

Tropical Geometry

Grigory Mikhalkin, Johannes Rau

November 16, 2018

Contents

1	Introduction	5
1.1	Algebraic geometry over what?	5
1.2	Tropical Numbers	8
1.3	Tropical monomials and integer affine geometry	11
1.4	Forgetting the phase leads to tropical numbers	15
1.5	Amoebas of affine algebraic varieties and their limits	17
1.6	Patchworking and tropical geometry	22
1.7	Concluding Remarks	28
2	Tropical hypersurfaces in \mathbf{R}^n	29
2.1	Polyhedral geometry dictionary I	30
2.2	Tropical Laurent polynomials and hypersurfaces	34
2.3	The polyhedral structure of hypersurfaces	40
2.4	The balancing condition	47
2.5	Planar Curves	56
2.6	Floor decompositions	69
3	Projective space and other tropical toric varieties	76
3.1	Tropical affine space \mathbf{T}^n	76
3.2	Tropical toric varieties	80
3.3	Tropical projective space	84
3.4	Projective hypersurfaces	88
3.5	Projective curves	95
4	Tropical cycles in \mathbf{R}^n	100
4.1	Polyhedral geometry dictionary II	100
4.2	Tropical cycles and subspaces in \mathbf{R}^n	102
4.3	Stable intersection	106
4.4	The divisor of a piecewise affine function	116
4.5	Tropical modifications	127

4.6	The projection formula	131
4.7	Examples and further topics	136
5	Projective tropical geometry	145
5.1	Tropical cycles in toric varieties	145
5.2	Projective degree	149
5.3	Modifications in toric varieties	153
5.4	Configurations of hyperplanes in projective space	166
5.5	Equivalence of rational curves	169
5.6	Equivalence of linear spaces	170
6	Tropical cycles and the Chow group	172
6.1	Tropical cycles	172
6.2	Push-forwards of tropical cycles	173
6.3	Linear equivalence of cycles	175
6.4	Algebraic equivalence of cycles	176
6.5	Linear systems	177
6.6	Recession Fans	177
7	Tropical manifolds	179
7.1	Tropical spaces	179
7.2	Regularity at Infinity	187
7.3	Tropical morphisms and structure sheaves	193
7.4	Smooth tropical spaces	196
7.5	Rational functions, Cartier divisors and modifications on abstract spaces	206
8	Tropical curves	210
8.1	Smooth tropical curves	210
8.2	Divisors and linear systems	213
8.3	Line bundles on curves	215
8.4	Riemann-Roch theorem	221
9	Tropical jacobians	229
9.1	(Coarse) differential forms, integration, and the Jacobian	229
9.2	Abelian varieties, polarization, and theta functions	234
9.3	The Abel-Jacobi theorem	238
9.4	Riemann's theorem	250

Acknowledgements

We benefited from numerous valuable conversations, in particular with Victor Batyrev, Benoît Bertrand, Erwan Brugallé, Ana Cannas da Silva, Jim Carlson, Yakov Eliashberg, Sergey Fomin, Andreas Gathmann, Mark Gross, Amihay Hanany, Ilia Itenberg, Michael Kapovich, Mikhail Kapranov, Ludmil Katzarkov, Yujiro Kawamata, Sean Keel, Viatcheslav Kharlamov, Askold Khovanskii, Barak Kol, Maxim Kontsevich, Lionel Lang, Conan Leung, Andrey Losev, Diane Maclagan, Hannah Markwig, Pierre Milman, Yong-Geun Oh, Mikael Passare, Jean-Jacques Risler, Hans Rullgård, Kristin Shaw, Eugenio Shustin, Yan Soibelman, David Speyer, Bernd Sturmfels, Louis Felipe Tabera, Ilya Tyomkin, Israel Vainsencher, Oleg Viro, Ilia Zharkov and Günter Ziegler on tropical geometry, especially during early stages of its development.

(To be continued.)

We would like to acknowledge support and encouragement from the Clay Mathematical Institute (USA) at the initial stage of this project. The project was also supported in part by the European Research Council (TROPGEO project), the Swiss National Science Foundation (projects no. 125070 and 140666). We would also like to thank the Mathematical Sciences Research Institute in Berkeley and the Max-Planck-Institut für Mathematik in Bonn for hospitality at different times during work on this project.

1 Introduction

1.1 Algebraic geometry over what?

The two words “Algebraic Geometry” may invoke different images for different people today. For some it might be something very abstract, and even somewhat formalistic, adapted to work in the most general setting possible. For many others, including the authors of this book, Algebraic Geometry is a branch of geometry in the first place — geometry of spaces defined by polynomial equations.

Note that (unlike many algebraic properties) the resulting geometry depends not only on the type of the defining equations, but also on the choice of the numbers where we look for solutions. The two most classical choices are the field \mathbf{R} of real numbers and the field \mathbf{C} of complex numbers. Both these fields come naturally enhanced with the so-called “Euclidean topology” induced by the metric $|x - y|$ between two points $x, y \in \mathbf{R}$ (or in \mathbf{C}). Furthermore, real algebraic varieties are differentiable manifolds (perhaps with singularities) from a topological viewpoint, and complex algebraic varieties are special kind of real algebraic varieties of twice the dimension.

To illustrate the two parallel classical theories let us recall the classical example of the so-called elliptic curve. Namely, consider a cubic curve in the complex projective plane \mathbf{CP}^2 . As long as the defining cubic polynomial is chosen generically, the resulting curve is topologically a 2-dimensional torus (see Figure 1.1). This torus is embedded in the complex projective plane \mathbf{CP}^2 . As \mathbf{CP}^2 is 4-dimensional such embedding is beyond our imagination tools.

Now consider the case of real coefficients. Even if the defining polynomial is chosen generically, the topological type of the curve is not fixed. But there are only two possible cases, see Figure 1.2.

As the ambient real projective plane \mathbf{RP}^2 is indeed 2-dimensional, we can actually draw how the curve is embedded there. Recall that topologically




Figure 1.1: A complex elliptic curve.




Figure 1.2: Two real elliptic curves.

\mathbf{RP}^2 can be obtained from a disc D^2 by identifying the antipodal points on its boundary circle, see Figure 1.3.




Figure 1.3: Real projective plane.




Figure 1.4: Elliptic curves in the real projective plane.

Figure 1.4 depicts embeddings of cubic curves in the real projective plane. Note that there might be different pictures inside D^2 before the self-identification of its boundary, but we get one of the two pictures above in $\mathbf{RP}^2 = D^2 / \sim$ for any smooth real curve. An example is given in Figure 1.5.




Figure 1.5: An elliptic curve in another presentation of \mathbf{RP}^2 by a disk.

Furthermore, the inclusion $\mathbf{R} \subset \mathbf{C}$ gives us an inclusion $\mathbf{RP}^2 \subset \mathbf{CP}^2$ as well as an inclusion of the real curve into the corresponding complex curve (see Figure 1.6).

We see that the same equation may yield quite different geometric spaces. At the same time, real and complex numbers may be the only examples of *fields* where algebraic geometry is that much geometric. There is also a continuation of this series with the quaternions \mathbf{H} and octonions \mathbf{O} which leads to very interesting geometry no longer based on fields as we loose




Figure 1.6: Complex elliptic curve with its real locus.

commutativity in the case of \mathbf{H} or even associativity in the case of \mathbf{O} .

In this book we study geometry based on a predecessor of the entire series \mathbf{R} , \mathbf{C} , \mathbf{H} , \mathbf{O} , called the *tropical numbers* \mathbf{T} .

1.2 Tropical Numbers

We consider the set

$$\mathbf{T} = [-\infty, +\infty) = \mathbf{R} \cup \{-\infty\}$$

enhanced with the arithmetic operations

$$"x + y" = \max\{x, y\}, \quad (1.1)$$

$$"xy" = x + y, \quad (1.2)$$

where we set $"(-\infty) + x" = "x + (-\infty)" = x$ and $"(-\infty)x" = "x(-\infty)" = -\infty$. These operations are called *tropical arithmetic operations*. We use quotation marks to distinguish them from the usual operations on \mathbf{R} .

Definition 1.2.1

The set \mathbf{T} enhanced with the arithmetic operations (1.1) and (1.2) is called the set of *tropical numbers*.

Remark 1.2.2

There are papers where $\min\{x, y\}$ is taken for tropical addition. In such case one has to modify the set of tropical numbers to include $+\infty$ and exclude $-\infty$. It is hard to say which choice is better. The choice of \max may be more natural from the *mathematical* viewpoint as we are more used to taking the logarithm whose base is greater than 1, cf. equation (1.4). Also when we add two numbers in this way the sum does not get smaller. However in some considerations in Computer Science and Physics (cf. [IM12]) taking the minimum is more natural. Clearly it does not really matter as the two choices are isomorphic under $x \mapsto -x$.

Remark 1.2.3

The term “tropical” was borrowed from Computer Science, where it was reportedly introduced to commemorate contributions of the Brazilian computer scientist Imre Simon. Simon introduced the semiring $(\mathbf{N}, \min, +)$ which was later baptized “tropical” following a suggestion of Christian Chofurt (see [Sim88; Cho92]). According to Jean-Eric Pin, the suggestion was (also) made by Dominique Perrin (see [Pin98]). Some years later, in the first days of what is now known as tropical geometry, the adjective “tropical” quickly outplayed other suggestions inspired by “non-archimedean amoebas”, “idempotent algebraic geometry” or “logarithmic limit sets”.

The set \mathbf{T} is a semigroup with respect to tropical addition. It is commutative, associative and admits the neutral element $0_{\mathbf{T}} = -\infty$. Nevertheless we do not get a group as there is no room for subtraction. Indeed, if “ $x + y$ ” = $0_{\mathbf{T}}$, then either $x = 0_{\mathbf{T}}$ or $y = 0_{\mathbf{T}}$. Thus the only element admitting an inverse with respect to tropical addition is the neutral element $0_{\mathbf{T}} = -\infty$.

We may note that tropical addition is idempotent, i.e. we have “ $x + x$ ” = x for all $x \in \mathbf{T}$. Idempotency makes tropical numbers non-Archimedean. Let us recall the Archimedes axiom (stated for the case of real numbers).

Axiom 1.2.4 (Archimedes)

For any positive real numbers $a, b \in \mathbf{R}$, $a, b > 0$ there exists a natural number $n \in \mathbf{N}$ such that

$$\underbrace{a + a + \dots + a}_n > b.$$

Clearly, the conventional linear order $>$ on $\mathbf{T} = \mathbf{R} \cup \{-\infty\}$ makes perfect sense tropically. Furthermore, we can express it in terms of tropical addition: we have $a \geq b$ if and only if “ $a + b$ ” = b . However, we have “ $\underbrace{a + a + \dots + a}_{n \text{ times}}$ ” = a independently of n and thus the Archimedes axiom does not hold for tropical numbers.

Remark 1.2.5

Note that the Euclidean topology on \mathbf{T} is determined by the linear order on \mathbf{T} : it is generated by the sets $U_a = \{x \in \mathbf{T} \mid x < a\}$ and $V_a = \{x \in \mathbf{T} \mid a < x\}$ (as a subbase) for all possible $a \in \mathbf{T}$.

The tropical non-zero numbers are $\mathbf{T}^\times = \mathbf{R}$. Of course, they form an honest group with respect to tropical multiplication as it coincides with the conventional addition. It is easy to check that the tropical arithmetic operations satisfy the distribution law

$$“(x + y)z” = “xz + yz”.$$

So we see that the only defect of tropical arithmetics is the missing subtraction which makes the tropical numbers a so-called *semifield* instead of a field. It does not stop us from defining polynomials. Let $A \subset \mathbf{N}^n$ be a finite set of integer vectors with non-negative entries. We denote the entries of $j \in A$ by $j = (j_1, \dots, j_n)$. The function

$$“\sum_{j \in A} a_j x_1^{j_1} \dots x_n^{j_n}” = \max_{j \in A} \{a_j + jx\} : \mathbf{T}^n \rightarrow \mathbf{T}$$

of $x = (x_1, \dots, x_n) \in \mathbf{T}^n$ is a *tropical polynomial* in n variables. Here, $a_j \in \mathbf{T}$ and jx denotes the standard scalar product $\sum_i j_i x_i$. If we also allow negative exponents, i.e. $A \subset \mathbf{Z}^n$, we get a *tropical Laurent polynomial*. If we even drop the finiteness condition of A , we obtain a *tropical Laurent series*.

Tropical polynomials are globally well-defined continuous (with respect to the Euclidean topology) functions $F : \mathbf{T}^n \rightarrow \mathbf{T}$ with $F(-\infty, \dots, -\infty) = a_0$. Laurent polynomials are always defined on $\mathbf{R}^n \subset \mathbf{T}^n$, but not necessarily on

$$\partial \mathbf{T}^n = \mathbf{T}^n \setminus \mathbf{R}^n = \{(x_1, \dots, x_n) \in \mathbf{T}^n \mid x_j = -\infty \text{ for some } j\}.$$

Indeed, if $x \in \mathbf{T}^n$ is such that $x_l = -\infty$, then all monomials “ $a_j x^j$ ” in a Laurent polynomial F have to satisfy $j_l \geq 0$ whenever $a_j \neq -\infty$ as otherwise the value of such monomial is $+\infty \notin \mathbf{T}$.

An infinite tropical Laurent series does not have to be well-defined even on \mathbf{R}^n . However, it is easy to check that the domain of a tropical Laurent series is convex (though not necessarily open or closed).

1.3 Tropical monomials and integer affine geometry

Each tropical monomial “ $a_j x^j$ ” = $a_j + jx$ with $a_j \neq -\infty$ is an affine function $\lambda : \mathbf{R}^n \rightarrow \mathbf{R}$, $x = (x_1, \dots, x_n)$, $j = (j_1, \dots, j_n)$. We may extend it to a part of $\partial \mathbf{T}^n$ by continuity. Namely, we define a domain D with

$$\mathbf{R}^n \subseteq D \subseteq \mathbf{T}^n$$

as \mathbf{T}^n minus all points $x = (x_1, \dots, x_n)$ with $x_l = -\infty$ whenever $j_l < 0$. Clearly, λ gets naturally extended to a continuous function $\bar{\lambda} : D \rightarrow \mathbf{T}$ and D is the maximal set in \mathbf{T}^n where such a continuous extension exists.

The *linear part* jx of the monomial λ is defined by an integer vector $j \in \mathbf{Z}^n$ of the dual vector space $(\mathbf{R}^n)^*$ (as we have the classical pairing $jx \in \mathbf{R}$, $x \in \mathbf{R}^n$). Thus we may invert our construction and define tropical monomials (and hereby also polynomials, series, etc.) starting from an arbitrary real vector space V of dimension n as long as we fix a lattice $N_V \subset V$. Here, by a lattice in a vector space V we mean a discrete free subgroup $N \subset V$ of rank n .

Indeed, once we fixed $N_V \subset V$ we have a canonical dual lattice $N_V^* \subset V^*$ which consists of linear functionals $V \rightarrow \mathbf{R}$ taking integer values on N_V . A tropical monomial on V is a function of the form $\lambda(x) + c$, $\lambda \in N_V^*$, $c \in \mathbf{R}$. Taking finite tropical sums of tropical monomials we get the notion of tropical polynomials on V . We call a real vector space V together with a lattice $N_V \subset V$ a *tropical vector space* of dimension n (or just a *tropical n -space*, not to be confused with tropical affine space \mathbf{T}^n) and the lattice N_V the *tropical lattice*.

Definition 1.3.1

An affine map $\Phi : V \rightarrow W$ between tropical vector spaces V, W is called *integer affine* if for any $x, y \in V$ with $x - y \in N_V$ we have $(\Phi(x) - \Phi(y)) \in N_W$. The map Φ is called an integer affine transformation of V if $V = W$, it is

invertible and the inverse function is an integer affine map, too. We define the *differential* $d\Phi : V \rightarrow W$ as the conventional differential (identifying $V = T_x V$ and $W = T_{\phi(x)} W$). Note that $d\Phi$ takes N_V to N_W . Clearly, an integer affine map $\Phi : V \rightarrow V$ is an integer affine transformation if and only if $d\Phi|_{N_V}$ is a bijection to N_V .

We can generalize these notions to local versions whenever the underlying space comes equipped with a so-called integer affine structure.

Definition 1.3.2

Let M be a smooth n -dimensional manifold. An *integer affine structure* on M is given by an open cover U_α and charts $\phi_\alpha : U_\alpha \rightarrow \mathbf{R}^n$ such that for each pair α, β the overlapping map $\phi_\beta \circ \phi_\alpha^{-1}$ is locally the restriction of an integer affine transformation $\Phi_{\beta\alpha} : \mathbf{R}^n \rightarrow \mathbf{R}^n$.

Two integer affine structures on M are considered equivalent if the union of their covers and charts gives an integer affine structure as well. A manifold M together with an equivalence class of integer affine structures is called an *integer affine manifold*.

Example 1.3.3

Let A be a real affine space with tangent vector space V . Obviously, by choosing a lattice $N_A \subset V$ (i.e. turning V into a tropical vector space) induces an integer affine structure on A . It is tempting to call such an integer affine manifold a tropical affine space. However, in this text we reserve this term for the space \mathbf{T}^n (following the algebro-geometric viewpoint here).

Example 1.3.4

Let (V, N) be a tropical vector space and let $\Lambda \subset V$ be an arbitrary lattice (unrelated to N). We declare points $x, y \in A$ equivalent if $x - y \in \Lambda$. Let $T := V/\Lambda$ denote the quotient space. Then T is a integer affine manifold diffeomorphic to $(S^1)^n$. Indeed, the quotient map $V \rightarrow T$ can be inverted locally and the usual atlas given by local inverse maps provides an integer affine structure on T . We call T a *tropical torus*.

We will see that the tropical structure on general tropical manifolds can be thought of as an extension of integer affine structures to polyhedral complexes. Let us preview this (without detailed explanation) by revisiting the example of a smooth cubic curve in the plane.

Now we do everything tropically borrowing the notions from the main part of the book. We will see that the tropical projective plane \mathbf{TP}^2 can be viewed as a compactification of the tropical 2-space \mathbf{R}^2 . Combinatorially, this compactification is a triangle whose interior is identified with the whole \mathbf{R}^2 .




Figure 1.7: Tropical projective plane \mathbf{TP}^2 .

A curve in \mathbf{TP}^2 might be represented as a picture in \mathbf{R}^2 which has to be compactified to get the full picture. Or we can consider already compact pictures in the triangle whose interior is equipped with an distorted integer affine structure inherited from a diffeomorphism with \mathbf{R}^2 . As we shall see in this book, Figure 1.8 depicts a smooth cubic curve in \mathbf{R}^2 before the compactification. It is the hypersurfaces in \mathbf{R}^2 defined by a tropical cubic polynomial in two variables. In the compactified view, the edges are no longer straight. Figure 1.9 provides a sketch (where the curvature of the edges is perhaps still not visible).

We may note that our cubic intersects each side of the triangle three times. Similarly to the situation over the real numbers, there is more than one “type” of a smooth tropical curve of given degree, but only finitely many.

To preview the relation between classical and tropical varieties we may look at Figure 1.10 which depicts the collapse of the complex elliptic curve from Figure 1.1 to the tropical cubic curve in \mathbf{TP}^2 . The left-hand side of Figure 1.10 is topologically a 2-torus $S^1 \times S^1$ minus nine points. These are the nine points (three from each side of the triangle, cf. Figure 1.9) that are




Figure 1.8: A tropical cubic curve in \mathbf{R}^2 (before compactification).




Figure 1.9: A tropical cubic curve embedded in \mathbf{TP}^2 .

attached after the compactification in \mathbf{TP}^2 . In the next section we look at the calculus that governs the collapse of complex varieties to tropical ones.




Figure 1.10: Collapse of a complex cubic curve to a tropical cubic curve.

1.4 Forgetting the phase leads to tropical numbers

To relate complex and tropical numbers consider a complex number

$$z = ae^{i\alpha}, \quad a, \alpha \in \mathbf{R}, a \geq 0.$$

The exponent α is called the argument, or the phase of z . It gets ignored when passing to tropical numbers. In order to do this consistently with addition and multiplication we measure the remaining parameter, the norm $a = |z|$, on logarithmic scale. We set

$$x = \log_t a = \log_t |z|,$$

where the base t of the logarithm is a large real number.

The map $z \mapsto \log_t |z|$ is a surjection $\mathbf{C} \rightarrow \mathbf{T}$ which we denote by Log_t (of course, we set $\log_t 0 = -\infty$). This map “forgets” the phase of z . Furthermore, it rescales the norm. We may use Log_t to induce arithmetic operations on \mathbf{T} from \mathbf{C} . However, since Log_t is not injective the resulting operations might be multivalued.

Indeed, addition on \mathbf{C} induces the following operation on \mathbf{T} .

$$x \vee_t y = \text{Log}_t(\text{Log}_t^{-1} x + \text{Log}_t^{-1} y) \tag{1.3}$$

We should stress here that $\text{Log}_t^{-1} x + \text{Log}_t^{-1} y$ denotes the set $\{z+w : \text{Log}_t(z) = x, \text{Log}_t(w) = y\}$ and we define $x \vee_t y$ to be the image of this set under Log_t . As $\text{Log}_t^{-1} x$ is a circle $\{z \in \mathbf{C} \mid |z| = t^x\}$, we get

$$x \vee_t y = [\log_t |t^x - t^y|, \log_t(t^x + t^y)] \subset \mathbf{T},$$

i.e. $x \vee_t y$ is an interval in \mathbf{T} instead of a specific number.

To get a unique number we define

$$“x +_t y” = \max\{x \vee_t y\} = \log_t(t^x + t^y). \quad (1.4)$$

This expression has a well-defined limit when $t \rightarrow +\infty$. We get

$$“x + y” = \lim_{t \rightarrow +\infty} “x +_t y” = \lim_{t \rightarrow +\infty} \log_t(t^x + t^y) = \max\{x, y\},$$

thus recovering tropical addition as a certain limit of addition of complex numbers with the help of rescaling by \log_t once we forgot the phase.

Remark 1.4.1

Note also that if $x \neq y$, then we have

$$“x + y” = \lim_{t \rightarrow +\infty} \min\{x \vee_t y\} = \lim_{t \rightarrow +\infty} \log_t |t^x - t^y| = \max\{x, y\}.$$

We see that in this case “ $x + y$ ” is the single limit of the multi-valued operation $x \vee_t y$ and therefore independent of the chosen phases for the preimages of x and y . If $x = y$ then $\min\{x \vee_t y\} = -\infty$, so it is independent of t . Hence, in the realm of multivalued operations, the limit of \vee_t for $t \rightarrow \infty$ is the *multivalued tropical addition*

$$a \vee b = \begin{cases} \max(a, b) & \text{if } a \neq b, \\ [\infty, a] & \text{if } a = b. \end{cases}$$

The operations \vee_t and \vee are examples of hypergroup additions, see [Vir10] for the relevant treatment in the context of tropical calculus (more explicitly, the definitions of \vee_t and \vee are given in sections 5.4 and 5.3). Viro suggests an another useful viewpoint on tropical calculus using the multi-valued addition \vee instead of the more conventional single-valued version given by \max .

Multiplication in \mathbf{C} induces a well-defined single-valued operation

$$“xy” = \text{Log}_t(\text{Log}_t^{-1} x \cdot \text{Log}_t^{-1} y) = \log_t(t^x t^y) = x + y,$$

as the norm of the product of two complex numbers is independent of their phases.

1.5 Amoebas of affine algebraic varieties and their limits

The map $\text{Log}_t : \mathbf{C} \rightarrow \mathbf{T}$ can be applied coordinatwise to generalize to the case of several variables.

$$\begin{aligned} \text{Log}_t : \mathbf{C}^n &\rightarrow \mathbf{T}^n, \\ (z_1, \dots, z_n) &\mapsto (\log_t |z_1|, \dots, \log_t |z_n|) \end{aligned}$$

Clearly, $\text{Log}_t((\mathbf{C}^\times)^n) = \mathbf{R}^n$. Images of algebraic subvarieties $V \subseteq (\mathbf{C}^\times)^n$ under Log_t are called *amoebas*. They were introduced in [GKZ08, Chapter 6]. The most well-known example is the amoeba

$$\mathcal{A}_t = \text{Log}_t(\{(z, w) \in (\mathbf{C}^\times)^2 \mid z + w + 1 = 0\}) \quad (1.5)$$

depicted in Figure 1.11.

It is the closed set in \mathbf{R}^2 bordered by three arcs

$$t^x + t^y = 1, \quad t^y + 1 = t^x, \quad 1 + t^x = t^y.$$

We add to the picture the three “asymptotics of the tentacles”: the negative part of the x -axis, the negative part of the y -axis and the diagonal ray $\{(x, x) \mid x \geq 0\}$. The union of these three rays, with the origin as vertex, is denoted by Γ . The tripod Γ separates the amoeba \mathcal{A}_t into three equal parts. To see that these parts are equal we note that the whole picture — without the boundary points — is symmetric with respect to the linear action on \mathbf{R}^2 by the symmetric group S_3 generated by $\begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$ and this action interchanges the three parts of the amoeba. Also, the action is volume-preserving as our S_3 is a subgroup of $\text{SL}_2(\mathbf{Z})$. Thus the area Z_t (which clearly depends on the parameter $t > 1$) of each part is the same. One can show that

$$Z_t = \frac{\pi^2}{6(\log t)^2} \quad (1.6)$$




Figure 1.11: The amoeba \mathcal{A}_t and three rays inside.

(in the following, the logarithm without index always refers to the natural base logarithm $\log := \log_e$) as we see in the next remark.

Remark 1.5.1

It was observed by Passare (see [Pas08]) that two different ways of computing the area Z_e produces yet another proof of Euler's formula $\zeta(2) = \pi^2/6$. Indeed, as $\log_t = \frac{\log}{\log(t)}$ coordinatewise it suffices to establish (1.6) for $t = e$. Let us compute the area of the part of \mathcal{A}_e in the negative quadrant. It is given by

$$\begin{aligned} \int_{-\infty}^0 \int_{\log(1-e^x)}^0 dx dy &= \int_{-\infty}^0 -\log(1-e^x) dx \\ &= \int_{-\infty}^0 \sum_{i=1}^{\infty} \frac{e^{ix}}{i} dx = \sum_{i=1}^{\infty} \int_{-\infty}^0 \frac{e^{ix}}{i} dx = \sum_{i=1}^{\infty} \frac{1}{i^2} = \zeta(2). \end{aligned}$$

On the other hand, one can show (using the fact that $\{z + w + 1 = 0\} \in \mathbf{C}^2$ is a holomorphic submanifold) that the area of \mathcal{A}_e is half of the area of the corresponding coamoeba, which is the image of $\{z + w + 1 = 0\} \in (\mathbf{C}^\times)^2$ under the argument map

$$\text{Arg} : (\mathbf{C}^\times)^2 \rightarrow (\mathbf{R}/2\pi\mathbf{Z})^2, (a_1 e^{i\alpha_1}, a_2 e^{i\alpha_2}) \mapsto (\alpha_1, \alpha_2).$$

Using elementary triangle geometry (for all $0 < \beta_1, \beta_2, \beta_1 + \beta_2 < \pi$ there exists a triangle with vertices $0, 1 \in \mathbf{C}$ and interior angles β_1 and β_2) one can

show that the coamoeba consists of two congruent triangles in $(\mathbf{R}/2\pi\mathbf{Z})^2$ (neglecting the boundary), each of area $\pi^2/2$. Hence $\text{area}(\mathcal{A}_e) = \pi^2/2$. But as $\zeta(2)$ represents a third of \mathcal{A}_e , it follows $\zeta(2) = \pi^2/6$.

If we consider a line $\{az + bw + c = 0\}$ with a, b, c non-zero, then the amoeba $\text{Log}_t\{az + bw + c = 0\}$ can be obtained from \mathcal{A}_t by the translation

$$\begin{aligned} x &\mapsto x + \text{Log}_t c - \text{Log}_t a, \\ y &\mapsto y + \text{Log}_t c - \text{Log}_t b, \end{aligned}$$

as $\{az + bw + c = 0\}$ can be obtained from $\{z + w + 1 = 0\}$ by the rescaling

$$z \mapsto \frac{c}{a}z, \quad w \mapsto \frac{c}{b}w.$$

Note that if $V \subseteq \mathbf{C}^n$ is fixed the only effect of varying t is the scaling of the target \mathbf{R}^n with the coefficient $\frac{1}{\log t}$, i.e.

$$\text{Log}_t(V) = \frac{1}{\log t} \text{Log}(V).$$

In particular, the limit of $\text{Log}_t\{az + bw + c = 0\}$ when $t \rightarrow +\infty$ does not depend on the coefficients (as long as they are non-zero) and is equal to Γ , the union of the three rays inside \mathcal{A}_t , see Figure 1.11.

The situation changes if we vary V simultaneously with varying the base t , i.e. if we consider a family of complex varieties $V_t \subseteq (\mathbf{C}^\times)^n$ with a real parameter t .

For example, let us take a family of lines $V_t = \{a(t)z + b(t)w + c(t) = 0\}$ in $(\mathbf{C}^\times)^2$, where $a(t) = \alpha t^A + o(t^A)$, $b(t) = \beta t^B + o(t^B)$, $c(t) = \gamma t^C + o(t^C)$, $\alpha, \beta, \gamma \in \mathbf{C}^\times$, $A, B, C \in \mathbf{R}$ are functions for large positive values of t with highest order terms as described. Then the limit L of $\text{Log}_t(V_t)$ for $t \rightarrow +\infty$ (which we may consider in the topology induced by the Hausdorff metric on neighbourhoods of compact sets in \mathbf{R}^n) depends only on A, B, C and is equal to the translation of Γ by

$$\begin{aligned} x &\mapsto x + C - A, \\ y &\mapsto y + C - B. \end{aligned}$$

Thus the numbers A, B, C define asymptotics of the amoebas $\text{Log}_t(V_t)$. Within the paradigm of tropical geometry we regard the limits L as geometric objects of their own, the so called *tropical lines*, and $A, B, C \in \mathbf{R} \subset \mathbf{T}$ as *tropical*




Figure 1.12: Tropical lines intersect in one point. They can be used to connect any pair of points in \mathbf{T}^2 .

coefficients defining this line. Indeed, we will later define the notion of a tropical hypersurface of a tropical polynomial, and in the case of the linear tropical polynomial “ $Ax + By + C$ ” this hypersurface will be equal to L . Note that the tropical lines we can get in this way only differ by translations in \mathbf{R}^2 .

Remark 1.5.2

Let us consider a special case when $b(t) = -1$, $a(t) = t^A$, $c(t) = t^C$. Then the curve V_t is a graph of the function $w(z) = t^A z + t^C$ while its amoeba can be written as

$$\text{Log}_t(V_t) = \{y = \text{“}Ax\text{”} \vee_t C\},$$

i.e. it can be thought of as the graph of the multivalued addition (1.3) of “ Ax ” = $A + x \in \mathbf{T}$ and $C \in \mathbf{R}$. The statement continues to hold in the limit case $t \rightarrow \infty$ if we use multivalued tropical addition \vee , i.e.

$$\lim_{t \rightarrow \infty} \text{Log}_t(V_t) = \{y = \text{“}Ax\text{”} \vee C\},$$

We easily get some familiar properties of lines for the new piecewise linear objects. Two *generic* lines intersect in a single point. And for two *generic* points in \mathbf{R}^2 there is a unique line connecting them. We get more tropical lines in \mathbf{R}^2 if we allow the coefficients A, B, C to be $-\infty = 0_{\mathbf{T}}$ as well: conventional horizontal lines, vertical lines and diagonal lines parallel to the vector $(1, 1)$, see Figure 2.21.

