal{Y}$ then

$$(\sigma_1 \cap \sigma_1') \cap (\sigma_2 \cap \sigma_2') = (\sigma_1 \cap \sigma_2) \cap (\sigma_1' \cap \sigma_2').$$

Then $\sigma_1 \cap \sigma_2$ is a face of both σ_1 and σ_2 and similarly $\sigma'_1 \cap \sigma'_2$ is a face of both σ'_1 and σ'_2 . It remains to show that if F is a face of $\sigma \in \mathcal{X}$ and F' is a face of $\sigma' \in \mathcal{Y}$ then $F \cap F'$ is a face of $\sigma \cap \sigma'$. However, this is precisely lemma 2.1.6.

We see in lemma 2.1.6 and the proof of lemma 2.1.7 that the faces of an intersection of two polyhedra $\sigma \cap \sigma'$ are exactly of the form $F \cap F'$ where F is a face σ and F' is a face of σ' .

Corollary 2.1.8. Let σ and σ' be polyhedra in \mathbb{R}^n then the faces of $\sigma \cap \sigma'$ are precisely of the form $F \cap F'$ where F is a face of σ and F' is a face of σ' .

Corollary 2.1.9. Any polyhedron σ in \mathbb{R}^n has finitely many faces.

Proof. Since a polyhedron is an intersection of finitely many half-spaces, using corollary 2.1.8, it suffices to show that a single half-space has finitely many faces. Let H_{κ} be a half-space and suppose that $H_{\kappa'}$ contains H_{κ} so that $H_{-\kappa'} \cap H_{\kappa}$ is a face of H_{κ} . Note that κ' must be some translation of κ otherwise $H_{\kappa'}$ could not possibly contain H_{κ} since it would cut through its interior. Therefore, every face of H_{κ} is of the form $H_{-\kappa-q} \cap H_{\kappa}$ for some $q \in \mathbb{R}$. Now, in order for $x \in H_{-\kappa-q} \cap H_{\kappa}$ and $H_{\kappa} \subset H_{\kappa+q}$ to be true, we must have $\kappa(x) = -q$, $q \geq 0$ and $\kappa(x) \geq 0$. This is only possible if q = 0 so the only faces of H_{κ} are \emptyset , H_{κ} and $H_{\kappa} \cap H_{-\kappa}$.

Definition 2.1.10 (Relative interior). Let σ be a polyhedron then we define the relative interior of σ , denoted by relint (σ) , to be the interior of σ when viewed as a subspace of $\operatorname{aff}(\sigma) \cong \mathbb{R}^{\dim(\sigma)}$.

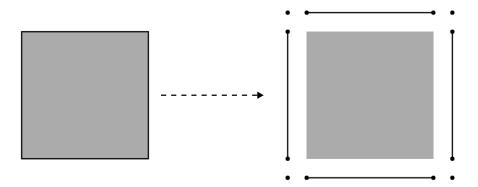


Figure 1: Illustration of the relative interior, which is coloured in gray, of a square and its faces, which are in black.

Figure 1 suggests that we can equivalently define the relative interior as the complement in a polyhedron of all its proper faces. This is indeed the case and will be the content of the following proposition.

Proposition 2.1.11. Let σ be a polyhedron in \mathbb{R}^n then we can equivalently define

$$\operatorname{relint}(\sigma) = \sigma \setminus \bigcup_{F \subsetneq \sigma \text{ face}} F$$

Proof. Without loss of generality we may assume that σ is full-dimensional so that we can work with the topological interior. Let $\operatorname{relint}(\sigma)$ be the relative interior of σ corresponding to the $\operatorname{aff}(\sigma)$ definition and $\operatorname{rfint}(\sigma)$ be the complement in σ of all its proper faces. Our goal is then to show that $\operatorname{relint}(\sigma) = \operatorname{rfint}(\sigma)$. Let $\sigma = H_{\kappa_1} \cap \ldots \cap H_{\kappa_t}$ be an irredundant representation of σ . Then by proposition 2.1.4 we have that

$$\operatorname{rfint}(\sigma) = \sigma \cap \left(\bigcup_{i=1}^{t} \sigma \cap H_{-\kappa_i}\right)^{c}$$

$$= \sigma \cap \bigcap_{i=1}^{t} \sigma^{c} \cup \{x \in \mathbb{R}^{n} : \kappa_i(x) > 0\}$$

$$= \bigcap_{i=1}^{t} \operatorname{int} H_{\kappa_i}$$

$$= \operatorname{int} \bigcap_{i=1}^{t} H_{\kappa_i}$$

$$= \operatorname{int} \sigma.$$

Notice that in the third line onward it should be $\{x \in \sigma : \kappa_i(x) > 0\}$ and not H_{κ_i} however, it is fine to replace σ by \mathbb{R}^n here because any x that lies in all these intersections is already in σ by definition.

A polyhedral set X in \mathbb{R}^n is a finite union of polyhedra in \mathbb{R}^n . Write $X = \sigma_1 \cup \cdots \cup \sigma_t$ then if each σ_i happens to be of dimension k then we say that X has pure dimension k. If Y is an open subset of X (with respect to the subspace topology) then we call Y an open polyhedral set. A polyhedral structure of X is a polyhedral complex \mathcal{X} such that $|\mathcal{X}| = X$. Given a polyhedral complex \mathcal{X} and any set $S \subset \mathbb{R}^n$ then we define $\mathcal{X}|_{S} := \{\sigma \in \mathcal{X} : \sigma \subset S\}$ which is clearly seen to be a polyhedral complex.

We state the following proposition and its proof due to [10, Proposition 4.1.1].

Proposition 2.1.12. Let $X = \sigma_1 \cup \cdots \cup \sigma_t$ be a polyhedral set in \mathbb{R}^n . Then there exists a polyhedral structure \mathcal{X} of X such that $\mathcal{X}|_{\sigma_i}$ is a subdivision of σ_i for each i.

Proof. Let $\kappa_1, \ldots, \kappa_m$ be a complete collection of integer affine functions such that any σ_i can be represented by intersecting half-spaces corresponding to appropriate choices of κ_j . We consider the subdivision of \mathbb{R}^n given by

$$\mathcal{G} = \mathcal{H}_{\kappa_1} \cap \cdots \cap \mathcal{H}_{\kappa_m}$$

(where \mathcal{H}_{κ_i} is as in example 2.1.5) which is a polyhedral complex by lemma 2.1.7. Then we claim that the polyhedral complex $\mathcal{X} := \mathcal{G}|_X$ satisfies our desired properties. To prove this it is enough to show that $\mathcal{X}|_{\sigma_i}$ is a polyhedral subdivision of σ_i since this would imply that $|\mathcal{G}|_X|$ includes each σ_i hence it also includes X by definition. Let $\sigma = \sigma_1$ and suppose that $\sigma = H_{\kappa_1} \cap \cdots \cap H_{\kappa_l}$. Then $\mathcal{G}|_{\sigma}$ includes (but is certainly not limited to) all elements of the form

$$\sigma \cap \bigcap_{i>l} H_{\pm \kappa_i}.$$

Then from this description, taking the union of all these elements it is clear that we achieve σ . The rest of the σ_i follow by an identical argument so we conclude that $\mathcal{G}|_X$ is a polyhedral structure for X.

Definition 2.1.13 (Balancing). Let \mathcal{X} be a pure m-dimensional polyhedral complex in \mathbb{R}^n . Then a weight function on \mathcal{X} is a function that assigns an integer $\omega(\sigma) \in \mathbb{Z}$ to every facet of $\sigma \in \mathcal{X}$. A weighted polyhedral complex (\mathcal{X}, ω) is a pair that consists of a polyhedral complex \mathcal{X} and a weight function ω on \mathcal{X} . If it is clear that we are talking about a weighted polyhedral complex, then we will just simply denote it by \mathcal{X} .

If \mathcal{X} has non-negative weights then we say that it is balanced if for every $\tau \in \mathcal{X}$ of dimension m-1 (or codimension one) the following holds:

$$\sum_{\substack{\sigma \text{ facet} \\ \tau \subset \sigma}} \omega(\sigma) v_{\sigma/\tau} \in L_{\mathbb{Z}}(\tau).$$

Here $v_{\sigma/\tau} \in \mathbb{Z}^n$ denotes an integer vector that points from τ towards σ and satisfies

$$L_{\mathbb{Z}}(\sigma) = L_{\mathbb{Z}}(\tau) + \mathbb{Z}v_{\sigma/\tau}.$$

This vector exists because $L(\sigma)/L(\tau)$ is a one-dimensional vector space, and there are only two choices for $v_{\sigma/\tau}$ and only one faces the correct direction.

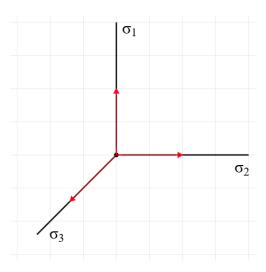


Figure 2: Illustration of the balancing condition.

Example 2.1.14. If we look at figure 2, then place weights one on each facet σ_1 , σ_2 and σ_3 , the result will be a balanced polyhedral complex. This is because if we call $\tau = \{(0,0)\}$ the vertex in the centre (the only codimension one cell), then $v_{\sigma_1/\tau} = (0,1)$, $v_{\sigma_2/\tau} = (1,0)$ and $v_{\sigma_3/\tau} = (-1,-1)$ whose sum is clearly seen to be $(0,0) \in L_{\mathbb{Z}}(\tau)$.

If \mathcal{X} is a weighted polyhedral complex, then we define its support $|\mathcal{X}|$ to be the union of facets in \mathcal{X} with non-zero weight. When talking about the support of a weighted polyhedral set, we will always assume this definition and not take the union of all facets as before.

Definition 2.1.15 (Equivalence of complexes). Let \mathcal{X} and \mathcal{Y} be two weighted polyhedral complexes in \mathbb{R}^n . We say that \mathcal{X} and \mathcal{Y} are equivalent if $|\mathcal{X}| = |\mathcal{Y}|$ and whenever relint $(\sigma) \cap \text{relint}(\tau) \neq \emptyset$ where $\sigma \in \mathcal{X}$ and $\tau \in \mathcal{Y}$ are facets, their weights agree.

A weighted polyhedral complex \mathcal{L} (of pure dimension the same as \mathcal{X}) is a refinement of \mathcal{X} if $|\mathcal{L}| = |\mathcal{X}|$ and for any $\sigma \in \mathcal{L}$ with $\sigma \subset |\mathcal{L}|$ (recall that this is not automatic because we exclude weight zero facets in the support), we have that σ is contained in some $\tau \in \mathcal{X}$. If σ is a facet (hence τ must be a facet since \mathcal{L} and \mathcal{X} are pure of the same dimension) then the weights of σ and τ agree.

Example 2.1.16. In figure 3, if we place a common weight of one on each facet of the polyhedral complexes \mathcal{X} and \mathcal{Y} , then we observe that they are equivalent. We note that one is not a refinement of the other because not all facets of one are contained in a facet of the other.

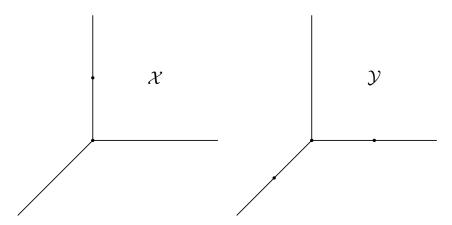


Figure 3: Example of two equivalent polyhedral complexes

Proposition 2.1.17. Let \mathcal{L} be a refinement of a weighted polyhedral complex \mathcal{X} . Then \mathcal{L} is balanced if and only if \mathcal{X} is balanced.

Proof. The direction where we first assume that \mathcal{X} is balanced is done in [9, Lemma 3.6.2]. Conversely, if \mathcal{X} is not balanced, then there is a codimension one cell $\sigma \in \mathcal{X}$ for which the balancing condition does not hold. Let $M_{\mathcal{L},\mathcal{X}}: \mathcal{L} \to \mathcal{X}$ be the function that sends a polyhedron $\tau \in \mathcal{L}$ to the smallest polyhedron in \mathcal{X} that contains it. Then, since \mathcal{L} is a refinement, there must exist a codimension one cell $\tau \in \mathcal{L}$ such that $M_{\mathcal{L},\mathcal{X}}(\tau) = \sigma$. We can then follow the argument in [9, Lemma 3.6.2] to see that \mathcal{L} is not balanced at τ .

Lemma 2.1.18. Let σ and τ be full-dimensional polyhedra in \mathbb{R}^n . Then

 $\operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau) \neq \emptyset \iff \sigma \cap \tau \text{ is full-dimensional in } \mathbb{R}^n.$

Proof. Suppose that $\operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau)$ is non-empty, then we note that $\operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau)$ is an open set in \mathbb{R}^n , hence it spans \mathbb{R}^n . However, we

also see that it is also contained inside $\sigma \cap \tau$ which implies that $\sigma \cap \tau$ must also be full dimensional. Conversely, if $\sigma \cap \tau$ is full-dimensional then

$$\operatorname{relint}(\sigma \cap \tau) = \operatorname{int}(\sigma \cap \tau) = \operatorname{int}(\sigma) \cap \operatorname{int}(\tau) = \operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau).$$

Looking at the extreme right and left we observe that $\operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau)$ is non-empty since $\operatorname{relint}(\sigma \cap \tau)$ is non-empty.