Remark 1.5.3

For an affine variety $V \subset \mathbf{C}^n$ we call $\text{Log}_t(V) \subset \mathbf{T}^n$ its affine amoeba. Note




Figure 1.13: Amoeba of a parabola.

that if $V \subset \mathbf{C}^n$ is irreducible and is not contained in a coordinate hyperplane $z_j = 0$, $j = 1, \dots, n$ then V is the topological closure of $V \cap (\mathbf{C}^\times)^n$ in \mathbf{C}^n . Furthermore, the amoeba $\text{Log}_t(V)$ is the topological closure of $\text{Log}_t(V \cap (\mathbf{C}^\times)^n)$ in \mathbf{T}^n . For example, the affine amoeba $\vec{\mathcal{A}}_t \subset \mathbf{T}^2$ of the line $\{z+w+1=0\} \subset \mathbf{C}^2$ can be obtained from \mathcal{A}_t (see (1.5)) by adding two points

$$\vec{\mathcal{A}}_t = \mathcal{A}_t \cup \{(-\infty, 0)\} \cup \{(0, -\infty)\}$$

at the far left and far lower apex of \mathcal{A}_t at Figure 1.11.

As a second example type of amoebas, let us consider the graph $G_t = \{(z, w) \in \mathbf{C}^2 \mid w = f_t(z)\}$ of a polynomial

$$f_t(z) = \sum_{j=0}^d a_j(t)z^j$$

in z whose coefficients $a_j(t)$ are \mathbf{C} -valued functions in the positive real variable t (defined for sufficiently large values of t) such that there exist $\alpha_j \in \mathbf{C}^\times$ and $A_j \in \mathbf{R}$

$$\lim_{t \rightarrow +\infty} \frac{a_j(t) - \alpha_j t^{A_j}}{t^{A_j}} = 0,$$

in other words, $a_j(t) = \alpha_j t^{A_j} + o(t^{A_j})$, $t \rightarrow +\infty$.

We may note that the amoeba $\text{Log}_t(G_t)$ is contained in the graph of the multivalued polynomial

$$F_t = A_0 \vee_t "A_1 x" \vee_t \cdots \vee_t "A_d x^d" \tag{1.7}$$




Figure 1.14: A tropical parabola as the limit of amoebas when $t \rightarrow +\infty$.

if $a_j(t) = \alpha_j t^{A_j}$. The case of $d = 2$, i.e. when G_t is a parabola is depicted on Figure 1.13.

Let us compare F_t against the corresponding tropical polynomial

$$F(x) = \left\langle \sum_{j=0}^d A_j x^j \right\rangle.$$

By Remark 1.4.1 the set $F_t(x) \subset \mathbf{T}$ has a one-point limit $\{F(x)\}$ when $t \rightarrow +\infty$ if x is such that the collection $\{A_d x^d\}$ has a unique maximum. If this maximum is not unique then this limit is $[-\infty, F(x)]$. Again, using multivalued tropical addition this can be reformulated by saying that the limit of $\text{Log}_t(G_t)$ is equal to the graph of $F_\infty = A_0 \vee A_1 x \vee \dots \vee A_d x^d$.

Clearly the same holds if we slightly perturb the coefficients $a_j(t)$ from $\alpha_j t^{A_j}$ to $\alpha_j t^{A_j} + o(t^{A_j})$. The limit of $\text{Log}_t(G_t)$ is shown on Figure 1.14.

1.6 Patchworking and tropical geometry

We saw that if we change the coefficients a, b, c in $az + bw + c = 0$, the amoeba only gets translated. If we choose a, b, c to be real, then the boundary of the amoeba \mathcal{A}_t will be the image of the real locus of the line L . In formulas, $\partial \mathcal{A}_t = \text{Log}_t(\mathbf{RL})$ where $\mathbf{RL} = \{(z, w) \in (\mathbf{R}^*)^2 \mid az + bw + c = 0\} = L \cap (\mathbf{R}^*)^2$.

If we assume $a, b, c \neq 0$, then the three arcs in the boundary of \mathcal{A}_t correspond to the three components of $\mathbf{RL} \cap (\mathbf{R}^*)^2$ which in turn correspond




Figure 1.15: Log_t maps the real line $z + w + 1 = 0$ to the three boundary arcs of the amoeba.

to three out of four quadrants of \mathbf{R}^2 . Which arc corresponds to which quadrant is determined by signs of the coefficients a, b, c . Furthermore, even if a, b, c are functions of a real positive parameter $t \rightarrow +\infty$ of the form $a(t) = \alpha t^{-A} + o(t^{-A})$ and the function is real we may still speak of the sign of $a(t)$ as the sign of the leading coefficient $\alpha \in \mathbf{R}$. Consider $\Gamma = \lim_{t \rightarrow +\infty} \text{Log}_t(L_t)$ where L_t is given by $a(t)z + b(t)w + c(t) = 0$ with the real coefficients $a(t), b(t), c(t)$. We saw already that Γ is a graph in \mathbf{R}^2 with one vertex. Now let us focus on \mathbf{RL} again. First, we split $(\mathbf{R}^*)^2$ into its four quadrants $\mathbf{R}_{>0}^2 \times \{+, -\}^2$ and compute the limits of \mathbf{RL} for each quadrant separately (see Figure 1.16). As described before, these components correspond to the three boundary arcs of \mathcal{A}_t , so the limit is easy to compute: For each quadrant, we just get a part of the tropical line consisting of two rays. In our example, the signs of $a(t), b(t), c(t)$ are all positive, i.e. $\alpha, \beta, \gamma > 0$.

We can summarize these pictures by drawing the whole limit as before, but now labeling the edges of Γ with the signs of the real quadrants whose part of \mathbf{RL} converge to the edge (see Figure 1.17).

Let us now construct \mathbf{RP}^2 (as a topological space) by gluing together four copies of \mathbf{TP}^2 along the sides at infinity as indicated in Figure 1.18. The four boundary sides of the picture are also glued together by identifying antipodal points. The inside of each triangle is homeomorphic to \mathbf{R}^2 .

Now we can just redraw the quadrant pictures from above in this representation of \mathbf{RP}^2 , see Figure 1.19. Equivalently, we could first draw a copy of Γ in each of the quadrants and then throw away those edges which are not labelled with the corresponding sign in the previous picture.




Figure 1.16: The logarithmic limit for each quadrant separately.




Figure 1.17: A tropical line whose edges are labelled with signs referring to the 4 quadrants.

As indicated, we denote the resulting set in \mathbf{RP}^2 by $\mathbf{R}\Gamma$. Note that $\mathbf{R}\Gamma$ is closed in \mathbf{RP}^2 and that topologically it is an embedded circle $S^1 = \mathbf{RP}^1 \subset \mathbf{RP}^2$. It can be considered as the limit of $\mathbf{R}L_t$ under a certain reparameterization (the so-called phase-tropical limit). Furthermore, $\mathbf{R}L_t$ is isotopic to $\mathbf{R}\Gamma$.¹ A similar construction works not only for lines but for smooth algebraic curves in \mathbf{RP}^2 . It was introduced by Viro in 1979 and is now known as Viro patchworking (cf. [Vir79]; for a list of references see [Vir06]). It is the most powerful construction tool currently known in Real Algebraic Geometry. One of the major breakthroughs obtained with the help of this technique was the construction by Itenberg of a counterexample to the so-

¹In our case of a line, this is true for any t (as long as the functions $a(t), b(t), c(t)$ are defined at t and do not vanish simultaneously). In more general cases we get similar isotopies for large values of t (whenever $\mathbf{R}\Gamma$ is smooth as a tropical variety).




Figure 1.18: \mathbf{RP}^2 obtained from gluing four copies of \mathbf{TP}^2 (antipodal points are also identified).




Figure 1.19: A real tropical line.

called Ragsdale conjecture standing open since 1906 (see [Rag06] for the conjecture and [IV96] for the counterexample). Let us review the relevant background of this conjecture.

There are only two homology types of circles embedded in \mathbf{RP}^2 . The line \mathbf{RL} is an example of the non-trivial class (as the complement $\mathbf{RP}^2 \setminus \mathbf{RL}$ is still connected). A circle which bounds a disc in \mathbf{RP}^2 is zero-homologous and is called an *oval*. If we consider a smooth real algebraic curve $\mathbf{RC} \subset \mathbf{RP}^2$, then its connected components are embedded circles. Note that non-trivially embedded circles must intersect by topological reasons and that their intersection points must be singular points of \mathbf{RC} . Thus if the degree of \mathbf{RC} is even then all its components are ovals; if it is odd, then all but one component are ovals.




Figure 1.20: A tropical curve of degree 10.

Each oval separates its complement into two connected components, the *interior* (homeomorphic to a disc) and the *exterior* (homeomorphic to a Möbius band). An oval is called even if it sits in the interior of an even number of other ovals (and odd otherwise). The Ragsdale conjecture stated that the number of even ovals, denoted by p , of a smooth real curve of even degree $2k$ is bounded by

$$p \leq \frac{3k(k-1)}{2} + 1.$$

It was noted by Viro [Vir80] that this inequality comes as a special case of the more general conjecture

$$b_1(\mathbf{R}X) \leq h^{1,1}(X), \quad (1.8)$$




Figure 1.21: The quasitropical limit of the counterexample to the Ragsdale conjecture from [Ite93].

where X is a smooth complex algebraic surface defined over \mathbf{R} and $\mathbf{R}X$ is its real locus (as usual, b_1 stands for the first Betti number and $h^{1,1}$ stands for the $(1, 1)$ -Hodge number). Furthermore, the inequality (1.8) implies the bound

$$n \leq \frac{3k(k-1)}{2} + 1.$$

for the number n of odd ovals that is weaker than the historical Ragsdale conjecture [Rag06] by 1.

In one of the first spectacular applications of the patchworking technique in [Vir80] Viro disproved the original Ragsdale conjecture for the number of odd ovals by that very one, leaving the more general conjecture (1.8) still plausible. Then the final counterexample was given by Itenberg [Ite93] in

yet another striking application of patchworking.

Figure 1.20 shows a very particular tropical curve of degree 10. As in the example of the line, this curve can be obtained as the limit of a family of real algebraic curves, given by a family of equations. After a choice of signs for the (leading terms of) the coefficients of these equations, we can again draw the quasitropical limit in the four quadrants (see Figure 1.21). The result corresponds to a smooth algebraic curve of degree 10 in \mathbf{RP}^2 with $p = 32$ even ovals (don't forget the big one which crosses the line at infinity several times), which exceeds the Ragsdale bound $\frac{3 \cdot 5 \cdot 4}{2} + 1 = 31$.

Tropical geometry can be viewed as a further development and generalization of patchworking. The first volume of this book takes an intrinsic point of view — Tropical Geometry per se. In some applications (e.g. some simple cases in Gromov-Witten theory) it can completely replace Complex Geometry. In the first volume we will look only at the most elementary instances of such applications (rational curves in toric varieties). In the second volume we plan to take a more detailed look at amoebas and at the so-called *phase-tropical varieties* linking tropical and complex geometries.

1.7 Concluding Remarks

2 Tropical hypersurfaces in \mathbf{R}^n

We start our presentation of tropical geometry with tropical hypersurfaces in \mathbf{R}^n . Recall from the introduction that \mathbf{R} can be regarded as the set of tropical non-zero numbers \mathbf{T}^\times . Hence, using the terminology from algebraic geometry, $\mathbf{R}^n = (\mathbf{T}^\times)^n$ is the tropical algebraic torus. Subvarieties of the algebraic torus are usually called very affine varieties in algebraic geometry. In this sense, we are dealing with tropical very affine geometry in this (and the next) chapter.

As many examples in the introduction showed, tropical algebraic geometry can typically be translated into notions from *ordinary* affine geometry (unfortunately in slight conflict with the meaning of “affine” in the previous paragraph) after unwrapping the definition of tropical arithmetics. Table 2.1 shows an incomplete dictionary between tropical algebro-geometric and ordinary affine language and the present chapters focuses on exploring this interplay in the case of hypersurfaces. It is therefore appropriate to start by recalling some basic notions from affine and polyhedral geometry.

Throughout the following, any reference to topology refers to the Euclidean topology of \mathbf{R}^n .

<i>algebraic geometry</i>	<i>affine geometry</i>
algebraic torus $(\mathbf{T}^\times)^n$	vectorspace \mathbf{R}^n
$(\mathbf{T}^\times)^n$ -torsor	affine space A
monomials “ x^j ”	integer linear forms jx
monomial maps $\Phi : (\mathbf{T}^\times)^n \rightarrow (\mathbf{T}^\times)^m$	integer linear maps $A : \mathbf{R}^n \rightarrow \mathbf{R}^m$
polynomials $f = \sum_j a_j x^j$	PL convex functions $f = \max_i \{a_i + jx\}$
hypersurfaces $V(f)$	$n - 1$ -dim’l polyhedral subspaces $X \subset \mathbf{R}^n$

Table 2.1: A dictionary between tropical algebraic and affine geometry of \mathbf{R}^n

2.1 Polyhedral geometry dictionary I

In this first section, we collect the very basic notions from polyhedral geometry that will be used throughout the rest of the text. In this way, we fix our terminology while avoiding unwanted interruptions in the subsequent sections. For more details and proofs of the theory needed here, we refer the reader to the (far more comprehensive) expositions in [GKZ08; Zie95; BG09].

2.1.1 Tropical vector spaces

Let V be a real vector space. A subgroup $N \subset V$ is called *lattice* in V if it is free of rank n and discrete (equivalently, free of rank n and spans V). A vector space V together with lattice N is called a *tropical vector space*. We use this terminology for lack of better alternatives. In particular, note that does not carry a \mathbf{T} -module structure, as one might expect in analogy with *real* and *complex* vector spaces.

If not specified otherwise, we always regard \mathbf{R}^n as tropical vector space with lattice \mathbf{Z}^n . A subspace $W \subset V$ is called a *rational subspace* if $\dim(W) = \text{rank}(W \cap N)$. It follows that $W \cap N$ resp. $N/(W \cap N)$ form lattices for W resp. V/W and turn these spaces into tropical vector spaces. The dual vector space V^* is a tropical vector space with dual lattice $N^* = \text{Hom}(N, \mathbf{Z})$ (where we identify $\lambda \in N^*$ with the function $V \rightarrow \mathbf{R}$ given by linear extension). The functions $\lambda \in N^*$ are called *integer linear forms* or *integer linear functions* on V .

A function $\kappa : V \rightarrow \mathbf{R}$ is called *integer affine* (or a *tropical monomial*) if it is the sum of an integer linear form and a *real* constant. In other words, they have the form

$$\kappa(x) = a + jx$$

for some $j \in N^*, a \in \mathbf{R}$.

A linear map $\Phi : V \rightarrow W$ between two tropical vector spaces $V = N \otimes \mathbf{R}$ and $W = M \otimes \mathbf{R}$ is called an *integer linear map* if it is induced by an linear map $N \rightarrow M$, i.e. if it sends lattice vectors to lattice vectors. A map $\Psi : V \rightarrow W$ is called *integer affine* if it is of the form $\Psi(x) = \Phi(x) + w$ with Φ integer linear and $w \in W$. The map Ψ is called an *integer affine isomorphism* or *tropical isomorphism* if there exists an integer affine inverse map. An




Figure 2.1: An unbounded and a bounded polyhedron

integer affine map $\Psi : \mathbf{R}^n \rightarrow \mathbf{R}^m$ can be written as $x \mapsto Ax + b$, where $A \in \text{Mat}(m \times n, \mathbf{Z})$ is a matrix with integer entries and $b \in \mathbf{R}^m$. The map Ψ is an isomorphism if and only if A is a square matrix with determinant ± 1 , i.e. if $A \in \text{GL}(n, \mathbf{Z})$.

2.1.2 Polyhedra

For each integer affine function κ on V , we define the *rational halfspace* of κ to be

$$H_\kappa := \{x \in V : \kappa(x) \geq 0\} = \{x \in V : jx \geq -a\}.$$

A subset $\sigma \subseteq V$ is called a *rational polyhedron* of X if it is the intersection of finitely many rational halfspaces. Let us emphasize that the adjective “rational” refers to the fact that the bounding inequalities have linear parts in N^* (or, equivalently, $N^* \otimes \mathbf{Q}$). As we will not work with other halfspaces or polyhedra, and also to avoid conflicts with other usage of the attribute “rational”, we mostly drop it in the following. Figure 2.1 depicts an unbounded and a compact polyhedron.

The *Minkowski sum* of two polyhedra

$$\sigma + \sigma' = \{x + y : x \in \sigma, y \in \sigma'\}$$

is a polyhedron again. A *face* τ of a polyhedron σ is given as $\sigma \cap H_{-\kappa}$, where κ is an integer affine-linear function such that $P \subset H_\kappa$. If $\tau = \sigma \cap H_{-\kappa} \neq \emptyset$, H_κ is called a *supporting halfspace* of σ . The *(relative) boundary* $\partial\sigma$ of σ is the union of all proper faces. The complement $\sigma^\circ := \sigma \setminus \partial\sigma$ is called the *(relative) interior* of σ . We denote by $L(\sigma)$ the real subspace of V spanned

by σ and called the *linear span* of σ . More precisely, $L(\sigma)$ is spanned by all differences $x - y$, $x, y \in \sigma$. We define the dimension of σ by $\dim(\sigma) := \dim(L(\sigma))$. We set $L_{\mathbf{Z}}(\sigma) = L(\sigma) \cap \mathbf{Z}^n$. If σ is rational, then $L_{\mathbf{Z}}(\sigma)$ is a lattice in $L(\sigma)$.

Let $\sigma \subset V$ be a polyhedron. The *recession cone* $\text{rc}(\sigma)$ of σ is the set of direction vectors of all rays contained in σ ,

$$\begin{aligned} \text{rc}(\sigma) &:= \{v \in V : \exists x \in \sigma \text{ such that } x + \mu v \in \sigma \forall \mu \geq 0\} \\ &= \{v \in V : x + \mu v \in \sigma \forall x \in \sigma, \mu \geq 0\}. \end{aligned}$$

Equivalently, if σ is given as the intersection of the affine halfspaces $j_i x \geq -a_i$, $i = 1, \dots, k$, then $\text{rc}(\sigma)$ is the cone obtained as the intersection of the linear halfspaces $j_i x \geq 0$, $i = 1, \dots, k$ (see [BG09, Proposition 1.23]).

Any polyhedral cone $\sigma \subset V$ is equal to

$$\sigma = \mathbf{R}_{\geq 0} v_1 + \dots + \mathbf{R}_{\geq 0} v_k$$

for suitable vectors $v_1, \dots, v_k \in V$, called *generators* of σ . Any bounded polyhedron $\sigma \subset V$ is the convex hull $\text{Conv}(A)$ of a finite set of points $A \subset V$, $|A| < \infty$. Finally, an arbitrary polyhedron $\sigma \subset V$ can be written as a Minkowski sum $\sigma = \tau + \rho$, where τ is bounded and ρ is a cone. Here, τ is not unique, but $\rho = \text{rc}(\sigma)$. See [BG09, Sections 1.B and 1.C].

2.1.3 Polyhedral complexes and fans

A collection $\mathcal{P} = \{\sigma_1, \dots, \sigma_m\}$ of polyhedra is called a *polyhedral complex* if for each σ_i all faces are also contained in \mathcal{P} and if each intersection $\sigma_i \cap \sigma_j$ produces a face of both σ_i and σ_j (if nonempty). The elements σ_i of \mathcal{P} are called the *cells* of \mathcal{P} . The *support* of \mathcal{P} is $|\mathcal{P}| := \bigcup_i \sigma_i$. If $|\mathcal{P}|$ is equal to a polyhedron P (e.g. \mathbf{R}^n) we call \mathcal{P} a *polyhedral subdivision* of σ (or \mathbf{R}^n , respectively). Given two polyhedral complexes \mathcal{X} and \mathcal{Y} , we call \mathcal{X} a *refinement* of \mathcal{Y} if each cell of \mathcal{X} is contained in a cell of \mathcal{Y} . In this case, the cells of \mathcal{X} contained in a fixed cell $\sigma \in \mathcal{Y}$ form a polyhedral subdivision of σ .

We say that \mathcal{P} is of *pure dimension* n if all maximal polyhedra in \mathcal{P} have dimension n . We will mostly deal with pure-dimensional complexes in the following. The cells of dimension n , 1 and 0 are called *facets*, *edges* and

vertices, respectively. An edge containing a single vertex is called a *ray*. The *k-skeleton* of a polyhedral complex is the set of polyhedra

$$\mathcal{P}^{(k)} := \{P \in \mathcal{P} : \dim(P) \leq k\}.$$

It forms a polyhedral complex again.

A polyhedron σ is called a *cone* if for each $x \in \sigma$ the whole ray $\mathbf{R}_{\geq 0}x$ is contained in σ . Equivalently, it is the finite intersection of rational half-spaces given by integer linear (not affine) functions. A polyhedral complex of cones is called a *polyhedral fan*. A fan is called *pointed* if it contains $\{0\}$. In this case, all of its cones are *pointed*, which means have $\{0\}$ as a face. Any polyhedron σ in V gives rise to a fan in dual space V^* which is called its *normal fan* and constructed as follows. For each face τ of σ , let C_τ be the cone in V^* consisting of those linear forms which are bounded on σ from below and whose minimum on σ is attained on τ ,

$$C_\tau := \{\lambda \in V^\vee : \lambda(x) \geq \lambda(y) \text{ for all } x \in \sigma, y \in \tau\}.$$

The collection of cones C_τ for all faces τ of σ forms a polyhedral fan which is the normal fan of σ , denoted by $\mathcal{N}(\sigma)$ (see [BG09, Proposition 1.67]).

Let \mathcal{X} be a polyhedral complex in V and fix a cell $\tau \in \mathcal{X}$. For any cell $\sigma \in \mathcal{X}$ containing τ , we define σ/τ to be the polyhedral cone in $V/L(\tau)$ generated by all differences $x - y, x \in \sigma, y \in \tau$. The collection of σ/τ for all $\sigma \supseteq \tau$ forms a pointed fan in $V/L(\tau)$ denoted by $\text{Star}_{\mathcal{X}}(\tau)$ and called the *star of \mathcal{X} at τ* .

Exercise 2.1.1

Let σ and σ' be polyhedra such that $\sigma \cap \sigma' \neq \emptyset$. Show that $\text{rc}(\sigma \cap \sigma') = \text{rc}(\sigma) \cap \text{rc}(\sigma')$.

Exercise 2.1.2

Prove that $\text{Star}_{\mathcal{X}}(\tau) = \{\sigma/\tau : \tau \subset \sigma\}$ is a pointed fan.

Exercise 2.1.3

Let \mathcal{X} and \mathcal{Y} be two polyhedral complexes in \mathbf{R}^n . Show that

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

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

Exercise 2.1.4

Let \mathcal{X} and \mathcal{Y} be two polyhedral complexes in \mathbf{R}^n and \mathbf{R}^m , respectively. Show that

$$\mathcal{X} \times \mathcal{Y} := \{\sigma \times \sigma' : \sigma \in \mathcal{X}, \sigma' \in \mathcal{Y}\}$$

is a polyhedral complex in \mathbf{R}^{n+m} and $|\mathcal{X} \times \mathcal{Y}| = |\mathcal{X}| \times |\mathcal{Y}|$.

2.2 Tropical Laurent polynomials and hypersurfaces

The natural class of functions on the tropical algebraic torus $\mathbf{R}^n = (\mathbf{T}^\times)^n$ are tropical Laurent polynomials (in contrast to tropical polynomials, which are the those Laurent polynomials that extend to \mathbf{T}^n). In this section, we develop the basic properties of these functions.

Recall that a *tropical monomial* κ on \mathbf{R}^n is a function of the form

$$\kappa(x_1, \dots, x_n) = "ax_1^{j_1} \dots x_n^{j_n}",$$

where $a \in \mathbf{R}$, $j_1, \dots, j_n \in \mathbf{Z}$ and x_1, \dots, x_n are the coordinates of \mathbf{R}^n . In standard arithmetic operations this can be expressed as

$$\kappa(x_1, \dots, x_n) = a + j_1 x_1 + \dots + j_n x_n.$$

As before we will use multi-index notation $j = (j_1, \dots, j_n) \in \mathbf{Z}^n$ resp. $x = (x_1, \dots, x_n) \in \mathbf{R}^n$ and write

$$\kappa(x) = "ax^j" = a + jx,$$

where jx is the standard scalar product.

A *tropical Laurent polynomial* $f : \mathbf{R}^n \rightarrow \mathbf{T}$ is a finite sum of tropical monomials, i.e., a function of the form

$$f(x) = " \sum_{j \in A} a_j x^j " = \max_{j \in A} \{a_j + jx\},$$

with $A \subseteq \mathbf{Z}^n$ finite and $a_j \in \mathbf{R}$ for all $j \in A$. By definition, a Laurent polynomial is a convex, piecewise integer affine function. We will presently give an example (see example 2.2.2) showing that the coefficients a_j are in

general not uniquely determined by the function f . In this book, we will adopt the convention that a Laurent polynomial is a function whereas an assignment of coefficients $\mathbf{Z}^n \rightarrow \mathbf{T}, j \mapsto a_j$ (with $a_j = -\infty$ for almost all j) is called a *representation* of f . We will say a few more words about the ambiguity later in this chapter.

Example 2.2.1

If a tropical Laurent polynomial f can be represented by only one term, then it is just an integer affine function, as discussed before. If f consists of two terms, say $f = "a_i x^i + a_j x^j"$, then it divides \mathbf{R}^n into two halfspaces along the hyperplane $a_i + ix = a_j + jx$. In one half, f equals $a_i + ix$, in the other f is $a_j + jx$. In particular, f is locally integer affine except for points in the hyperplane, where f is non-differentiable and strictly convex.

Example 2.2.2

We take three polynomials $f_1 = "0 + 1x + x^2"$, $f_2 = "0 + x + x^2"$ and $f_3 = "0 + (-1)x + x^2"$ in one variable x . First, let us point out some possibly confusing facts. Note that neither the constant term 0 nor the coefficient 1 of $1x$ can be omitted here, because $0 \neq 0_{\mathbf{T}} = -\infty$ and $1 \neq 1_{\mathbf{T}} = 0$. In fact, $x^2 = "0x^2"$, but $"1x" \neq x = "0x"$. After this piece of warning, let us just draw the graphs of the three functions (see Figure 2.2). The important




Figure 2.2: Three tropical polynomials

observation is that f_2 and f_3 are actually the same functions coming from different representations as a sum of monomials. For f_3 , the linear term is never maximal. For f_2 , the linear term is maximal at 0, but only together with the other two terms. In fact, the "shortest" representation of both functions is given by $f_2 = f_3 = "0 + x^2"$. Apart from that, note that again the functions are non-differentiable and strictly convex only at a finite number

of breaking points. Away from these points, only one monomial attains the maximum and hence the function is locally affine linear.

Example 2.2.3

The expression

$$f(x, y) = “(-1)x^2 + (-1)y^2 + 1xy + x + y + 0”$$

defines a tropical polynomial of degree 2 in two variables. Its graph is displayed in Figure 2.3.




Figure 2.3: The graph of $\max\{(-1) + 2x, (-1) + 2y, 1 + x + y, x, y, 0\}$

Classical algebraic geometry is the study of zero-sets of classical polynomials. Consequently, the object of study in this chapter should be zero-sets of tropical polynomials. However, note that the naive definition $\{x \in \mathbf{R}^n : f(x) = 0_{\mathbf{T}} = \infty\}$ produces the empty set whenever $f \not\equiv -\infty$ (i.e., $a_j \neq \infty$ for at least one $j \in A$) and therefore we obviously need some alternative definition. We will give this definition now and afterwards provide some first, intrinsically tropical evidence that our definition is indeed the “right” one. The extrinsic, but probably more convincing motivation for our notion of tropical hypersurfaces is the fact that they describe limits of amoebas of hypersurfaces as exemplified to in the introduction. We will not pursue this line of thought here.

Definition 2.2.4

Let f be a tropical Laurent polynomial in n variables. Then we define the

hypersurface $V(f) \subseteq \mathbf{R}^n$ to be the set of points in \mathbf{R}^n where f is not differentiable.

In classical arithmetics, the zero element of a group is distinguished by the property of being idempotent, i.e. $0 + 0 = 0$. In tropical arithmetics, this is true for any number as “ $x + x$ ” = $\max\{x, x\} = x$. However, we might consider this as a hint that a tropical sum should be called “zero” if the maximum is attained by at least two terms. This leads to an alternative definition of the tropical hypersurface of a tropical polynomial. The following proposition shows that both definitions coincide.

Proposition 2.2.5

Let $f = \sum_{j \in A} a_j x^j$ be a representation of tropical polynomial. Then the hypersurface $V(f)$ is equal to the set of points $x \in \mathbf{R}^n$ where the maximum $f(x)$ is attained by at least two monomials, i.e.

$$V(f) = \{x \in \mathbf{R}^n : \exists i \neq j \in A \text{ such that } f(x) = a_i x^i = a_j x^j\}.$$

Proof. If the maximum is attained by only one monomial, then f is locally affine and thus differentiable. If two monomials attain the maximum, then f is strictly convex at x and therefore cannot be differentiable. \square

Example 2.2.6

For the three (in fact, two) polynomials from example 2.2.2, the hypersurfaces are just finite sets of points (see Figure 2.4).




Figure 2.4: The “zeros” of tropical polynomials

The polynomial f from example 2.2.3 produces a connected hypersurface $V(f)$ consisting of 4 vertices, 3 bounded edges and 6 unbounded rays (see Figure 2.5).




Figure 2.5: The hypersurface of “ $(-1)x^2 + (-1)y^2 + 1xy + x + y + 0$ ”

Remark 2.2.7

As mentioned in the introduction, there is an alternative definition of the tropical sum of two numbers given by *multivalued addition*. This was suggested by Viro (cf. [Vir10]) and is given, for each two numbers $a, b \in \mathbf{T} = \mathbf{R} \cup \{-\infty\}$, by

$$a \curlyvee b = \begin{cases} \max(a, b) & \text{if } a \neq b, \\ [-\infty, a] & \text{if } a = b. \end{cases} \quad (2.1)$$

Note that when $a = b$, the sum $a \curlyvee b$ is not just a single element but the set $[-\infty, a] := \{-\infty\} \cup \{x \in \mathbf{R} : x \leq a\}$. The set \mathbf{T} equipped with the multivalued addition \curlyvee and the uni-valued multiplication “ \cdot ” = + forms a structure which is called a *hyperfield* (in our case, the *tropical hyperfield*). One reason to consider such generalizations of ordinary algebra is that they allow to describe “zero-sets” in a natural way. Of course, when considering a multi-valued sum, “being zero” should be replaced by “containing zero”. Note that a sum $a_1 \curlyvee \cdots \curlyvee a_n$ in the tropical hyperfield contains $-\infty$ if and only if we find $i \neq j$ such that $a_i = a_j \geq a_k, k = 1, \dots, n$, i.e., if the maximum of the summands occurs at least twice. Hence, if we replace a tropical polynomial

$$f(x) = \max(a_{j_1} + j_1x, \dots, a_{j_n} + j_nx)$$

by the hyperfield version

$$f^\curlyvee(x) = (a_{j_1} + j_1x) \curlyvee \cdots \curlyvee (a_{j_n} + j_nx),$$

we find

$$V(f) = \{x \in \mathbf{R}^n : -\infty \in f^\curlyvee(x)\} =: V(f^\curlyvee).$$

In summary, sacrificing uni-valued addition leads to a more natural definition of tropical zero-sets.

Exercise 2.2.8

Let $f : \mathbf{R} \rightarrow \mathbf{R}$ be a univariate tropical polynomial of degree d and with non-vanishing constant term. Show that \mathbf{T} is *algebraically closed* in the following sense. There exist $c \in \mathbf{R}$ and $\alpha_1, \dots, \alpha_n \in \mathbf{R}$ (unique up to reordering) such that

$$f(x) = "c \prod_{i=1}^n (x + \alpha_i)"$$

for all $x \in \mathbf{R}$ (as functions, not as polynomial representations). Show that $V(f) = \{\alpha_1, \dots, \alpha_n\}$.

Exercise 2.2.9

Let $f, g : \mathbf{R}^n \rightarrow \mathbf{R}$ be tropical Laurent polynomials. Show that $V("fg") = V(f) \cup V(g)$.

Exercise 2.2.10

Given a set H , we denote by $\text{Pot}(H)$ the set of subsets of R . A *hyperfield* is a set H together with binary operations $\gamma : H \times H \rightarrow \text{Pot}(H)$ and $\cdot : H \times H \rightarrow H$ and two elements $0_H, 1_H \in H$ such that for all $a, b, c \in H$

$$\begin{aligned} a \gamma (b \gamma c) &= (a \gamma b) \gamma c, & 0_H \gamma a &= \{a\}, \\ \exists! -a \in H \text{ s.t. } 0_H &\in a \gamma (-a), & a \gamma b &= b \gamma a, \\ a \in b \gamma c &\Rightarrow b \in a \gamma -c, & 0_H \cdot a &= 0_H, \\ (H \setminus \{0_H\}, \cdot) &\text{ is a commutative group,} & a \cdot (b \gamma c) &= (a \gamma b) \cdot (a \gamma c). \end{aligned}$$

Note that some of the axiom (e.g. associativity) require to consider the canonical extension of $\gamma : H \times H \rightarrow \text{Pot}(H)$ to $\text{Pot}(H) \times \text{Pot}(H) \rightarrow \text{Pot}(H)$.

- (a) With γ defined as in Equation 2.1, show that $(\mathbf{T}, \gamma, +, -\infty, 0)$ forms a hyperfield.
- (b) With γ_t defined as in Equation 1.3, show that $(\mathbf{T}, \gamma_t, +, -\infty, 0)$ forms a hyperfield for all $t > 0$.
- (c) Show that $\lim_{t \rightarrow \infty} \gamma_t = \gamma$ (cf. Remark 1.4.1).

2.3 The polyhedral structure of hypersurfaces

Of course, tropical hypersurfaces are not just sets but carry much more structure. In this section, we describe their structure as polyhedral complexes. For each monomial parameterized by $j \in A$, set

$$\sigma_j := \{x \in \mathbf{R}^n : f(x) = "a_j x^j"\}$$

to be the locus of points where the chosen monomial is maximal. The sets σ_j subdivide \mathbf{R}^n into the domains of linearity of f . If σ_j is a neighbourhood of a point x , then f is obviously differentiable at x , with differential $df_x = j$ (regarding j as a covector). Note that σ_j is a rational polyhedron in \mathbf{R}^n as it is the intersection of the halfspaces $a_j + jx \geq a_i + ix$ for all $i \in A$. Moreover, the intersection of two polyhedra σ_i and σ_j is either empty or a common face (given by intersecting with the plane $a_i + ix = a_j + jx$). It follows that the collection of polyhedra

$$\mathcal{S}(f) := \{\text{faces } \tau \text{ of } \sigma_j \text{ for some } j \in A\}$$

forms a polyhedral subdivision of \mathbf{R}^n . While the definition of σ_j may depend on the representation of f (see Example 2.3.1), we prove in Lemma 2.3.4 that $\mathcal{S}(f)$ only depends on the function f .

Example 2.3.1

Our polynomials from the previous examples give subdivisions as indicated in Figure 2.6. Note that in the second and third case the subdivisions $\mathcal{S}(f_2)$ and $\mathcal{S}(f_3)$ are identical, but the definition of σ_1 depends on the chosen representation. For the polynomial in two variables, the subdivision $\mathcal{S}(f)$ is shown in Figure 2.7.




Figure 2.6: The subdivision induced by polynomials




Figure 2.7: Another subdivision

Obviously, a point $x \in \mathbf{R}^n$ lies in $n - 1$ -cell of $\mathcal{S}(f)$ if and only if at least two monomials attain the maximum $f(x)$ at x . We get the following corollary.

Corollary 2.3.2

The hypersurface $V(f)$ is equal to the $n - 1$ -skeleton of $\mathcal{S}(f)$. In particular, $V(f)$ is canonically equipped with the structure of a rational polyhedral complex of pure dimension $n - 1$.

It remains to show that $\mathcal{S}(f)$ does not depend on the representation of f . On the way, we define the reduced representation of f given by throwing away all monomials “ $a_j x^j$ ” for which σ_j is empty or not full-dimensional.

Definition 2.3.3

Let $f = “\sum_{j \in A} a_j x^j”$ be a tropical Laurent polynomial. The set

$$A^{\text{red}} := \{j \in A : \dim(\sigma_j) = n\} = \{df_x : x \in \mathbf{R}^n \setminus V(f)\}$$

is called the *reduced support* of f . The truncated polynomial

$$f^{\text{red}} = “\sum_{j \in A^{\text{red}}} a_j x^j”$$

is called the *reduced representation* of f .

Lemma 2.3.4

Let f be a tropical Laurent polynomial. On the level of functions on \mathbf{R}^n , we have $f \equiv f^{\text{red}}$. Moreover, two tropical Laurent polynomials describe the same

function if and only if their reduced representations agree (as abstract polynomials). In particular, the subdivision $\mathcal{S}(f)$ is independent of the representation of f .

Proof. As $\mathcal{S}(f)$ forms a subdivision of \mathbf{R}^n , every point $x \in \mathbf{R}^n$ is contained in some σ_j with $\dim(\sigma_j) = n$, so

$$f(x) = "a_j x^j" = f^{\text{red}}(x).$$

Moreover, if two polynomials describe the same function, they have the same reduced support (by the second description of A^{red}). The value $f(x)$ for any $x \in \mathbf{R}^n$ with $df_x = i$ determines the coefficient of x^i uniquely. Thus the second claim follows. Finally, since obviously $\mathcal{S}(f) = \mathcal{S}(f^{\text{red}})$ (by the first description of A^{red}), it follows that $\mathcal{S}(f)$ is completely determined by the underlying function. \square

We have seen that each tropical Laurent polynomial subdivides \mathbf{R}^n into its domains of linearity $\mathcal{S}(f)$. Next, we want to describe the dual subdivision to $\mathcal{S}(f)$ which is a certain subdivision of the Newton polytope of f . The *Newton polytope* $\text{NP}(f)$ of $f = \sum_j a_j x^j$ is given by

$$\text{NP}(f) = \text{Conv}\{j \in \mathbf{Z}^n : a_j \neq -\infty\},$$

i.e. the convex hull of all appearing exponents. Note that $\text{NP}(f)$ naturally lives in the dual space of \mathbf{R}^n .

Definition 2.3.5

Let $f = \sum_{j \in A} a_j x^j$ be a tropical Laurent polynomial. We set

$$\tilde{A} := \{(j, -a_j) \in \mathbf{Z}^n \times \mathbf{R} : j \in A \text{ and } a_j \neq -\infty\}$$

and $\tilde{P} = \text{Conv}(\tilde{A})$. The projection of the lower faces of \tilde{P} (i.e. those which are also faces of $\tilde{P} + \rho$ with half-ray $\rho := \{0\} \times \mathbf{R}_{\geq 0}$) produces a subdivision $\text{SD}(f)$ of $\text{NP}(f)$, which we call the *dual subdivision* of f .

Example 2.3.6

Our polynomials in one variable from example 2.2.2 all have the same Newton polytope $\text{Conv}\{0, 2\}$. Only for the first polynomial, $\text{SD}(f_1)$ is non-trivial,




Figure 2.8: The dual subdivision

i.e. the segment is divided into two unit segments with vertex $\{1\}$ (see Figure 2.8).

For the polynomial f of example 2.2.3, the dual subdivision is again more interesting. In Figure 2.9, only the lower faces of \tilde{P} are drawn solidly. We




Figure 2.9: The dual subdivision of “ $(-1)x^2 + (-1)y^2 + 1xy + x + y + 0$ ”

obtain a subdivision of the triangle of size 2 into 4 triangles of size 1. Moreover, the subdivision consists of 3 internal edges and 6 edges in the boundary of the big triangle.

In our examples, we can easily observe an inclusion-reversing duality between the cells of $\mathcal{S}(f)$ and $SD(f)$. Let us formulate this systematically. For each cell $\sigma \in \mathcal{S}(f)$ we define

$$A_\sigma := \{j \in A : \sigma \subseteq \sigma_j\}.$$

In other words, A_σ is the set of monomials which are maximal on σ . We find

$$\sigma = \bigcap_{j \in A_\sigma} \sigma_j. \tag{2.2}$$

Analogously, for a point $x \in \mathbf{R}^n$ we set $A_x := \{j \in A : x \in \sigma_j\}$. We have $x \in \sigma^\square$ if and only if $A_x = A_\sigma$.

Theorem 2.3.7

The subdivisions $\mathcal{S}(f)$ of \mathbf{R}^n and $\text{SD}(f)$ of $\text{NP}(f)$ are dual in the following sense. There is an inclusion-reversing bijection of cells given by

$$\begin{aligned} \mathcal{S}(f) &\rightarrow \text{SD}(f) \\ \sigma &\mapsto D_\sigma := \text{Conv}(A_\sigma) \end{aligned} \tag{2.3}$$

such that $\dim \sigma + \dim D_\sigma = n$ and $L(\sigma)^\perp = L(D_\sigma)$. Moreover, $\text{SD}(f)$ is independent of the representation of f and the set of vertices of $\text{SD}(f)$ is equal to the reduced support of f .

Proof. First let us show that D_σ is indeed a cell of $\text{SD}(f)$. To do so, we pick a point x in the relative interior of σ . Let us consider the linear form $(x, -1)$ on $\tilde{P} \subset \mathbf{R}^{n+1}$. For each vertex $(j, -a_j)$ of \tilde{P} we find $(x, -1)(j, -a_j) = a_j + jx$. Thus the face of \tilde{P} on which $(x, -1)$ is maximal is exactly the convex hull of the points \tilde{A}_σ with

$$\tilde{A}_\sigma := \{(j, -a_j) : j \in A_\sigma\}.$$

The projection of this face to \mathbf{R}^n is D_σ , and therefore D_σ is indeed a face of $\text{SD}(f)$. Next we see that all lower faces of \tilde{P} are obtained in this way for a suitable $(x, -1)$. It follows that equation (2.3) indeed describes a well-defined bijection whose inverse map is given by

$$D \mapsto \bigcap_{j \in D \cap A} \sigma_j.$$

The previous arguments also provide another way to construct $\mathcal{S}(f)$. Let $\mathcal{N}(\tilde{P})$ be the normal fan of \tilde{P} , then $\mathcal{S}(f)$ is the subdivision obtained from intersecting $\mathcal{N}(\tilde{P})$ with the plane $\mathbf{R}^n \times \{-1\} \cong \mathbf{R}^n$ (see Figure 2.10). Using this description the orthogonality and dimension statements follow directly from the corresponding statements for the dual cells of a polyhedron and its normal fan.

Finally, the equality $(x, -1)(j, -a_j) = a_j + jx$ also implies that the vertices of \tilde{P} detected by linear forms of type $(x, -1)$ (which correspond to the vertices of $\text{SD}(f)$) are in bijection to A^{red} . Hence the lower faces of \tilde{P} and hence $\text{SD}(f)$ only depend on the function f and the last claims follow. \square



Figure 2.10: Dual subdivisions via normal fans

Example 2.3.8

Let $A \in \mathbf{Z}^n$ be a finite set and let $f = \sum_{j \in A} x^j$ be the Laurent polynomial with only trivial coefficients $a_j = 0, j \in A$. Then $SD(f)$ is just the trivial subdivision of $NP(f) = \text{Conv}(A)$ (i.e. $SD(f)$ contains $NP(f)$ and all its faces) and $\mathcal{S}(f)$ is just the normal fan of $NP(f)$.

Remark 2.3.9

Let \mathcal{S} be a subdivision of a polyhedron $P \subseteq \mathbf{R}^n$. \mathcal{S} is called a *regular subdivision* if it can be obtained by projecting down the lower faces of some polyhedron $\tilde{P} \subseteq \mathbf{R}^{n+1}$. Equivalently, \mathcal{S} is regular if there exists a convex function on P which is affine on each cell of \mathcal{S} (to get the polyhedron in \mathbf{R}^{n+1} , we take the convex hull of the graph of the function; the other way around, the union of lower faces of the polyhedron describes the graph of a suitable function). Note that both $\mathcal{S}(f)$ and $SD(f)$ are regular subdivisions. Such subdivisions are sometimes also called *convex* or *coherent*. They play a prominent role in the context of the Viro’s patchworking technique (see [Vir06] for references) as well as the study of discriminants and resultants (see [GKZ08]).

Note that not all subdivisions are regular. An example is given in Figure 2.11. Assume there is convex function g inducing this subdivision. We can assume that g is constant zero on the inner square. Then if we fix one $g(v_1)$ on an outer vertex, on the next vertex v_2 in clockwise direction, we need $g(v_2) \geq g(v_1)$, due to the diagonal edge subdividing the corresponding trapezoid. Going around the square completely gives a contradiction.

Another way to describe the duality of $\mathcal{S}(f)$ and $SD(f)$ can be formulated in terms of Legendre transforms. Let $g : S \rightarrow \mathbf{R}$ be a function on an arbitrary set $S \subseteq \mathbf{R}^n$. The *Legendre transform* g^* is a function on covectors

of \mathbf{R}^n given by

$$g^*(w) = \sup_{x \in S} \{wx - g(x)\}.$$

It is easy to check that g^* takes finite values on a convex set of \mathbf{R}^n (possibly empty) and that g^* is convex. Moreover, if g is a convex function in the beginning, then $(g^*)^* = g$. By definition, a tropical Laurent polynomial $f = \sum_{j \in A} a_j x^j$ is equal to the Legendre transform of its inverted coefficient map $A \rightarrow \mathbf{R}, j \mapsto -a_j$.

Lemma 2.3.10

Let f be a Laurent polynomial and let g be the convex function on $\text{NP}(f)$ whose graph is the lower hull of

$$\tilde{P} = \text{Conv}(\{(j, -a_j) : j \in A\}).$$

Then $f^* = g$ and equivalently $f = g^*$.

Proof. We set $g' : A \rightarrow \mathbf{R}, j \mapsto -a_j$. As $f = g'^*$, it suffices to show $g'^* = g^*$. For a fixed $x \in \mathbf{R}^n$, we can compute $g'^*(x)$ resp. $g^*(x)$ as the maximum value of $(x, -1)$ on $\{(j, -a_j) : j \in A\}$ resp. \tilde{P} . This maximum is always attained on at least one vertex of \tilde{P} and the vertices of \tilde{P} are contained in $\{(j, -a_j) : j \in A\}$. Hence the claim follows. \square

Exercise 2.3.11

Let f be a tropical Laurent polynomial. Show that the Newton polytope of f can be described as

$$\text{NP}(f) = \{w \in \mathbf{R}^n : wx - f(x) \text{ is bounded from above}\}.$$




Figure 2.11: A non-regular subdivision

Exercise 2.3.12

Let $f, g : \mathbf{R}^n \rightarrow \mathbf{R}$ be tropical Laurent polynomials. Show that $\text{NP}("fg") = \text{NP}(f) + \text{NP}(g)$.

Exercise 2.3.13

Let f be a tropical Laurent polynomial. In this exercise, we want to prove $\mathcal{N}(\text{NP}(f)) = \{\text{rc}(\sigma) : \sigma \in \mathcal{S}(f)\}$. Recall that $\text{rc}(\sigma)$ denotes the recession cone of σ .

- (a) Let $\sigma \subset \mathbf{R}^n \times \mathbf{R}$ be a (closed) cone. For $k = 0, 1$ set $\sigma_k := \sigma \cap (\mathbf{R}^n \cap \{-k\}) \subset \mathbf{R}^n$. Show that $\text{rc}(\sigma_1) = \sigma_0$.
- (b) For $D \in \text{SD}(f)$, let $F \subset \text{NP}(f)$ be the unique face of $\text{NP}(f)$ such that $D^\square \subset F^\square$. Let σ_D and σ_F be the dual cell/cone of D and F in $\mathcal{S}(f)$ and $\mathcal{N}(\text{NP}(f))$, respectively. Show that $\text{rc}(\sigma_D) = \text{rc}(\sigma_F)$.
- (c) Conclude that $\mathcal{N}(\text{NP}(f)) = \{\text{rc}(\sigma) : \sigma \in \mathcal{S}(f)\}$.

2.4 The balancing condition

So far, we have seen that to each tropical Laurent polynomial f we can associate a subdivision $\mathcal{S}(f)$ of \mathbf{R}^n and a dual subdivision $\text{SD}(f)$ of $\text{NP}(f)$ which is induced by the Legendre transform of f . We described the hypersurface $V(f)$ as a set — the points where f is not differentiable — and as polyhedral complex of pure dimension $n - 1$ — the $n - 1$ -skeleton of $\mathcal{S}(f)$. We now add yet another layer to our description, namely weights of points. We will later say that smooth points of $V(f)$ are those with weight 1. Moreover, these weights are necessary to formulate the most fundamental structure property of tropical objects, the balancing condition.

For each polytope P , we define the *volume* $\text{Vol}(P)$ to be the volume of σ measured in the affine space spanned by P and normalized such that $\text{Vol}(S_n) = 1$, where S_n denotes the standard simplex in \mathbf{R}^n given by $x_1 + \dots + x_n \leq 1$ and $x_i \geq 0, i = 1, \dots, n$. In other words, the volume of a full-dimensional P is the usual volume of P in \mathbf{R}^n divided by $n!$. Let P be a lattice polytope, i.e. all vertices are integer points. We call P a *minimal simplex* if $\text{Vol}(P) = 1$. This is the case if and only if σ can be mapped to the standard simplex S_n by an integer affine transformation. For any lattice

polytope P we have $\text{Vol}(P) \in \mathbf{N}$. This follows from the fact that we can always triangulate P into lattice simplices (not necessarily minimal ones). The volume of a simplex with vertices v_0, \dots, v_n is given by the determinant of the vectors $v_1 - v_0, \dots, v_n - v_0$, which is obviously integer for lattice simplices. For one-dimensional lattice polytopes E (i.e., edges) we find $\text{Vol}(E) = \#\{E \cap \mathbf{Z}^n\} - 1$. For two-dimensional lattice polytopes P , Pick's theorem states $\text{Vol}(P) = 2i + b - 1$, where $i := \#\{\sigma^\square \cap \mathbf{Z}^n\}$ is the number of interior lattice points and $b := \#\{\partial \sigma \cap \mathbf{Z}^n\}$ is the number of lattice points in the boundary.

Definition 2.4.1

Let f be a tropical Laurent polynomial. We turn its hypersurface $V(f)$ into a weighted polyhedral complex as follows. The *weight* of each cell σ of $V(f)$ is the volume of the corresponding cell D_σ in the dual subdivision $\text{SD}(f)$,

$$\omega(\sigma) := \text{Vol}(D_\sigma) \in \mathbf{N}.$$

For each point $x \in V(f)$ in the relative interior of σ we define the *weight* $\omega(x) \in \mathbf{N}$ (or *multiplicity* of x) to be $\omega(\sigma)$.

A point resp. cell of $V(f)$ is called *smooth* or *singular* point/cell of $V(f)$ if its weight is equal to 1 or greater than 1, respectively. If $V(f)$ does not have any singular points, it is called a *smooth hypersurface*. This is equivalent to the condition that all cells of $\text{SD}(f)$ are minimal simplices, in which case we call $\text{SD}(f)$ a *unimodular subdivision*. Of course, it suffices to check this condition for the maximal cells of $\text{SD}(f)$, as faces of minimal simplices are minimal simplices again.

Remark 2.4.2

Let σ be a facet of $V(f)$. The dual cell D_σ is an edge whose endpoints, say i and j , are the exponents of the two monomials in the reduced representation of f which attain the maximum at F . Thus the weight $\omega(F)$ can be considered as a measure of the change of slope when crossing from the linearity domain P_i to P_j through F . This is precisely what happens for polynomials in one variable. In general, one should think of the change of slope relative to F .

Example 2.4.3

Let us revisit our running examples. The polynomials from example 2.2.2




Figure 2.12: Multiplicities of tropical zeros




Figure 2.13: $V((-1)x^2 + (-1)y^2 + 1xy + x + y + 0)$ is smooth.

describe hypersurfaces of either two points with multiplicity 1 or one single point with multiplicity 2. Correspondingly, the dual subdivision divides the interval $[0, 2]$ in either two segments of volume 1 or one segment of volume 2 (see Figure 2.12). In the case of the planar conic from example 2.2.3, the dual subdivision only consists of four minimal triangles of volume 1 and therefore the conic is smooth (see Figure 2.13).

Remark 2.4.4

For a smooth hypersurface $V(f)$, the reduced support of f is $\text{NP}(f) \cap \mathbf{Z}^n$. This follows from the fact that a minimal simplex does not contain integer points other than its vertices and thus each point in $\text{NP}(f) \cap \mathbf{Z}^n$ must be a




Figure 2.14: A tropical 2-dimensional hyperplane

vertex of $\text{SD}(f)$.

Example 2.4.5

The *tropical hyperplane* in \mathbf{R}^n is the hypersurface $V(f)$ with $f = “x_1 + \dots + x_n + 0”$. It is obviously smooth, as $\text{NP}(f)$ is the standard simplex itself. Figure 2.14 depicts the 2-dimensional hyperplane in \mathbf{R}^3 .

A hypersurface $V(f)$, equipped with multiplicities as above, satisfies the so-called balancing condition — a fundamental property of tropical varieties.

Definition 2.4.6

Let \mathcal{X} be a pure-dimensional polyhedral complex in \mathbf{R}^n . A *weight function* on \mathcal{X} is given by a weight $\omega(\sigma) \in \mathbf{Z}$ for each facet σ of \mathcal{X} . In this case, we say \mathcal{X} is a *weighted polyhedral complex*. We define its support $|\mathcal{X}|$ as the union of facets of non-zero weight.

Let \mathcal{X} be a positively weighted polyhedral complex in \mathbf{R}^n . We say \mathcal{X} is *balanced* or satisfies the *balancing condition* if for every cell $\tau \in \mathcal{X}$ of codimension one the following equation holds:

$$\sum_{\substack{\sigma \text{ facet} \\ \tau \subset \sigma}} \omega(\sigma) v_{\sigma/\tau} \in L(\tau)$$

Here $v_{\sigma/\tau} \in \mathbf{Z}^n$ denotes a *primitive generator of σ modulo τ* , i.e. an integer vector that points from τ to the direction of σ and satisfies

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




Figure 2.15: The balancing condition

The balancing condition requires that the weighted sum of primitive vectors around a codimension one cell τ is parallel to τ . An equivalent formulation is as follows. Consider the ray $\sigma/\tau \in \text{Star}_{\mathcal{X}}(\tau)$ let $u_{\sigma/\tau} \in \mathbf{Z}^n / L_{\mathbf{Z}}(\tau)$ denote its unique primitive generator. By definition, $v_{\sigma/\tau} + L(\tau) = u_{\sigma/\tau}$ and hence the balancing condition can be equivalently expressed as

$$\sum_{\substack{\sigma \text{ facet} \\ \tau \subset \sigma}} \omega(\sigma) u_{\sigma/\tau} = 0 \in \mathbf{R}^n / L(\tau).$$

Allowing non-zero weights will be convenient later, but for all practical purposes (cf. definition of support) we may always throw away all facets of weight 0 and hence assume that all weights are non-zero. Obviously, the balancing condition is not affected by this operation.

Theorem 2.4.7

Any hypersurface $V(f)$ of a Laurent polynomial $f = \sum_{j \in A} a_j x^j$ forms a balanced polyhedral complex.

Proof. Let us first consider the two-dimensional case, i.e. assume f is a polynomial in two variables and hence $V(f)$ is a piecewise linear graph in the plane. We have to check the balancing condition around each vertex ν of $V(f)$. Let D_ν be the 2-cell in the dual subdivision. By duality we see that locally around ν the subdivision $\mathcal{S}(f)$ looks like the normal fan of D_ν . In particular, for each edge σ containing ν the primitive generator $v_{\sigma/\nu}$ is

orthogonal to the dual edge D_σ in D_ν , and $\omega(\sigma)$ is by definition just the integer length of D_σ . Thus, when we concatenate all the vectors $\omega(\sigma)v_{\sigma/\nu}$




Figure 2.16: The balancing condition in two dimensions

in say clockwise order, the prescribed chain of edges is equal to a right angle rotation of the boundary of D_ν . The fact that this boundary closes up is thus equivalent to the fact

$$\sum_{\substack{\sigma \text{ edge} \\ \nu \subset \sigma}} \omega(\sigma)v_{\sigma/\nu} = 0.$$

This finishes the proof for \mathbf{R}^2 , and essentially the same argument can be applied in the general case. Let τ be a codimension one face of $V(f)$ and let D_τ be the dual 2-cell of $SD(f)$. Then $L(D_\tau)$ is canonically the dual space of $\mathbf{R}^n/L(\tau)$ and for each facet σ containing τ , the primitive generator $u_{\sigma/\tau} \in \mathbf{R}^n/L(\tau)$ is orthogonal to the corresponding edge $D_\sigma \in SD(f)$. Thus, once again the “closedness” of D_τ implies the balancing condition. \square

Remark 2.4.8

Note that the balancing condition only involves the multiplicity of the facets of $V(f)$ or, in other words, of generic points of $V(f)$. Since we defined multiplicities for all cells of $\mathcal{S}(f)$, one might ask whether $\mathcal{S}(f)^{(k)}$ is balanced for all k . Indeed, we will prove this in Theorem 4.7.6.

Remark 2.4.9

Let us one more time revisit our question of defining the “zero set” of a tropical polynomial. Now that we know that the balancing condition is a fundamental property of tropical hypersurfaces, we might be bothered by

the fact that the graph of a tropical Laurent polynomial

$$\Gamma = \{(x, y) \in \mathbf{R}^n \times \mathbf{R} : y = f(x)\}$$

does *not* satisfy the balancing condition (of course, Γ carries a natural polyhedral structure by lifting $\mathcal{S}(f)$, and we use multiplicity 1 for all facets here). Note that in classical algebraic geometry, the graph of a polynomial




Figure 2.17: The graph of $\max\{(-1) + 2x, (-1) + 2y, 1 + x + y, x, y, 0\}$

is a hypersurface again given by an algebraic equation, namely $y = f(x)$. Tropically, however, the graph of a polynomial is obviously a polyhedral complex, but the balancing condition fails exactly at the points over $V(f)$ where the function is strictly convex. We can repair this by using the *multivalued* tropical addition \curlyvee from remark 2.2.7 again. Let us consider the multivalued function f^\curlyvee and its “graph”

$$\Gamma^\curlyvee = \{(x, y) \in \mathbf{R}^n \times \mathbf{T} : y \in f^\curlyvee(x)\}.$$

We find that Γ^\curlyvee contains the ordinary graph Γ , but additionally, over each point $x \in V(f)$, contains the half-infinite interval $\{(x, y) : y \leq f(x)\}$. Moreover, Γ^\curlyvee is canonically balanced. To that end, we equip each facet of Γ^\curlyvee which is not already contained in Γ with the multiplicity of its projection to \mathbf{R}^n which is a facet of $V(f)$. It is easy to see that Γ^\curlyvee is the only balanced completion of Γ if we only allow to add facets in direction $(0, -1) \in \mathbf{R}^n \times \mathbf{R}$. This is why Γ^\curlyvee is sometimes called the completed graph of f . Note that, in contrast to Γ , the completed version Γ^\curlyvee intersects the “zero-section” $\mathbf{R}^n \times \{-\infty\}$ and this intersection gives back $V(f)$ — yet another reason for our




Figure 2.18: The graph of $((-1) + 2x) \vee ((-1) + 2y) \vee (1 + x + y) \vee x \vee y \vee 0$

definition of the zeros of a tropical polynomial. We will meet the completed graph again when introducing the concept of tropical modifications later.

We close this section by proving an inverse statement to theorem 2.4.7. This emphasizes the importance of the balancing condition by showing it is not only a necessary but also sufficient condition for characterizing tropical hypersurfaces (at least if we extend the notion to quotients of tropical polynomials).

Theorem 2.4.10

Let \mathcal{X} be a balanced polyhedral complex in \mathbf{R}^n which is of dimension $n - 1$ and all of whose weights are positive. Then there exist tropical polynomials f such that \mathcal{X} is a weighted refinement of $V(f)$. Moreover, f is unique up to adding an integer affine function.

Here, \mathcal{X} is a *weighted refinement* of $V(f)$ if $|\mathcal{X}| = V(f)$, for each $\sigma \in \mathcal{X}$ there exists $\sigma' \in \mathcal{S}(f)^{(n-1)}$ with $\sigma \subset \sigma'$, and if σ a facet, $\omega_{\mathcal{X}}(\sigma) = \omega_{V(f)}(\sigma')$.

Proof. Let C_1, \dots, C_l be the connected components of $\mathbf{R}^n \setminus |\mathcal{X}|$. Our plan is to find a monomial f_i for each component C_i such that $f = \sum f_i$ and $f|_{C_i} = f_i$.

Let us consider the graph Γ with vertices v_1, \dots, v_l corresponding to the connected components and edges corresponding to facets of \mathcal{X} . An edge connects v_i and v_j if the associated cell separates C_i from C_j . Note that Γ

is connected since $\mathcal{X}^{(n-2)}$ is of codimension 2 and hence $\mathbf{R}^n \setminus |\mathcal{X}^{(n-2)}|$ is connected.

We pick one vertex, say v_1 , and set $f_1 \equiv 0$. Assume that v_2 is connected to v_1 by an edge and let $\sigma \in \mathcal{X}$ be the corresponding facet. Then we define

$$f_2 = f_1 + \omega(\sigma)h_\sigma,$$

where h_σ is the integer affine function such that $h_\sigma|_\sigma = 0$ and $dh_\sigma(v_{C_2/\sigma}) = 1$ (where $v_{C_2/\sigma}$ is a primitive generator of C_2 modulo σ). Using this procedure recursively, we may define integer affine functions f_i for all $i = 1, \dots, l$.