Proposition 2.1.19. Let \mathcal{X} and \mathcal{Y} be two weighted polyhedral complexes. Then \mathcal{X} and \mathcal{Y} are equivalent if and only if they have a common refinement.

Proof. First, assume that \mathcal{X} and \mathcal{Y} are equivalent, then consider the polyhedral complex $\mathcal{Z} := \mathcal{X} \cap \mathcal{Y}$. Note that the facets of \mathcal{Z} are of the form $\sigma \cap \tau$ where $\sigma \in \mathcal{X}$ and $\tau \in \mathcal{Y}$ are facets such that $\operatorname{relint}(\sigma) \cap \operatorname{relint}(\tau) \neq \emptyset$ according to lemma 2.1.18. Therefore, we may place weights on the facets of \mathcal{Z} that match the weights of the facets of \mathcal{X} and \mathcal{Y} that make up its composition. Then it is clear that under this weighting, \mathcal{Z} is a common refinement of \mathcal{X} and \mathcal{Y} .

On the other hand, let \mathcal{L} be a common refinement of \mathcal{X} and \mathcal{Y} then we have

$$|\mathcal{X}| = |\mathcal{L}| = |\mathcal{Y}|$$

so the first condition is clear. Next, let $\sigma \in \mathcal{X}$ and $\tau \in \mathcal{Y}$ with $\mathrm{relint}(\sigma) \cap \mathrm{relint}(\tau) \neq \emptyset$. Taking $p \in \mathrm{relint}(\sigma) \cap \mathrm{relint}(\tau)$, there exists a facet $\rho \in \mathcal{L}$ that contains p with $\rho \subset \sigma'$ and $\rho \subset \tau'$ for some facets $\sigma' \in \mathcal{X}$ and $\tau' \in \mathcal{Y}$. However, note that the relative interiors of polyhedra in a polyhedral complex are disjoint because the intersection of two polyhedra in a complex is a face of both by definition. Therefore, we must have $\sigma = \sigma'$ and $\tau = \tau'$ and since \mathcal{L} is a common refinement, the weight of ρ is the same as those of σ and τ .

We obtain two corollaries for free from combining both proposition 2.1.19 and proposition 2.1.17.

Corollary 2.1.20. Let \mathcal{L} be a common refinement of a weighted polyhedral complex \mathcal{X} . Then \mathcal{X} and \mathcal{L} are equivalent complexes.

Corollary 2.1.21. If the balancing condition holds for a weighted polyhedral complex \mathcal{X} then any equivalent complex is also balanced.

Proposition 2.1.22. Let σ be a polyhedron in \mathbb{R}^n then given a point $p \in \sigma$ we consider the set

$$\operatorname{Star}_{\sigma}(p) := \{ v \in \mathbb{R}^n : p + \varepsilon \cdot v \in \sigma \text{ for } \varepsilon > 0 \text{ sufficiently small} \}.$$

This set is a polyhedron in \mathbb{R}^n .

Proof. If p lies in the interior of σ then this set is simply \mathbb{R}^n so let us assume that p lies on the boundary of σ . Let $\sigma = H_{\kappa_1} \cap \cdots \cap H_{\kappa_l}$ with $\kappa_i(x) = \langle m_i, x \rangle + c_i$ and set $\widetilde{\kappa}_i(x) = \kappa_i(x) - c_i$. Let $\{\kappa_1, \ldots, \kappa_t\}$ be the set of all functions where $\kappa_i(p) = 0$ (this set is non-empty because by assumption p lies on the boundary) then we claim that

$$\operatorname{Star}_{\sigma}(p) = \bigcap_{i=1}^{t} H_{\widetilde{\kappa}_i}.$$

Let $v \in \operatorname{Star}_{\sigma}(p)$ then there exists some $\varepsilon > 0$ such that $\kappa_i(p + \varepsilon \cdot v) \geq 0$. This means that

$$\langle m_i, p \rangle + c_i + \varepsilon \langle m_i, v \rangle \ge 0.$$

Then by assumption we have that $\kappa_i(p) = 0$ so that $\langle m_i, v \rangle \geq 0$ which implies that $\widetilde{\kappa}_i(v) \geq 0$ which further implies that $\operatorname{Star}_{\sigma}(p) \subset H_{\widetilde{\kappa}_1} \cap \cdots \cap H_{\widetilde{\kappa}_t}$. Conversely, let $v \in H_{\widetilde{\kappa}_1} \cap \cdots \cap H_{\widetilde{\kappa}_t}$ and let $\varepsilon > 0$ to be determined then for $1 \leq i \leq t$ we have

$$\kappa_i(p + \varepsilon \cdot v) = \langle m_i, p \rangle + c_i + \varepsilon \langle m_i, v \rangle = \varepsilon \langle m_i, v \rangle \ge 0$$

for any choice of $\varepsilon > 0$, so the issue arises when we consider κ_i for i > t. Let i > t and set $\kappa_i(p) = \delta_i > 0$ and then

$$\kappa_i(p+\varepsilon\cdot v)=\delta_i+\varepsilon\langle m_i,v\rangle.$$

If $\langle m_i, v \rangle \geq 0$ then we are done since any choice of ε will do the job, otherwise let ε be small enough such that

$$\varepsilon \le \frac{\delta_i}{|\langle m_i, v \rangle|}.$$

Then we have that

$$\kappa_i(p+\varepsilon\cdot v) = \delta_i + \varepsilon \langle m_i, v \rangle = \delta_i - \varepsilon |\langle m_i, v \rangle| \ge 0.$$

It then follows that by choosing ε small enough that $\kappa_i(p+\varepsilon \cdot v) \geq 0$ for all i which implies that $p+\varepsilon \cdot v \in \sigma$. This then shows the desired equality, which implies that $\operatorname{Star}_{\sigma}(p)$ is a polyhedron.

Remark 2.1.23. In fact, we see that $\operatorname{Star}_{\sigma}(p)$ is a cone, which means that it is a polyhedron that contains all the rays $\mathbb{R}_{\geq 0}x$ where $x \in \operatorname{Star}_{\sigma}(p)$ since we can always choose a sufficiently small ε .

We call $\operatorname{Star}_{\sigma}(p)$ the star of σ at p. If X is an open polyhedral set we can construct a similar set which we shall also call the star of X at a point $p \in X$.

Definition 2.1.24. Let X be an open polyhedral set of the polyhedral set $Y = \sigma_1 \cup \cdots \cup \sigma_l$ and p a point in X. Then we define

$$\operatorname{Star}_X(p) := \operatorname{Star}_{\sigma_1}(p) \cup \cdots \cup \operatorname{Star}_{\sigma_l}(p)$$

to be the star of X at p.

Note that this definition only depends on the ambient polyhedral set. Indeed, when constructing the star of an open polyhedral set, we only care about what happens locally at p so it does not matter whether we are looking at X or Y to construct it. Furthermore, if $p \notin \sigma_i$ for some i then we note that $\operatorname{Star}_{\sigma_i}(p) = \emptyset$ therefore contributes nothing and if p lies in the relative interior of some polyhedron σ_i then $\operatorname{Star}_{\sigma_i}(p) \cong \mathbb{R}^{\dim(\sigma)}$ as one would expect.

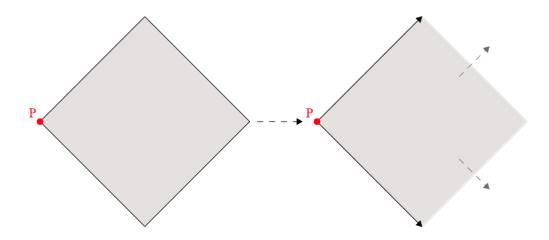


Figure 4: The star of the polyhedron at the point labeled P.

3 Tropical Spaces

The strategy we will proceed with is to define what a tropical space should look like locally, by means of an atlas of charts. Therefore, we shall start with defining exactly what these local building blocks actually are. All definitions are again from [10], however all theorems, lemmas, and propositions in this section are either said in passing in [10, Section 6], or set as an exercise, with a few exceptions. We make those statements precise by giving a precise statement and their proof.

3.1 Building Blocks

Definition 3.1.1 (Open polyhedral sets). Let $\mathbb{T} = \mathbb{R} \cup \{-\infty\}$, topologized to be homeomorphic to a half-line. Define $[n] := \{1, \dots, n\}$ then the space \mathbb{T}^n has a natural decomposition as

$$\mathbb{T}^n = \coprod_{I \subset [n]} \mathbb{R}_I,$$

where

$$\mathbb{R}_I := \{ (x_i) \in \mathbb{T}^n : x_i = -\infty \iff i \in I \}.$$

can be identified with $\mathbb{R}^{n-|I|}$ by removing each x_i for $i \in I$. We also define

$$\mathbb{T}_I := \{(x_i) \in \mathbb{T}^n : x_i = -\infty \text{ for } i \in I\}$$

which can similarly be naturally identified with $\mathbb{T}^{n-|I|}$. Then a rational polyhedron in \mathbb{T}^n is the closure (in \mathbb{T}^n) of a usual polyhedron in $\mathbb{R}_I \cong \mathbb{R}^{n-|I|}$ for some $I \subset [n]$ and we call I the sedentarity of the polyhedron. A finite union of polyhedra in \mathbb{T}^n is a polyhedral set in \mathbb{T}^n and an open set of a polyhedral set is said to be an open polyhedral set in \mathbb{T}^n . Let X be an open polyhedral set, then we say that a point $x \in X$ is of sedentarity I if $x \in X \cap \mathbb{R}_I$. We denote the sedentarity of a point $x \in X$ by $\operatorname{sed}(x)$ and similarly for a polyhedron $\overline{\sigma}$ by $\operatorname{sed}(\overline{\sigma})$.

Example 3.1.2. Figure 5 depicts a polyhedron in \mathbb{T}^n that is induced by the polyhedron in $\mathbb{R}_{\emptyset} \cong \mathbb{R}^2$ that is the intersection of the two half-spaces $\{y \leq 1\}$ and $\{x \leq 1\}$.

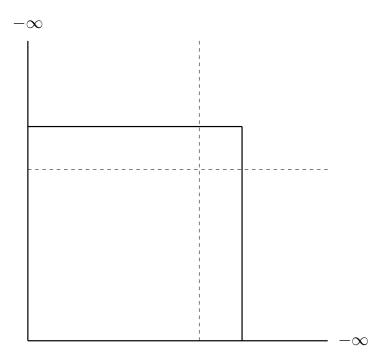


Figure 5: Example of a polyhedron in \mathbb{T}^2 .

If X is an open polyhedral set in \mathbb{T}^n or \mathbb{R}^n then there is a dense open subset of points in X that have an open neighbourhood that is homeomorphic to an open in \mathbb{R}^m for some m. These points are called generic points and the set of them is denoted by X^{gen} . If p is a generic point that has an open neighbourhood that is homeomorphic to an open in \mathbb{R}^m for some m, then we call m the dimension of X at p, denoted by $\dim_X(p)$. Note that this is well-defined, since open subsets of \mathbb{R}^n cannot be homeomorphic to ones in \mathbb{R}^m if $n \neq m$. Now we get a locally constant function $\dim_X : X^{\text{gen}} \to \mathbb{Z}_{\geq 0}$ and we call its maximum the dimension of X, denoted by $\dim(X)$. If \dim_X is constant, then we say that X is of pure dimension $\dim(X)$.

If σ is a polyhedron in \mathbb{R}^n then we have two notions of dimension of σ . The one imposed as above and the dimension of the linear subspace that it generates. We show that these two definitions coincide.

Proposition 3.1.3. Let σ be a polyhedron in \mathbb{R}^n then the dimension of σ when considered as an open polyhedral set coincides with the dimension of the linear subspace spanned by σ .

Proof. Let k_1 be the dimension of σ when considered as an open polyhedral set and k_2 be the dimension of $L(\sigma)$. We firstly have $k_2 \geq k_1$ because if $x \in \sigma$ has an open neighbourhood that is isomorphic to an open of \mathbb{R}^{k_1} then $L(\sigma)$ must at least have dimension k_1 because it contains this open neighbourhood.

Conversely, we see that $\operatorname{relint}(\sigma) \subset \operatorname{aff}(\sigma) \cong \mathbb{R}^{k_2}$ is open in σ hence is isomorphic to an open subset of \mathbb{R}^{k_2} so we have $k_1 \geq k_2$ since k_1 is the maximum of the function $\dim_{\sigma} : \sigma^{\operatorname{gen}} \to \mathbb{Z}$. We conclude that $k_1 = k_2$ as desired.

Definition 3.1.4 (Weights). A weight function on an open polyhedral set X in \mathbb{T}^n is a locally constant function $\omega: X^{\text{gen}} \to \mathbb{Z}$, or in other words a choice of integer for every connected component of X^{gen} . An open polyhedral set in \mathbb{T}^n equipped with a weight function is said to be a weighted open polyhedral set in \mathbb{T}^n , and if all weights are non-negative, then it is said to be effective. The same can be said for open polyhedral sets in \mathbb{R}^n .

We shall henceforth assume that all weights are non-zero.