We need to check that the definition of f_i does not depend on the chosen path from v_1 to v_i . Equivalently, we show that for any loop in Γ adding up the various functions $\omega(\sigma)h_\sigma$ produces 0. Let $\tau \in \mathcal{X}$ be a $n - 2$ -cell. Going around τ in a small loop gives rise to a simple loop l_τ in Γ . Note that $\mathbf{R}^n \setminus |\mathcal{X}^{(n-3)}|$ is simply connected since $\mathcal{X}^{(n-3)}$ is of codimension 3 in \mathbf{R}^n . It follows that it suffices to check the addition to zero for the loops l_τ .

Let $\sigma_1, \dots, \sigma_h$ be the facets adjacent to τ , ordered according to l_τ , and let h_{σ_j} be the integer affine functions defined above. Then the condition

$$\sum \omega(\sigma_j)h_{\sigma_j} = 0$$

is just the dual version (and thus equivalent to) the balancing condition of \mathcal{X} at R . We have seen this in more details in the proof of theorem 2.4.7.

So far we proved that the above procedure produces a well-defined monomial f_i for each connected component C_i . Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be the unique continuous function with $f|_{C_i} := f_i$ for all i . It is clear from the construction of the f_i that f is well-defined and convex on each line segment in $\mathbf{R}^n \setminus \mathcal{X}^{(n-2)}$. By continuity, f is convex on all of \mathbf{R}^n and therefore equal to twice its Legendre dual “ $\sum_{i=1}^l f_i$ ”. Hence f is a tropical polynomial. It is now easy to check that \mathcal{X} is a weighted refinement of $V(f)$.

To verify uniqueness, let g be another tropical polynomial satisfying the hypothesis. Then g is also integer affine on each component of $\mathbf{R}^n \setminus |\mathcal{X}| = \mathbf{R}^n \setminus V(f)$. Moreover, since the weights of \mathcal{X} and $V(g)$ agree, g needs to satisfy the same transition formula than f when traversing a facet $\sigma \in \mathcal{X}$. It follows that $g = f + g_1$, where g_1 is the integer affine functions which describes g on C_1 . □

Exercise 2.4.11

Let $f_1, f_2, f_3 : \mathbf{R}^2 \rightarrow \mathbf{R}$ be the tropical polynomials given by

$$\begin{aligned} f_1 &= "(-1) + x + 1y + x^2 + xy + (-1)y^2", \\ f_2 &= "(-2) + (-2)x^3 + (-2)y^3 + x + y + x^2 + y^2 + x^2y + xy^2 + 1xy", \\ f_3 &= "0 + (-1)x + (-3)x^2 + y + 1xy + x^2y + (-2)y^2 + xy^2 + x^2y^2". \end{aligned}$$

In order to compute the associated tropical curves $V(f_i)$, proceed as follows.

- (a) Compute the Newton polytopes $\text{NP}(f_i)$ and the dual subdivisions $\text{SD}(f_i)$.
- (b) Compute (some of) the vertices of $V(f_i)$. (Each triangle in $\text{SD}(f_i)$ singles out three terms of f_i . The corresponding vertex is the point where these three terms attain the maximum simultaneously.)
- (c) Draw the curves $V(f_i) \subset \mathbf{R}^2$ by adding the edges.

Exercise 2.4.12

Let $f : \mathbf{R} \rightarrow \mathbf{R}$ be a univariate polynomial and $\alpha \in \mathbf{R}$. Let $\text{ord}_\alpha(f)$ be the maximal number k such that $(x - \alpha)^k | f$ (cf. Exercise 2.2.8). Show that $\omega_{V(f)}(\alpha) = \text{ord}_\alpha(f)$.

Exercise 2.4.13

Let \mathcal{X} be a weighted polyhedral complex of pure dimension k . For any cell τ the fan $\text{Star}_{\mathcal{X}}(\tau)$ carries induced weights given by $\omega(\sigma/\tau) = \omega(\sigma)$. Show that the following statements are equivalent.

- (a) The polyhedral complex \mathcal{X} is balanced.
- (b) The stars $\text{Star}_{\mathcal{X}}(\tau)$ are balanced for any $\tau \in \mathcal{X}$.
- (c) The stars $\text{Star}_{\mathcal{X}}(\tau)$ are balanced for any $k - 1$ -cell τ .

2.5 Planar Curves

This section is devoted to the study planar curves as a particular example of tropical hypersurfaces. While they are easily visualizable, they still exhibit




Figure 2.19: A tropical planar line

interesting combinatorial, but also geometric properties with have counterparts in higher dimensions.

A tropical *planar curve* is the hypersurface $C = V(f)$ of a tropical Laurent polynomial

$$f = \sum_{j \in \mathbf{Z}^2} a_j x^{j_1} y^{j_2}$$

in two variables x, y . Let $\Delta_d, d \in \mathbf{N}$ denote the convex hull of the points $(0, 0), (d, 0), (0, d)$. We call Δ_d the *d-fold standard simplex*. Mostly, we will consider polynomials whose Newton polytope is $\text{NP}(f) = \Delta_d$ for a suitable d . In this case we call $V(f)$ a curve of *degree d*. We may extend this definition to arbitrary polynomials by saying $V(f)$ is of degree d if d is the smallest natural number such that (a shift of) Δ_d contains $\text{NP}(f)$. If $\text{NP}(f)$ is not equal to this minimal (shift of) Δ_d , we call $V(f)$ *degenerated*.

Of course, we start with *planar lines* $L = V(f)$, where f is a polynomial of degree 1. The general form of such a polynomial is

$$f(x, y) = "ax + by + c".$$

Let us start with $a = b = c = 0$ (see Figure 2.19). Then $L = V(f)$ consists of three rays $\mathbf{R}_{\geq}(1, 1), \mathbf{R}_{\geq}(-1, 0)$ and $\mathbf{R}_{\geq}(0, -1)$. The point $(0, 0)$ is the single vertex of L . Each ray corresponds to two of the three terms of f being maximal, while $(0, 0)$ is the single point where all three terms attain the maximum. What happens if we change a, b, c ? There is still a unique point where all three monomials are maximal, $(c - a, c - b)$. Therefore

$L_{a,b,c} = V("ax + by + c")$ is just the translation of L to this point. If we decrease one coefficient at a time, we move L in one of the directions $(-1, -1)$, $(1, 0)$ or $(0, 1)$. In the limit, i.e. when one coefficient becomes $-\infty$, we end up with an ordinary classical line whose Newton polytope is just a segment (see Figure 2.20). In total, there are a single non-degenerated and 3 degenerated types of tropical lines, and two lines of the same type are translations of each other. Let us stress that the only classical lines which show up as "tropical lines" are the lines of slope $(1, 1)$, $(-1, 0)$ or $(0, -1)$. Let us now decrease a second coefficient of our linear polynomial, such that in the limit two coefficients are $-\infty$. Geometrically, the line vanishes at infinity and $V(f)$ is empty. This reminds us of the fact that we work with the tropical algebraic torus $\mathbf{R}^n = (\mathbf{T}^\times)^n$ here, a non-compact space on which single monomials are "non-vanishing" functions. Later on, we will consider various compactifications of \mathbf{R}^n , in particular tropical projective space \mathbf{TP}^2 . In this space, when decreasing two of the coefficients, the corresponding moving line will attain one of the coordinate lines $x = -\infty$, $y = -\infty$ or $z = -\infty$ as limit.¹

For now, let us stick to the non-compact picture. There are two elementary properties of classical planar lines which we want to study tropically now: Two generic lines intersect in a single point, and, dually, through two different points in the plane there passes a unique line.

First, let us consider the intersection of two tropical lines. Indeed, for most pairs of lines, we get exactly one point of intersection, but there are two notable exceptions, as illustrated in the Figure 2.22. While in the first

¹In particular, we see that the dual space of lines in \mathbf{TP}^2 is itself a copy of \mathbf{TP}^2 and inherits a stratification into affine spaces corresponding to the seven types of lines.




Figure 2.20: Moving a tropical line




Figure 2.21: The types of tropical lines




Figure 2.22: Intersections of tropical lines

two pictures we get a perfectly nice single point of intersection, the third picture shows two parallel degenerated lines with no intersection. As in classical geometry, we will get rid of this special case by compactifying \mathbf{R}^2 . However, a classical geometer would very likely compactify to the real projective plane \mathbf{RP}^2 to make the lines intersect at infinity. This is different from the tropical compactifications we will study later, e.g. \mathbf{TP}^2 . The fourth picture shows a more interesting way of failure. Here, the two tropical lines have a whole ray in common. This type of abnormality does not have a classical counterpart (except for taking the same line twice) and will later encourage us to introduce the concept of *stable intersection*. The stable intersection of two curves consists only of those points in the set-theoretic intersection which are stable under small deformations — small translations, in our case. E.g. for lines, a small translation of one of them will yield a unique intersection point close to the apex of the common ray. When moving the translation back to the original line, the limit of intersection points

is just this single point, which we therefore call the stable intersection of the two lines. The dual problem of finding a line through two points shows a completely similar behaviour.

Proposition 2.5.1

Any pair of points $p_1, p_2 \in \mathbf{R}^2$ can be joined by a tropical line. Furthermore, this line is unique if and only if the points do not lie on an ordinary classical line of slope $(1, 1)$, $(-1, 0)$ or $(0, -1)$.

Proof. Even though the statement is completely elementary, let us give a short proof here. First, we set L to be the tropical line whose single vertex is p_1 . If p_2 lies on one of the rays of L , we are already done (with existence). If not, p_2 is contained in one of the sectors of $\mathbf{R}^2 \setminus L$ corresponding to a certain monomial being maximal. We move L into this sector by decreasing the coefficient of the corresponding monomial. While doing so, the two rays of L that bound the sector will sweep it out completely. Hence we just stop when the moving line meets p_2 (see Figure 2.24). The uniqueness statement follows from our previous discussion of the intersection of two lines. Namely, if we can find two different lines containing p_1 and p_2 , then the two lines have a ray of slope $(1, 1)$, $(-1, 0)$ or $(0, -1)$ in common, and the statement follows. \square

We now consider curves of degree 2 given by polynomials of the form

$$f(x, y) = "ax^2 + bxy + cy^2 + dx + ey + f".$$




Figure 2.23: The stable intersection point of two lines




Figure 2.24: Finding a line passing through two points

Contrary to the case of lines, where we found only one non-degenerated example up to translations, the variety of conics is already more interesting: There exist several combinatorial types of non-degenerated conics and two curves of the same type are not necessarily just translation of each other. Figure 2.25 gives a list of the four “smooth” combinatorial types of smooth non-degenerated conics, given by the four unimodular subdivisions of Δ_2 . In all four cases the complement $\mathbf{R}^2 \setminus V(f)$ consists of six




Figure 2.25: Smooth (non-degenerated) conics

connected components — the linearity domains of the six monomials in f . Let us again deform a given conic by changing the coefficient of just one monomial. This will enlarge or shrink the corresponding connected component, depending on whether we increase or decrease the coefficient. More precisely, the deformation will move *all* edges adjacent to this component while all other edges rest in the sense that they stay in the same affine line (they might change length, however). This follows from the fact that the position of each of these edges is given by an equation of the form

$a_i + i_1x + i_2y = a_j + j_1x + j_2y$ with $i, j \in \Delta_2 \cap \mathbf{Z}^2$, and this affine line moves if and only if a_i or a_j is the coefficient being changed (see Figure 2.26). This describes small changes of a single coefficient. However, if we keep




Figure 2.26: Increasing the x -coefficient

deforming, a couple of things can eventually happen. Of particular interest is the case when the combinatorial type of the conic changes. This happens if one of the edges shrinks to length zero and vanishes. In this case also the corresponding edge in the dual subdivision vanishes which merges together two of the triangles to a 2-cell of volume 2. Such a 2-cell can either be a parallelogram or a triangle with one side of length 2. Both situations can be illustrated by deforming the conic from the previous picture. Increasing, as before, the coefficient of the monomial “ x ”, we run into a conic with a singular point whose dual cell is a square. Increasing the coefficient further, we get a smooth conic again, but the combinatorial type has changed (see Figure 2.27). If we instead increase the monomial “ x^2 ”, we run into a singular conic with a ray of multiplicity 2 (see Figure 2.28). Note that in this case, increasing the coefficient further does not remove the weight 2 edge. The parallelogram degeneration corresponds to a reducible conic which decomposes into two lines. The second degeneration may be interpreted as a conic that is tangent to the coordinate axis $y = -\infty$ at infinity.

For higher degree curves, the number of combinatorial types becomes large very quickly. Already, for curves of degree 3, there are 79 smooth non-degenerate combinatorial types. Figure 2.29 depicts particular examples of a smooth and a singular cubic. In general, a planar curve is smooth if and only if each of its vertices has exactly 3 adjacent edges (we say all




Figure 2.27: Passing through a singular conic

vertices are 3-valent) and each pair of the primitive integer vectors v_1, v_2, v_3 spanning the 3 edges of a vertex form a \mathbf{Z} -basis of \mathbf{Z}^2 . We now present a few special cases of combinatorial types for arbitrary degrees.

The *honeycomb triangulation* of Δ_d is defined by the property that all its edges are parallel to one of the three boundary edges of Δ_d . In other words, it is obtained from the collection of lines $x = i, y = j, x + y = k$ with $i, j, k \in \{0, \dots, d\}$. Our figure depicts the case $d = 5$. The curves of this combinatorial type are called *honeycombs*. They proved to be useful in




Figure 2.28: A conic with a “fat” edge




Figure 2.29: A smooth and a singular cubic

the context of the Horn problem (e.g. [KT01]).

The *bathroom tiling curves* form a similar type of curves. Their dual triangulation of Δ_d is given by $x = i, y = j$ with $i, j \in \{0, \dots, d\}$, $x + y = k, x - y = l$ with $k, l \equiv d \pmod{2}$. Again, our picture shows an example of degree 5. These curves are of interest for example when constructing real




Figure 2.30: A honeycomb curve of degree 5




Figure 2.31: A bathroom tiling curve of degree 5

algebraic curves.

Let us also recall an example of a special tropical curve which we mentioned in the introduction. Figure 2.32 shows the curve of degree 10 together with its subdivision. It is this curve (together with the extra data of suitable signs) that was used by Itenberg (cf. [IV96]) to disprove the famous Ragsdale conjecture (cf. [Rag06]; the conjecture was an inequality involving the numbers of “odd” and “even” ovals of a real planar curve of given degree).

A fundamental theorem in the study of classical planar curves is Bézout’s




Figure 2.32: The Itenberg-Ragsdale curve of degree 10




Figure 2.33: The 6 intersection points of a conic and a cubic

theorem which states that two projective curves of degree d and e have de points of intersections (in various meanings and under various assumptions). Previously we discussed the case of lines which are supposed to intersect in a single point. It is easy to convince oneself that a more general statement should also be true tropically. For example, a smooth non-degenerate curve of degree d has exactly d unbounded rays in each of the directions $(1, 1)$, $(-1, 0)$ and $(0, -1)$ corresponding to the d dual segments in the boundary of Δ_d . Thus two such tropical curves of degree d and e can be translated in such a way that they only intersect in such rays of fixed direction (for each curve) — in this case there are exactly de intersection points (see Figure 2.33). We can actually prove a general version of Bézout’s theorem based on our understanding of dual subdivisions. For this purpose, let us assume that C and D are non-degenerate tropical planar curves of degree d resp. e and given by the polynomials $f = \sum_{j \in \Delta_d} a_j x^{j_1} y^{j_2}$ resp. $g = \sum_{j \in \Delta_e} b_j x^{j_1} y^{j_2}$. The non-degeneracy condition ensures that no intersections “at infinity” occur, so it is enough to work in \mathbf{R}^n here. Let us furthermore assume that C and D intersect *transversally*. This means that that $C \cap D$ is finite and each intersection point lies in the relative interior of an edge of C and an edge of D . We count each such point with an *intersection multiplicity* defined as follows. Assume the intersecting edges are σ and τ , then we define

$$\begin{aligned} \text{mult}(p) &:= \omega(\sigma)\omega(\tau)[\mathbf{Z}^2 : L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\tau)] \\ &= \omega(\sigma)\omega(\tau)|\det(v, w)|, \end{aligned}$$

where v and w are primitive integer vectors describing the slope of σ resp. τ . This multiplicity is one if and only if both weights are one and v, w form a lattice basis of \mathbf{Z}^2 .



Figure 2.34: Intersection points of multiplicity 1 and 2

Theorem 2.5.2

Let $C = V(f)$ and $D = V(g)$ be two non-degenerate tropical curves of degree d and e , respectively, which intersect transversally. The number of intersection points, counted with multiplicities, is equal to de .

$$de = \sum_{p \in C \cap D} \text{mult}(p)$$

Proof. The main trick is to consider the union of the two curves $B := C \cup D = V("fg")$. Note that each intersection point $p \in C \cap D$ is a vertex of $C \cup D$ and thus has a corresponding dual cell σ in $\text{SD}("fg")$. This dual cell must be a parallelogram whose pairs of parallel edges correspond to the edges of C resp. D that intersect. Moreover, we see that the multiplicity of p can be computed in terms of the volume of σ , namely

$$\text{mult}(p) = \frac{\text{Vol}(\sigma)}{2}.$$

If a vertex of B is not an intersection point, it is just a vertex of C or D . Note that the corresponding dual triangle in $\text{SD}("fg")$ is just a shift of the corresponding triangle in $\text{SD}(f)$ resp. $\text{SD}(g)$ (see Figure 2.35). Indeed, assume that v is a vertex of C and therefore by assumption $v \notin D$. Then locally around v the function g is affine-linear, say $b_j + jx$, and therefore " fg " is locally equal to $f + b_j + jx = "fb_jx^j"$, which corresponds to a shift of the dual picture by j . Hence the maximal cells of $\text{SD}("fg")$ are

either triangles, which are in (volume-preserving) bijection to the triangles of $SD(f)$ and $SD(g)$, or parallelograms, which are in bijection to $C \cap D$. This implies

$$\begin{aligned} \text{Vol}(\text{parallelograms of } SD("fg")) \\ &= \text{Vol}(\Delta_{d+e}) - \text{Vol}(\Delta_d) - \text{Vol}(\Delta_e) \\ &= (d+e)^2 - d^2 - e^2 = 2de. \end{aligned}$$

Together with the above formula, this proves the result. □

Our proof is adapted from a paper of Vigeland (cf. [Vig09]).

Exercise 2.5.3

In this exercise we give an alternative proof of tropical Bézout’s theorem (cf. [RST05]). Let $C = V(f)$ and $D = V(g)$ be two non-degenerate tropical curves of degree d and e , respectively. We denote $\deg(C \cdot D) = \sum_{p \in C \cap D} \text{mult}(p)$ whenever C and D intersect transversally.

- (a) Show that the set of vectors $v \in \mathbf{R}^2$ such that C and $D + v$ do not have a vertex in common is connected.
- (b) Assume that C is a fan and Y is an (usual) affine line. Show that $\deg(C \cdot (D + v))$ is constant for all v such that the intersection is transversal.
- (c) Using the previous two items, show that $\deg(X \cdot (Y + v)) = \text{const.}$ also holds in the general case (for all “generic” v).




Figure 2.35: Bézout’s theorem in the case of two tropical conics

- (d) Find a particular v such that C and $D + v$ intersect transversally in de points, all of multiplicity 1 (Hint in Figure 2.33).

Exercise 2.5.4

Let $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ be a tropical polynomial and $C = V(f)$ the associated planar tropical curve. The (*embedded*) genus $g(C)$ of C is the first Betti number $b_1(C) = \dim H_1(C, \mathbf{Z})$ (i.e. the number of independent 1-cycles in C). Show the following statements.

- (a) The genus $g(C)$ is equal to the number of vertices of $SD(f)$ contained in $NP(f)^\square$.
- (b) Assume that C is a non-degenerated smooth curve of degree d . Then

$$g(C) = \frac{(d-1)(d-2)}{2}.$$

Exercise 2.5.5

Let $C \subset \mathbf{R}^2$ be a planar tropical curve of degree d (possibly degenerated). For each ray ρ of C with primitive generator $v = (v_1, v_2) \in \mathbf{Z}^2$, we set $\text{ord}_\infty(\rho) := \omega(\rho) \max\{v_1, v_2, 0\}$. Show that $d = \sum_\rho \text{ord}_\infty(\rho)$, where the sum runs through all rays of C .

2.6 Floor decompositions

Floor decomposition is a combinatorial tool allowing to construct/describe a hypersurface in \mathbf{R}^n using several hypersurfaces in \mathbf{R}^{n-1} . The hypersurfaces which admit such a description (as well as their dual subdivisions) are of rather special type and are called *floor-decomposed*. Particular types of floor-decomposed subdivisions appeared in the context of Viro's patchworking method, see [IV96]. See Figure 2.39 for an example. In tropical geometry, they are particularly useful in the context of enumerative geometry [BM07; BM09]. We present the basic ideas here.

Floor decompositions can be defined with respect to any surjective linear map of tropical vector spaces $V \rightarrow W$ of dimension n and $n - 1$, respectively. For simplicity, we choose explicit coordinates here and only consider the fixed projection $\pi : \mathbf{R}^n \rightarrow \mathbf{R}^{n-1}$ forgetting the last coordinate. We will




Figure 2.36: Floor-decomposed subdivisions in dimension 2 and 3

also need the dual projection $q : \mathbf{R}^n \rightarrow \mathbf{R}$ which forgets all except the n -th coordinate.

Definition 2.6.1

Let $V(f) \subset \mathbf{R}^n$ be a hypersurface with dual subdivision $SD(f)$. The $V(f)$ is called *floor-decomposed* (with respect to π) if for each cell $D \in SD(f)$ the image $q(\sigma)$, which is an integer interval in \mathbf{R} , is of length at most 1 (i.e. of length 0 or 1).

Examples of floor-decomposed subdivisions are given in Figure 2.36. The surface corresponding to the three-dimensional subdivision is displayed in Figure 2.37.

A floor-decomposed hypersurface splits naturally into floors and elevators as follows. We assume $q(\text{NP}(f)) = [0, m]$. For a cell σ of $V(f)$, we denote the dual cell in $SD(f)$ by D_σ . For $i = 1, \dots, m$, we define

$$F_i := \bigcup_{\substack{\sigma \text{ cell of } V(f) \\ q(D_\sigma)=[i-1,i]}} \sigma,$$

and call this closed set a *floor* of $V(f)$. For $i = 0, \dots, m$, we define

$$E_i := \bigcup_{\substack{\sigma \text{ cell of } V(f) \\ q(D_\sigma)=\{i\}}} \sigma^\square,$$




Figure 2.37: Two views of a floor-decomposed surface

and call this set an *elevator* of $V(f)$. (In a more refined setting, one might call the connected components of the E_i elevators.) The floors and elevators in Figure 2.37 are shown in gray and red, respectively.

Accordingly, we can split the polynomial f , i.e. we can write f as a polynomial of the last coordinate x_n

$$f = \sum_{i=0}^m f_i x_n^i,$$

with $f_i \in \mathbf{T}[x_1, \dots, x_{n-1}]$. We have the following relation.

Proposition 2.6.2

The projection $\pi(E_i)$ of an each elevator is a tropical hypersurface in \mathbf{R}^{n-1} given by $V(f_i)$.

Each floor F_i projects one-to-one to \mathbf{R}^{n-1} and is equal to the graph of the function “ f_{i-1}/f_i ” = $f_{i-1} - f_i$ on \mathbf{R}^{n-1} .

The floors and elevator projections of our example surface are displayed in Figure 2.38.

Proof. For the first part, let us fix an elevator E_i . We want to show that $\pi(E_i) = V(f_i)$ as sets. First, take $x \in E_i$. By definitions, the dual cell D_x




Figure 2.38: The elevators and floors of our example surface

of x in $\text{SD}(f)$ satisfies $q(D_x) = \{i\}$. It follows that locally at x we have $f \equiv f_{D_x} \equiv "f_i x_n^i"$. In particular, $f_i(\pi(x))$ attains its maximum at least twice and thus $\pi(x) \in V(f_i)$. For the other direction, pick $y \in \mathbf{R}^{n-1}$. We may consider the polynomial $g = "\sum_i f_i(y)x_n^i"$ in one variable x_n obtained by restricting f to $\pi^{-1}(y)$. Note that g is of degree m and its subdivision $\text{SD}(g)$ can be computed as

$$\text{SD}(g) = \{q(D_\sigma) : \sigma \in \mathcal{S}(f) \text{ such that } \sigma \cap \pi^{-1}(y) \neq \emptyset\}.$$

As f is floor-decomposed, all cells of $\text{SD}(g)$ have length 1, i.e. $[0, m]$ is maximally subdivided. It follows that for any i , there exists x_n such that the only maximal term in $g(x_n)$ is " $f_i(y)x_n^i$ ". If we further assume $y \in V(f_i)$, the point $x = (y, x_n) \in \mathbf{R}^n$ is in $V(f)$ and the maximum of $f(x)$ is only attained on the i -th level. It follows $x \in E_i$ and therefore $y \in \pi(E_i)$. This proves $\pi(E_i) = V(f_i)$. Note that also the multiplicities of the two sets are compatible in the sense that $\omega_{V(f)}(x) = \omega_{V(f_i)}(\pi(x))$. This is true because we defined the multiplicity to be the volume of the dual cell in $\text{SD}(f)$ resp. $\text{SD}(f_i)$, which is the same (up to being embedded in different ambient spaces).

Now let us prove the second part of the assertion. First note that $x \in F_i$ is equivalent to

$$f(x) = "f_i(\pi(x))x_n^i" = "f_{i-1}(\pi(x))x_n^{i-1}."$$

Transforming the second equality according to the rules of tropical arithmetics, we get $x_n = f_{i-1}(\pi(x)) - f_i(\pi(x))$. So F_i is contained in the graph

of the function $f_{i-1} - f_i$ and therefore $\pi|_{F_i}$ is injective. For surjectivity, we use again the fact that, for any $y \in \mathbf{R}^{n-1}$ and corresponding one variable polynomial $g = \sum f_i(y)x_n^i$, $SD(g)$ maximally subdivides $[0, m]$. Hence, there exists a (unique) “zero” z of g with $g(z) = f_i(y)z^i = f_{i-1}(y)z^{i-1}$ and therefore $(y, z) \in F_i$. \square

The nice feature of floor-decomposed hypersurfaces is that they are completely described by (the projection of) their elevators and the “height” of their floors. This is the statement of the following proposition. For simplicity, we restrict ourselves to the case of non-degenerated hypersurfaces, i.e. $NP(f) = \Delta_d$, the simplex of size d . Of course, for other toric varieties and different shapes of Newton polytopes, one can proceed similarly.

We make the following convention. For a non-degenerated hypersurface V in \mathbf{TP}^n , we denote by f_V the unique tropical polynomial (not Laurent polynomials) in n variables such that $V = V(f_V)$ and the constant term of f_V is 0 (not $-\infty$).

Proposition 2.6.3

Let V_0, \dots, V_{d-1} be non-degenerated tropical hypersurfaces in \mathbf{R}^{n-1} of degree $\deg(V_i) = d - i$. Moreover, let $h_1 < h_2 < \dots < h_d$ be real numbers and set $f_i := f_{V_i} - \sum_{0 < j \leq i} h_j$ for all $i = 0, \dots, d$ (with and $f_{V_d} \equiv 0$). Under the assumption

$$(2f_i - f_{i-1} - f_{i+1})(x) > 0$$

for all $x \in \mathbf{R}^{n-1}$ and $i = 1, \dots, m - 1$, there exists a unique generic floor-decomposed hypersurface $V \subset \mathbf{TP}^n$ such that $\pi(E_i) = V_i$ for all elevators and each floor F_i intersects the x_n -coordinate axis at height h_i . This hypersurface is given by $V = V(f)$, where $f = \sum_{i=0}^m f_i x_n^i$. In particular, V has degree d .

Proof. Let us study $V := V(f)$ with $f := \sum_{i=0}^m f_i x_n^i$. The assumption $(2f_i - f_{i-1} - f_{i+1})(x) > 0$ makes sure that the various floors (given by the graphs of $f_{i-1} - f_i$) do not intersect and therefore V is indeed floor-decomposed. By proposition 2.6.2, the elevators of V satisfy $\pi(E_i) = V(f_i) = V_i$. Moreover, again by proposition 2.6.2, the height of the intersection of F_i with the x_n -coordinate axis is given by the the difference of the constant terms of f_{i-1} and f_i , which is $-h_1 - \dots - h_{i-1} + h_1 + \dots + h_{i-1} + h_m = h_m$. For the uniqueness of V , we note that the condition $\pi(E_i) = V_i$ already implies $f_i - f_{V_i} \equiv \text{const}$. These constants are uniquely fixed (up to adding a global constant) by the “heights” of the floors F_i . \square




Figure 2.39: The Itenberg-Viro subdivision in two dimensions

Example 2.6.4

Figure 2.39 shows an example of a floordecomposed curve that appears in the context of Viro’s patchworking method. A simple example for the construction of a stacking floors on top of each other is given in Figure 2.40.

References

The main themes of this chapter — tropical polynomials, tropical hypersurfaces, dual subdivisions and the balancing condition — appear in many of the early papers on tropical geometry and we do not attempt to describe the history of the development of these notions and its diverse predecessors in other areas of mathematics. Among the earliest “tropical” papers treating (or at least mentioning) some of this material, let us mention [Mik01; Mik04; Stu02; SS04; Mik05; EKL06]. Note, in particular, that [EKL06] is based on an earlier unpublished preprint by Kapranov treating tropical hypersurfaces (under the name “non-archimedean amoebas”).




Figure 2.40: Stacking floors on top of each other

3 Projective space and other tropical toric varieties

In the previous chapter we studied the tropical space \mathbf{R}^n and its hypersurfaces. Let us recall once more that $\mathbf{R} = \mathbf{T}^\times$ is the set of units for tropical multiplication on $\mathbf{T} = \mathbf{R} \cup \{-\infty\}$. Hence, $\mathbf{R}^n = (\mathbf{T}^\times)^n$ is the tropical algebraic torus with classical analogue $(\mathbf{C}^\times)^n$ and we treated what is sometimes called very affine varieties so far.

The goal of this chapter is to study tropical geometry for certain compactifications of \mathbf{R}^n , in particular, tropical projective space. In fact, most of the new features that will show up are already present in tropical affine space \mathbf{T}^n , a partial compactification of \mathbf{R}^n . The set of new points $\mathbf{T}^n \setminus \mathbf{R}^n$ at “infinity” behaves differently from the points in \mathbf{R}^n , even on a topological level. In particular, the number of coordinates equal to $-\infty$, called the *sedentarity* of a point, is an invariant under tropical automorphisms of \mathbf{T}^n . This is of course in contrast to the classical situation, with \mathbf{C}^n and \mathbf{CP}^n being fully homogeneous spaces.

A natural class of compactifications of $\mathbf{R}^n = (\mathbf{T}^\times)^n$ are toric varieties. The main feature of toric varieties is that they can be constructed by monomial maps and binomial equations, respectively, which is to say, without using *addition*. Since the peculiar idempotent nature of tropical arithmetics occurs only in the context of tropical additions, it turns out that toric varieties can be pushed to the tropical world instantly.

3.1 Tropical affine space \mathbf{T}^n

The natural first step when compactifying \mathbf{R}^n is, of course, to consider *tropical affine space* \mathbf{T}^n — with classical analogue \mathbf{C}^n . Here, we always equip \mathbf{T} with the topology generated by intervals of the type $[-\infty, a)$, (b, c) , $a, b, c \in \mathbf{R}$, and use the product topology on \mathbf{T}^n .

An element of \mathbf{T}^n is an n -tuple of numbers $x = (x_1, \dots, x_n)$, but now some of these coordinates might be equal to $x_i = -\infty$. If all coordinate entries x_i are finite, i.e. $x \in \mathbf{R}^n$, we call x a *points at finite distance*. The additional points $x \in \mathbf{T}^n \setminus \mathbf{R}^n$, with one or more coordinate entries equal to $-\infty$, are called *points at infinity*. Of course, we can refine this distinction by setting for each subset $I \subseteq [n]$

$$\mathbf{R}_I := \mathbf{R}^{n \setminus I} := \{x \in \mathbf{T}^n : x_i = -\infty \forall i \in I, x_i \neq -\infty \forall i \notin I\}.$$

This gives us a natural stratification

$$\mathbf{T}^n = \bigsqcup_{I \subseteq [n]} \mathbf{R}_I.$$

Note that for all I we have an identification $\mathbf{R}_I = \mathbf{R}^{n \setminus I} \cong \mathbf{R}^{n-|I|}$, which means that all the strata in $\mathbf{T}^n \setminus \mathbf{R}^n$ are tropical algebraic tori of smaller dimension (see Figure 3.1). Analogously, we set

$$\mathbf{T}_I := \mathbf{T}^{n \setminus I} := \{x \in \mathbf{T}^n : x_i = -\infty \forall i \in I\}.$$

Each \mathbf{T}_I is again a tropical affine space $\mathbf{T}_I = \mathbf{T}^{n \setminus I} \cong \mathbf{T}^{n-|I|}$.




Figure 3.1: \mathbf{T}^2 and its stratification

Definition 3.1.1

Let x be a point in \mathbf{T}^n . The *sedentarity* $\text{sed}(x)$ is the number of coordinates of x equal to $-\infty$. Namely, if $x \in \mathbf{R}_I$, then $\text{sed}(x) = |I|$.

How do the points of different sedentarity interact? How are the various strata glued together geometrically? To understand this better, let $x, v \in \mathbf{R}^n$ be vectors consider the ray $R = x + \mathbf{R}_{\geq 0}v \subseteq \mathbf{R}^n$. We are interested in the closure $\bar{R} \in \mathbf{T}^n$. A few examples of rays in \mathbf{T}^2 are depicted in Figure 3.2 (the vectors labelling each ray describe the particular choice of v).




Figure 3.2: Rays in \mathbf{T}^2 labelled by their direction vectors

We have to distinguish two cases. Whenever v contains a strictly positive entry, no limit point exists, i.e. $\bar{R} = R$. Conversely, if all entries of v are non-positive, a limit point $\bar{x} = \bar{x}(x, v)$ is added to $\bar{R} = R \cup \{\bar{x}\}$. However, figure 3.2 might give us a wrong impression of where this point \bar{x} is located. Its coordinates are

$$\bar{x}_i = \begin{cases} x_i & \text{if } v_i = 0, \\ -\infty & \text{if } v_i < 0. \end{cases}$$

In particular, whenever $v_i < 0$, the coordinate \bar{x}_i does not depend on the starting point x . In our examples, it follows that two of the rays, though having different directions, run into the same limit point $(-\infty, -\infty)$, independent of their exact starting point (see Figure 3.3). Only the two rays with direction $(-1, 0)$ resp. $(0, -1)$ have limit points different from $(-\infty, -\infty)$.

Let $\sigma \subseteq \mathbf{R}^n$ be the cone spanned by the negative standard basis vectors $-e_1, \dots, -e_n$. Its faces are of the form σ_I , where $I \subseteq [n]$ and σ_I is the cone spanned by $-e_i, i \in I$. Then the above discussion can be formalized and summarized as follows. If $v \notin \sigma$, then no limit point is added and $\bar{R} = R$. Conversely, if $v \in \sigma$, let σ_I be the minimal face of σ containing v . Then




Figure 3.3: The closures of the rays in \mathbf{T}^2

$\bar{x} \in \mathbf{R}_I$ and its finite coordinates are just given by the projection

$$\begin{aligned} \mathbf{R}^n &\rightarrow \mathbf{R}^{n-|I|} \cong \mathbf{R}_I, \\ (x_i)_{1 \leq i \leq n} &\mapsto (x_i)_{i \notin I}. \end{aligned}$$

We see that the stratum containing \bar{x} is determined by the minimal face containing v . For example, the codimension one strata of \mathbf{T}^n can only be reached along the n special directions $-e_1, \dots, -e_n$.




Figure 3.4: Tropical affine space \mathbf{T}^3

Exercise 3.1.2

Let $A = (a_{ij}) \in \text{Mat}(m \times n, \mathbf{Z})$ be an integer matrix and let $\phi : \mathbf{R}^n \rightarrow \mathbf{R}^m, x \mapsto Ax$ be the associated integer linear map. Set $I(A) := \{i : a_{ki} \geq 0 \forall k\}$ and $U(A) = \mathbf{T}^n \setminus \bigcup_{i \notin I(A)} \mathbf{T}_{\{i\}} = \bigcup_{I \subset I(A)} \mathbf{R}_I$. Show that the following statements.

- (a) There exists a unique continuous extension of ϕ_A to a map $\phi : \mathbf{T}^n \supset U(A) \rightarrow \mathbf{T}^m$.
- (b) For $I \subset I(A)$, we set $K = K(I) = \{k : a_{ki} > 0 \text{ for some } i \in I\}$. Then the extension ϕ maps $\mathbf{R}_I \subset \mathbf{T}^n$ to $\mathbf{R}_K \subset \mathbf{T}^m$.
- (c) For $K = K(I)$, the map $\phi : \mathbf{R}_I \rightarrow \mathbf{R}_K$ is the integer linear map associated to the matrix A with K -rows and I -columns removed.

3.2 Tropical toric varieties

A complex monomial map is a map of

$$\begin{aligned} \Phi_A : (\mathbf{C}^\times)^n &\rightarrow (\mathbf{C}^\times)^m, \\ (z_1, \dots, z_n) &\mapsto \left(\prod_i z_i^{a_{1i}}, \dots, \prod_i z_i^{a_{mi}} \right), \end{aligned}$$

where $A = (a_{ij}) \in \text{Mat}(m \times n, \mathbf{Z})$ is an integer matrix. Complex toric varieties are glued via (invertible) monomial maps. Obviously, the concept can be transferred to any other group, and in particular, to (\mathbf{T}, \cdot) . Since “ $\prod_i x_i^{a_{ki}}$ ” = $\sum a_{ki} x_i$, a *tropical monomial map* is nothing but an integer linear map

$$\begin{aligned} \phi_A : \mathbf{R}^n &\rightarrow \mathbf{R}^m, \\ x = (x_1, \dots, x_n) &\mapsto Ax = \left(\sum_i a_{1i} x_i, \dots, \sum_i a_{mi} x_i \right). \end{aligned}$$

We are mostly interested in invertible maps, i.e., in the case $m = n$ and $\det(A) = \pm 1$ (again, independent of the group). In other words: Tropical monomial transformations of \mathbf{R}^n are \mathbf{Z} -invertible integer linear maps of \mathbf{R}^n .

We would like to extend ϕ_A to (parts of) \mathbf{T}^n — the problem that may occur is that $a_{ki} \cdot (-\infty)$ is not defined when $a_{ki} < 0$. Let $\mathbf{T}^n = \bigsqcup \mathbf{R}_I$ be the stratification of \mathbf{T}^n as before. Then ϕ_A can be extended to the stratum \mathbf{R}_I if

and only if $a_{ki} \geq 0$ for all $i \in I$ (see Exercise 3.1.2). Again, the analogous statements hold in the complex setting, which means that the rules for gluing affine patches together are exactly the same. For simplicity, we restrict to smooth toric varieties here.

Recall that a pointed fan Ξ in \mathbf{R}^n is called *unimodular* if every cone σ of Ξ can be generated by $\dim(\sigma)$ integer vectors that form a lattice basis of $L_{\mathbf{Z}}(\sigma)$. Let Ξ be a unimodular pointed fan in \mathbf{R}^n . For any cone $\sigma \in \Xi$ of dimension k , fix a lattice basis $e_1^\sigma, \dots, e_n^\sigma$ of \mathbf{Z}^n such that $-e_1^\sigma, \dots, -e_k^\sigma$ generate σ and set

$$U_\sigma := \mathbf{T}^k \times (\mathbf{T}^\times)^{n-k} = \mathbf{T}^k \times \mathbf{R}^{n-k}.$$

By convention, we use the standard basis for $\sigma = \{0\}$, i.e., $e_i^{\{0\}} = e_i$. Given two cones $\sigma, \sigma' \in \Xi$, let $A^{\sigma, \sigma'} \in \text{GL}(n, \mathbf{Z})$ be the coordinate change matrix from basis $\{e_i^{\sigma'}\}$ to basis $\{e_i^\sigma\}$. Explicitly, if $M^\sigma \in \text{GL}(n, \mathbf{Z})$ is the matrix whose columns are given by e_i^σ , then $A^{\sigma, \sigma'} = (M^\sigma)^{-1}M^{\sigma'}$. We obviously have $A^{\sigma, \sigma'}A^{\sigma', \sigma''} = A^{\sigma, \sigma''}$ and $(A^{\sigma, \sigma'})^{-1} = A^{\sigma', \sigma}$. It is straightforward to check that the corresponding maps $\phi_{A^{\sigma, \sigma'}}$ extend to boundary strata as summarized in the following lemma.

Lemma 3.2.1

Let τ be a face of $\sigma \in \Xi$. Then the map $\phi_{A^{\sigma, \tau}} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ extends to an open embedding $\phi_{\sigma, \tau} = \phi_{A^{\sigma, \tau}} : U_\tau \rightarrow U_\sigma$. We denote the image by $U_\sigma^\tau \subset U_\sigma$.

Let $\sigma, \sigma' \in \Xi$ two cones with $\tau = \sigma \cap \sigma' \neq \emptyset$. Then the map $\phi_{A^{\sigma, \sigma'}} : \mathbf{R}^n \rightarrow \mathbf{R}^n$ extends to a homeomorphism $\phi_{\sigma, \sigma'} = \phi_{A^{\sigma, \sigma'}} : U_{\sigma'}^\tau \rightarrow U_\sigma^\tau$. Moreover, $\phi_{\sigma, \sigma'} \circ \phi_{\sigma', \sigma''} = \phi_{\sigma, \sigma''}$ whenever defined.

Exercise 3.2.2

Prove the preceding lemma (e.g. using Exercise 3.1.2).

Definition 3.2.3

Let Ξ be a unimodular pointed fan in \mathbf{R}^n . The *tropical smooth toric variety* associated to Ξ is the topological space

$$X_\Xi := \bigcup_{\sigma \in \Xi} U_\sigma / \sim,$$

where \sim is given by $x \sim \phi_{\sigma, \sigma'}(x)$ for all $\sigma, \sigma' \in \Xi$, $x \in U_{\sigma'}^\tau$.

The following properties follow easily from the previous discussion.

Proposition 3.2.4

Let X_{Ξ} be a tropical smooth toric variety. Then the following properties hold.

- (a) $X(\Xi)$ contains $\mathbf{R}^n = U_{\{0\}}$ as an open dense subset and addition on \mathbf{R}^n extends to an action of \mathbf{R}^n on $X(\Xi)$,
- (b) The orbits of the tropical torus action are in bijection to the cones of Ξ . For each $\sigma \in \Xi$ of dimension k the corresponding orbit is given by

$$\mathbf{R}_{\sigma} = \{-\infty\} \times \mathbf{R}^{n-k} \subset \mathbf{T}^k \times \mathbf{R}^{n-k} = U_{\sigma}, \quad (3.1)$$

or rather, its image in X_{Ξ} . It can be canonically identified with the quotient $\mathbf{R}_{\sigma} \cong \mathbf{R}^n / L(\sigma)$.

- (c) Let X_{τ} denote the closure of \mathbf{R}_{τ} in X_{Ξ} , $\tau \in \Xi$. Then X_{τ} is the tropical smooth toric variety associated to $\text{Star}_{\Xi}(\tau)$ and can be written as disjoint union of orbits

$$X_{\tau} = X_{\text{Star}_{\Xi}(\tau)} = \bigsqcup_{\tau \subset \sigma} \mathbf{R}_{\sigma}.$$

- (d) Let $x \in \mathbf{R}^n$ be a point, and let $v \in \mathbf{R}^n$ be a direction vector. Then $x + \lambda v$ converges for $\lambda \in \mathbf{R}, \lambda \rightarrow \infty$ to a point \bar{x} if and only if v is contained in the support of Ξ . Moreover, if v lies in the relative interior of σ , then $\bar{x} \in \mathbf{R}_{\sigma}$ and is equal to the image of x under the projection $\pi_{\sigma} : \mathbf{R}^n \rightarrow \mathbf{R}^n / L(\sigma) \cong \mathbf{R}_{\sigma}$ (using the canonical identification mentioned previously).
- (e) The topological space X_{Ξ} is compact if and only if the fan Ξ is complete, i.e., $|\Xi| = \mathbf{R}^n$.

There is a corresponding complex smooth toric variety $\mathbf{C}X(\Xi)$ which is obtained from exact same gluing procedure, but using complex monomial maps on $(\mathbf{C}^{\times})^n$ instead. Properties (a), (b), (c), (e) can be literally translated to the complex case (replacing \mathbf{R}, \mathbf{T} by $\mathbf{C}^{\times}, \mathbf{C}$). Property (d) reflects the corresponding statement about limits of (translations of) one-parameter subgroups in $(\mathbf{C}^{\times})^n$ (which are of the form $(z_1 t^{\nu_1}, \dots, z_n t^{\nu_n})$ with $z \in (\mathbf{C}^{\times})^n, t \in \mathbf{C}^{\times}, \nu \in \mathbf{Z}^n$).

Remark 3.2.5

Property (d) together with the requirement that the embeddings $\mathbf{R}_\sigma \hookrightarrow X_\Xi$ are continuous fixes the topology on X_Ξ uniquely. Hence an alternative construction of X_Ξ could be based on defining a topology on $\bigsqcup_{\sigma \in \Xi} \mathbf{R}/L(\sigma)$ satisfying these two properties. Note that this description extends to arbitrary (not necessarily unimodular) pointed fans Ξ , providing (not necessarily smooth) tropical toric varieties. Proposition 3.2.4 still holds for such general toric varieties, except for the description of orbits given in Equation (3.2.4) (in the sense that the open charts U_σ are no longer of the form $\mathbf{T}^k \times \mathbf{R}^{n-k}$ in general). The description of toric varieties in terms of semi-group homomorphisms (see for example [Ful93]) can also be immediately “tropicalized”.

For later use, we extend the discussion of limit points of rays to arbitrary polyhedra $P \subset \mathbf{R}^n$. Recall that the *recession cone* $\text{rc}(P)$ of P is the set of direction vectors of all rays contained in P .

Proposition 3.2.6

Let X_Ξ be a tropical toric variety, $\sigma \in \Xi$ a cone and $P \subset \mathbf{R}^n$ a polyhedron. Let \bar{P} denote the closure of P in X_Ξ . Then the following holds.

- (a) $\bar{P} \cap \mathbf{R}_\sigma = \emptyset \iff \text{rc}(P) \cap \sigma^\square = \emptyset$.
- (b) If $\bar{P} \cap \mathbf{R}_\sigma \neq \emptyset$, then $\bar{P} \cap \mathbf{R}_\sigma = \pi_\sigma(P)$. In particular, $\bar{P} \cap \mathbf{R}_\sigma$ is a rational polyhedron in \mathbf{R}_σ and $\dim(P) = \dim(\bar{P} \cap \mathbf{R}_\sigma) + \dim(L(P) \cap L(\sigma))$.

Example 3.2.7

Before treating “serious” examples of tropical toric varieties in the next section, let us look back at tropical affine space from Section 3.1. Let $\sigma_I, I \subset \{1, \dots, n\}$ be the cones generated by the negative standard basis vectors $-e_i, i \in I$ and let Ξ be the fan formed by all such cones. Then obviously the tropical toric variety X_Ξ is canonically isomorphic to $U_{\sigma_{[n]}} = \mathbf{T}^n$. Indeed, choosing the lattice basis e_1, \dots, e_n globally (i.e., $M^\sigma = \text{Id}$ for all σ), all gluing maps are identity maps on the subsets $U_{\sigma_I} \subset \mathbf{T}^n$. Moreover, the torus orbit \mathbf{R}_{σ_I} agrees with the sedentarity stratum \mathbf{R}_I , and the closure X_{σ_I} is equal to \mathbf{T}_I .

Exercise 3.2.8

Prove Proposition 3.2.4, for example, by copying the proofs of the corresponding facts from your favorite textbook on classical toric varieties (e.g. [Ful93]).

Exercise 3.2.9

Prove Proposition 3.2.6.

3.3 Tropical projective space

The most basic compact toric variety is, of course, projective space. We start with the symmetric description of \mathbf{TP}^n as a quotient and then recover the fan by using property (d) above.

Definition 3.3.1

Tropical projective space is

$$\mathbf{TP}^n := \mathbf{T}^{n+1} \setminus \{-\infty\} / \sim,$$

where the equivalence relation \sim is given by $x \sim “\lambda \cdot x”$ for all $\lambda \in \mathbf{T}^\times$. In other words,

$$(x_0, \dots, x_n) \sim (x_0 + \lambda, \dots, x_n + \lambda)$$

for all $\lambda \in \mathbf{R}$. We equip \mathbf{TP}^n with the quotient topology.

As usual, projective coordinates are denoted by $(x_0 : \dots : x_n)$. Projective space \mathbf{TP}^n can be covered by affine patches $U_i = \{x \in \mathbf{TP}^n : x_i \neq -\infty\}$ which can be identified with \mathbf{T}^n via

$$(x_0 : \dots : x_n) \mapsto (x_0 - x_i, \dots, x_{i-1} - x_i, x_{i+1} - x_i, \dots, x_n - x_i).$$

On the overlaps $U_i \cap U_j$ we get induced maps glueing two copies of \mathbf{T}^n via

$$\begin{aligned} \phi_{ij} : (y_0, \dots, y_{i-1}, y_{i+1}, \dots, y_n) \mapsto \\ (y_0 - y_j, \dots, y_{i-1} - y_j, -y_j, y_{i+1} - y_j, \dots, \widehat{y_j - y_j}, \dots, y_n - y_j). \end{aligned} \quad (3.2)$$

Thus \mathbf{TP}^n is obtained by glueing $n + 1$ copies of \mathbf{T}^n along these integer linear transformations (whenever defined). The gluing process is visualized for the tropical projective plane \mathbf{TP}^2 in Figure 3.5. Note that \mathbf{TP}^n contains




Figure 3.5: Glueing \mathbf{TP}^2 from its three affine patches

$\mathbf{R}^{n+1}/\mathbf{R}(1, \dots, 1) \cong \mathbf{R}^n$ as an open dense subset and this torus acts on \mathbf{TP}^n by coordinatewise addition. The orbits of this action are of the form

$$\mathbf{R}_I = \{(x_0 : \dots : x_n) \in \mathbf{TP}^n : x_i = -\infty \forall i \in I, x_i \neq -\infty \forall i \notin I\}$$

with $I \subsetneq \{0, \dots, n\}$.

Let us review the construction in the language of toric variety. For any proper subset $I \subsetneq \{0, \dots, n\}$, we define $\sigma_I \subset \mathbf{R}^n$ as the cone generated by the vector $-e_i, i \in I$. Here, e_1, \dots, e_n denotes the standard basis of \mathbf{R}^n and $-e_0 = e_1 + \dots + e_n = (1, \dots, 1)$. Let Ξ be the fan consisting of all such cones. We denote the maximal cones by $\sigma_i := \sigma_{\{0, \dots, \hat{i}, \dots, n\}}, i = 0, \dots, n$. For each such cone, we take as lattice basis the one given by $e_0, \dots, \hat{e}_i, \dots, e_n$. Then it is straightforward to check that linear maps described by A^{σ_j, σ_i} are equal to the gluing maps ϕ_{ij} from above. Indeed, since $\sum_{k=0}^n y_k e_k = 0$, by setting $y_i = 0$ we change coordinates via the computation

$$\sum_{k \neq i} y_k e_k = \sum_{k=0}^n y_k e_k = \sum_{k=0}^n (y_k - y_j) e_k = \sum_{k \neq j} (y_k - y_j) e_k,$$

which is in agreement with ϕ_{ij} . Hence $X_\Xi = \mathbf{TP}^n$. Figure 3.6 shows the fan and orbit stratification for \mathbf{TP}^2 .

Remark 3.3.2

Note that the fans we associate to tropical toric varieties are the reflection at




Figure 3.6: The stratification and the fan of \mathbf{TP}^2

the origin of fans commonly used in classical toric geometry. This because it seems more natural tropically to follow v in positive direction, i.e. $\lambda \rightarrow \infty$, whereas in the classical case one usually performs $t \rightarrow 0$ (for $t \in \mathbf{C}^\times$).

By Proposition 3.2.4 part (d), the torus $\mathbf{R}^n \subset \mathbf{TP}^n$ is equipped with $n + 1$ *divisorial directions* $(-1, \dots, 0), (0, -1, \dots, 0), \dots, (0, \dots, -1), (1, \dots, 1)$. In homogeneous coordinates, we can rewrite them more symmetrically as $(-1 : \dots : 0), (0 : -1 : \dots : 0), \dots, (0 : \dots : -1)$. Among all rays in \mathbf{R}^n , only those whose direction vector is one of the divisorial directions end up at point in an $(n - 1)$ -dimensional stratum of $\mathbf{TP}^n \setminus \mathbf{R}^n$. All other rays have limit points in strata of higher codimension. Examples for \mathbf{TP}^2 are given in Figure 3.7.

Note, in particular, that our illustration of \mathbf{TP}^2 as a triangle, while reflecting correctly the stratification into torus orbits, does not exhibit the metric properties of \mathbf{TP}^n : Of course, the points on the boundary of the triangle are infinitely far away from all points in the interior, and points in the interior of edges can only be reached via the corresponding divisorial direction. Note also that the divisorial directions showed up before, for example, in our treatment of non-degenerated planar curves in Section 2.5. In the same way, \mathbf{TP}^n is topologically an n -simplex (with matching stratifications). The $n + 1$ divisorial directions are in correspondence to the $n + 1$ maximal faces of the simplex (see Figure 3.8).




Figure 3.7: Some rays and their limit points in \mathbf{TP}^2




Figure 3.8: \mathbf{TP}^3 and its distinguished directions

3.4 Projective hypersurfaces

In chapter 2, we used tropical Laurent polynomials to describe hypersurfaces in \mathbf{R}^n . We will now see how a homogeneous tropical polynomial describes a hypersurface in \mathbf{TP}^n . Let us start with \mathbf{T}^n first. Let

$$f(x) = \sum_{j \in \mathbf{N}^n} a_j x^j$$

be a tropical polynomial in n variables (with multi-index notation). The coefficients $a_j \in \mathbf{T}$ are tropical numbers and the set $A = \{j : a_j \neq -\infty\}$ is finite (and non-empty). As all exponents are positive now, this describes a function

$$f : \mathbf{T}^n \rightarrow \mathbf{T},$$

and similar to the case of \mathbf{R}^n , the locus of points where two terms in f attain the maximum simultaneously is independent of the polynomial representation.

Definition 3.4.1

The *affine hypersurface* $V(f) \subset \mathbf{T}^n$ given by the tropical polynomial $f \in \mathbf{T}[x_1, \dots, x_n]$ is (as a set)

$$V(f) := \{x \in \mathbf{T}^n : \exists i \neq j \in \mathbf{N}^n \text{ such that } f(x) = a_i x^i = a_j x^j\}. \quad (3.3)$$

Note that since we allow any $i, j \in \mathbf{N}^n$, the hypersurface $V(f)$ contains the “honest” zero-set $\{x : f(x) = -\infty\}$ of f , a possibly non-empty set now.

Clearly, $V(f) \cap \mathbf{R}^n$ is equal to the very affine hypersurface $V(f|_{\mathbf{R}^n})$. Since $V(f)$ is closed, the closure of $V(f|_{\mathbf{R}^n})$ is also contained in $V(f)$. We denote it by $V(f)^{\text{fin}}$ and call it the *finite part* or *sedentarity zero part* of $V(f)$. Beyond $V(f)^{\text{fin}}$, we claim that $V(f)$ consists only of coordinate hyperplanes

$$H_i := V(x_i) = \{x \in \mathbf{T}^n : x_i = -\infty\}.$$

Let $\text{ord}_{x_i}(f) \geq 0$ denote the maximal integer k such that x_i divides f (i.e. f can be written as $f = x_i^k f'$ for some other tropical polynomial f'). Note that k only depends on the Newton polytope of f and hence on the function f (not on the polynomial representation). We say f is *monomialfree* if $\text{ord}_{x_i}(f) = 0$ for all $i = 1, \dots, n$ (equivalently, $x_i \nmid f$).

Theorem 3.4.2

Let $f \in \mathbf{T}[x_1, \dots, x_n]$ be a non-zero tropical polynomial. Then

$$V(f) = V(f)^{\text{fin}} \cup H_{i_1} \cup \dots \cup H_{i_k},$$

where $i_1, \dots, i_k \in \{1, \dots, n\}$ are exactly those indices for which $\text{ord}_{x_i}(f) > 0$.

Proof. Given two polynomials $g, h \in \mathbf{T}[x_1, \dots, x_n]$, we can check easily that $V("gh") = V(g) \cup V(h)$. Let us write $f = "x^j g"$, where $j = (j_1, \dots, j_n) \in \mathbf{N}^n$ is the exponent vector of a monomial and g is monomialfree. Then $V(f) = V(x^j) \cup V(g)$. Clearly, $V(x^j) = \bigcup_{i: j_i > 0} H_i$. Moreover, note that $f|_{\mathbf{R}^n}$ and $g|_{\mathbf{R}^n}$ only differ by the integer linear function $x \mapsto jx$ and thus $V(f|_{\mathbf{R}^n}) = V(g|_{\mathbf{R}^n})$ and $V(f)^{\text{fin}} = V(g)^{\text{fin}}$. So it suffices to prove the statement in the monomialfree case.

So from now on we assume f is monomialfree and it remains to show that $V(f) = V(f)^{\text{fin}}$. The inclusion " \supset " is obvious since $V(f)$ is closed and $V(f|_{\mathbf{R}^n}) \subset V(f)$. For the other inclusion, let $p \in V(f)$ be point, say, in the stratum $\mathbf{R}_I \subset \mathbf{T}^n$. Without loss of generality, we may assume $I = \{1, \dots, s\}$, $s = |I|$. Let us first consider the case $f(p) \neq -\infty$. Then there exist two terms " $a_j x^j$ ", " $a_{j'} x^{j'}$ " attaining the (finite) maximum at p . In particular, this implies $j_i = j'_i = 0$ for all $i = 1, \dots, s$. Considering the points $p(\lambda) := (-\lambda, \dots, -\lambda, p_{s+1}, \dots, p_n)$, we see that for sufficiently large $\lambda \in \mathbf{R}$

$$f(p(\lambda)) = "a_j p(\lambda)^j" = "a_{j'} p(\lambda)^{j'}" = f(p).$$

Hence $p(\lambda) \in V(f|_{\mathbf{R}^n})$ for large λ and therefore $p \in V(f)^{\text{fin}}$.

Let us consider the case $f(p) = \infty$. We set $\mathbf{T}_p^s := \mathbf{T}^s \times \{(p_{s+1}, \dots, p_n)\} \subset \mathbf{T}^n$ and consider the restriction

$$f_p := f|_{\mathbf{T}_p^s} = f(x_i = p_i : i = s+1, \dots, n) \in \mathbf{T}[x_1, \dots, x_s].$$

Obviously, f_p is monomialfree, $f_p(p) = f(p) = -\infty$ and

$$V(f_p) \times \{(p_{s+1}, \dots, p_n)\} \subset V(f) \cap \mathbf{T}_p^s.$$

Hence the claim follows from the following lemma. □

Lemma 3.4.3

Let $f \in \mathbf{T}[x_1, \dots, x_n]$ be a non-zero monomialfree tropical polynomial such that $f((-\infty)^n) = -\infty$, $(-\infty)^n = (-\infty, \dots, -\infty) \in \mathbf{T}^n$. Then $(-\infty)^n \in V(f)^{\text{fin}}$.

Proof. Let $P := \text{NP}(f)$ be the Newton polytope of f and let $\mathcal{N}(P)$ be its normal fan. By our assumptions, P lies in the positive orthant of \mathbf{R}^n but is touching all the hyperplanes $\{j_i = 0\}, i = 1, \dots, n$. Consider the open negative orthant $N := (\mathbf{R}_{<0})^n$. We want to show that N intersects one of the non-maximal cells of $\mathcal{N}(P)$. Assuming the contrary, it follows that N is contained in a n -cone of $\mathcal{N}(P)$ corresponding to a vertex V of P . Choosing a sequence of rays in N converging to $-e_i$, it follows by continuity that V is contained in $\{j_i = 0\}$. Running through all $i = 1, \dots, n$, we find $V = (0, \dots, 0) \in P$, which is a contradiction to $f((-\infty)^n) = -\infty$. Hence N intersects a non-maximal cone σ_D of $\mathcal{N}(P)$ dual to a positive-dimensional face F of P . Let D be a positive-dimensional cell of $\text{SD}(f)$ such that $D^\square \subset F^\square$. Let $\sigma_D \subset V(f|_{\mathbf{R}^n})$ be the cell of $\mathcal{S}(f)^{(n-1)}$ dual to D . By Exercise 2.3.13 we have $\text{rc}(\sigma_D) = \sigma_F$, which intersects N . Hence by Proposition 3.2.6 we see that $(-\infty)^n \in \overline{\sigma_D} \subset \mathbf{T}^n$, which proves the claim. \square

So far, we described $V(f)$ as a set. Clearly, as in the case of \mathbf{R}^n , we would like to add a weighted polyhedral structure to $V(f)$. Postponing the definition of polyhedral complexes in the compactified setting to Chapter 5, we restrict ourselves here to an adhoc description of the weight function based on the description of $V(f)$ as a set.

Definition 3.4.4

Let X_Ξ be a tropical toric variety with torus \mathbf{R}^n . A *tropical divisor* or *tropical $n - 1$ -cycle* D is a formal sum of the form

$$D = D_0 + k_1 D_1 + \dots + k_l D_l,$$

where

- D_0 is a (single) $n - 1$ -cycle in \mathbf{R}^n ,
- $k_1, \dots, k_n \in \mathbf{Z}$,
- $D_i = X_{\rho_i}$, where ρ_1, \dots, ρ_l denote the rays of Ξ .

The *support* of D is the union of $\overline{|D_0|}$ with D_i for all i such that $k_i \neq 0$.