Proposition 3.1.5. Let σ be a polyhedron in \mathbb{R}_J for some $J \subset [n]$ and $\overline{\sigma}$ its closure in \mathbb{T}^n . Then $\overline{\sigma} \cap \mathbb{R}_I$ for any $I \subset [n]$ is a polyhedron in $\mathbb{R}^{n-|I|}$. Additionally, if $\operatorname{sed}(\overline{\sigma}) \subset I$ we have $\overline{\overline{\sigma} \cap \mathbb{R}_I} = \overline{\sigma} \cap \mathbb{T}_I$.

Proof. If it is not the case that $\operatorname{sed}(\overline{\sigma}) \subset I$ then $\overline{\sigma} \cap \mathbb{R}_I$ is empty (which is a polyhedron) since all points in $\overline{\sigma}$ will have $-\infty$ in coordinates outside of I, so we henceforth assume that $\operatorname{sed}(\overline{\sigma}) \subset I$. When we take the closure of a polyhedron in \mathbb{R}^m then the only points we introduce by taking the closure in \mathbb{T}^n are points with $-\infty$ so $\overline{\sigma} \cap \mathbb{R}_I$ is just the projection of σ to $\mathbb{R}^{n-|I|}$. In particular, it is enough to show that the projection of a polyhedron in \mathbb{R}^n to \mathbb{R}^m is again a polyhedron for which, by induction, it is enough to show this for m = n - 1. Let σ be given by the inequalities

$$\langle m_i, x \rangle + b_i x_n \le a_i$$

where $m_i \in \mathbb{Z}^m$, $b_i \in \mathbb{Z}$, and $a_i \in \mathbb{R}$. Depending on the sign of b_i we get inequalities of the form

$$\langle m'_i, x \rangle + a_i/b_i \ge x_n \text{ if } b_i > 0;$$

 $\langle m'_j, x \rangle + a_j/b_j \le x_n \text{ if } b_j < 0;$
 $\langle m_k, x \rangle \le a_k \text{ if } b_k = 0;$

where $m'_i = -m_i/b_i$ for all i. We can combine the first two to get the set of inequalities

$$\langle m_j' - m_i', x \rangle \le a_i/b_i - a_j/b_j.$$

Then these sets of inequalities and the third type define the projection of σ . Moving onto the second claim, we see that $\overline{\sigma} \cap \mathbb{R}_I \subset \overline{\sigma} \cap \mathbb{T}_I$ since the closure of an intersection is contained in the intersection of the closures. Conversely, if $x \in \overline{\sigma} \cap \mathbb{T}_I$ then there is a sequence $(x_k)_{k \geq 0}$ in $\sigma \subset \mathbb{R}_{\text{sed}(\overline{\sigma})}$ that converges to x. We can then project this sequence down to $\overline{\sigma} \cap \mathbb{R}_I$ which still converges to x since the Ith coordinates of x are all $-\infty$ by assumption.

Corollary 3.1.6. Let $\sigma \subset \mathbb{R}_I$ and $\tau \subset \mathbb{R}_J$ be polyhedra for some $I, J \subset [n]$, then we have

$$\overline{\sigma} \cap \overline{\tau} = \overline{\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}}.$$

In particular, the intersection of polyhedra in \mathbb{T}^n is a polyhedron in \mathbb{T}^n , and if I = J then

$$\overline{\sigma} \cap \overline{\tau} = \overline{\sigma \cap \tau}.$$

Proof. Suppose that we know the proposition when at least one of σ or τ is a half-space, then we know it is true for all polyhedra. Indeed, write $\tau = H_{\kappa_1} \cap \cdots \cap H_{\kappa_t}$ then we have

$$\overline{\sigma} \cap \overline{\tau} = \overline{\sigma} \cap \overline{H_{\kappa_1}} \cap \dots \cap \overline{H_{\kappa_t}}$$

$$= \overline{\sigma} \cap \overline{H_{\kappa_1}} \cap \mathbb{R}_{I \cup J} \cap \overline{H_{\kappa_2}} \cap \dots \cap \overline{H_{\kappa_t}}$$

$$= \overline{\sigma} \cap \overline{H_{\kappa_1}} \cap \overline{H_{\kappa_2}} \cap \mathbb{R}_{I \cup J} \cap \overline{H_{\kappa_3}} \cap \dots \cap \overline{H_{\kappa_t}}$$

$$\vdots$$

$$= \overline{\sigma} \cap \overline{H_{\kappa_1}} \cap \dots \cap \overline{H_{\kappa_t}} \cap \mathbb{R}_{I \cup J}$$

$$= \overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}.$$

The first and last equality uses the case I = J for the decomposition of τ and the rest is a repeated use of proposition 3.1.5 and the special case of the statement of this proposition. Now assume that $\tau = H_{\kappa}$ is a half space then $\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J} \subset \overline{\sigma} \cap \overline{\tau}$ and the right-hand side is closed in \mathbb{T}^n so $\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J} \subset \overline{\sigma} \cap \overline{\tau}$ $\overline{\sigma} \cap \overline{\tau}$. Conversely, letting $x \in \overline{\sigma} \cap \overline{\tau}$ then there exists a sequence $(x_i)_{i \geq 0}$ in $\overline{\sigma} \cap \mathbb{R}_{I \cap J}$ that converges to x, and if $x \in \mathbb{R}_K$ for some $K \supset I \cup J$ then without loss of generality we may assume that the K coordinates of x_i are strictly decreasing towards $-\infty$ by choosing a subsequence of $(x_i)_{i>0}$. Then looking at the proof of proposition 3.1.5 we see that $\overline{\tau} \cap \mathbb{R}_{I \cap J}$ is a half-space in $\mathbb{R}_{I \cap J}$ which we will denote by $H_{\kappa'}$. Then necessarily the linear part of κ' has all coordinates in K negative and if we focus on the $[n]\setminus K$ coordinates of $\kappa'(x_i)$ as $i\to\infty$ then it is bounded while the K coordinate approach ∞ and so will eventually dominate the bounded portion for large enough i. Therefore, for large enough i we will have $\kappa'(x_i) \geq 0$ which implies that $x_i \in H_{\kappa'}$ for sufficiently large i. Now we have that for sufficiently large i every $x_i \in \overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}$ so that $x \in \overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}$.

Now for the other claim we note that

$$\overline{\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}} = \overline{(\overline{\sigma} \cap \mathbb{R}_{I \cup J}) \cap (\overline{\tau} \cap \mathbb{R}_{I \cup J})}$$

which is a closure of the intersection of two polyhedra in $\mathbb{R}_{I\cup J}$. Then for the other claim we see that

$$\overline{\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_I} = \overline{(\overline{\sigma} \cap \mathbb{R}_I) \cap (\overline{\tau} \cap \mathbb{R}_I)} = \overline{\sigma \cap \tau}$$

since for polyhedra $\sigma \subset \mathbb{R}_I$ we always have $\sigma = \overline{\sigma} \cap \mathbb{R}_I$.

Corollary 3.1.7. Let X be a polyhedral set in \mathbb{T}^n then we can decompose X as $\bigcup_{I\subset [n]}X_I$ where each X_I is a polyhedral set in \mathbb{T}_I .

Proof. Define $X_I := \overline{X \cap \mathbb{R}_I}$ (where the closure is taken in \mathbb{T}_I) then if $X = \sigma_1 \cup \cdots \cup \sigma_l$ where each σ_i is a polyhedron in \mathbb{T}^n then

$$\overline{X \cap \mathbb{R}_I} = \overline{\sigma_1 \cap \mathbb{R}_I} \cup \dots \cup \overline{\sigma_l \cap \mathbb{R}_I}$$

is a polyhedral set in $\underline{\mathbb{T}_I}$ since each $\sigma_i \cap \mathbb{R}_I$ is a polyhedron in $\mathbb{R}^{n-|I|}$ by proposition 3.1.5, therefore $\overline{\sigma_i \cap \mathbb{R}_I}$ is a polyhedron in \mathbb{T}_I . Additionally, we note that $\sigma_i \cap \mathbb{T}_I$ is closed in \mathbb{T}_I hence $\overline{\sigma_i \cap \mathbb{R}_I} \subset \sigma_i \cap \mathbb{T}_I$. Therefore,

$$\overline{X \cap \mathbb{R}_I} \subset X \cap \mathbb{T}_I \subset X.$$

Moreover, any point $x \in X$ lies in some \mathbb{R}_I hence $X = \bigcup_{I \subseteq [n]} X_I$ as desired. \square

Corollary 3.1.8. Let X be an open polyhedral set in \mathbb{T}^n then we can decompose X as $\bigcup_{I \subset [n]} X_I$ where each X_I is an open polyhedral set in \mathbb{T}_I .

Proof. Suppose that X is an open subset of the polyhedral set Y in \mathbb{T}^n . Then it is enough to show that $X_I := \overline{X \cap \mathbb{R}_I}$ is an open subset of $Y_I = \overline{Y \cap \mathbb{R}_I}$ since Y_I is a polyhedral set in \mathbb{T}_I by corollary 3.1.7. To that end, let $x \in X_I$ then if $x \in X \cap \mathbb{R}_I$ then we can find an open containing it that is fully contained inside $X \cap \mathbb{R}_I$ since $X \cap \mathbb{R}_I$ is open in $Y \cap \mathbb{R}_I$. Then the only problem occurs if x is a limit point of $X \cap \mathbb{R}_I$, in which case there are some coordinates of x that are $-\infty$ apart from those in I. Let $J \subset [n]$ be sedentarity of x and denote the jth coordinates by x_j for $j \in J$. Then there exists some set of the form $[-\infty, l_j)$ such that $\{x_1\} \times \cdots \times \{x_{j-1}\} \times [-\infty, l_j) \times \{x_{j+1}\} \times \cdots \times \{x_n\} \subset X_I$. It cannot be the case that no matter how small l_j we make that this is not fully contained in X_I because this would necessarily subdivide a region around x into infinitely many parts in Y_I which is impossible since Y_I is only a finite union of polyhedra. We do this for every x_j for $j \in J$ and for the remaining ones we just take a usual open in $X \cap \mathbb{R}_I$ as it is already open in $Y \cap \mathbb{R}_I$. Then taking the product of all of these opens will be an open around x contained in X_I .

Definition 3.1.9. Let $\overline{\sigma}$ be a polyhedron in \mathbb{T}^n where σ is a polyhedron in \mathbb{R}_I for some $I \subset [n]$. We have three types of faces of σ , the faces that are the closures in \mathbb{T}^n of the original faces of σ , the closures of non-empty sets of the form $\overline{\sigma} \cap \mathbb{R}_J$ for some $J \supset I$ and lastly the sets of the form $\overline{\tau} \cap \mathbb{R}_J$ where $\tau \subset \sigma$ is a face and $J \supset I$. Faces of the first kind are called finite faces, whereas those of the second and third kinds are called infinite faces (these are all polyhedra by proposition 3.1.5).

A polyhedral complex in \mathbb{T}^n is defined analogously to the case of \mathbb{R}^n just replaced with polyhedra in \mathbb{T}^n and their faces. If X is an open polyhedral set then a polyhedral structure for it is a structure for the ambient polyhedral set.

We take inspiration from proposition 2.1.11 and define the relative interior of a polyhedron in \mathbb{T}^n as the complement in said polyhedron of all its proper faces.

Remark 3.1.10. We note that [10] excluded faces of the third type; however, we found it much more natural, especially when proving the following lemmas, to include these faces.

Example 3.1.11. Observe figure 6, it depicts the polyhedron given in figure 5 with all its faces labelled red.

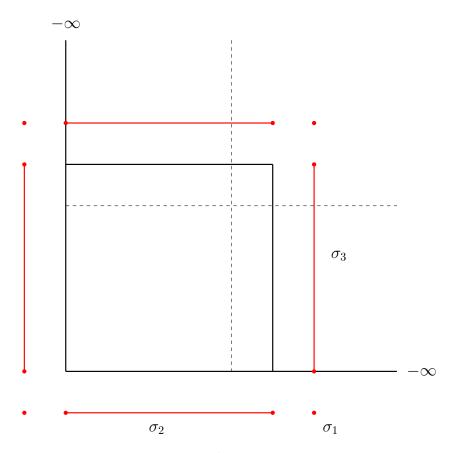


Figure 6: A polyhedron in \mathbb{T}^2 with all its faces labelled in red.

The face labelled σ_3 is an example of a face of the first type, the face labelled σ_2 is a face of the second type, and lastly σ_1 is a face of the third type.

Lemma 3.1.12. Let $\overline{\sigma}$ be a polyhedron in \mathbb{T}^n . Then a face of a face of $\overline{\sigma}$ is a face of $\overline{\sigma}$.

Proof. Let $\overline{\tau'} \subset \overline{\tau} \subset \overline{\sigma}$ be such that $\overline{\tau}$ is a face of $\overline{\sigma}$ and $\overline{\tau'}$ is a face of $\overline{\tau}$.

Let τ be a face of $\overline{\sigma}$ so that $\overline{\tau}$ is a finite face of $\overline{\sigma}$. Then if $\overline{\tau'}$ is a finite face of $\overline{\tau}$ then $\overline{\tau'}$ is a face of $\overline{\sigma}$ because a face of a face of a polyhedron is a face of the polyhedron when we are in \mathbb{R}^n . If $\overline{\tau'}$ is an infinite face, say $\overline{\tau'} = \overline{\overline{\tau''}} \cap \mathbb{R}_I$ for some finite face $\overline{\tau''}$ of $\overline{\tau}$. Then we observe that $\overline{\tau''}$ is a finite face of $\overline{\sigma}$ and by definition an infinite face of $\overline{\tau''}$ is also a face of $\overline{\sigma}$.

Let $\overline{\tau} = \overline{\overline{\tau''} \cap \mathbb{R}_I}$ be an infinite face of $\overline{\sigma}$ with $\overline{\tau''}$ a finite face of $\overline{\sigma}$. Then an infinite face of $\overline{\tau}$ is of the form $\overline{\overline{\tau'''} \cap \mathbb{R}_J}$ where $\overline{\tau'''}$ is a finite face of $\overline{\tau''}$ hence a finite face of $\overline{\sigma}$ which implies that $\overline{\overline{\tau'''} \cap \mathbb{R}_J}$ is a face of $\overline{\sigma}$. Then a finite face of $\overline{\tau}$ is just the same but with J = I so again a face of $\overline{\sigma}$.

Remark 3.1.13. Much like in the case of \mathbb{R}^n any polyhedron in \mathbb{T}^n has finitely many faces because any finite face of $\overline{\sigma}$ is induced from faces of σ of which there are finitely many by corollary 2.1.9. Then there are finitely many infinite faces because there are only finitely many choices for $I \subset [n]$.

Lemma 3.1.14. Let σ be a polyhedron in \mathbb{R}^n and $\pi : \mathbb{R}^n \to \mathbb{R}^m$ be a projection, then any face of $\pi(\sigma)$ is a projection of some face of σ .

Proof. Firstly, by induction, it is enough to consider the case where m = n - 1. Let $\tau = \pi(\sigma) \cap H_{-\kappa}$ be a face of $\pi(\sigma)$ then $\pi^{-1}(H_{\kappa})$ is a half space, since we may consider the function $\kappa'(x) = \langle m, x \rangle + c$ where m is equal to the linear part of κ in the first n - 1 coordinates and the last coordinate we set to zero while c takes the same value. Then any point $p \in \sigma$ lies in this new half space since $\pi(p)$ lies in $\pi(\sigma)$ and the last coordinate of p does not matter. It follows that $\tau' = \sigma \cap H_{-\kappa'}$ projects to τ .

Remark 3.1.15. Note that projections of faces are not necessarily faces of the projection. Consider a square-based pyramid in \mathbb{R}^3 whose apex is a 0-dimensional face, then the projection down to the base is not a face of the projection since the projection of the apex will lie in the interior of the base.

Proposition 3.1.16. Let $\mathcal{X} = \{\sigma_1, \dots, \sigma_t\}$ be a polyhedral complex in \mathbb{R}_I for some $I \subset [n]$. Then the collection

$$\overline{\mathcal{X}} := \{ \overline{\sigma_i} : 1 \le i \le t \} \cup \{ \text{infinite faces of } \overline{\sigma_i} : 1 \le i \le t \}$$

is a polyhedral complex in \mathbb{T}^n .

Proof. Let $\sigma_i, \sigma_j \in \mathcal{X}$ then by proposition 3.1.6 we have $\overline{\sigma_i} \cap \overline{\sigma_j} = \overline{\sigma_i \cap \sigma_j}$ is a face of both. Then we note that if $J \subset K \subset [n]$ then $\overline{\overline{\sigma_i} \cap \mathbb{R}_J} \cap \mathbb{R}_K = \overline{\sigma_i} \cap \mathbb{R}_K$ since the left-hand side is the projection to \mathbb{R}_J followed by the one to \mathbb{R}_K while the right-hand side is the projection directly to \mathbb{R}_K . Now let $J_1, J_2 \subset [n]$ be arbitrary, then

$$\overline{\overline{\sigma_i} \cap \mathbb{R}_{J_1}} \cap \overline{\sigma_j} \cap \mathbb{R}_{J_2} = \overline{\overline{\overline{\sigma_i} \cap \mathbb{R}_{J_1}} \cap \overline{\overline{\sigma_j} \cap \mathbb{R}_{J_2}} \cap \mathbb{R}_{J_1 \cup J_2}} = \overline{\overline{\sigma_i} \cap \sigma_j} \cap \mathbb{R}_{J_1 \cup J_2}$$

The right-hand side is a face of both the inputs on the left-hand side, since by lemma 3.1.14 the faces of $\overline{\sigma_i} \cap \mathbb{R}_{J_1}$ are of the form $\overline{\tau} \cap \mathbb{R}_K$ where τ is a face of σ_i and $K \supset J_1$ which is exactly the form the right hand side takes. The same argument follows for $\overline{\sigma_j} \cap \mathbb{R}_{J_2}$. This shows that the intersection of any two elements of $\overline{\mathcal{X}}$ is a face of both.

Now for the second condition, let $\overline{\sigma_i} \in \overline{\mathcal{X}}$ then any face of it already lies inside $\overline{\mathcal{X}}$ by construction. Similarly, any face of $\overline{\sigma_i} \cap \mathbb{R}_J$ is of the form $\overline{\tau} \cap \mathbb{R}_K$ for some $K \supset J$ and τ a face of σ_i which also lies inside $\overline{\mathcal{X}}$ by construction. \square

Proposition 3.1.17. Let \mathcal{X} and \mathcal{Y} be two complexes in \mathbb{T}^n , then the intersection

$$\mathcal{X} \cap \mathcal{Y} := \{ \overline{\sigma} \cap \overline{\tau} : \overline{\sigma} \in \mathcal{X}, \overline{\tau} \in \mathcal{Y}, \overline{\sigma} \cap \overline{\tau} \neq \emptyset \}$$

is a polyhedral complex in \mathbb{T}^n .

Proof. Let $\overline{\sigma} \in \mathcal{X}$ and $\overline{\tau} \in \mathcal{Y}$ and suppose that $\sigma \subset \mathbb{R}_I$ and $\tau \subset \mathbb{R}_J$. Then a finite face of $\overline{\sigma} \cap \overline{\tau}$ is the closure of a face of $\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}$ by proposition 3.1.6. Then by the proof of lemma 2.1.7 faces of $\overline{\sigma} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}$ are of the form $F \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}$ where F is a face of $\overline{\sigma} \cap \mathbb{R}_{I \cup J}$ or $\overline{\sigma} \cap F' \cap \mathbb{R}_{I \cup J}$ where F' is a face of $\overline{\tau} \cap \mathbb{R}_{I \cup J}$. Assume we are in the first case, then the argument is identical for the second. Then we may write $F = \overline{\widetilde{F}} \cap \mathbb{R}_{I \cup J}$ where \widetilde{F} is a face of σ by lemma 3.1.14. Then it follows that

$$\overline{F \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}} = \overline{\overline{\widetilde{F}} \cap \overline{\tau} \cap \mathbb{R}_{I \cup J}} = \overline{\widetilde{F}} \cap \overline{\tau}$$

which is an element of $\mathcal{X} \cap \mathcal{Y}$. Next, for the other condition it is enough to show that if F is a face of $\overline{\sigma} \in \mathcal{X}$ and F' is a face of $\overline{\tau} \in \mathcal{Y}$ then $F \cap F'$ is a face of $\overline{\sigma} \cap \overline{\tau}$. Suppose both are finite faces, then write $F = \overline{\sigma_F}$ where σ_F is a face of σ and $F' = \overline{\tau_{F'}}$ where $\tau_{F'}$ is a face of τ then

$$F \cap F' = \overline{\overline{\sigma_F} \cap \overline{\tau_{F'}} \cap \mathbb{R}_{I \cup J}}.$$

We then observe that $\overline{\sigma_F} \cap \mathbb{R}_{I \cup J}$ and $\overline{\tau_{F'}} \cap \mathbb{R}_{I \cup J}$ is a face of σ and τ respectively which implies by the proof of lemma 2.1.7 that $F \cap F'$ is a finite face of $\overline{\sigma} \cap \overline{\tau}$ which we see by using the description of $\overline{\sigma} \cap \overline{\tau}$ given in proposition 3.1.6. Now suppose that $F' = \overline{\overline{\tau_{F'}}} \cap \mathbb{R}_K$ is an infinite face of $\overline{\tau}$, then

$$F \cap F' = \overline{\sigma_F} \cap \overline{\tau_{F'}} \cap \mathbb{R}_K = \overline{\sigma_F} \cap \overline{\tau_{F'}} \cap \mathbb{R}_{I \cup K}$$

and the right-hand side is an infinite face of $\overline{\sigma} \cap \overline{\tau}$ by again using a similar description. The case of two infinite faces is identical as the previous case except I is replaced by some $K' \supset I$ in the final equality which is still an infinite face.

Corollary 3.1.18. Every polyhedral set has a polyhedral structure.

Proof. We can follow the same proof as proposition 3.1.7 using proposition 3.1.17 and the complexes $\overline{\mathcal{H}_{\kappa_1}}, \ldots, \overline{\mathcal{H}_{\kappa_t}}$ where $\kappa_1, \ldots, \kappa_t$ are a complete set of functions that represent every polyhedron in a given polyhedral set. This works because if $\overline{\sigma}$ and $\overline{\tau}$ are polyhedra in \mathbb{T}^n of the same sedentarity, then by corollary 3.1.6

$$\overline{\sigma} \cap \overline{\tau} = \overline{\sigma \cap \tau}$$
.

In particular if $\overline{\sigma}$ is a polyhedron that makes up our polyhedral set then the intersection $\overline{\mathcal{H}_{\kappa_1}} \cap \cdots \cap \overline{\mathcal{H}_{\kappa_t}}$ contains elements of the form $\overline{\sigma}$ intersected with extra items for which the union of all such elements give $\overline{\sigma}$ which implies the desired result.

Let X be a pure n dimensional open polyhedral set and P be the ambient polyhedral set. Then any polyhedron in P that has a non-empty intersection with X necessarily is n dimensional by proposition 3.1.3 so we may force P to be pure n dimensional by removing any polyhedra that do not intersect with X. This forms a new polyhedral set P' of which X is still an open subset, and any polyhedral structure of P' is necessarily purely n dimensional. When talking about the structure of a pure dimensional open polyhedral set, then we will always assume this structure to be pure of the same dimension.

With this we wish to define the codimension of a point p in a pure dimensional open polyhedral set X, denoted by $\operatorname{codim}_X(p)$, as the minimum of the codimension of the cells whose relative interior contains p for the structure constructed in corollary 3.1.18. This comes with the question if every point p lies in the relative interior of some cell of the structure. We shall address this in proposition 3.1.19.

Proposition 3.1.19. Let X be a pure n dimensional open polyhedral set and \mathcal{X} a structure of X. Then for any point $p \in X$ there exists a cell $\overline{\sigma} \in \mathcal{X}$ that contains p in its relative interior.

Proof. Let $\overline{\sigma_0} \in \mathcal{X}$ be a polyhedron that contains p. If p is not already in the relative interior of $\overline{\sigma_0}$ then p lies inside a proper face of $\overline{\sigma_0}$, say $\overline{\sigma_1}$. We again check if p lies in the relative interior, if not then it lies in a proper face of $\overline{\sigma_2}$. We get a chain of faces of $\overline{\sigma_0}$ which eventually must terminate since any polyhedron in \mathbb{T}^n has finitely many faces. Therefore, there must eventually exist some $\overline{\sigma_i}$ that contains p in its relative interior.

Remark 3.1.20. According to [10] the definition of the codimension is independent of the choice of polyhedral structure. It is unknown whether or not this is true, as there was no proof given nor could we come up with one or a counterexample, therefore, we comprise by choosing a specific polyhedral structure.

Note that $p \in X^{\text{gen}}$ if and only if $\operatorname{codim}_X(p) = 0$. Indeed, if $p \in X^{\text{gen}}$ then any polyhedron that contains it in its relative interior must be of dimension $\dim(X)$. On the other hand, if $\operatorname{codim}_X(p) = 0$ then there is a polyhedron whose relative interior contains p and is of dimension $\dim(X)$.

Lemma 3.1.21. Let $\sigma \subset \mathbb{R}_I$ and $\tau \subset \mathbb{R}_J$ be two polyhedra with $\overline{\sigma} = \overline{\tau}$. Then I = J and $\sigma = \tau$.

Proof. First, suppose that there is some $m \in I$ such that $m \notin J$, and let $x \in \overline{\sigma}$. Then, by assumption $x \in \overline{\tau}$ and since $m \notin J$, there must be some sequence $(x_k)_{k \geq 0}$ with $x_k \in \tau$ whose mth coordinates are all finite. In particular, we have $x_k \in \tau \subset \overline{\sigma}$ but the mth coordinates of any $x \in \overline{\sigma}$ must be $-\infty$ since $\overline{\sigma}$ is of sedentarity I and $m \in I$. This shows that we must have $I \subset J$ and we can prove the reverse inclusion in exactly the same way, so this implies that I = J. Then we achieve $\sigma = \tau$ be noting that $\sigma = \overline{\sigma} \cap \mathbb{R}_I$ and similarly for τ , then intersecting both sides of $\overline{\sigma} = \overline{\tau}$ with \mathbb{R}_I .