Definition 3.4.5

The *affine hypersurface* $V(f) \subset \mathbf{T}^n$ defined by the tropical polynomial $f \in$

$\mathbf{T}[x_1, \dots, x_n]$ is (as a divisor)

$$V(f) := V(f|_{\mathbf{R}^n}) + \sum_{i=1}^n \text{ord}_{x_i}(f)H_i.$$

By Theorem 3.4.2 the support of $V(f)$ agrees with the set-theoretic definition in Equation (3.3).

Once again, given our treatment of the case \mathbf{T}^n , we can easily extrapolate to arbitrary tropical toric varieties. Since the procedure is completely analogous to the case of classical toric varieties, we will only discuss the case \mathbf{TP}^n in more details and restrict ourselves to a brief outline of the general case in Remark 3.4.8.

Let

$$F(x) = \sum_{j \in \mathbf{N}^{n+1}} a_j x^j$$

be a tropical homogeneous polynomial of degree $d \in \mathbf{Z}$ in $n + 1$ variables. In more details, the coefficients $a_j \in \mathbf{T}$ are tropical numbers, the set $A = \{j : a_j \neq -\infty\}$ is finite and for all $j \in A$ we have $|j| = j_0 + \dots + j_n = d$. In projective coordinates, this does not quite define a function to \mathbf{T} as we have $F(\lambda \cdot x) = \lambda^d F(x) = \lambda d + F(x)$. However, the set

$$V(F) = \{x \in \mathbf{TP}^n : \exists i \neq j \in \mathbf{N}^n \text{ such that } F(x) = a_i x^i = a_j x^j\} \quad (3.4)$$

is still well-defined. Fixing an affine chart $U_i \cong \mathbf{T}^n \subset \mathbf{TP}^n$ corresponds to the dehomogenization

$$f_i = F(y_0, \dots, 1, \dots, y_n) \in \mathbf{T}[y_0, \dots, \hat{y}_i, \dots, y_n].$$

Under the coordinate change (3.2), these polynomial follow the rule

$$f_i(y) - f_j(\phi_{ij}(y)) = dy_j.$$

In particular, since the difference is an integer linear function very affine hypersurface $V(f_i|_{\mathbf{R}^n}) \subset \mathbf{R}^n \subset U_i$ is independent on the affine chart (under the coordinate changes ϕ_{ij}). We call its closure in \mathbf{TP}^n the *finite part* or *sedentarity zero part* of $V(F)$, denoted by $V(F)^{\text{fin}}$.

Definition 3.4.6

The *projective hypersurface* $V(F) \subset \mathbf{TP}^n$ defined by the tropical homogeneous polynomial F is the tropical divisor

$$V(F) := V(F)^{\text{fin}} + \sum_{i=1}^n \text{ord}_{x_i}(F)H_i,$$

where $H_i = \{x \in \mathbf{TP}^n : x_i = -\infty\}$ is the i -th coordinate hyperplane.

Remark 3.4.7

Note that $\text{ord}_{x_k}(F) = \text{ord}_{y_k}(f_i)$ for all $i \neq k$. In particular, defining $V(F) \cap U_i$ in the obvious way, we have $V(F) \cap U_i = V(f_i)$ for all $i = 0, \dots, n$. Together with Theorem 3.4.2, this also shows that the support of $V(F)$ agrees with the set-theoretic definition in Equation (3.4). Moreover, we may define the dual subdivision $\text{SD}(F)$ of $\text{NP}(F)$ as before. The canonical identification $\text{NP}(F) \cong \text{NP}(f_i)$ for all i also identifies the subdivisions $\text{SD}(F) \cong \text{SD}(f_i)$. Hence we get a duality between $V(F)^{\text{fin}}$ and $\text{SD}(F)$ as in the very affine case.

Remark 3.4.8

The definitions can be immediately generalized to arbitrary tropical toric varieties following the standard representation of Cox rings for (classical) toric varieties. More precisely, let Ξ be a pointed fan in \mathbf{R}^n and let v_1, \dots, v_l be the primitive generators of its rays. Let A be the kernel of the map $\mathbf{Z}^l \rightarrow \mathbf{Z}^n, e_i \mapsto v_i$. Set $B = A^*$ and let $\delta : \mathbf{Z}^l \rightarrow B$ be the map dual to $A \hookrightarrow \mathbf{Z}^l$. The *toric degree* of a monomial x^j in the variables x_1, \dots, x_l is given by $\delta(j) \in B$. Fixing $d \in B$, a *homogeneous polynomial of degree d* is a tropical sum

$$F(x) = \left\langle \sum_{j \in \mathbf{N}^l} a_j x^j \right\rangle \in \mathbf{T}[x_1, \dots, x_l]$$

such that $A = \{j : a_j \neq -\infty\}$ is finite and for all $j \in A$ we have $\delta(j) = d$. Let $c : \mathbf{Z}^n \rightarrow \mathbf{Z}^l$ be the map dual to $\mathbf{Z}^l \rightarrow \mathbf{Z}^n$ from above. Fixing j_0 with $\delta(j_0) = d$, we can define the *dehomogenization*

$$f(x) = \left\langle \sum_{i \in \mathbf{N}^n} a_{c(i)+j_0} y^i \right\rangle \in \mathbf{T}[y_1^\pm, \dots, y_n^\pm].$$

Choosing different j_0 only results in multiplying f by an monomial y^i , so the very affine hypersurface $V(f) \in \mathbf{R}^n$ is well-defined. We call its closure

in \mathbf{TP}^n the *finite part* or *sedentarity zero part* of $V(F)$, denoted by $V(F)^{\text{fin}}$. Finally, we define the *toric hypersurface* $V(F) \subset X_{\Xi}$ as the tropical divisor

$$V(F) := V(F)^{\text{fin}} + \sum_{i=1}^n \text{ord}_{x_i}(F)D_i,$$

where D_i is the toric boundary divisor spanned by the vector v_i . Again, for a given affine chart $U_i \cong \mathbf{T}^n$, there exists a unique choice of j_0 such that $V(F) \cap U_i = V(\tilde{f})$, where \tilde{f} is obtained from f via the corresponding coordinate change. The statements about the dual subdivision $\text{SD}(F)$ of $\text{NP}(F)$ and its duality with $V(F)^{\text{fin}}$ extend without difficulties.

Proposition 3.4.9

Let X_{Ξ} be a tropical toric variety and let

$$D = D_0 + k_1D_1 + \cdots + k_lD_l$$

be a tropical divisor in the notation of Definition 3.4.4. Assume that all weights of D_0 are positive and that $k_i \geq 0$ for all $i = 1, \dots, l$. Then there exists a Ξ -homogeneous polynomial F such that $V(F) = D$.

Proof. We content ourselves to discuss the projective case, the general case can be easily adapted. First, note that by Theorem 2.4.10 there exists a polynomial $g \in \mathbf{T}[y_1, \dots, y_n]$ such that $V(g) = D_0$, and we may assume that g is monomialfree. Let $G \in \mathbf{T}[x_0, \dots, x_n]$ denote the homogenization of g (still monomialfree). Then the polynomial $F = x^j G$ with $j_i = k_i$ satisfies $V(F) = D$. \square

Let us close this section with a short digression on the homotopy type of tropical hypersurfaces.

Proposition 3.4.10

Let $V(f)$ be a hypersurface in \mathbf{R}^n whose Newton polytope $\text{NP}(f)$ is full-dimensional. Then $V(f)$ is homotopy-equivalent to a bouquet of $k(n-1)$ -spheres. Here k is the number of vertices of the dual subdivision $\text{SD}(f)$ which are interior points of $\text{NP}(f)$.

Proof. Recall that each vertex of $\text{SD}(f)$ corresponds to a connected component of $\mathbf{R}^n \setminus V(f)$. For vertices in the boundary of $\text{NP}(f)$, the boundary of

this component is contractible and the homotopy type after removing the component is unchanged. For vertices in the interior of $\text{NP}(f)$, the boundary of the corresponding component is homeomorphic to an $(n-1)$ -sphere. Up to homotopy equivalence, removing such a component is equivalent to removing a point of \mathbf{R}^n . But \mathbf{R}^n minus k points is a bouquet of k $(n-1)$ -spheres, as claimed. \square

Remark 3.4.11

Assume $\text{NP}(f)$ is not full-dimensional, but generates the affine subspace $A \subset \mathbf{R}^n$ in dual space. Then $V(f)$ is translation-invariant along A^\perp and the projection to \mathbf{R}^n/A^\perp provides a contraction of $V(f)$ to a hypersurface in a smaller-dimensional space whose Newton polytope is full-dimensional (namely $\text{NP}(f) \subset A$). Therefore the proposition is basically valid in this case, too. The only difference is that now we count vertices in the relative interior of $\text{NP}(f)$ and $V(f)$ is a bouquet of $(m-1)$ -spheres, where m is the dimension of $\text{NP}(f)$.

When considering hypersurfaces in compact toric varieties, the statement remains nearly unchanged. All we have to do is to use the right notion of being an interior point now. We state the result for projective hypersurfaces here, although the generalization to other (compact) toric varieties is straightforward.

Proposition 3.4.12

Let F be a homogeneous polynomial of degree d such that $x_i \nmid F$ for all i and define Δ_d to be the simplex obtained as convex hull of $(d, 0, \dots, 0), \dots, (0, \dots, d) \in \mathbf{N}^{n+1}$. Let k denote the number of vertices of $\text{SD}(F)$ contained in Δ_d^\square . Then the projective hypersurface $V(F) \subset \mathbf{TP}^n$ is homotopy-equivalent to a bouquet of k $(n-1)$ -spheres.

Proof. Note that \mathbf{TP}^n is homeomorphic to a closed ball of dimension n . As in the proof for \mathbf{R}^n , each vertex of $\text{SD}(f)$ corresponds to a connected component of $\mathbf{TP}^n \setminus V(f)$. Again, these components differ, depending on whether the vertex is in the boundary of Δ_d or not. Accordingly, the boundary of the connected component is contractible or homeomorphic to an $(n-1)$ -sphere. Thus again, $V(F)$ is homotopy equivalent to an n -ball minus k interior points, which is a bouquet of k $(n-1)$ -spheres. \square




Figure 3.9: The homotopy type of a plane projective cubic ($k = 1$)

Exercise 3.4.13

Let $f : \mathbf{T}^n \rightarrow \mathbf{T}$ be a tropical polynomial. Show that $V(f) \subset \mathbf{T}^n$ is closed. Conclude that projective hypersurfaces $V(F) \subset \mathbf{TP}^n$ are closed.

Exercise 3.4.14

Let $f = \sum_{j \in A} a_j x^j$ be a tropical polynomial. For $I \subset \{1, \dots, n\}$, set $A_I \subset A$ to be the subset of exponent vectors $j = (j_1, \dots, j_n)$ with $j_i = 0$ whenever $i \in I$. We call

$$f_I := \sum_{j \in A_I} a_j x^j \in \mathbf{T}[x_i : i \notin I]$$

the *truncation* of f to I . Show that $V(f) \cap \mathbf{T}_I = V(f_I)$, including the special case $\mathbf{T}_I \in V(f) \Leftrightarrow f_I = -\infty$, and an agreement of the weights otherwise.

3.5 Projective curves

A tropical homogeneous polynomial in 3 variables describes a tropical curve in \mathbf{TP}^2 . Such a curve always splits into its sedentarity zero part and possibly a union of coordinate lines. We focus our attention to sedentarity zero curves here. Such curves may equivalently be described by a not necessarily homogeneous polynomial f in 2 variables which is not divisible by monomials. Here we work with the fixed (very) affine chart $\mathbf{R}^2 = \{0\} \times \mathbf{R}^2 \subset \mathbf{TP}^2$. Then f defines the sedentarity zero projective curve $\overline{V(f)} \subset \mathbf{TP}^2$. In consistency with section 2.5, the degree of $V(f)$ is the smallest number such that




Figure 3.10: A generic and a non-generic projective conic in \mathbf{TP}^2

$\text{NP}(f) \subseteq \Delta_d$. Here Δ_d is the d -fold standard simplex, i.e. the convex hull of the points $(0, 0), (d, 0), (0, d)$. In section 2.5, we mostly considered non-degenerated curves with $\text{NP}(f) = \Delta_d$. In projective world, they are distinguished by not containing any of the torus fixed points $p_0 = (0, -\infty, -\infty)$, $p_1 = (-\infty, 0, -\infty)$, $p_2 = (-\infty, -\infty, 0) \in \mathbf{TP}^2$. Vice versa, any curve $V(f)$ with $\text{NP}(f) \neq \Delta_d$ contains at least one of these points (see Figure 3.10).

Our goal is to extend Bézout theorem from Section 2.5 to projective curves. To do so, we need to define intersection multiplicities for curves intersecting at torus fixed points. We will do this in the following adhoc manner following [BS11]. Let X, Y be two curves in \mathbf{T}^2 and let ρ_1, ρ_2 be (unbounded) rays of X_1 resp. X_2 . These rays intersect at $(-\infty, -\infty)$ if and only if both primitive generators v_1, v_2 have only negative entries, say $v_i = (x_i, y_i), x_i, y_i \in \mathbf{Z}_{<0}$ (see Figure 3.11). In this case, the contribution of the intersection of these two rays to the intersection multiplicity at $(-\infty, -\infty)$ is defined to be

$$\text{mult}(\rho_1, \rho_2) = \omega(\rho_1)\omega(\rho_2) \min\{x_1y_2, x_2y_1\}.$$

The general definition is then just a summation over all such pairs of rays.

Definition 3.5.1




Figure 3.11: Two rays intersecting at $(-\infty, -\infty)$

Let X, Y be two affine tropical curves in \mathbf{T}^2 . The *intersection multiplicity* of two X and Y at $(-\infty, -\infty)$ is

$$\text{mult}((-\infty, -\infty)) = \sum_{\rho_1, \rho_2} \omega(\rho_1)\omega(\rho_2) \min\{x_1y_2, x_2y_1\},$$

where the sum is taken over all pairs of rays ρ_1 and ρ_2 of X and Y , respectively, with primitive generators $(x_i, y_i) \in \mathbf{Z}_{<0}^2$, $i = 1, 2$.

Let $X, Y \subset \mathbf{TP}^2$ be two projective tropical curves which intersect transversally in \mathbf{R}^2 . For the three torus fix points p_0, p_1, p_2 , we use the multiplicity $\text{mult}((-\infty, -\infty))$ in the corresponding charts to define multiplicities $\text{mult}(p_0)$, $\text{mult}(p_1)$, $\text{mult}(p_2)$. The *intersection product* $X \cdot Y$ is the formal sum of points

$$X \cdot Y = \sum_{p \in \mathbf{R}^2} \text{mult}(p)p + \text{mult}(p_0)p_0 + \text{mult}(p_1)p_1 + \text{mult}(p_2)p_2,$$

where $\text{mult}(p)$ denotes intersection multiplicity in \mathbf{R}^2 from Section 2.5 (for the curves $X \cap \mathbf{R}^2$ and $Y \cap \mathbf{R}^2$).

Our justification for this ad-hoc definition is that Bézout’s theorem holds. In the course of the proof, we will deform the rays going to the torus fixed points to “non-degenerated” rays, providing more precise evidence as to why this is the correct definition (see Figure 3.12).

Theorem 3.5.2 (Tropical Bézout’s theorem)

Let X and Y be two tropical sedentarity zero curves in \mathbf{TP}^2 of degree d and e ,

respectively. Then

$$\deg(X \cdot Y) = \deg(X) \cdot \deg(Y)$$

holds.

Proof. First note that in the non-degenerated case, i.e. when both X and Y do not contain torus fixed points, the statement is already proven in Theorem 2.5.2.

In the degenerated case, let σ be a ray of X containing a torus fixed point. In the corresponding affine chart \mathbf{T}^2 , let $v = (x_1, y_1)$ be the primitive generator of this ray, $x_1, y_1 < 0$. Now we perform the following “deformation” of σ . We pick a point $p \in \sigma \cap \mathbf{R}^2$ such that the remaining unbounded part $p + \mathbf{R}_{\geq 0}v$ does not intersect Y . We replace this part of σ by the two rays $p + \mathbf{R}_{\geq 0}(-1, 0)$ and $p + \mathbf{R}_{\geq 0}(0, -1)$ with weights $\omega(\sigma)x_1$ and $\omega(\sigma)y_1$, respectively. By construction, this deformed curve X' is still balanced, hence a tropical curve. By Exercise 2.5.5, it has the same degree as X .

By choosing p close enough to $(-\infty, -\infty)$, we can ensure that the two new rays only intersect Y in rays also going to $(-\infty, -\infty)$. Then, each such ray τ of Y , with primitive generator (x_2, y_2) , intersects exactly one of the two rays. Moreover, by our assumption that $p + \mathbf{R}_{\geq 0}v$ does not intersect Y , it follows that each τ intersects the $(-1, 0)$ -ray only if $x_1y_2 \leq x_2y_1$, and vice versa for the $(0, -1)$ -ray. Hence the intersection multiplicity at this new intersection point is

$$\omega(\sigma)\omega(\tau) \min\{x_1y_2, x_2y_1\}.$$

In other words, the new finite intersection points of $X' \cdot Y$ exactly compensate the loss of intersection multiplicities at $(-\infty, -\infty)$, compared to $X \cdot Y$, and hence $\deg(X \cdot Y) = \deg(X' \cdot Y)$. Repeating this procedure for each “degenerated” ray of X and Y , we eventually end up with two non-degenerated tropical curves and the assertion follows. \square

Exercise 3.5.3

Let $X, Y \subset \mathbf{TP}^2$ be two tropical curves of sedentarity zero such that $X \cap \mathbf{R}^2$ and $Y \cap \mathbf{R}^2$ intersect transversally. Show that X and Y intersect only in \mathbf{R}^2 and some of the torus fixed points.

Exercise 3.5.4

Show that any two tropical projective curves $X, Y \subset \mathbf{TP}^2$ intersect.




Figure 3.12: Replacing a degenerated ray by two standard ends

Exercise 3.5.5

Let $(\mathbf{TP}^2)^*$ be the set of tropical lines in \mathbf{TP}^2 . Show that $(\mathbf{TP}^2)^*$ can be canonically identified with \mathbf{TP}^2 . We call $(\mathbf{TP}^2)^*$ the *dual projective plane*.

4 Tropical cycles in \mathbf{R}^n

In this chapter, we jump back to case of very affine tropical geometry in \mathbf{R}^n . We will define tropical cycles as balanced polyhedral sets of any codimension and introduce the concept of stable intersection of such cycles. We also extend the construction of associating a hypersurface to a tropical polynomial to arbitrary piecewise affine functions on tropical cycles. This is strongly related to the concept of tropical modifications, an important construction in tropical geometry which does not have a direct counterpart in classical geometry.

4.1 Polyhedral geometry dictionary II

A *polyhedral set* X (of pure dimension k) in \mathbf{R}^n is a finite union of polyhedra $X = \sigma_1 \cup \dots \cup \sigma_l$ (with $\dim(\sigma_i) = k$). Obviously, the support of a polyhedral complex \mathcal{X} is a polyhedral set. Vice versa, a *polyhedral structure* of X is a polyhedral complex \mathcal{X} such that $|\mathcal{X}| = X$. Let us collect a few useful facts concerning polyhedral sets. We use the following notations. Given an integer affine function $\kappa : \mathbf{R}^n \rightarrow \mathbf{R}$, the *halfspace complex* associated to κ is the polyhedral subdivision of \mathbf{R}^n

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

Given a polyhedral complex \mathcal{X} in \mathbf{R}^n and an arbitrary set $S \subset \mathbf{R}^n$, the *restriction of \mathcal{X} to S* is $\mathcal{X}|_S := \{\sigma \in \mathcal{X} : \sigma \subset S\}$. It is clear that $\mathcal{X}|_S$ is a polyhedral complex.

Proposition 4.1.1

The following statements hold true.

- (a) *Let $X = \sigma_1 \cup \dots \cup \sigma_l$ be a polyhedral set. Then there exists a polyhedral structure \mathcal{X} of X such that $\mathcal{X}|_{\sigma_i}$ is a subdivision of σ_i for all $i = 1, \dots, l$. In particular, any polyhedral set admits a polyhedral structure.*

(b) Let \mathcal{X} and \mathcal{Y} be two polyhedral complexes in \mathbf{R}^n and set $X = |\mathcal{X}|$, $Y = |\mathcal{Y}|$. Then there exists a polyhedral structure \mathcal{Z} of $X \cup Y$ such that $\mathcal{Z}|_X$ and $\mathcal{Z}|_Y$ are refinements of \mathcal{X} and \mathcal{Y} , respectively. In particular, if $X = Y$, then \mathcal{X} and \mathcal{Y} have a common refinement.

(c) In both previous items, the promised polyhedral structure can be completed to a polyhedral subdivision of \mathbf{R}^n .

Proof. To prove (a), let $\kappa_1, \dots, \kappa_m : \mathbf{R}^n \rightarrow \mathbf{R}$ be a collection of integer affine forms such that each σ_i , $i = 1, \dots, l$ can be represented by intersecting some of the halfspaces $H_{\kappa_1}, \dots, H_{\kappa_m}$. Consider the subdivision of \mathbf{R}^n

$$\mathcal{S} = \mathcal{G}_{\kappa_1} \cap \dots \cap \mathcal{G}_{\kappa_m}$$

(see Exercise 2.1.3). Then the polyhedral complex $\mathcal{X} = \mathcal{S}|_X$ satisfies the conditions in part (a).

To prove (b), apply the construction for part (a) to the union of polyhedra

$$X \cup Y = \bigcup_{\sigma \in \mathcal{X}} \sigma \cup \bigcup_{\sigma \in \mathcal{Y}} \sigma.$$

Part (b) is obvious, the completion is given by \mathcal{S} . □

In other words, polyhedral sets and polyhedral complexes up to common refinements are essentially the same thing.

Let us turn to fan sets now. A *fan set* is polyhedral set X such that for all $x \in X$ the ray $\mathbf{R}_{\geq 0}x$ is contained in X . It follows from Exercise 4.1.3 that such an X can be written as the union of polyhedral cones. A *fan structure* of X is a polyhedral fan \mathcal{X} such that $X = |\mathcal{X}|$. Clearly, Proposition 4.1.1 and its proof can be adapted such that all involved sets, polyhedra, complexes are fan sets, cones, fans, respectively. In particular, any fan set admits a fan structure and two such structures have a common (fan) refinement.

For later use, let us recall a theorem about simplicial and unimodular refinements of fans. A cone σ of dimension k is *simplicial* if it can be generated by k vectors v_1, \dots, v_k . This implies that σ is pointed. Moreover, σ is *unimodular* if it can be generated by a part of a lattice basis (equivalently, if $\langle v_1, \dots, v_k \rangle_{\mathbf{Z}} = L_{\mathbf{Z}}(\sigma)$). A fan \mathcal{F} is *simplicial* or *unimodular* if all its cones are simplicial or unimodular, respectively. A fan \mathcal{F} is *complete* if $|\mathcal{F}| = \mathbf{R}^n$.

Theorem 4.1.2

Let \mathcal{F} be a fan in \mathbf{R}^n . Then the following statements are true.

- (a) There exists a simplicial fan \mathcal{G} which is a refinement of \mathcal{F} containing all its simplicial cones.
- (b) There exists a unimodular fan \mathcal{G} which is a refinement of \mathcal{F} containing all its unimodular cones.
- (c) There exists a complete fan \mathcal{E} containing all cones of \mathcal{F} .

Proof. For (a) and (b), see for example [CLS11, Theorems 11.1.7 and 11.1.9]. For (c), see [Ewa96, Theorem III.2.8] (or [EI06]). \square

Exercise 4.1.3

Let $\sigma \in \mathbf{R}^n$ be a (rational) polyhedron. Show that $\mathbf{R}_{\geq 0}\sigma = \{\lambda x : \lambda \in \mathbf{R}_{\geq 0}, x \in \sigma\}$ is a (polyhedral) cone.

4.2 Tropical cycles and subspaces in \mathbf{R}^n

In chapter 2 we explained that a weighted polyhedral complex of pure dimension k in \mathbf{R}^n is called balanced if it satisfies the balancing condition

$$\sum_{\substack{\sigma \text{ facet} \\ \tau \subset \sigma}} \omega(\sigma)v_{\sigma/\tau} = 0 \pmod{L(\tau)}$$

for all $k-1$ cells τ (see Definition 2.4.6). Tropical hypersurfaces are examples of balanced polyhedral complexes of dimension $n-1$, but note that the definition makes sense in any codimension. Indeed, such balanced polyhedral complexes of arbitrary codimension are the basic objects in tropical geometry.

Compared to the case of hypersurfaces, we will slightly shift our point of view by focusing on the underlying support sets and not insisting on a fixed subdivision into cells. Here are the relevant definitions. Recall that for a weighted polyhedral complex \mathcal{X} , the support $|\mathcal{X}|$ is the union of cells of non-zero weight.

Definition 4.2.1

Two weighted polyhedral complexes \mathcal{X} and \mathcal{Y} in \mathbf{R}^n are *equivalent* if $|\mathcal{X}| = |\mathcal{Y}|$ and for any pair of facets $\sigma \in \mathcal{X}, \sigma' \in \mathcal{Y}$ with $\sigma^\square \cap \sigma'^\square \neq \emptyset$ the weights agree. A weighted polyhedral complex \mathcal{Z} is a (*weighted*) *refinement* of \mathcal{X} if $|\mathcal{Z}| = |\mathcal{X}|$, any $\sigma \in \mathcal{Z}, \sigma \subset |\mathcal{Z}|$, is contained in some $\sigma' \in \mathcal{X}$, and if σ is a facet, then $\omega_{\mathcal{Z}}(\sigma) = \omega_{\mathcal{X}}(\sigma')$.

Clearly, any (unweighted) polyhedral complex \mathcal{Z} such that $\mathcal{Z}|_{|\mathcal{X}|}$ is a refinement of $\mathcal{X}|_{|\mathcal{X}|}$ inherits a unique weight function such that \mathcal{Z} is a weighted refinement of \mathcal{X} . It is easy to check that a weighted refinement of \mathcal{X} is balanced if and only if \mathcal{X} is balanced (see Exercise 4.2.9). Moreover, using Proposition 4.1.1 we see that \mathcal{X} and \mathcal{Y} are equivalent if and only if they have common weighted refinement (see Exercise 4.2.10). Consequently, if the balancing condition holds for \mathcal{X} , then also for all equivalent complexes. This leads to the following definition.

Definition 4.2.2

A *tropical k -cycle* X in \mathbf{R}^n is an equivalence class of balanced polyhedral complexes in \mathbf{R}^n of pure dimension k . A *tropical subspace* is a tropical cycle that can be represented by a complex with only positive weights. A *tropical fan cycle* is a tropical cycle X that can be represented by a fan \mathcal{X} .

A reformulation emphasizing the underlying support set can be given as follows. Let $X \subset \mathbf{R}^n$ be a polyhedral set. A point $x \in X$ is called *generic* if there exists a polyhedral structure for X such that x is contained in the relative interior of a facet. Equivalently, there exists an affine subspace $A \subset \mathbf{R}^n$ and a neighbourhood $x \in U \subset \mathbf{R}^n$ such that $X \cap U = A \cap U$. The set of generic points is an open dense subset of X and is denoted by X^{gen} . A *weighted polyhedral set* is a polyhedral set X equipped with a locally constant function $\omega : X^{\text{gen}} \rightarrow \mathbf{Z} \setminus \{0\}$, called *weight function* of X .

Any polyhedral structure \mathcal{X} of X inherits a weight function for the facets by setting $\omega(\sigma) = \omega(x)$ for any $x \in \sigma^\square$. Clearly, two polyhedral structures equipped with the inherited weights are equivalent in the sense of definition 4.2.1. Vice versa, a balanced polyhedral complex \mathcal{X} also induces a weight function on $|\mathcal{X}|^{\text{gen}}$ by setting $\omega(x) = \omega(\sigma)$ for $x \in \sigma^\square$ and extending the function to all of $|\mathcal{X}|^{\text{gen}}$ (cf. Exercise 4.2.11). This leads to the following equivalent reformulation of Definition 4.2.2.

Definition 4.2.3

A tropical k -cycle X (or *balanced polyhedral subset*) in \mathbf{R}^n is a weighted polyhedral set in \mathbf{R}^n of pure dimension k such that for any polyhedral structure of X , equipped with the inherited weights, the balancing condition is satisfied. A *tropical subspace* is a tropical cycle with positive weight function. A *tropical fan cycle* is a tropical cycle X supported on a fan set.

In the following we will take the freedom to take both viewpoints simultaneously. For example, will regard a tropical cycle X as a subset of \mathbf{R}^n and use the notation $|X|$ only when we want to emphasize that we are forgetting the weight function.

The set of all k -cycles of \mathbf{R}^n is denoted by $Z_k(\mathbf{R}^n)$ and forms a group under taking unions and adding weights. That is to say, to obtain $X + Y$ we first equip $(X \cup Y)^{\text{gen}} \subset X^{\text{gen}} \cup Y^{\text{gen}}$ with the weight function $\omega_X(x) + \omega_Y(x)$ (for $x \in (X \cup Y)^{\text{gen}} \subset X^{\text{gen}} \cup Y^{\text{gen}}$), extending ω_X and ω_Y by zero if necessary, and secondly we restrict to the closure of the set of points with non-zero weight. In particular, $X + Y \subset X \cup Y$, but in general equality does not hold. For a subset $X \subset \mathbf{R}^n$, we denote by $Z_k(X)$ groups of tropical k -cycles whose support is contained in X .

Note that in tropical geometry the distinction between between algebraic subvarieties as honest geometric objects and algebraic cycles as formal sums of irreducible subvarieties does not really exist. In particular, our definition of tropical subspaces as “effective” tropical cycles is somewhat arbitrary. Rather, one should think of tropical cycles as a mixture between the two concepts, carrying features of honest geometric subobjects (weights already occur for hypersurfaces) as well as algebraic properties of cycles (such as they form a group). For later use, let us prove the following lemma.

Example 4.2.4

Let $A \subset \mathbf{R}^n$ be a rational affine subspace of dimension k . We denote by $[A]$ the k -cycle supported on A with constant weight function 1.

Lemma 4.2.5

Any k -cycle X in \mathbf{R}^n can be written as a difference $X = X_1 - X_2$ of tropical k -subspaces X_1, X_2 .

Proof. We set $X_2 = \sum_{\sigma: \omega(\sigma) < 0} -\omega(\sigma)[A_\sigma]$. Here, the sum runs through all facets of (a polyhedral structure of) X with negative weights and A_σ denotes

the affine space spanned by σ . By definition, X_2 and $X + X_2$ are tropical subspaces, and hence the claim follows. \square

Example 4.2.6

Choose non-negative integers $k \leq n$. Let e_1, \dots, e_n denote the standard basis of \mathbf{R}^n and set $e_0 = -\sum_{i=1}^n e_i = (-1, \dots, -1)$. For any subset $I \subset \{0, \dots, n\}$, let C_I denote the rational polyhedral cone in \mathbf{R}^n spanned by the vectors $-e_i, i \in I$. We consider the polyhedral fan

$$\mathcal{L}_k := \{C_I : I \subset \{0, \dots, n\}, |I| \leq k\}.$$

Equipped with trivial weights all equal to one, \mathcal{L}_k is balanced. Indeed, fix $I \subset \{0, \dots, n\}$ with $|I| = k-1$, the the corresponding codimension one cone $C := C_I$ has adjacent facets $C_j := C_{I \cup \{j\}}, j \in \{0, \dots, n\} \setminus I$. We can use as primitive generators

$$v_{C_j/C} := -e_j.$$

Hence the balancing condition is equivalent to

$$\sum_{j \in \{0, \dots, n\} \setminus I} -e_j \in L(C) = \langle e_i : i \in I \rangle_{\mathbf{R}},$$

which is obvious since $e_0 + \dots + e_n = 0$. The associated k -cycle L_k is called the *standard tropical k -plane*.

For later use, we close this section with a brief discussion of recession fans in the one-dimensional case (see Section 6.6 for the general treatment).

Definition 4.2.7

Let $C \in Z_1(\mathbf{R}^n)$ be a one-dimensional tropical cycle with polyhedral structure \mathcal{C} . The *recession fan* of \mathcal{C} is the fan $\text{RF}(\mathcal{C}) = \{\text{rc}(\sigma) : \sigma \in \mathcal{C}\}$ equipped with the following weights. For a ray σ of $\text{RF}(\mathcal{C})$, let $\sigma_1, \dots, \sigma_k$ be the rays of \mathcal{C} with $\text{rc}(\sigma_i) = \sigma$. Then $\omega(\sigma) := \sum_{i=1}^k \omega(\sigma_i)$. The *recession fan cycle* $\text{RF}(C)$ is the weighted polyhedral set represented by $\text{RF}(\mathcal{C})$.

It is clear that $\text{RF}(C)$ is independent of the chosen polyhedral structure \mathcal{C} .

Lemma 4.2.8

The weighted fan $\text{RF}(\mathcal{C})$ is balanced. Hence $\text{RF}(C) \in Z_1(\mathbf{R}^n)$ is a tropical fan cycle.

Proof. We sum up the balancing conditions for all vertices of \mathcal{C} . Every bounded edge σ with endpoints p_1, p_2 contributes twice, and the contributions cancel since $v_{\sigma/p_1} = -v_{\sigma/p_2}$. The remaining terms are exactly the terms in the balancing condition for $\text{RF}(\mathcal{C})$ at the origin. \square

Exercise 4.2.9

Let \mathcal{X} be a weighted polyhedral complex and let \mathcal{Z} be a weighted refinement. Show that \mathcal{Z} is balanced if and only if \mathcal{X} is balanced

Exercise 4.2.10

Let \mathcal{X} and \mathcal{Y} be two weighted polyhedral complexes. Show that \mathcal{X} and \mathcal{Y} are equivalent if and only if they have common weighted refinement.

Exercise 4.2.11

Let \mathcal{X} be a polyhedral complex of dimension k . A cell $\tau \in \mathcal{X}^{(k-1)}$ is called *redundant* if $\tau^\square \subset |\mathcal{X}|^{\text{gen}}$. Show that the weight functions for $|\mathcal{X}|$ are in bijection with the non-zero weight functions for \mathcal{X} which are balanced for all redundant $(k-1)$ -cells τ .

Exercise 4.2.12

Let L_k be the standard k -plane in \mathbf{R}^n defined in Example 4.2.6. Let C_I be a cone of \mathcal{L}_k , the standard fan structure of L_k . Show that $\text{Star}_{\mathcal{L}_k}(C_I)$ is again (the fan structure of) a standard plane, namely the standard $(k|I|)$ -plane in $\mathbf{R}^n/L(C_I)$ (where the standard directions in $\mathbf{R}^n/L(C_I)$ are given by the images of $e_j, j \notin I$).

4.3 Stable intersection

In this Section we treat the concept of stable intersection of tropical cycles alluded to in Section 2.5. The adjective “stable” refers to the fact that this type of intersection behaves continuously under small deformations of the two cycles. In fact, we will use this property in order to define stable intersections. The main feature of the construction is that it always produces as output a specific cycle of the correct codimension, even in the case of say self-intersections (without passing to equivalence relations such as rational equivalence).

Given two cycles X and Y in \mathbf{R}^n of pure dimension k resp. l , our goal is to construct a cycle $X \cdot Y$ of pure dimension $k + l - n$ which is supported on $X \cap Y$. Note that $X \cap Y$ is obviously a polyhedral set again, but might have parts of dimension bigger than $k + l - n$. So our approach is as follows: First we define $X \cdot Y$ in nice cases, namely when the intersection is transversal. Then, in the general case, we translate one of the cycles slightly. For generic translations, the resulting cycles are transversal and we can define the stable intersection as the limit of the transversal intersections when moving the translated cycles back to the original one.

4.3.1 Transversal intersections

Let us start with the definition of the transversal case. In the following, we always assume that $X \cap Y$ is equipped with the polyhedral structure $\mathcal{X} \cap \mathcal{Y}$ for two (sufficiently nice) polyhedral structures of X and Y , respectively (see Exercise 2.1.3).

Definition 4.3.1 (Transversal intersection)

Let X and Y be two cycles in \mathbf{R}^n of pure dimension k resp. l . We say X and Y intersect transversally if $X \cap Y$ is of dimension $k + l - n$ and if every facet τ of $X \cap Y$ can be written uniquely as $\tau = \sigma \cap \sigma'$ with facets σ, σ' of X resp. Y (for suitable polyhedral structures \mathcal{X}, \mathcal{Y}).

In this case, we define the *transversal intersection* $\mathcal{X} \cdot \mathcal{Y}$ to be the polyhedral complex $\mathcal{X} \cap \mathcal{Y}$ with weights

$$\omega(\tau) = \omega(\sigma) \cdot \omega(\sigma') \cdot [\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\sigma')],$$

(where τ, σ, σ' are as before). The *transversal intersection* $X \cdot Y$ is the weighted polyhedral set represented by $\mathcal{X} \cdot \mathcal{Y}$.

Proposition 4.3.2

In the transversal case, the transversal intersection $X \cdot Y$ is balanced and hence forms a tropical $(k + l - n)$ -cycle.

Proof. As \mathcal{X} and \mathcal{Y} intersect transversally, any codimension one cell of $\mathcal{X} \cap \mathcal{Y}$ lies in the codimension one skeleton of either \mathcal{X} or \mathcal{Y} . We may assume the former, i.e. the codimension one cell can be written uniquely as $\tau \cap \sigma'$,




Figure 4.1: A transversal and two non-transversal intersections

where τ is a codimension one cell of \mathcal{X} and σ' is a facet of \mathcal{Y} . Let $\sigma \supset \tau$ be a facet of \mathcal{X} . We may compare primitive generators of $\sigma \cap \sigma'$ (modulo $\tau \cap \sigma'$) and σ (modulo τ). By definition we have

$$\begin{aligned} \mathbf{Z}v_{\sigma \cap \sigma' / \tau \cap \sigma'} + L_{\mathbf{Z}}(\tau) &= (L_{\mathbf{Z}}(\sigma) \cap L_{\mathbf{Z}}(\sigma')) + L_{\mathbf{Z}}(\tau), \\ \mathbf{Z}v_{\sigma / \tau} + L_{\mathbf{Z}}(\tau) &= L_{\mathbf{Z}}(\sigma), \end{aligned}$$

and therefore

$$v_{\sigma \cap \sigma' / \tau \cap \sigma'} = [L_{\mathbf{Z}}(\sigma) : (L_{\mathbf{Z}}(\sigma) \cap L_{\mathbf{Z}}(\sigma')) + L_{\mathbf{Z}}(\tau)] \cdot v_{\sigma / \tau} \pmod{L_{\mathbf{Z}}(\tau)}.$$

Plugging in the identity

$$[L_{\mathbf{Z}}(\sigma) : (L_{\mathbf{Z}}(\sigma) \cap L_{\mathbf{Z}}(\sigma')) + L_{\mathbf{Z}}(\tau)] = [L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\sigma') : L_{\mathbf{Z}}(\tau) + L_{\mathbf{Z}}(\sigma')].$$

and multiplying by $[\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\sigma')]$ we get

$$[\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\sigma')] \cdot v_{\sigma \cap \sigma' / \tau \cap \sigma'} = [\mathbf{Z}^n : L_{\mathbf{Z}}(\tau) + L_{\mathbf{Z}}(\sigma')] \cdot v_{\sigma / \tau} \pmod{L_{\mathbf{Z}}(\tau)}. \quad (4.1)$$

Now let $\sigma_1, \dots, \sigma_m$ be the collection of facets containing τ (with primitive vectors v_i). Then the facets of $\mathcal{X} \cap \mathcal{Y}$ containing $\tau \cap \sigma'$ are $\sigma_i \cap \sigma'$ (with primitive vectors w_i) and the above relation gives

$$\begin{aligned} \sum_{i=1}^m \omega(\sigma_i \cap \sigma') w_i &= \omega(\sigma') \sum_{i=1}^m \omega(\sigma_i) [\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma_i) + L_{\mathbf{Z}}(\sigma')] w_i \\ &= \omega(\sigma') \cdot [\mathbf{Z}^n : L_{\mathbf{Z}}(\tau) + L_{\mathbf{Z}}(\sigma')] \cdot \sum_{i=1}^m \omega(\sigma_i) v_i = 0, \end{aligned} \quad (4.2)$$

i.e. the balancing condition around τ implies the balancing condition around $\tau \cap \sigma'$. \square




Figure 4.2: A non-transversal intersection and two small deformations

4.3.2 The Moving Lemma

To extend the intersection of two cycles to the non-transversal case, we need the following “moving lemma”. As usual, the degree of a zero-dimensional cycle X , which is just a weighted formal sum of points $X = \sum \omega_i p_i$, is the sum of all weights $\deg(X) = \sum \omega_i$.

Proposition 4.3.3

Let X and Y be two cycles in \mathbf{R}^n of pure dimensions k resp. l . Then the following holds.

- (a) For a generic vector $v \in \mathbf{R}^n$, the intersection of X and the translation $Y + v$ is transversal.
- (b) Assume $n = k + l$. Then $\deg(X \cdot (Y + v))$ is the same for all such generic $v \in \mathbf{R}^n$.

Remark 4.3.4

In the following, a set of generic vectors will always be the complement of a polyhedral subset of dimension $n - 1$ (or even a union of affine subspaces).

Proof. Choose arbitrary polyhedral structures \mathcal{X} and \mathcal{Y} . For the first statement, we consider all pairs of cells $\sigma \in \mathcal{X}$, $\sigma' \in \mathcal{Y}$. If $L(\sigma) + L(\sigma') \neq \mathbf{R}^n$, then for any $v \in \mathbf{R}^n \setminus (L(\sigma) + L(\sigma'))$, we have $L(\sigma) \cap (L(\sigma') + v) = \emptyset$. It

follows that $\sigma \cap (\sigma' + \nu) = \emptyset$ for all ν not contained in a certain translation of the linear subspace $L(\sigma) + L(\sigma')$. Hence for any ν not contained in any of these affine subspaces, all intersections $\sigma \cap (\sigma' + \nu)$ are either empty or of “expected dimension”, which implies that X and $Y + \nu$ intersect transversally.

For the second statement, let us first refine the previous consideration. For any vector $\nu \in \mathbf{R}^n$, the data of pairs of cells $\sigma \in \mathcal{X}$, $\sigma' \in \mathcal{Y}$ such that $\sigma \cap (\sigma' + \nu) \neq \emptyset$ is called the intersection type of ν . It is easy to check that the set of vectors with given intersection type form (the interior of) a polyhedron, and this subdivides \mathbf{R}^n into a complete polyhedral complex. \mathcal{X} and $\mathcal{Y} + \nu$ intersect transversally if ν is contained in the interior of a maximal cell of this subdivision.

Now assume we are given two such generic vectors. If they are contained in the same cell, the degree is obviously constant (as it is determined by \mathcal{X} , \mathcal{Y} and the intersection type). If not, we can connect the two vectors by passing through at most codimension one cells of the subdivision of \mathbf{R}^n . In other words, it suffices to study intersection types where $\sigma \cap (\sigma' + \nu) = \emptyset$ whenever the sum of the codimensions of σ and σ' is greater than one (otherwise, the intersection type is of higher codimension).

Therefore let τ be a codimension one cell, say of \mathcal{X} , such that there is a (unique) facet σ' of \mathcal{Y} with $\tau \cap (\sigma' + \nu) \neq \emptyset$. It is enough to show the invariance of the degree locally at τ . That is to say, we may assume that $\mathcal{Y} = \{\sigma'\}$ is a linear space and $\mathcal{X} = \{\tau, \sigma_1, \dots, \sigma_m\}$ is a fan with exactly one codimension one cell τ (a linear space of dimension $k - 1$) and facets σ_i . If $\dim(\tau + \sigma') < n - 1$, for generic ν we have $X \cap (Y + \nu) = \emptyset$ and hence $\deg(X \cdot (Y + \nu)) = 0$ is invariant. So we may assume that $H = \tau + \sigma'$ is a hyperplane. Let ν_i be primitive generators for $\sigma_i \supset \tau$. The balancing condition states

$$\sum_{i=1}^m \omega(\sigma_i) \nu_i = 0 \pmod{\tau}.$$

Choose a primitive generator w of $\mathbf{Z}^n / (H \cap \mathbf{Z}^n)$ and write $\nu_i = \lambda_i w \pmod{H}$ for (unique) $\lambda_i \in \mathbf{Z}$. It follows

$$\sum_{i=1}^m \omega(\sigma_i) \lambda_i = 0. \tag{4.3}$$




Figure 4.3: Passing a codimension one case

On the other hand, using the shorthand

$$[H : \tau + \sigma'] := [(H \cap \mathbf{Z}^n) : (\tau \cap \mathbf{Z}^n) + (\sigma' \cap \mathbf{Z}^n)]$$

we may write

$$[\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma_i) + L_{\mathbf{Z}}(\sigma')] = |\lambda_i| \cdot [H : \tau + \sigma'].$$

Now, if we move Y towards w , (i.e. v is a translation vector which points in the same direction as w modulo H), then $Y + v$ intersects the facets σ_i with $\lambda_i > 0$ and

$$\deg(X \cdot (Y + v)) = \omega(\sigma') [H : \tau + \sigma'] \sum_{\substack{i \\ \lambda_i > 0}} \omega(\sigma_i) \lambda_i.$$

Contrary, when moving Y towards $-w$, we must sum over i with $\lambda_i < 0$ instead. By equation (4.3) both sums are equal. \square

4.3.3 Stable intersections

We are now ready to define the stable intersection of two cycles in the general case.

Definition 4.3.5

Let X and Y be two tropical cycles in \mathbf{R}^n of pure dimension k resp. l . The *stable intersection* $X \cdot Y$ is

$$X \cdot Y := \lim_{\epsilon \rightarrow 0} X \cdot (Y + \epsilon v),$$

where $v \in \mathbf{R}^n$ is a vector such that X and $(Y + \epsilon v)$ intersect transversally for small $\epsilon > 0$.

More precisely, let \mathcal{X} and \mathcal{Y} be polyhedral structure for X and Y , respectively. The stable intersection $\mathcal{X} \cdot \mathcal{Y}$ is weighted polyhedral complex $\mathcal{X} \cap \mathcal{Y}^{(k+l-n)}$ with weights

$$\omega(\tau) = \sum_{\sigma, \sigma'} \omega(\sigma)\omega(\sigma')[\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma) + L_{\mathbf{Z}}(\sigma')], \quad (4.4)$$

where the sum runs through the pairs of facets $\sigma \in \mathcal{X}, \sigma' \in \mathcal{Y}$ with $\tau = \sigma \cap \sigma'$ and $\sigma \cap (\sigma' + \epsilon v) \neq \emptyset$ for small $\epsilon > 0$. The weighted polyhedral set associated to $\mathcal{X} \cdot \mathcal{Y}$ is $X \cdot Y$.

We have to show that $X \cdot Y$ does not depend on v and is balanced. It is convenient to prove these facts together with a fact that stable intersections can be computed locally. This is summarized in the following Proposition.

Proposition 4.3.6

Let X and Y be two cycles in \mathbf{R}^n of pure dimension k resp. l with polyhedral structures \mathcal{X} and \mathcal{Y} . Then the following holds.

- (a) *The stable intersection $X \cdot Y$ is well-defined, i.e. does not depend on the choice of $v \in \mathbf{R}^n$.*
- (b) *The stable intersection can be computed locally. In formulas, we have*

$$\omega(\tau) = \deg(\text{Star}_{\mathcal{X}}(\tau) \cdot \text{Star}_{\mathcal{Y}}(\tau)),$$

for any $k + l - n$ -cell τ of $X \cap Y$ and, more general, for any cell ρ

$$\text{Star}_{\mathcal{X} \cdot \mathcal{Y}}(\rho) = \text{Star}_{\mathcal{X}}(\rho) \cdot \text{Star}_{\mathcal{Y}}(\rho).$$

- (c) *The stable intersection $X \cdot Y$ is balanced and therefore defines a tropical $(k + l - n)$ -cycle in \mathbf{R}^n .*
- (d) *Assume $k + l = n$. Then $\deg(X \cdot Y) = \deg(X \cdot (Y + v))$ for all $v \in \mathbf{R}^n$.*
- (e) *Assume $k + l = n + 1$ and X, Y fan cycles. Then $X \cdot Y = \text{RF}(X \cdot (Y + v))$ for all $v \in \mathbf{R}^n$.*

Proof. We start by noting that the weight of τ in $\mathcal{X} \cdot \mathcal{Y}$, computed with respect to a generic ν , is

$$\omega(\tau) = \deg(\text{Star}_{\mathcal{X}}(\tau) \cdot (\text{Star}_{\mathcal{Y}}(\tau) + [\nu])).$$

This follows directly from Equation 4.4 and the observation that dividing by $L(\tau)$ does not affect any of the lattice indices. With the help of our earlier invariance result 4.3.3, this proves (a) and, using the definition of stable intersection again, (b) and (d).

Similarly, to prove (e) it suffices to consider generic vectors ν . We choose a generic ν and set $\mathcal{C} = \mathcal{X} \cdot (\mathcal{Y} + \nu)$. In order to prove $\text{RF}(\mathcal{C}) = \mathcal{X} \cdot \mathcal{Y}$, we note that by Exercise 2.1.1 we have $\text{rc}(\sigma \cap (\sigma' + \nu)) = \text{rc}(\sigma \cap \sigma')$ for all facets $\sigma \in \mathcal{X}$, $\sigma' \in \mathcal{Y}$ such that $\sigma \cap (\sigma' + \nu) \neq \emptyset$. Hence for a given ray $\tau \in \mathcal{X} \cap \mathcal{Y}$, the terms in the sum 4.4 correspond exactly to the rays in \mathcal{C} whose recession cone is equal to τ . It follows that $\tau \in \text{RF}(\mathcal{C})$ if $\omega_{\mathcal{X} \cdot \mathcal{Y}}(\tau) \neq 0$ and moreover $\omega_{\mathcal{X} \cdot \mathcal{Y}}(\tau) = \omega_{\text{RF}(\mathcal{C})}(\tau)$. This proves (e).

It remains to prove (c). By (b), it suffices to consider the case where \mathcal{X} and \mathcal{Y} are fans and $k + l = n + 1$. By (e), in this case we have $\mathcal{X} \cdot \mathcal{Y} = \text{RF}(\mathcal{X} \cdot (\mathcal{Y} + \nu))$ for all ν . For generic ν , we know that $\mathcal{X} \cdot (\mathcal{Y} + \nu)$ is balanced by Proposition 4.3.2. Hence $\mathcal{X} \cdot \mathcal{Y}$ is balanced by Lemma 4.2.8. \square

Stable intersection provides a map

$$\cdot : Z_k(\mathbf{R}^n) \times Z_l(\mathbf{R}^n) \rightarrow Z_{k+l-n}(\mathbf{R}^n).$$

Let us emphasize that this intersection product is constructed on the non-compact space \mathbf{R}^n and without passing to equivalence classes of any kind — even in the case of non-transversal (or self-)intersections. The product \cdot turns $Z_*(\mathbf{R}^n)$ into a graded commutative \mathbf{R} -algebra with unit $\mathbf{1} = \mathbf{R}^n$, as the following proposition shows.

Proposition 4.3.7

Stable intersection is associative, commutative, bilinear and its neutral element is given by the “fundamental” cycle $[\mathbf{R}^n]$. We have $X \cdot Y \subseteq X \cap Y$.

Proof. Commutativity, Bilinearity, the neutral element assertion and $X \cdot Y \subseteq X \cap Y$ follow directly from the definitions.

It remains to show associativity. We first study the situation where all of the involved intersections are transversal. In this case, each facet τ of

$\mathcal{X} \cap \mathcal{Y} \cap \mathcal{Z}$ can be written uniquely as $\tau = \sigma_1 \cap \sigma_2 \cap \sigma_3$ for facets of \mathcal{X} , \mathcal{Y} and \mathcal{Z} respectively. The weight of τ in $(\mathcal{X} \cdot \mathcal{Y}) \cdot \mathcal{Z}$ is given by

$$\begin{aligned} \omega(\tau) &= \omega(\sigma_1)\omega(\sigma_2)\omega(\sigma_3) \\ &\quad \cdot [\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma_1) + L_{\mathbf{Z}}(\sigma_2)] [\mathbf{Z}^n : (L_{\mathbf{Z}}(\sigma_1) \cap L_{\mathbf{Z}}(\sigma_2)) + L_{\mathbf{Z}}(\sigma_3)]. \end{aligned}$$

After plugging in the identity

$$\begin{aligned} &[\mathbf{Z}^n : (L_{\mathbf{Z}}(\sigma_1) \cap L_{\mathbf{Z}}(\sigma_2)) + L_{\mathbf{Z}}(\sigma_3)] \\ &= [\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma_2) + L_{\mathbf{Z}}(\sigma_3)] [L_{\mathbf{Z}}(\sigma_2) : (L_{\mathbf{Z}}(\sigma_1) \cap L_{\mathbf{Z}}(\sigma_2)) + (L_{\mathbf{Z}}(\sigma_2) \cap L_{\mathbf{Z}}(\sigma_3))], \end{aligned}$$

we obtain an expression which is symmetric in σ_1 and σ_3 and thus is equal to the weight in $X \cdot (Y \cdot Z)$. For the general case, we just have to unwind our definition and observe that a triple stable intersection can be computed by moving two of the cycles simultaneously, i.e.

$$(X \cdot Y) \cdot Z = \lim_{\substack{v \rightarrow 0 \\ u \rightarrow 0}} (X \cdot (Y + v)) \cdot (Z + u).$$

Therefore the previous argument is sufficient. \square

Example 4.3.8

Let us consider the stable intersection of the standard planes L_k defined in Example 4.2.6. We claim

$$L_k \cdot L_l = L_{n-k-l}.$$

We set $m = n - k - l$. By definitions, the m -skeleton of $L_k \cap L_l$ is exactly the support of L_m , so we just have to show that all facets appear with weight 1 in $L_k \cdot L_l$. To do this, we use locality and compute the weight of a facet by intersecting the corresponding stars instead. We mentioned before that these stars are standard planes again (cf. example 4.2.12), so in fact it suffices to show that

$$L_k \cdot L_{n-k} = \{0\}$$

with weight 1. This may be verified explicitly by choosing the translation vector

$$v := (\underbrace{-1, \dots, -1}_{k \text{ times}}, \underbrace{1, \dots, 1}_{n-k \text{ times}}),$$




Figure 4.4: Two planes in \mathbf{R}^3 which intersect in a line with a bounded edge

for example. Indeed, in this case, the only facets of L_k and $L_{n-k} + v$ that intersect are $C_{\{1,\dots,k\}}$ and $C_{\{k+1,\dots,n\}} + v$. The intersection point

$$\sum_{i=1}^k e_i = \sum_{i=k+1}^n e_i + v.$$

has weight 1, as $L_{\mathbf{Z}} C_{\{1,\dots,k\}} + L_{\mathbf{Z}} C_{\{k+1,\dots,n\}} = \mathbf{Z}e_1 + \dots + \mathbf{Z}e_n = \mathbf{Z}^n$. So the claim follows. In particular, if we denote by $H := L_{n-1}$ the *standard hyperplane* in \mathbf{R}^n , then

$$L_{n-k} = H^k = H \cdots H \text{ (} k \text{ times)}.$$

Figure 4.4 show the intersection of two copies of $H \subset \mathbf{R}^3$ moved to transversal position.

Exercise 4.3.9

Let X and Y be tropical cycles in \mathbf{R}^n that intersect transversally. Show that $X \cdot Y$ does not depend on the chosen polyhedral structures \mathcal{X} and \mathcal{Y} .

4.4 The divisor of a piecewise affine function

In Chapter 2 we discussed how a tropical polynomial gives rise to a tropical hypersurface which, in turn, represents a tropical subspace of dimension $n - 1$. In this section, we want to extend this construction to more general functions on tropical cycles and describe its relationship to stable intersections. In the following, we use the same letter X for the support of a tropical cycle $X \subset \mathbf{R}^n$ (forgetting the weights).

4.4.1 Piecewise integer affine, regular and rational functions

Definition 4.4.1

Let $X \subset \mathbf{R}^n$ be a tropical cycle. A function $f : X \rightarrow \mathbf{R}$ is a *piecewise integer affine function* if there exists a polyhedral structure \mathcal{X} for X such that for any cell $\sigma \in \mathcal{X}$, the restriction $f|_{\sigma}$ is the restriction of an integer affine function on \mathbf{R}^n to σ . Any polyhedral structure satisfying this property is called *sufficiently fine for f* . We denote the set of piecewise integer affine functions on X by $\text{PA}_{\mathbf{Z}}(X)$. The function f is *piecewise integer linear* if furthermore \mathcal{X} can be chosen to be a fan and $f(0) = 0$.

Note that a piecewise integer affine function is always continuous. It is clear that $(\text{PA}_{\mathbf{Z}}(X), +)$ forms an abelian group under ordinary addition, i.e., tropical multiplication. It is equally straightforward to check that $\text{PA}_{\mathbf{Z}}(X)$ is closed under tropical addition. In particular, after adding the constant function $-\infty$, we obtain a semigroup $(\text{PA}_{\mathbf{Z}}(X) \cup \{-\infty\}, \max)$.

Example 4.4.2

Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a tropical polynomial. Then f is a piecewise integer affine function. A sufficiently fine polyhedral structure is given by $\mathcal{S}(f)$.

Let $g : \mathbf{R}^n \rightarrow \mathbf{R}$ be a second polynomial. Then the tropical quotient “ f/g ” : $x \mapsto f(x) - g(x)$ is a piecewise integer affine function as well, with sufficiently fine subdivision given by the intersection of the subdivisions $\mathcal{S}(f)$ and $\mathcal{S}(g)$.

Note that our definition includes a finiteness condition since polyhedral structures are assumed to be finite here. For example, a tropical Laurent

series in n variables is in general not included in our definition since its domains of linearity may produce an infinite polyhedral structure in \mathbf{R}^n .

We may turn these two basic examples into general definitions.

Definition 4.4.3

Let $X \subset \mathbf{R}^n$ be a tropical cycle. A *polynomial function* $f : X \rightarrow \mathbf{R}$ is the restriction of a tropical polynomial $\mathbf{R}^n \rightarrow \mathbf{R}$ to f . A *(tropical) regular function* $f : X \rightarrow \mathbf{R}$ is a piecewise integer affine function which is locally polynomial. We denote the set of regular functions by $\mathcal{O}(X)$.

A *(tropical) rational function* $f : X \rightarrow \mathbf{R}$ is the restriction of a tropical quotient of two tropical polynomials “ g/h ” to X . We denote the set of rational functions by $\mathcal{R}(X)$. A *locally rational function* $f : X \rightarrow \mathbf{R}$ is a piecewise integer affine function which is locally rational.

Note that the definitions, even though imposing a local condition, still contain the finiteness condition inherited from our definition of piecewise integer affine function. In particular, with the given definition neither regular functions nor locally rational functions form a sheaf on X .

4.4.2 The divisor construction

Definition 4.4.4

Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on a tropical cycle $X \in Z_m(\mathbf{R}^n)$ and let \mathcal{X} be sufficiently fine polyhedral structure. The *divisor of f on \mathcal{X}* is the weighted polyhedral complex $\text{div}_{\mathcal{X}}(f)$ (or $f \cdot \mathcal{X}$) constructed as follows.

- The underlying polyhedral complex is \mathcal{X}^{m-1} .
- For a facet $\tau \in \mathcal{X}^{m-1}$, the weight of τ in $\text{div}_{\mathcal{X}}(f)$ is given by

$$\omega(\tau) := \left(\sum_{i=1}^k \omega(\sigma_i) df|_{\sigma_i}(v_{\sigma_i/\tau}) \right) - df|_{\tau} \left(\sum_{i=1}^k \omega(\sigma_i) v_{\sigma_i/\tau} \right). \quad (4.5)$$

Here, $\sigma_1, \dots, \sigma_k$ denote the facets of \mathcal{X} containing τ , $v_{\sigma_i/\tau}$ is a choice of primitive generator for facet, and $df|_{\sigma_i}$ denotes the differential of the (affine) function $f|_{\sigma_i}$.

Remark 4.4.5

A few remarks are in order.

- First note that by the balancing condition for \mathcal{X} around τ , the sum $\sum_{i=1}^k \omega(\sigma_i)v_{\sigma_i/\tau}$ is a vector in $L(\tau)$ and its value under $df|_\tau$ is well-defined. Second, let us check that the definition of weights is independent of the chosen primitive generators $v_{\sigma_i/\tau}$. Indeed, for an alternative choice $v'_{\sigma_i/\tau}$ we have $w_i := v'_{\sigma_i/\tau} - v_{\sigma_i/\tau} \in L(\tau)$. It follows

$$df|_{\sigma_i}(v'_{\sigma_i/\tau}) = df|_{\sigma_i}(v_{\sigma_i/\tau}) + df|_\tau(w_i).$$

Hence we get a correction term $\sum_{i=1}^k \omega(\sigma_i)df|_\tau(w_i)$ in the first sum which cancels, by linearity of $df|_\tau$, with the correction term $df|_\tau(\sum_{i=1}^k \omega(\sigma_i)w_i)$.

- We can always choose primitive vectors such that $\sum_{i=1}^k \omega(\sigma_i)v_{\sigma_i/\tau} = 0$. In this case the weight formula (4.5) simplifies to

$$\omega(\tau) := \sum_{i=1}^k \omega(\sigma_i)df|_{\sigma_i}(v_{\sigma_i/\tau}).$$

To verify that such a choice of primitive vectors is possible, we set $g := \gcd\{\omega(\sigma_i)\}$ and write $g = \sum_{i=1}^k a_i\omega(\sigma_i)$, $a_i \in \mathbf{Z}$. Then for any choice of primitive vectors we have

$$w := \frac{1}{g} \sum_{i=1}^k \omega(\sigma_i)v'_{\sigma_i/\tau} \in \text{PA}_{\mathbf{Z}}(\tau).$$

Setting $v_{\sigma_i/\tau} := v'_{\sigma_i/\tau} - a_i w$ we obtain a collection of primitive vectors whose sum up to zero.