Proposition 3.1.22. Let X be a polyhedral set in \mathbb{T}^n , \mathcal{X} be a polyhedral structure of X and $I \subset [n]$ be such that $X \cap \mathbb{R}_I$ is non-empty. Then the collection

$$\mathcal{X} \cap \mathbb{R}_I := \{ \sigma : \overline{\sigma} \in \mathcal{X} \text{ and } \sigma \subset \mathbb{R}_I \}$$

is a polyhedral structure for $X \cap \mathbb{R}_I$.

Proof. Firstly, let $\sigma, \tau \in \mathcal{X} \cap \mathbb{R}_I$ then $\overline{\sigma} \cap \overline{\tau} = \overline{\sigma \cap \tau}$ so $\sigma \cap \tau \in \mathcal{X} \cap \mathbb{R}_I$ and we wish to show that $\sigma \cap \tau$ is a face of σ and τ . To that end, suppose that $\overline{\sigma} \cap \overline{\tau} = \overline{F} \cap \mathbb{R}_J$ for some face F of σ (an infinite face of σ) for some $J \supset I$. Then by lemma 3.1.21 we must have I = J which implies that $\overline{\sigma} \cap \overline{\tau} = \overline{F}$ hence \overline{F} is a face of $\overline{\tau}$. Now \overline{F} cannot be an infinite face of sedentarity $J \supsetneq I$ of $\overline{\tau}$ since \overline{F} is of sedentarity I. Therefore, \overline{F} is a face of sedentarity I of $\overline{\tau}$ which implies it is a finite face hence F is a face of both σ and τ . Then by lemma 3.1.21 we have $\sigma \cap \tau = F$ so in this case $\sigma \cap \tau$ is a face of both. If F is instead a face of $\overline{\tau}$ then we can make a near identical argument. Next if $\overline{\sigma} \cap \overline{\tau}$ is a finite face of $\overline{\sigma}$ and $\overline{\tau}$ say $\overline{\sigma} \cap \overline{\tau} = \overline{F}$ then we achieve $\sigma \cap \tau = F$ which is a face of both.

The last two arguments imply that $\mathcal{X} \cap \mathbb{R}_I$ is a polyhedral structure so it remains to show that $|\mathcal{X} \cap \mathbb{R}_I| = X \cap \mathbb{R}_I$. It is clear by definition that $|\mathcal{X} \cap \mathbb{R}_I| \subset X \cap \mathbb{R}_I$ since \mathcal{X} is a structure for X and $\mathcal{X} \cap \mathbb{R}_I$ only considers those polyhedra of sedentarity I so it is contained in \mathbb{R}_I . Conversely, let $x \in X \cap \mathbb{R}_I$ then there exists a polyhedron $\overline{\sigma} \in \mathcal{X}$ that contains x since \mathcal{X} is a structure for X. Now since \mathcal{X} is a complex we must have $\overline{\overline{\sigma} \cap \mathbb{R}_I} \in \mathcal{X}$ hence $\overline{\sigma} \cap \mathbb{R}_I \in \mathcal{X} \cap \mathbb{R}_I$ and it also contains x which implies that $x \in |\mathcal{X} \cap \mathbb{R}_I|$. We conclude that $\mathcal{X} \cap \mathbb{R}_I$ is a polyhedral structure of $X \cap \mathbb{R}_I$ as desired. \square

Definition 3.1.23 (Star). Let $x \in X$ be a point of sedentarity I then by proposition 3.1.5 we have that $X \cap \mathbb{R}_I$ is an open polyhedral set in \mathbb{R}_I that contains x. We can then form the star of $X \cap \mathbb{R}_I$ at x, that is a union of finitely many cones, which is often called a fan, in \mathbb{R}_I from which we define

$$\operatorname{Star}_X(x) := \operatorname{Star}_{X \cap \mathbb{R}_I}(x).$$

The m-dimensional part of $\operatorname{Star}_X(x)$ is the closure of $\dim_{\operatorname{Star}_X(x)}^{-1}(m)$ in \mathbb{R}_I (which may be empty) and is denoted by $\operatorname{Star}_X(x)^{(m)}$. If non-empty, this is an m-dimensional polyhedral set in \mathbb{R}_I and it inherits a weight function from the weights around the point x since a weight function must be locally constant.

Definition 3.1.24. Any polyhedral structure \mathcal{L} of $\operatorname{Star}_X(x)^{(m)}$ inherits its weight function by assigning $\omega(\sigma) = \omega(p)$ for any $p \in \operatorname{relint}(\sigma)$ and $p \in X$. Any two weighted polyhedral structures obtained this way are equivalent, so one being balanced is equivalent to the other being balanced. We say that $\operatorname{Star}_X(x)^{(m)}$ is balanced if any choice of polyhedral structure of it is balanced. If X is pure m-dimensional and $\operatorname{Star}_X(x)^{(m)}$ is either empty or balanced for every $x \in X$, then we say that X is a balanced open polyhedral set.

Definition 3.1.25 (Morphisms of open polyhedral sets in \mathbb{T}^n). A map $F: \mathbb{R}^n \to \mathbb{R}^m$ is said to be affine \mathbb{Z} -linear if it is of the form $x \mapsto Ax + b$ where A is a matrix with integer entries and $b \in \mathbb{R}^m$. Then a map $f: \mathbb{R}^n \to \mathbb{T}^m$ is said to be affine \mathbb{Z} -linear if it is the composition of an affine \mathbb{Z} -linear map $F: \mathbb{R}^n \to \mathbb{R}_I$ followed by the inclusion $\mathbb{R}_I \hookrightarrow \mathbb{T}^m$ for some $I \subset [m]$.

If $X \subset \mathbb{T}^n$ and $Y \subset \mathbb{T}^m$ are two open polyhedral sets, then a continuous map $f: X \to Y$ is a tropical morphism if it is locally affine \mathbb{Z} -linear in the following sense: For each point $x \in X$ there is an open neighbourhood U containing x such that for all $I \subset [n]$ with $U \cap \mathbb{R}_I$ non-empty, the map $f|_{U \cap \mathbb{R}_I}$ is the restriction of an affine \mathbb{Z} -linear map from $\mathbb{R}_I \to \mathbb{T}^m$.

If X and Y are weighted open polyhedral sets, then a tropical isomorphism $f: X \to Y$ is a tropical morphism with an inverse tropical morphism $g: Y \to X$ and the weight of x and f(x) match for all $x \in X^{\text{gen}}$. Note that the fact that f is a homeomorphism implies that x is generic if and only if f(x) is generic.

Example 3.1.26. We have that $\mathbb{R} = \mathbb{T} \setminus \{-\infty\}$ and \mathbb{T} are open polyhedral sets in \mathbb{T} and a morphism $f : \mathbb{R} \to \mathbb{T}$ consists of a choice of base point in \mathbb{T} and the slope of f. There can be no slope changes because if there were, then locally around where this change happens f cannot be affine \mathbb{Z} -linear. Similarly, if we look at a tropical morphism $f : \mathbb{R}^n \to \mathbb{T}$ then this is equivalent to choosing a base point in \mathbb{T} and choosing the slope of f along the standard basis vectors of \mathbb{R}^n .

Definition 3.1.27 (Tropical structures). Let X be a topological space. A tropical structure on X is a collection of maps $\varphi_i : U_i \to V_i$ with the following properties:

- The collection $\{U_i\}_i$ forms an open covering of X.
- Each V_i is an effective and balanced open polyhedral set in \mathbb{T}^n for some n.
- Each φ_i is a homeomorphism.
- For every pair i, j with $U_i \cap U_j \neq \emptyset$ the map

$$\varphi_i \circ \varphi_i^{-1} : \varphi_j(U_i \cap U_j) \to \varphi_i(U_i \cap U_j)$$

is a tropical isomorphism of balanced open polyhedral sets in \mathbb{T}^n .

Two tropical structures on a topological space X are said to be equivalent if their union is also a tropical structure for X. A topological space X equipped with an equivalence class of tropical structures is said to be a tropical space, and if φ is a map in the structure that contains a point $x \in X$ then φ is said to be a chart at x. If X is non-empty and all V_i happen to be pure of dimension n, then we say that X is of dimension n.

Example 3.1.28. We give a few examples of tropical spaces.

- Every effectively weighted and balanced open polyhedral set X in \mathbb{T}^n can be equipped with the global chart id : $X \to X$, therefore, every effective weighted open polyhedral set in \mathbb{T}^n is a tropical space. Furthermore, if X is pure of dimension n, then so is its tropical space.
- If X is a tropical space and $U \subset X$ is open, then U is a tropical space by equipping it with the charts of X restricted to U.

Definition 3.1.29 (Morphisms of tropical spaces). Let X and Y be tropical spaces. Then a continuous map $f: X \to Y$ is said to be a tropical morphism if for any pair of charts $\varphi: U \to V$ for X and $\psi: U' \to V'$ for Y the composition

$$\psi \circ f \circ \varphi^{-1} : \varphi(f^{-1}(U') \cap U) \to V'.$$

is a tropical morphism of open polyhedral sets. It is an isomorphism if it is a homeomorphism and these maps are tropical isomorphisms onto their image.

3.2 Tropical polynomials

The vanishing set of a tropical polynomial is an important example of a tropical space, similar to algebraic geometry. We place two binary operations on the set \mathbb{T} :

$$x \oplus y := \max\{x, y\}$$
 and $x \odot y := x + y$,

where \oplus plays the role of addition and \odot the role of multiplication. These turn $\mathbb T$ into a semiring, which we call the tropical numbers. Note that the tropical numbers has no additive inverses except the additive neutral element $-\infty$ because

$$x \oplus x = \max\{x, x\} = x,$$

so the existence of an additive inverse implies $x = -\infty$. Now we shall consider tropical polynomials, which are functions $f: \mathbb{R}^n \to \mathbb{R}$ of the form

$$f(x) = \bigoplus_{u \in I} c_u \odot x^u = \max\{c_u + \langle u, x \rangle\}_{u \in I}$$
 (1)

where I is some finite subset of \mathbb{Z}^n and $c_u \in \mathbb{R}$ and x^u means $x_1^{u_1} \odot \cdots \odot x_n^{u_n}$. Unlike algebraic geometry, the vanishing set of a tropical polynomial f is defined as the set of all $x \in \mathbb{R}^n$ such that f(x) is attained by at least two monomials in equation 1. If $f(x) = a_0 x^n + \cdots + a_n$ is a polynomial over an algebraically closed valued field K with valuation ν , not tropical, then we can tropicalize it by considering the tropical polynomial

$$trop(f)(x) := \max \{ \nu(a_i) + (n-i) \cdot x \}_{0 \le i \le n}.$$

Then it turns out that V(f) and $V(\operatorname{trop}(f))$ are closely related, which one can see more details of in [9]. We can extend this to multivariate polynomials in the exact same way. For example, the multivariate polynomial f(x,y) = x + y + c over \mathbb{C} with the trivial valuation tropicalizes to the tropical line given in the second example of example 3.2.1.

Example 3.2.1.

- Consider $f(x) = x^2 \oplus 2 \odot x \oplus 0 = \max\{2x, x+2, 0\}$. The maximum occurs twice at x = -2, and x = 2.
- Consider $f(x,y) = x \oplus y \oplus 0 = \max\{x,y,0\}$. The maximum occurs at least twice in the sets, $\{x = 0, y \leq 0\}$, $\{y = 0, x \leq 0\}$, $\{x = y, x, y \geq 0\}$. We call this the tropical line, and it is illustrated in figure 10.

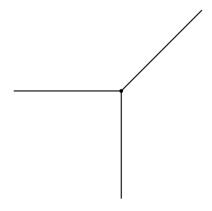


Figure 7: Illustration of the tropical line.

Notice that each example can be given the structure of a polyhedral complex, this is not a coincidence, and in fact the zero set of any tropical polynomial can be assigned a polyhedral complex. As a matter of fact, according to [9, Proposition 3.3.2], the following theorem is true:

Theorem 3.2.2. Let f be a tropical polynomial in n variables. Then V(f) is the support of a effectively weighted balanced polyhedral complex pure of dimension n-1.

We note that V(f) is not initially a complex in \mathbb{T}^n but only in \mathbb{R}^n , however, by taking the closure and inheriting the weights it naturally becomes one in \mathbb{T}^n . It remains balanced because if we take the star around any point of higher sedentarity (in other words, points that get added by taking the closure) then the result is empty since it must be lower dimensional than V(f). In particular, if f is a tropical polynomial then $\overline{V(f)}$ (where the closure is taken in \mathbb{T}^n) is a balanced open polyhedral set, hence is automatically a tropical space.

3.3 Regularity

When we later talk about tropical curves, regularity will act as our notion of a smooth curve. We remark that [10] has a separate notion of smoothness for arbitrary tropical spaces, but this requires a large detour, fortunately, regularity implies smoothness for curves (according to [10]) so in this way we bypass said detour. We cite the following proposition and its proof, which one can find at [10, Proposition 7.2.3].

Proposition 3.3.1. Let X be an open polyhedral set in \mathbb{T}^N of pure dimension n. Then the following conditions are equivalent:

• For every $I \subset [n]$ such that $X \cap \mathbb{R}_I \neq \emptyset$ we have

$$\dim(X \cap \mathbb{R}_I) = n - |I|.$$

• For each stratum \mathbb{T}_I such that $X \cap \mathbb{T}_I \neq \emptyset$ we have

$$\dim(X \cap \mathbb{T}_I) \leq n - |I|.$$

• For each point $p \in X$ we have

$$|\operatorname{sed}(p)| \le \operatorname{codim}_X(p).$$

Proof. See [10, Proposition 7.2.3].

Definition 3.3.2 (Regularity at infinity). Let X be an open polyhedral set pure of dimension n. Then X is called regular at infinity if the conditions of proposition 3.3.1 hold. A tropical space X is called regular at infinity if there exists a tropical structure $\{\varphi_i: U_i \to V_i\}_i$ where each V_i is regular at infinity.

Example 3.3.3. The following examples are taken from [10, Example 7.2.5].

- Any bounded pure m dimensional polyhedron in $\mathbb{R}_{\emptyset} \cong \mathbb{R}^n$ induces a polyhedron that is regular at infinity in \mathbb{T}^n . This is because the sedentarity of any point is \emptyset which is certainly has fewer elements than the codimension.
- Figure 8 depicts a polyhedron that is regular at infinity because the sedentarity of the single point in $\mathbb{R}_{[2]}$ is [2] while the codimension is two. The sedentarity of the rest of the points is \emptyset so condition three of proposition 3.3.1 is automatically satisfied for these points.

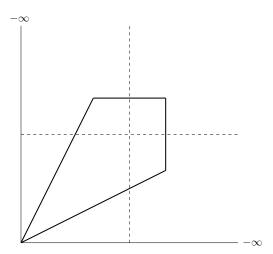


Figure 8: Polyhedron in \mathbb{T}^2 that is regular at infinity.

- Figure 9 depicts a polyhedron in \mathbb{T}^2 that is irregular at infinity because the sedentarity of the point in $\mathbb{R}_{[2]}$ is [2] while the codimension is only 1. However, if we instead consider that it is a tropical space, we find that it is isomorphic to \mathbb{T}^1 , which means that it is regular as a tropical space but not as an open polyhedral set.
- If σ is a bounded m dimensional polyhedron in \mathbb{R}_I for some $I \neq \emptyset$ then it will not be regular because this would imply that

$$m = \dim(\sigma \cap \mathbb{R}_I) = m - |I|.$$

However if we consider it as a tropical subspace of \mathbb{T}_I then it will be regular since we are now in the \mathbb{R}_{\emptyset} case from the first example.

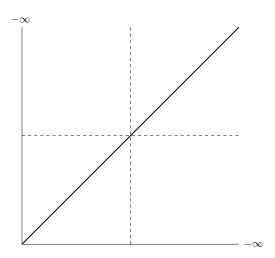


Figure 9: Polyhedron in \mathbb{T}^2 that is not regular at infinity.

3.4 Tropical Curves

An important class of tropical spaces are tropical curves, and in most literature one will find these to be defined as a metric graph. It turns out that this also happens in our case, however one must impose the notion of regularity.

Definition 3.4.1. A tropical curve is a connected tropical space of dimension one.

We shall adopt the following definition of a CW-complex from [7].

Definition 3.4.2 (CW-Complex). We define a CW-complex inductively.

- 1. Let X_0 be a discrete set of points that we call 0-cells.
- 2. We form the *n*-skeleton X_n from X_{n-1} by attaching *n*-cells $\{e_{\alpha}^n\}_{\alpha}$ (each of which is homeomorphic to the interior of the *n*-disk) through a continuous map $\varphi_{\alpha}: \partial e_{\alpha}^n \to X_{n-1}$ for each α . In particular, $X_n := X_{n-1} \bigsqcup_{\alpha} e_{\alpha}^n / \sim$ where \sim is the equivalence relation $x \sim \varphi_{\alpha}(x)$ for every $x \in \partial e_{\alpha}^n$.
- 3. This inductive process can either terminate at some n, in which case we set $X = X_n$ for some non-negative integer n, or this can go on indefinitely for which we set $X = \bigcup_n X_n$. In the second case we give it the topology that $U \subset X$ is open (resp. closed) if and only if $U \cap X_n$ is open (resp. closed) for every n.

If we terminate the process in definition 3.4.2 early, say $X = X_n$, then we say that the CW-complex X has dimension n. We shall only concern ourselves with one-dimensional CW-complexes as they will be important when we go on to define a metric graph.

Remark 3.4.3. Note that if X is non-empty, then X_0 must be non-empty. Indeed, let n > 0 be the minimal integer such that $X_n \neq \emptyset$ and $X_{n-1} = \emptyset$. Then there must exist some gluing map $\varphi_{\alpha} : \partial e_{\alpha}^n \to X^{n-1}$, however such a map cannot exist if X^{n-1} is empty.

On the other hand, the empty set is a CW-complex with no cells.

Example 3.4.4. A 1-dimensional CW-complex can be thought of as a graph. The 0-cells are the vertices and the 1-cells are the edges. Indeed, by definition 1 cells are open intervals, and the gluing action is precisely choosing to which vertices the end points of this open interval should be attached to.

We say that a 1-dimensional CW-complex is finite if there are finitely many vertices and edges. The valency of a vertex is the size of its equivalence class minus one (since the vertex itself is in its own equivalence class), and if a point is in the interior of an edge then we say it is of valency two.

Definition 3.4.5 (Metric graph). A topological space G is said to be a metric graph if it is a connected one-dimensional finite CW-complex that is equipped with a complete metric on $G \setminus \{1\text{-valent vertices}\}$.

A metric graph with open ends is a metric graph with some 1-valent vertices removed. We exclude the case of a single edge with both endpoints removed; however, we still allow a single edge that is subdivided by a vertex. Henceforth, we will always assume that when referring to metric graphs, we are always talking about one that may have open ends.

Two metric graphs are said to be isomorphic if there exists a homeomorphism between the two, which becomes an isometry after removing all 1-valent vertices. Note that edges can be identified with $(-\infty, 0]$, $[-\infty, 0]$ and [0, l] depending on if it is an open end, closed end or an inner edge in which case l is its length.

According to [10, Proposition 8.15] every tropical curve that is regular at infinity can be realized as a metric graph. We will leave this statement as a black box, however the details of how one can go about proving this can be found in [10, Exercise 8.1.6].

Theorem 3.4.6. There is a bijection between isomorphism classes of regular tropical curves and isomorphism classes of metric graphs.

The discussion in [10, p.211] gives a very concrete discussion on how one can go from a metric graph Γ to its tropical curve under the bijection. Then for the other direction we shall give a concrete example.

Example 3.4.7. We go through how to turn the tropical line in figure 10 into a metric graph.

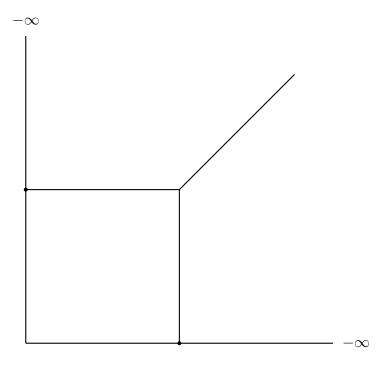


Figure 10: The tropical line in \mathbb{T}^2 .

According to [10, p.211], any regular tropical curve is necessarily smooth, which means that all weights are one so we in fact do not need to consider weights at all. Now, denoting the standard basis vectors in \mathbb{R}^2 by e_i , we define the length of $-e_1$, $-e_2$ and e_1+e_2 to all be one. This gives us a unique complete metric on the tropical line (excluding the vertices at $-\infty$), hence turns it into a metric graph.

This process is very much the same for all tropical curves as they all locally look like the fan generated by $-e_1, \ldots, -e_n, e_1 + \cdots + e_n$ for some n which we denote by L(n+1), or \mathbb{T} in the case of 1-valent vertices which, for convenience, we denote by L(1). We again define the metric to give a length of 1 for each e_i which gives the tropical curve a unique complete metric.

4 Metric Graphs

All definitions, unless otherwise stated, will henceforth be taken from [2]. However, all proofs are done by myself.

Let Γ be a metric graph, then we set $V(\Gamma)$ and $E(\Gamma)$ to be its vertex and edge sets respectively. Note that we can always subdivide any edge by adding a vertex to its interior, or remove any 2-valent vertex by completing the edge it lies on.

Definition 4.0.1 (Morphisms of metric graphs). Let Γ and Γ' be two metric graphs and $\varphi : \Gamma' \to \Gamma$ be a continuous map. Let $V(\Gamma)$ and $V(\Gamma')$ be vertex sets for each metric graph. Then φ is said to be a $(V(\Gamma'), V(\Gamma))$ -morphism if:

- We have $\varphi(V(\Gamma')) \subset V(\Gamma)$.
- We have $\varphi^{-1}(E(\Gamma)) \subset E(\Gamma')$.
- The restriction of φ to any edge e' of Γ' is a dilation by some factor $d_{e'}(\varphi) \in \mathbb{Z}_{>0}$. This dilation is called the degree of φ along e'.

The map φ is simply called a morphism if there exists a vertex set of Γ' and Γ such that the above holds. It is called finite if all dilations are non-zero.

Remark 4.0.2. Note that if φ were a morphism of metric graphs which had an inverse morphism φ^{-1} then the dilation factor must have been one which makes φ an isometry, hence matches the definition in the previous section.

This then raises the question of whether or not the category of metric graphs, which we denote by \mathcal{M} , and the category of regular tropical curves, which we denote by \mathcal{T}^{reg} , are equivalent. In the case that we make a functor $F: \mathcal{M} \to \mathcal{T}^{\text{reg}}$ that is the identity on both objects and morphisms, this will not result in an equivalence of categories. Take for example, the metric graphs L(3) and L(4) except we cut off the rays so that they are a length of one. The morphism (of metric graphs) mapping the points $(-1,0,0) \mapsto (-1,0)$, $(0,-1,0) \mapsto (-1,0)$, $(0,0,-1) \mapsto (0,-1)$ and $(1,1,1) \mapsto (0,-1)$ is perfectly valid, since all lengths are one. However, if we want to transfer this to a morphism of tropical curves then we will want an affine transformation that reflects exactly this map. In other words, if we denote the map by $x \mapsto Ax + b$ where

$$A = \begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix} \text{ and } b = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

then we would require the following equalities to be satisfied:

$$-a + b_1 = -1, -d + b_2 = 0,$$

 $-b + b_1 = -1, -e + b_2 = 0,$

$$-c + b_1 = 0, -f + b_2 = -1,$$

 $a + b + c + b_1 = 0, d + e + f + b_2 = -1.$

This then forces $b_1 = -1/2$, which implies that a, b, c are non-integers which is already a problem. Unfortunately, this means that our choice of F is not an equivalence, whether or not there exists some choice of functor that works is unknown, however more than likely an equivalence would require the change up in the definitions of morphisms of both metric graphs and tropical curves. In particular, if we allow real dilations and real entries then we do get an equivalence. Every metric graph is given charts to L(n) for some n, so it is enough to look at maps from an open around (0,0) in L(n) to an open around (0,0) in L(m) for some $n,m\geq 1$. These maps must send (0,0) in L(n) to (0,0) in L(m), otherwise we just choose L(2) for our chart. Any morphism of metric graphs with real dilations is a morphism of tropical cures with real entries, because we get a near identical line of inequalities as in the previous example, from which it is easy to extract the required entries of A and b. This implies that F is at least well-defined, then faithfulness is clear from the fact that F is the identity on morphisms and fullness is satisfied because every possible continuous map takes the form of a morphism of metric graphs with real dilations. Therefore, in the case that we allow real values, F is an equivalence of categories. Whether or not having real dilations and real entries is the right thing to impose is hard to say, since almost all sources impose their integer variants. Nonetheless, there is still some value in exploring the integer variants since this is how morphisms on metric graphs are defined in most sources, and tropical curves are often defined to be metric graphs to begin with.

Definition 4.0.3. Let $r \in \mathbb{Z}_{>0}$ then consider the convex hull in \mathbb{R}^2 of (0,0) and a point different from (0,0), we shall call this a branch. The star with r branches, denoted by S_r , is a union of r branches, no two of which are parallel.

If Γ is a metric graph and $p \in \Gamma$ is a point, then there exists a neighbour-hood U_p around p that is homeomorphic to some S_r where r is the valency of p. The set of tangent directions at p is given by $T_p(\Gamma) := \varinjlim_{U_p} \pi_0(U_p \setminus \{p\})$, where π_0 takes all path components of its input and the limit is taken over all the neighbourhoods of p. The size of $T_p(\Gamma)$ is naturally the valency of p.

Let $\varphi : \Gamma' \to \Gamma$ be a morphism of metric graphs, $p' \in \Gamma'$, $v' \in T_{p'}(\Gamma')$ and $e' \in E(\Gamma')$ be the edge in the direction of v'. Then we define $d_{v'}(\varphi) := d_{e'}(\varphi)$ and this induces a map

$$d\varphi(p'): \{v' \in T_{p'}(\Gamma'): d_{v'}(\varphi) \neq 0\} \to T_p(\Gamma)$$

in the obvious way. The reason we exclude those that have dilation zero is because this implies that an edge gets mapped to a vertex.

Definition 4.0.4 (Harmonic morphisms). Let $\varphi : \Gamma' \to \Gamma$ be a morphism of metric graphs and $p' \in \Gamma'$. Then we say that the morphism is harmonic at p' if for any choice $v \in T_{\varphi(p')}(\Gamma)$ the quantity

$$d_{p'}(\varphi) := \sum_{v' \in d\varphi(p')^{-1}(v)} d_{v'}(\varphi)$$

evaluates to the same value. If φ is surjective and harmonic at all $p' \in \Gamma'$ then we say that φ is a harmonic morphism.

Proposition 4.0.5. Let $\varphi : \Gamma' \to \Gamma$ be a harmonic morphism then the quantity

$$\sum_{p' \in \varphi^{-1}(p)} d_{p'}(\varphi)$$

is independent of the choice of $p \in \Gamma$. We call this quantity the degree of φ , denoted by $\deg(\varphi)$.

Proof. It is enough to show this for points along an edge $e \in E(\Gamma)$ since we can inductively extend it to all points of Γ . Let $p \in e$ be a vertex and $q \in e$ be an interior point, then it is enough to show that the degree of φ is the same for both of these. We will temporarily introduce the notation $\deg_p(\varphi)$ for the degree of φ with respect to p. To that end, let p_1, \ldots, p_n be vertices in Γ' that map to p and e_{11}, \ldots, e_{1m_1} be all edges incident to p_1 that map to e, e_{21}, \ldots, e_{2m_2} be all the edges incident to p_2 that map to e, and so on. Then by choosing the direction $v \in T_p(\Gamma)$ (which is allowed since φ is harmonic) pointing towards e we may write

$$\deg_p(\varphi) = \sum_{i=1}^n \sum_{j=1}^{m_i} d_{e_{ij}}(\varphi)$$

where we take it with respect to the point p. This matches the degree of φ when considering the point q because e_{ij} are by definition all the edges that map to e, and no matter which direction we choose, we will get the dilations as above since it is an interior point.

Definition 4.0.6 (Rational functions). Let Γ be a metric graph and $U \subset \Gamma$ an open subset, then a continuous function $f: U \to \mathbb{R}$ is said to be a rational function if it is piecewise affine with integer slopes. Recall that every edge can be identified with $[-\infty, 0], (-\infty, 0]$ or [0, l], so we can equivalently say piecewise affine with integer slopes along every edge. We denote the set of rational functions of Γ by $\text{Rat}(\Gamma)$.

Note that the term piecewise affine implies that there are only finitely many slope changes. The definition of rational functions mimics the idea of tropicalizing rational functions on an algebraic variety. If f/g is a quotient of two polynomials on an algebraic variety over an algebraically closed and valued field, then we can tropicalize it to $\operatorname{trop}(f) - \operatorname{trop}(g)$ which is piecewise affine with integer slopes.

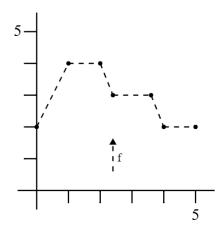


Figure 11: Example of a rational function f on the interval [0, 5].

Let Γ be a metric graph; then we define $\mathrm{Div}(\Gamma)$ to be the free abelian group of Γ with its elements being called divisors. If $D \in \mathrm{Div}(\Gamma)$ is a divisor, then $\deg(D)$ is the sum of all its coefficients.

The order of a rational function f of a metric graph Γ is the sum of its outgoing slopes from the point p, which we denote by $\operatorname{ord}_p(f)$. The divisor associated with f, denoted by $\operatorname{div}(f)$, is defined by

$$\operatorname{div}(f) := \sum_{p \in \Gamma} \operatorname{ord}_p(f)(p).$$

Note that we can equivalently write

$$\operatorname{div}(f) = \sum_{e \in E(\Gamma)} \operatorname{div}(f|_e).$$

If we take the rational function in figure 11 then the order of f at any point that lies on a linear part of f is zero. Then for the other points we have $\operatorname{ord}_0(f) = 2$, $\operatorname{ord}_1(f) = -2$, $\operatorname{ord}_2(f) = 2$, $\operatorname{ord}_{2.5}(f) = -2$, $\operatorname{ord}_{3.5}(f) = -2$, $\operatorname{ord}_4(f) = 2$ and lastly $\operatorname{ord}_5(f) = 0$. This gives us

$$\operatorname{div}(f) = 2(0) - 2(1) + 2(2) - 2(2.5) - 2(3.5) + 2(4).$$

Notice that in this case deg(div(f)) = 0. This is not a coincidence, and as it turns out, if Γ is a compact metric graph this will always be the case.

Proposition 4.0.7. Let Γ be a compact metric graph and $f:\Gamma\to\mathbb{R}$ a rational function. Then we have $\deg(\operatorname{div}(f))=0$.

Proof. Without loss of generality, we may assume that f is linear along every edge of Γ by adding vertices whenever f changes slope. Then it is enough to show that $\deg(\operatorname{div}(f|_e))$ for any edge $e \in E(\Gamma)$ is zero. However, by assumption f is linear on every edge so the only points that contribute to $\operatorname{div}(f|_e)$ are the endpoints of the edge e. Since Γ is compact, every edge will have two different end points or is a loop. In the case of a loop, since we subdivided Γ so that f has no slope changes along an edge, so f maps this loop to a single point which means $\operatorname{div}(f|_e) = 0$ in this case. In the case that e has two different end points, they both have opposite outgoing slopes, so their sum vanishes. \square

Definition 4.0.8. Let $\varphi : \Gamma' \to \Gamma$ be a harmonic morphism then we define the push-forward $\varphi_* : \operatorname{Div}(\Gamma') \to \operatorname{Div}(\Gamma)$ and the pull-back $\varphi^* : \operatorname{Div}(\Gamma) \to \operatorname{Div}(\Gamma')$ by

$$\varphi_*(p') = (\varphi(p'))$$
 and $\varphi^*(p) = \sum_{p' \in \varphi^{-1}(p)} d_{p'}(\varphi)(p')$

where we extend it linearly to all divisors. We also note that if $D \in \text{Div}(\Gamma)$ and $D' \in \text{Div}(\Gamma')$ then $\deg(\varphi^*(D)) = \deg(\varphi) \deg(D)$ and $\deg(\varphi_*(D)) = \deg(D)$ which matches the classical case. We may also define the push-forward of rational functions $\varphi_* : \text{Rat}(\Gamma') \to \text{Rat}(\Gamma)$ and the pull-back $\varphi^* : \text{Rat}(\Gamma) \to \text{Rat}(\Gamma')$ by

$$\varphi_*(f)(p) = \sum_{p' \in \varphi^{-1}(p)} d_{p'}(\varphi) f(p') \text{ and } \varphi^*(f) = f \circ \varphi.$$

Note that the push-forward of rational functions is well-defined because $d_{p'}(\varphi)$ is constant if $p' \in \varphi^{-1}(p)$ where p varies along an edge of Γ . Since the sum of piecewise integral affine functions is again piecewise integral affine, the push-forward is piecewise integral affine on each edge, hence the whole metric graph. Additionally, the pull-back is well-defined because by definition it is a dilation on edges. The definition for push-forward and pull-back of divisors are taken from [2], however, the definition for push-forward and pull-back of functions were made to satisfy the desired properties in the next section.

5 Weil-Reciprocity

In classical algebraic geometry, the Weil-reciprocity law is stated (due to [11]) as follows:

Theorem 5.0.1 (Weil-reciprocity law). Let C be a smooth projective curve over a field K and let $f, g \in \overline{K}(C)^*$ be functions such that $\operatorname{div}(f)$ and $\operatorname{div}(g)$ have disjoint support. Then we have

$$f(\operatorname{div}(g)) = g(\operatorname{div}(f))$$

where if $D = \sum_{p \in C} n_p p$ we define

$$f(D) := \prod_{p \in C} f(p)^{n_p}. \tag{2}$$

A possible approach to proving this is as in [11, Exercise 2.11] where one defines a morphism $g: C \to \mathbb{P}^1$ from the function g and use that the push forward and pullback 'play nicely' with (2). That is if $\varphi: C \to C'$ is a morphism of smooth curves then

$$f(\varphi^*D') = (\varphi_*f)(D')$$
 and $f'(\varphi_*D) = (\varphi^*f')(D)$

where φ^* and φ_* are the classical pull-back and push-forward respectively. We will follow this same approach to prove an analogue of Weil-reciprocity in the tropical world. This, of course, comes with many things that one must define, most of which we have done in previous sections.

Although our definition of a tropical curve will be slightly more general than in [8], there exists a statement and a proof of a tropical analogue of Weil-reciprocity given in it, however the proof does not provide much insight into the interplay between the classical and tropical worlds. We hope that by tropicalizing the classical proof of Weil-reciprocity, we will remedy this.

We will start by stating what we aim to prove, the statement of which is inspired by [8, Theorem 2]. In the tropical world, multiplication is replaced by addition, so a natural analogue to theorem 5.0.1 is the following theorem.

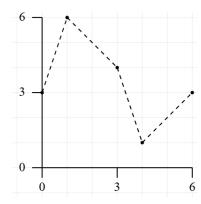
Theorem 5.0.2 (Tropical Weil-reciprocity). Let Γ be a compact metric graph and f, g be rational functions on Γ . Then

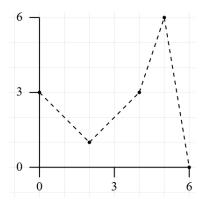
$$f(\operatorname{div}(g)) = g(\operatorname{div}(f))$$

where if $D = \sum_{p \in \Gamma} n_p p$ we define

$$f(D) := \sum_{p \in \Gamma} f(p) \cdot n_p.$$

Example 5.0.3. We begin with an example of the tropical Weil-reciprocity law in action.





$$\operatorname{div}(f) = 3(0) - 4(1) - 2(3) + 4(4) - (6) \operatorname{div}(g) = -(0) + 2(2) + 2(4) - 9(5) + 6(6)$$

Figure 12: Two rational functions on the same metric graph, and their divisors.

We now compute $f(\operatorname{div}(g))$ and $g(\operatorname{div}(f))$ and see that they are indeed equal.

- $f(\operatorname{div}(g)) = -f(0) + 2f(2) + 2f(4) 9f(5) + 6f(6) = 9$
- $g(\operatorname{div}(f)) = 3g(0) 4g(1) 2g(3) + 4g(4) g(6) = 9$

Unless otherwise stated, the discussions that follow in this section are entirely my own work. We shall now go through with establishing the remaining tools that we require in order to prove theorem 5.0.2 here.

Lemma 5.0.4. Let Γ, Γ' be metric graphs and $\varphi : \Gamma' \to \Gamma$ a harmonic morphism. Then for any rational functions $f \in \text{Rat}(\Gamma)$, $g \in \text{Rat}(\Gamma')$ we have

$$\varphi^* \operatorname{div}(f) = \operatorname{div}(\varphi^* f) \text{ and } \varphi_* \operatorname{div}(g) = \operatorname{div}(\varphi_* g).$$

Proof. Let $p' \in \Gamma'$ then by definition we have

$$\varphi^* \operatorname{div}(f)(p') = d_{p'}(\varphi) \operatorname{ord}_{\varphi(p')}(f).$$

On the other hand, let e_1, \ldots, e_m be all edges incident to p' with non-zero degrees with respect to φ . If we then denote the slope of f along some edge $e \in \Gamma$ by $\Delta_e(f)$, then we may write

$$\operatorname{div}(f \circ \varphi)(p') = d_{e_1}(\varphi) \cdot \Delta_{\varphi(e_1)}(f) + \dots + d_{e_m}(\varphi) \cdot \Delta_{\varphi(e_m)}(f).$$

Then group up the edges e_1, \ldots, e_m by declaring that each group maps to a unique edge. In other words, e_1, \ldots, e_{m_1} map to the same edge, $e_{m_1+1}, \ldots, e_{m_2}$

all map to the same edge, until the last group e_{m_t+1}, \ldots, e_m . Then we may rewrite

$$\operatorname{div}(f \circ \varphi)(p') = (d_{e_1}(\varphi) + \dots + d_{e_{m_1}}(\varphi)) \cdot \Delta_{\varphi(e_1)}(f) + \dots + (d_{e_{m_t+1}}(\varphi) + \dots + d_{e_m}(\varphi)) \cdot \Delta_{\varphi(e_{m_t+1})}(f).$$

However by harmonicity we may rewrite this to

$$\operatorname{div}(f \circ \varphi)(p') = d_{p'}(\varphi) \Delta_{\varphi(e_1)}(f) + \dots + d_{p'}(\varphi) \Delta_{\varphi(e_{m_{+}+1})}(f) = d_{p'}(\varphi) \operatorname{ord}_{\varphi(p')}(f).$$

The last equality is because harmonic morphisms are surjective so we achieve all edges incident to $\varphi(p')$, hence all outgoing slopes of f from $\varphi(p')$. Then for the next claim we have

$$\varphi_* \operatorname{div}(g)(p) = \sum_{p' \in \varphi^{-1}(p)} \operatorname{ord}_{p'}(g).$$
(3)

On the other hand, the slope of φ_*g along an edge $e \in \Gamma$ is the sum of the slopes of g along the edges that map to e under φ . This is because on the interior of an edge $d_{p'}(\varphi)$ is the degree of the said edge, and under the dilation, we divide the slopes of g by the degree of the edge to get the new slope. Since the push-forward requires multiplying by the degree of the edge, we recover the slopes of g. Then it follows that the order at a point is the sum of the orders as in (3).

Lemma 5.0.5. Let $\varphi : \Gamma' \to \Gamma$ be a harmonic morphism of metric graphs. Then for any rational functions $f \in \operatorname{Rat}(\Gamma)$, $g \in \operatorname{Rat}(\Gamma')$ and divisors $D \in \operatorname{Rat}(\Gamma)$, $D' \in \operatorname{Div}(\Gamma')$ we have

$$g(\varphi^*D) = (\varphi_*g)(D)$$
 and $f(\varphi_*D') = (\varphi^*f)(D')$.

Proof. Note that as functions on divisors, φ^* , φ_* and any rational function is linear, so we may assume that D = (p) and D' = (p') are single points. We first prove the equality on the left, and to that end we have

$$g(\varphi^*(p)) = g\left(\sum_{p' \in \varphi^{-1}(p)} d_{p'}(\varphi)(p')\right)$$
$$= \sum_{p' \in \varphi^{-1}(p)} g(p') \cdot d_{p'}(\varphi)$$

which is exactly the evaluation of φ_*g at p. Next we have

$$f(\varphi_*(p')) = (f \circ \varphi)(p') = (\varphi^*f)(p')$$

which proves the desired result.

Next, we construct a harmonic morphism out of a rational function, and we will illustrate how to do this in general in the following excursion.

Figure 13 illustrates the construction of a morphism associated to the rational function f whose domain is a line segment, which we shall call Γ . Here we place vertices on Γ on the end points of every slope change and the points that map to the same value as these slope changes, in particular the vertices occur at $f^{-1}(\{0, \pm 1, -2\})$ which can be seen in the second image.

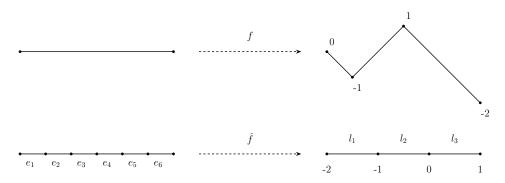


Figure 13: Illustration of the construction of a morphism from a rational function.

In this case there were no slopes that are 0, however if this happened to be the case then we collapse this line segment down to a single point (in other words we set its degree to 0). The lengths of e_1, \ldots, e_6 may vary depending on how steep we choose the slopes of f to be. In any case, the edge e_1 corresponds to the segment going from 0 to -1 in the codomain of f, the edge e_2 corresponds to the segment going from -1 to 0, the edge e_3 the segment from 0 to 1, and so forth. The edges are then mapped to the corresponding edge in the codomain of \hat{f} , for example, we map e_1 to the edge e_2 . All that remains is to choose the lengths of e_1 , e_2 and e_3 so that e_3 is harmonic. If we denote the lengths of an edge e_3 by e_4 and similarly e_2 for the others then this is equivalent to asking

$$\frac{d_{l_1}}{d_{e_6}} = \frac{d_{l_2}}{d_{e_2}} + \frac{d_{l_2}}{d_{e_5}} = \frac{d_{l_3}}{d_{e_3}} + \frac{d_{l_3}}{d_{e_4}}.$$

Then we require that each fraction is a positive integer, since they represent the dilations of the morphism \hat{f} . However, this is not always possible, for example if f happened to force $d_{e_2} = 1$ and $d_{e_5} = \sqrt{2}$ then we must have that $d_{e_5}/d_{e_2} = \sqrt{2}$ must be rational, which is not the case. Nevertheless, none of the lemmas that were proven in this section explicitly required the harmonic morphism to have integer dilations. This leads to the following definition.

Definition 5.0.6. Let Γ' , Γ be metric graphs and $\varphi : \Gamma' \to \Gamma$ a continuous map. Then φ is said to be a pseudo-morphism if it is a morphism in the regular

sense, except the dilations are allowed to be any nonnegative real number. Additionally, φ is said to be pseudo-harmonic if is it a pseudo-morphism and is harmonic in the sense of definition 4.0.4. We also call a function $f: \Gamma \to \mathbb{R}$ a pseudo-rational function if it is a rational function, possibly with real slopes.

Now, given any rational function f, it is always possible to construct a pseudo-harmonic morphism \hat{f} from it. For instance, if we have a look at the earlier example, then we can choose

$$d_{l_1} = 1$$
, $d_{l_2} = \frac{1}{d_{e_6}} \left(\frac{1}{d_{e_2}} + \frac{1}{d_{e_5}} \right)^{-1}$ and $d_{l_3} = \frac{1}{d_{e_6}} \left(\frac{1}{d_{e_3}} + \frac{1}{d_{e_4}} \right)^{-1}$.

This will always be possible given any metric graph and rational function f because we will always get a chain of equalities in the shape of the above example. Now if we denote the codomain of \hat{f} by Γ , then there exists a pseudorational function $i_f:\Gamma\to\mathbb{R}$ such that $\hat{f}^*i=f$ which essentially undoes the dilations. The last ingredient to prove is that Weil-reciprocity is true for any line segment, since given any metric graph Γ and an $f\in\mathrm{Rat}(\Gamma)$ then the codomain of \hat{f} is a line segment. The following proof is taken from [8].

Lemma 5.0.7. Let Γ be a compact line segment, and $f, g \in \text{Rat}(\Gamma)$ then

$$f(\operatorname{div}(g)) = g(\operatorname{div}(f)).$$

Proof. Let $x_1 < \cdots < x_t$ be a complete list of points in Γ of the poles and zeroes of f and g. Then on each interval $[x_i, x_{i+1}]$ for $1 \le i \le t-1$ we have that f and g are both linear, so we denote the slope of f and g on said interval by α_i and β_i respectively. Then we have

$$f(\operatorname{div}(g)) = \beta_1 f(x_1) + \sum_{i=2}^{t-1} (\beta_i - \beta_{i-1}) f(x_i) - \beta_{t-1} f(x_t)$$
$$= \sum_{i=1}^{t-1} \beta_i (f(x_i) - f(x_{i+1}))$$

and similarly

$$g(\operatorname{div}(f)) = \sum_{i=1}^{t-1} \alpha_i (g(x_i) - g(x_{i+1})).$$

Now we can rewrite

$$\alpha_i = \frac{f(x_{i+1}) - f(x_i)}{x_{i+1} - x_i}$$
 and $\beta_i = \frac{g(x_{i+1}) - g(x_i)}{x_{i+1} - x_i}$.

Plugging these into $f(\operatorname{div}(g))$ and $g(\operatorname{div}(g))$ we see that they are equal by changing the position of the $x_{i+1} - x_i$ in the denominator.

Remark 5.0.8. Note that lemma 5.0.7 is not necessarily true if Γ is not compact. For example, we follow the same notation as in the proof with $\Gamma = (a, b]$ for some $-\infty \leq a, b < \infty$ with slopes α_0 , β_0 for f, g on the interval $(a, x_1]$. Then by following the proof we get

$$f(\operatorname{div}(g)) = -\beta_0 f(x_1) + \sum_{i=1}^{t-1} \beta_i (f(x_i) - f(x_{i+1}))$$
$$= -\beta_0 f(x_1) + \sum_{i=1}^{t-1} \alpha_i (g(x_i) - g(x_{i+1}))$$
$$= -\beta_0 f(x_1) + \alpha_0 g(x_1) + g(\operatorname{div}(f)).$$

In particular, if $f(\operatorname{div}(g)) = g(\operatorname{div}(f))$ then we must have $\beta_0 f(x_1) = \alpha_0 g(x_1)$. However, this is certainly not true for all rational functions $f, g \in \operatorname{Rat}(\Gamma)$. This is the reason we must restrict ourselves to compact metric graphs in theorem 5.0.2.

Now we are in a position to prove the general Weil-reciprocity law for compact metric graphs. For convenience, we shall state it again and then prove it.

Theorem 5.0.9 (Tropical Weil-reciprocity). Let Γ be a compact metric graph and f, g be rational functions on Γ . Then

$$f(\operatorname{div}(g)) = g(\operatorname{div}(f))$$

where if $D = \sum_{p \in \Gamma} n_p p$ we define

$$f(D) := \sum_{p \in \Gamma} f(p) \cdot n_p.$$

Proof. Let Γ be a compact metric graph and $f, g \in \text{Rat}(\Gamma)$. Then we have

$$f(\operatorname{div}(g)) = f(\operatorname{div}(\hat{g}^*i_g))$$

$$= (\hat{g}_*f)\operatorname{div}(i_g)$$

$$= i_g(\operatorname{div}(\hat{g}_*f))$$

$$= (\hat{g}^*i_g)\operatorname{div}(f)$$

$$= q(\operatorname{div}(f)).$$

These equalities follow from lemma 5.0.7 and the properties of the push-forward and pull-back that we proved in the earlier portions of this section. \Box

6 Conclusion and Discussion

In this paper, we set out to transfer the classical Weil-reciprocity law into the tropical setting, aiming for a proof that closely parallels the algebraic-geometric argument and thereby illustrating the interplay between classical and tropical worlds. After laying the groundwork in polyhedral geometry and formalizing tropical spaces via charts and balancing conditions, we specialized to tropical curves realized as metric graphs. We then developed notions of divisors, rational functions, and harmonic morphisms on metric graphs, and defined push-forward and pull-back operations in analogy with the classical case. The centrepiece is the proof of the tropical Weil-reciprocity theorem, which verifies that for any two rational functions f, g on a compact metric graph Γ one has

$$f(\operatorname{div}(g)) = g(\operatorname{div}(f))$$

which is identical to the classical case. Our approach mirrors the classical push and pull argument by constructing, for a rational function, a suitable (pseudo-)harmonic morphism to a line segment and then invoking reciprocity on said segment using lemma 5.0.7.

This viewpoint hopefully offers conceptual clarity compared to the proof given in [8], emphasizing how tropical analogues of push-forwards and pullbacks behave much like their algebraic counterparts under appropriate definitions. It also highlights why compactness is crucial; as noted in remark 5.0.8, the statement fails for non-compact graphs, reflecting subtle boundary behaviour in the tropical context.

There remain natural avenues for further exploration. One may investigate whether some variation of the tropical Weil-reciprocity law exists for non-compact metric graphs. Additionally, we may investigate concepts that are reliant on Weil-reciprocity, for example, the Deligne pairing (as seen in [3, pp.366-368]) since we know what tropical line bundles are due to [1].

Ultimately, this work contributes to the growing corpus of tropical analogues of classical theorems and concepts (e.g., Riemann–Roch, vector bundles, ψ -classes, etc., as pointed out in the introduction), illustrating that tropical geometry not only provides combinatorial counterparts but often recapitulates the structural essence of algebraic phenomena. We hope that the methods and insights developed here will serve as a template for further tropicalization of classical results, enriching both worlds, the tropical and classical perspectives.

References

- [1] L. Allermann, Chern classes of tropical vector bundles, arXiv preprint: arXiv:0911.2909, 2009.
- [2] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff, *Lifting Harmonic Morphisms I: Metrized Complexes and Berkovich Skeleta*, arXiv preprint: arXiv:1303.4812v3, 2013.
- [3] E. Arbarello, M. Cornalba, and P.A. Griffiths, *Geometry of Algebraic Curves*, Vol. 2, Springer-Verlag Berlin Heidelberg 2011, ISBN: 978-3-540-42688-2
- [4] R. Cavalieri, A. Gross, H. Markwig, *Tropical* ψ *classes*, arXiv preprint: arXiv:2009.00586, 2020.
- [5] A. Gathmann, M. Kerber, A Riemann-Roch theorem in tropical geometry, arXiv preprint: arXiv:math/0612129, 2007.
- [6] B. Grünbaum, Convex Polytopes, 2nd ed., Springer New York, NY, 2003.
 ISBN: 978-0-387-00424-2, DOI: 10.1007/978-1-4613-0019-9.
- [7] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002. Available at: https://pi.math.cornell.edu/~hatcher/AT/AT.pdf. ISBN: 978-0-521-79540-1.
- [8] N. Kalinin and M. Magin, Tropical Weil's Reciprocity Law and Weil's Pairing, arXiv preprint: arXiv:2408.06372v1, 2024.
- [9] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*, Graduate Studies in Mathematics, Vol. 161, American Mathematical Society, 2015. ISBN: 978-0-8218-5198-2.
- [10] G. Mikhalkin and J. Rau, Tropical Geometry, (Unpublished draft), November 16, 2018.
- [11] J.H. Silverman, *The Arithmetic of Elliptic Curves*, Springer-Verlag, New York, 2009. ISBN: 978-0-387-09494-6, DOI: 10.1007/978-0-387-09494-6.
- [12] G.M. Ziegler, Lectures on Polytopes, Springer-Verlag, New York, 1995.
 ISBN: 978-1-4613-8431-1, DOI: 10.1007/978-1-4613-8431-1.