- Note that the construction may produce zero weights $\omega(\tau) = 0$. In this case, the (weighted) support $|\text{div}_{\mathcal{X}}(f)|$ is strictly contained in $|\mathcal{X}^{m-1}|$. In particular, $\omega(\tau) = 0$ if f can be written as the restriction of an integer affine function g in a neighbourhood of τ^\square , since then $df|_{\sigma_i} = dg$ for all i and the two terms cancel by linearity of dg (as above).

Example 4.4.6

Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a tropical polynomial with associated subdivision $\mathcal{S}(f)$. Then $\operatorname{div}_{\mathcal{S}(f)}(f) = V(f)$. Indeed, let P be an $(n-1)$ -cell of $\mathcal{S}(f)$ and let h_1, h_2 the integer linear functions associated to the endpoints of the dual cell σ_P . Let h be the integer linear form such that $h_2 - h_1 = \operatorname{Vol}(\sigma_P) \cdot h$ and let F_1, F_2 be the facets of $\mathcal{S}(f)$ on which h_1, h_2 attain the maximum, respectively. The primitive generators v of F_2 modulo P are characterized by the property $(v) = 1$. Hence

$$\omega(\tau) = h_2(v) + h_1(-v) = \operatorname{Vol}(\sigma_P)h(v) = \operatorname{Vol}(\sigma_P).$$

Proposition 4.4.7

The weighted polyhedral complex $\operatorname{div}_{\mathcal{X}}(f)$ is balanced.

Proof. We need to check the balancing condition for each $m-2$ -cell $\rho \in \mathcal{X}$. Let $f_\rho : \mathbf{R}^n \rightarrow \mathbf{R}$ be an integer affine function restricting to $f|_\rho$ on ρ . Then the function $f - f_\rho$ descends to a piecewise integer linear function \tilde{f} on $\mathcal{S} := \operatorname{Star}_{\mathcal{X}}(\rho)$, a polyhedral fan in $\mathbf{R}^n/L(\rho)$. By the previous remark, the weights of $\operatorname{div}_{\mathcal{X}}(f)$ are invariant under adding a (globally) integer affine function. Hence, the balancing condition at ρ is equivalent to the balancing condition of $\operatorname{div}_{\mathcal{S}}(\tilde{f})$ at the origin. In other words, we may restrict to the case where $m = 2$ and \mathcal{X} is a balanced polyhedral fan (with 0-cell $\{0\}$).

We may additionally assume that \mathcal{X} is unimodular, that is to say, each 2-cone $\sigma \in \mathcal{X}$ is generated by two integer vectors which form a lattice basis of $L_{\mathbf{Z}}(\sigma)$. To see this, recall that any \mathcal{X} admits a unimodular refinement and note that refining \mathcal{X} does not affect the balancing condition. Indeed, any newly added ray gets assigned weight 0 by the previous remark. So let us assume that \mathcal{X} is unimodular.

Let τ_1, \dots, τ_l denote the rays of \mathcal{X} , and let v_1, \dots, v_l be their primitive generators. If a 2-cone σ is spanned by τ_i and τ_j , then unimodularity implies that we can use v_j as a primitive generator for σ modulo τ_i (and vice versa). We use the following notation. For $i \neq j \in \{1, \dots, l\}$, we set $\omega_{ij} = \omega(\sigma_{ij})$ if τ_i and τ_j span a 2-cone σ_{ij} of \mathcal{X} , and $\omega_{ij} = 0$ otherwise. Note that $\omega_{ij} = \omega_{ji}$. The balancing condition of \mathcal{X} implies that there is an integer $\alpha_i \in \mathbf{Z}$ such that

$$\sum_{j=1}^l \omega_{ij} v_j = \alpha_i v_i \in L(\tau_i). \tag{4.6}$$

Then the weight of τ_i in $\operatorname{div}_{\mathcal{X}}(f)$ can be computed as

$$\omega(\tau_i) = \left(\sum_{j=1}^l \omega_{ij} f(v_j) \right) - \alpha_i f(v_i).$$

Finally, the balancing condition of $\operatorname{div}_{\mathcal{X}}(f)$ at the origin can be checked by computing

$$\begin{aligned} \sum_{i=1}^l \omega(\tau_i) v_i &= \left(\sum_{i,j=1}^l \omega_{ij} f(v_j) v_i \right) - \sum_{i=1}^l \alpha_i f(v_i) v_i \\ &= \left(\sum_{i,j=1}^l \omega_{ij} f(v_i) v_j \right) - \sum_{i=1}^l \alpha_i f(v_i) v_i \\ &= \sum_{i=1}^l f(v_i) \left(\left(\sum_{j=1}^l \omega_{ij} v_j \right) - \alpha_i v_i \right). \end{aligned}$$

But the term in big brackets in the last line is zero by Equation (4.6). \square

Definition 4.4.8

Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on a tropical cycle $X \in Z_m(\mathbf{R}^n)$ and let \mathcal{X} be sufficiently fine polyhedral structure. The *divisor of f* is the tropical cycle of dimension $m - 1$ defined by $\operatorname{div}_{\mathcal{X}}(f)$. We denote it by $\operatorname{div}_X(f)$ or $f \cdot X$.

Remark 4.4.9

By the previous proposition, $\operatorname{div}_{\mathcal{X}}(f)$ is balanced and hence defines a tropical cycle. Note that this tropical cycle does not depend on the chosen polyhedral structure \mathcal{X} . Indeed, let \mathcal{X}' be a refinement of \mathcal{X} and let τ be $m - 1$ -cell of \mathcal{X}' which is not contained in any $m - 1$ -cell of \mathcal{X} . Then f is integer affine in a neighbourhood of τ^\square and therefore the weight of τ in $\operatorname{div}_{\mathcal{X}'}(f)$ is zero. For all other $m - 1$ -cells of \mathcal{X} , obviously the weight agrees with the weight of the (unique) $m - 1$ -cell of \mathcal{X} containing τ . Hence $\operatorname{div}_{\mathcal{X}'}(f)$ and $\operatorname{div}_{\mathcal{X}}(f)$ are equivalent as weighted polyhedral complexes and induce the same tropical cycle.

Example 4.4.10

Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a tropical polynomial. Then $\operatorname{div}(f) = V(f)$.

Example 4.4.11

Let $X = L_2 = V("0 + x + y + z") \subset \mathbf{R}^3$ be the standard plane in \mathbf{R}^3 from Example 4.2.6. Note that X contains the (honest) line $L = \mathbf{R}(e_0 + e_3) = \mathbf{R}(e_1 + e_2)$ as a subcycle. We want to describe a function $f : X \rightarrow \mathbf{R}$ such that $f \cdot X = L$.

We start by subdividing the cones $\sigma_{\{0,3\}}$ and $\sigma_{\{1,2\}}$ by intersecting rays in the direction $-(e_0 + e_3)$ and $-(e_1 + e_2)$, respectively. Note that the refined fan \mathcal{X} is (still) unimodular. In particular, there exists a unique function g which is linear on the cones of \mathcal{X} and maps the primitive generators according to

$$-e_i \mapsto 0 \ (i = 0, 1, 2, 3), \quad -(e_0 + e_3), -(e_1 + e_2) \mapsto -1.$$

The divisor of this function carries the weights

$$\begin{aligned} \omega(\mathbf{R}_{\leq 0}e_0) &= g(-e_1) + g(-e_2) + g(-(e_0 + e_3)) - g(0) &= -1, \\ \omega(\mathbf{R}_{\leq 0}(e_0 + e_3)) &= g(-e_0) + g(-e_3) - g(-(e_0 + e_3)) &= 1, \end{aligned}$$

and symmetrically for all other rays. In other words, $g \cdot X = L - L_1$, where L_1 denotes the standard line in \mathbf{R}^3 . But $h \cdot X = L_1$ for $h = "0 + x + y + z"$ (see Examples 4.3.8 and 4.4.20, or just calculate it directly). Therefore the function $f = g + h$ satisfies $f \cdot X = L$ (see Remark 4.4.15).

We may generalize Theorem 2.4.10 to arbitrary tropical $n - 1$ -cycles allowing negative weights.

Corollary 4.4.12

Let $X \subset \mathbf{R}^n$ be a tropical $n - 1$ -cycle. Then there exists a tropical rational function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ such that $X = \text{div}(f)$. The function f is unique up to adding an integer affine function (among all piecewise integer affine functions). In particular, any piecewise integer affine function on \mathbf{R}^n is rational.

Proof. By Lemma 4.2.5, we can write $X = X_1 - X_2$ with X_1, X_2 tropical $n - 1$ -subspaces. By Theorem 2.4.10 there exist tropical polynomials f_1, f_2 such that $X_i = V(f_i) = \text{div}(f_i)$, $i = 1, 2$. Hence the rational function $f = f_1 - f_2$ satisfies $V(f) = X$. To show uniqueness and the last claim, let f be a piecewise affine function such that $V(f) = 0$ (as a cycle). Let \mathcal{X} be a sufficiently fine subdivision of \mathbf{R}^n . Then for any $n - 1$ -cell $\tau \in \mathcal{X}$ the condition $\omega_{\text{div}(f)}(\tau) = 0$ implies that the two affine functions on both sides of τ agree. Hence f is a globally integer affine function. \square

Corollary 4.4.13

Any piecewise integer affine function $f : X \rightarrow \mathbf{R}$ is locally rational.

Proof. Since this is a local statement, we may restrict to the case where f is piecewise linear function on a fan cycle X . Let \mathcal{X} be a sufficiently fine polyhedral fan structure for f on X . By refining \mathcal{X} further if necessary, we may assume that $\mathcal{X} \subset \mathcal{F}$ where \mathcal{F} is unimodular complete fan of \mathbf{R}^n . Then by choosing arbitrary integer values on the primitive generators of rays not contained in \mathcal{X} and extending by linearity, we may construct an extension $\tilde{f} : \mathbf{R}^n \rightarrow \mathbf{R}$ of f which is integer linear on cones of \mathcal{F} . Then by Corollary 4.4.12, \tilde{f} is a rational function, and the claim follows. \square

Exercise 4.4.14

$X = \mathbf{R}_{\geq 0}(-1, 2) \cup \mathbf{R}_{\geq 0}(-1, -2) \cup 2 \cdot \mathbf{R}_{\geq 0}(1, 0)$, $f : X \rightarrow \mathbf{R}$, $(x, y) \mapsto \frac{y}{2}$. Show that f is piecewise integer linear on X and $\text{div}(f) = 0$. Show that f is not regular. In particular, it is not true that any locally convex piecewise integer affine function is regular.

The statement holds, however, if we allow Laurent polynomials with *rational* exponents: Show that every locally convex piecewise integer affine function $f : X \rightarrow \mathbf{R}$ is locally the restriction of a function $x \mapsto \max_{j \in A} \{a_j + jx\}$ with $A \in \mathbf{Q}^n$ finite and $a_j \in \mathbf{R}$ (for arbitrary X).

4.4.3 Further properties of the divisor construction

Remark 4.4.15

For further reference, let us collect some of the properties of the divisor construction (most of which appeared in the previous proofs). Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on a tropical cycle $X \in Z_m(\mathbf{R}^n)$. We say f is *locally affine* at $x \in X$ if there is a neighbourhood $x \in U \subset X$ and a affine function $F : \mathbf{R}^n \rightarrow \mathbf{R}$ such that $f|_U = F|_U$. We denote by $V(f) \subset X$ the set of points where f is *not* locally affine. The following holds.

- (a) The support of $\text{div}(f)$ is contained in $V(f)$. In particular, if f is the restriction of an integer affine function, then $\text{div}(f) = 0$.
- (b) Let \mathcal{X} be a sufficiently fine polyhedral structure for f and let τ be a cell of $\mathcal{X}^{(m-1)}$. Then df induces piecewise integer linear function

$f^\tau : \text{Star}_x(\tau) \rightarrow \mathbf{R}$ which is unique up to adding an integer affine function. Moreover, we have

$$f^\tau \cdot \text{Star}_x(\tau) = \text{Star}_{f \cdot x}(\tau).$$

In other words, $f \cdot X$ can be computed locally.

(c) By the linearity property of Equation (4.5) it follows that the map

$$\begin{aligned} \text{PA}_Z(X) \times Z_k(X) &\rightarrow Z_{k-1}(X), \\ (f, Z) &\mapsto f \cdot Z := f|_Z \cdot Z \end{aligned}$$

is linear in both arguments.

Recall that a tropical subspace is a tropical cycle with only positive weights.

Proposition 4.4.16

Let $f : X \rightarrow \mathbf{R}$ be a regular function on a tropical subspace $X \in Z_m(\mathbf{R}^n)$. Then $f \cdot X$ is a tropical subspace and its support is equal to $V(f)$.

In fact, the statement can be generalized to locally convex functions. A function f is called locally convex at x if it is equal to the restriction of a convex function on \mathbf{R}^n in a neighbourhood of x . Clearly, regular functions are locally convex functions.

Proof. By locality we can restrict to the case where X is a one-dimensional fan and f is linear on each ray. Let ρ_1, \dots, ρ_l be the rays with weights $\omega_i > 0$ and primitive generators v_i . Let ω denote the weight of the origin in $\text{div}(f)$. Then

$$\omega = \sum_{i=1}^l \omega_i f(v_i) \geq f\left(\sum_{i=1}^l \omega_i v_i\right) = 0,$$

where the inequality follows from the local convexity of f . This proves the first claim.

For the second claim, it remains to show that if $\omega = 0$ then f is locally affine. Assuming the contrary, there exists a linear combination $\sum_{i=1}^l \lambda_i v_i = 0$, $\lambda_i \in \mathbf{Z}$, such that $\sum_{i=1}^l \lambda_i f(v_i) \neq 0$. We may assume $\sum_{i=1}^l \lambda_i f(v_i) < 0$ by replacing λ_i with $-\lambda_i$, if necessary. Moreover, replacing λ_i by $\lambda_i + m\omega_i$

for sufficiently large $m \in \mathbf{N}$ we may assume $\lambda_i > 0$. Hence the λ_i define another tropical subspace X' (with $|X'| \subset |X|$) such that $f \cdot X'$ carries a negative weight the origin. This is a contradiction to the first claim of the proposition, and therefore proves the second claim. \square

Example 4.4.17

Let $X = L_2 \subset \mathbf{R}^3$ be the standard plane and $L \in X$ the (honest) line from Example 4.4.11. Note that L is rigid in X : No non-trivial translation of L lies in X . Philosophically, we might hope to detect this by computing a negative self-intersection of L in X . Since in Example 4.4.11 we constructed a function $f : X \rightarrow \mathbf{R}$ such that $f \cdot X = L$, we can interpret this self-intersection as $f^2 \cdot X$. By definition, this 0-cycle is supported $\{0\}$ and the weight of the origin is given by

$$f(-(e_0 + e_3)) + f(-(e_1 + e_2)) = 0 - 1 = -1.$$

So indeed, the self-intersection $f^2 \cdot X$ is negative. In particular, f is not a regular (or locally convex) function.

Proposition 4.4.18

Let $f, g : X \rightarrow \mathbf{R}$ be a two piecewise integer affine functions on a tropical cycle $X \in Z_m(\mathbf{R}^n)$. Then

$$f \cdot (g \cdot X) = g \cdot (f \cdot X).$$

Proof. The proof is very similar to the proof of 4.4.7. Again by locality, we may restrict to the case of a two-dimensional unimodular fan \mathcal{X} and f, g piecewise integer linear functions on \mathcal{X} . Let ω and ω' be the weight of the origin in $f \cdot (g \cdot X)$ and $g \cdot (f \cdot X)$, respectively. Let τ_1, \dots, τ_l be the rays of \mathcal{X} and let ω_i and ω'_i be the weights of ρ_i in $g \cdot X$ and $f \cdot X$, respectively. Then

$$\begin{aligned} \omega &= \sum_{i=1}^l \omega_i f(v_i) = \left(\sum_{i,j=1}^l \omega_{ij} g(v_j) f(v_i) \right) - \sum_{i=1}^l \alpha_i g(v_i) f(v_i) \\ &= \left(\sum_{i,j=1}^l \omega_{ij} f(v_j) g(v_i) \right) - \sum_{i=1}^l \alpha_i f(v_i) g(v_i) \\ &= \sum_{i=1}^l \omega'_i g(v_i) = \omega', \end{aligned}$$

where $\omega_{ij}, v_i, \alpha_i$ are defined as before in 4.4.7. □

4.4.4 Compatibility with stable intersection

Our next goal is to show that the divisor construction is compatible with stable intersection. More precisely, we want to prove the following theorem.

Theorem 4.4.19

Let X and Y be tropical cycles in \mathbf{R}^n . Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on X . Then

$$f \cdot (X \cdot Y) = (f \cdot X) \cdot Y.$$

Example 4.4.20

Let $X \subset \mathbf{R}^n$ be a tropical cycle and let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a tropical polynomial. Then

$$f \cdot X = V(f) \cdot X.$$

Let f_1, \dots, f_k be a collection of tropical polynomials. Then

$$f_1 \cdots f_k \cdot \mathbf{R}^n = V(f_1) \cdots V(f_k).$$

We need the following lemma.

Lemma 4.4.21

Let X be a tropical fan cycle, $f : X \rightarrow \mathbf{R}$ a piecewise integer linear function on X and $C \subset X$ a one-dimensional tropical cycle. Then

$$\deg(f \cdot C) = \deg(f \cdot \text{RF}(C)).$$

Proof. First note that $\deg(f \cdot C) = \sum_{i=1}^l \omega(\sigma_i) df|_{\sigma_i}(v_i)$, where $\sigma_1, \dots, \sigma_l$ are the rays of \mathcal{C} (for some polyhedral structure) and v_i is a primitive generator for σ_i . As in the proof of Lemma 4.2.8, this follows easily from the fact that each bounded edge of \mathcal{C} produces two contributions to the degree which cancel each other. Since X is a fan and f is linear on a polyhedral fan structure for X , we have $f(v_i)$. Since the v_i are also primitive generators for $\text{rc}(\sigma_i)$, we see that $\deg(f \cdot C)$ is equal to $\deg(f \cdot \text{RF}(C)) = \sum_{i=1}^l \omega(\sigma_i) f(v_i)$. □

Proof of Theorem 4.4.19. By locality of both sides, it suffices to restrict to the case where X, Y are tropical fan cycles of dimension k, l , respectively, $k + l = n + 1$, and f is piecewise linear on X . Let $Y + v$ be a generic translation of Y which is transverse both to X and $f \cdot X$. Then by Lemma 4.4.21 and Proposition 4.3.6 (d) and (e), it suffices to show $f \cdot (X \cdot (Y + v)) = (f \cdot X) \cdot (Y + v)$, i.e., we may restrict to the transversal case.

So let us assume that $\mathcal{Y} = \{L\}$ is an ordinary linear space of dimension l and that \mathcal{X} is a polyhedral fan of dimension k with a single $(k - 1)$ -cell Q , linear space transversal to L , and facets $\sigma_1, \dots, \sigma_l$. We denote by v_i and w_i the primitive generators of σ_i modulo Q and $\sigma_i \cap L$ modulo $Q \cap L = \{0\}$, respectively. We may assume $\sum_{i=1}^l \omega(\sigma_i)v_i = 0$ for simplicity. With the notations $a_i := [\mathbf{Z}^n : L_{\mathbf{Z}}(\sigma_i) + (L \cap \mathbf{Z}^n)]$ and $a = [\mathbf{Z}^n : (Q \cap \mathbf{Z}^n) + (L \cap \mathbf{Z}^n)]$, we get

$$a_i w_i = a v_i \pmod{Q}$$

for all i (see Equation 4.1). Finally, let ω and ω' denote the weight of the origin in $f \cdot (\mathcal{X} \cdot \mathcal{Y})$ and $(f \cdot \mathcal{X}) \cdot \mathcal{Y}$, respectively. Then

$$\begin{aligned} \omega &= \sum_{i=1}^l \omega_{\mathcal{X} \cdot \mathcal{Y}}(\sigma_i \cap L) f(w_i) = \omega_{\mathcal{Y}}(L) \sum_{i=1}^l \omega_{\mathcal{X}}(\sigma_i) a_i f(w_i) \\ &= \omega_{\mathcal{Y}}(L) \cdot a \cdot \sum_{i=1}^l \omega(\sigma_i) f(v_i) = \omega_{\mathcal{Y}}(L) \cdot a \cdot \omega_{f \cdot \mathcal{X}}(Q) = \omega'. \end{aligned}$$

□

Exercise 4.4.22

Let X be a tropical 1-cycle in \mathbf{R}^n and $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function. Let \mathcal{X} be a sufficiently fine (pointed) polyhedral structure and let ρ_1, \dots, ρ_l denote the rays (i.e. unbounded edges) of \mathcal{X} , with primitive generators v_1, \dots, v_l . Then

$$\deg(f \cdot X) = \sum_{i=1}^l \omega(\rho_i) df|_{\rho_i}(v_i).$$

4.5 Tropical modifications

In classical algebraic geometry, it is a basic feature of any regular function $f : X \rightarrow K$ that its graph $\Gamma(X, f) \subset X \times K$ is Zariski closed. In Remark 2.4.9 we constructed a balanced completion of the graph of a tropical polynomial f which served as a motivation for the definition of $V(f)$. We will extend this graph completion to arbitrary piecewise integer affine functions now. Beyond giving some justification of the divisor construction, these graph completions are also of interest in their own right and are called *tropical modifications*.

The graph of a function $f : X \times \mathbf{R}$ is denoted by $\Gamma(X, f)$ or just $\Gamma(f)$. If $X \subset \mathbf{R}^n$ is a polyhedral set and f is a piecewise integer affine function, then $\Gamma(X, f) \subset \mathbf{R}^n \times \mathbf{R}$ is a polyhedral set. Let $\pi : \mathbf{R}^n \times \mathbf{R} \rightarrow \mathbf{R}^n$ denote the projection. If X is weighted, we equip $\Gamma(X, f)$ with weights by setting $\omega(\tilde{x}) = \omega(\pi(\tilde{x}))$ for all $\tilde{x} \in \Gamma(X, f)^{\text{gen}}$. Note that the weighted polyhedral set $\Gamma(X, f)$ is in general not balanced. However, in a sense to be made precise, there is a unique balanced completion $\Gamma(X, f)$. Let us discuss this in more details.

Let \mathcal{X} be a polyhedral structure of X sufficiently fine for f . For $\sigma \in \mathcal{X}$, we denote by $\tilde{\sigma} := \Gamma(\sigma, f|_{\sigma})$ the lift of σ under f . The polyhedra $\tilde{\sigma}$ form a polyhedral structure $\tilde{\mathcal{X}}$ for $\Gamma(X, f)$. Let τ be a $k-1$ -cell of \mathcal{X} , let $\sigma_1, \dots, \sigma_l$ be the facets of \mathcal{X} adjacent to τ and let v_1, \dots, v_l be corresponding primitive generators in \mathbf{Z}^n . For simplicity, we choose them such that $\sum_{i=1}^l \omega(\sigma_i)v_i = 0$. Then the facets around $\tilde{\tau}$ in $\tilde{\mathcal{X}}$ are $\tilde{\sigma}_1, \dots, \tilde{\sigma}_l$, and as primitive generators we can use the lifts

$$\tilde{v}_i := (v_i, df|_{\sigma_i}(v_i)) \in \mathbf{Z}^n \times \mathbf{Z}.$$

The weighted sum of primitive generators around $\tilde{\tau}$ gives

$$\begin{aligned} \sum_{i=1}^l \omega(\tilde{\sigma}_i)\tilde{v}_i &= \sum_{i=1}^l \omega(\sigma_i)(v_i, df|_{\sigma_i}(v_i)) \\ &= (0, \sum_{i=1}^l \omega(\sigma_i)df|_{\sigma_i}(v_i)) \\ &= (0, \omega(\tau)) \in \mathbf{Z}^n \times \mathbf{Z}, \end{aligned} \tag{4.7}$$




Figure 4.5: The modification of \mathbf{R}^2 along “ $(-1)x^2+1xy+(-1)y^2+x+y+0$ ”

where $\omega(\tau)$ denotes the weight of τ in $\text{div}_{\mathcal{X}}(f)$ defined in Equation (4.5). Hence a natural way to restore the balancing condition around $\tilde{\tau}$ is to add another facet

$$\tau_{\leq} = \tilde{\tau} + (\{0\} \times \mathbf{R}_{\leq 0}) = \{(x, y) \in \mathbf{R}^n \times \mathbf{R} : x \in \tau, y \leq f(x)\}$$

equipped with weight $\omega(\tau_{\leq}) = \omega(\tau)$ (see Figure 4.5). Let us summarize the discussion in the following definition.

Definition 4.5.1

Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on tropical k -cycle and let \mathcal{X} be a sufficiently fine polyhedral structure. The *(open) modification of \mathcal{X} along f* is the weighted polyhedral complex

$$\text{Mod}(\mathcal{X}, f) := \{\tilde{\sigma} : \sigma \in \mathcal{X}\} \cup \{\tau_{\leq} : \tau \in \text{div}_{\mathcal{X}}(f)\},$$

where $\tilde{\sigma} = \Gamma(\sigma, f|_{\sigma})$ and $\tau_{\leq} = \tilde{\tau} + (\{0\} \times \mathbf{R}_{\leq 0})$. The weights are given by $\omega(\tilde{\sigma}) = \omega_{\mathcal{X}}(\sigma)$ (for facets $\sigma \in \mathcal{X}$) and $\omega(\tau_{\leq}) = \omega_{\text{div}_{\mathcal{X}}(f)}(\tau)$ (for facets $\tau \in \text{div}_{\mathcal{X}}(f)$).

Lemma 4.5.2

The weighted polyhedral complex $\text{Mod}(\mathcal{X}, f)$ in $\mathbf{R}^n \times \mathbf{R}$ is balanced.

Proof. The $k-1$ -cells of $\text{Mod}(\mathcal{X}, f)$ are of two types. For $\tilde{\tau}, \tau$ an $k-1$ -cell of \mathcal{X} , the balancing condition follows from the computation in Equation (4.7). Note that $(0, -1) \in \mathbf{Z}^n \times \mathbf{Z}$ is a primitive generator for τ_{\leq} modulo $\tilde{\tau}$.

The second type of codimension cells is of the form ρ_{\leq} for a $k - 2$ -cell $\rho \in \mathcal{X}$. In this case, obviously the equality

$$\text{Star}_{\text{Mod}(\mathcal{X}, f)}(\rho_{\leq}) = \text{Star}_{\text{div}_{\mathcal{X}}(f)}(\rho)$$

holds. Hence the balancing condition around ρ_{\leq} is equivalent to the balancing condition around ρ in $\text{div}_{\mathcal{X}}(f)$, which was proven in Proposition 4.4.7. \square

Definition 4.5.3

Let $f : X \rightarrow \mathbf{R}$ be a piecewise integer affine function on tropical k -cycle X . The (open) modification $\text{Mod}(X, f)$ of X along f is the tropical cycle in $\mathbf{R}^n \times \mathbf{R}$ defined by $\text{Mod}(\mathcal{X}, f)$ for any sufficiently fine polyhedral structure \mathcal{X} . The map $\pi : \text{Mod}(\mathcal{X}, f) \rightarrow X$ is the contraction of the modification.

Remark 4.5.4

Let us make a few remarks here.

- It follows from Equation (4.7) that $\text{Mod}(\mathcal{X}, f)$ is the *unique* balanced completion of the graph $\Gamma(X, f)$ under the condition that we only allow adding cells in the “downward” direction (more precisely, only polyhedra P such that $(0, -1) \in \text{rc}(P)$).
- Let us, for the moment, consider the closure $\overline{\text{Mod}(X, f)}$ of $\text{Mod}(X, f)$ in $\mathbf{R}^n \times \mathbf{T}$. Let us denote by $H_{-\infty} = \mathbf{R}^n \times \{-\infty\} \cong \mathbf{R}^n$ the hyperplane added at infinity. Then by construction the intersection $\overline{\text{Mod}(X, f)} \cap H_{-\infty}$ is equal to $\text{div}_X(f)$, and this equality is respecting the weights (in an obvious sense). See Figure 4.6 for an example. Hence, in agreement with Remark 2.4.9, the modification construction justifies to some extent our definition of $\text{div}(f)$ as locus of (tropical) zeros and poles of f . Here, poles are actually described in terms of zeros with negative weight. It might be tempting to use facets in upward direction $(0, 1)$ instead of negative weights. Note, however, that for a general piecewise integer affine function there is no globally consistent construction which does so.
- Let $X \subset (\mathbf{C}^\times)^n$ be a classical very affine variety and $f : X \rightarrow \mathbf{C}$ a regular (or rational) function. Then taking the very affine graph $\Gamma(X, f) \cap$




Figure 4.6: The modification of \mathbf{R}^2 along the standard line is the standard plane.

$(\mathbf{C}^\times)^{n+1}$ is isomorphic X after removing the zero locus $V(f)$ (or the divisor of zeros and poles of f , respectively). Hence, the process of tropical open modifications can be regarded as the result of a very affine embedding of a principally open subset of X .

Assume that $f = "g/h" : X \rightarrow \mathbf{R}$ is a rational function for two polynomial functions g, h on X . In this case there exists an alternative version of the modification along f which *does* distinguish between zeros and poles, in contrast to the previous remark.

Definition 4.5.5

Let $g, h : X \rightarrow \mathbf{R}$ be two polynomial functions on a tropical subspace X . Let \mathcal{X} be a polyhedral structure sufficiently fine for both g and h . The *(open) modification* of \mathcal{X} along the quotient of g by h is the weighted polyhedral complex

$$\text{Mod}(\mathcal{X}, g, h) := \{\tilde{\sigma} : \sigma \in \mathcal{X}\} \cup \{\tau_{\leq} : \tau \in \text{div}_{\mathcal{X}}(g)\} \cup \{\tau_{\geq} : \tau \in \text{div}_{\mathcal{X}}(h)\},$$

where $\tilde{\sigma} = \Gamma(\sigma, (g-h)|_{\sigma})$, $\tau_{\leq} = \tilde{\tau} + (\{0\} \times \mathbf{R}_{\leq 0})$ and $\tau_{\geq} = \tilde{\tau} + (\{0\} \times \mathbf{R}_{\geq 0})$. The weights are given by $\omega(\tilde{\sigma}) = \omega_{\mathcal{X}}(\sigma)$, $\omega(\tau_{\leq}) = \omega_{\text{div}_{\mathcal{X}}(g)}(\tau)$ and $\omega(\tau_{\geq}) = \omega_{\text{div}_{\mathcal{X}}(h)}(\tau)$, respectively. The *(open) modification* $\text{Mod}(\mathcal{X}, g, h)$ of X along the quotient of g by h is the tropical cycle defined by $\text{Mod}(\mathcal{X}, g, h)$.

Remark 4.5.6

Again, a few remarks.

- It is straightforward to adapt the proof of Lemma 4.5.2 to show that $\text{Mod}(\mathcal{X}, g, h)$ is balanced and hence $\text{Mod}(\mathcal{X}, g, h)$ is well-defined as a tropical cycle.
- The construction does depend on the choice of g, h and not only on “ g/h ”. Also, in general no canonical minimal choice for g, h given “ g/h ” might be available.
- Let us identify the open torus in \mathbf{TP}^1 with \mathbf{R} by mapping $(x : y) \mapsto x - y$. Then the map $g - h : X \rightarrow \mathbf{R}$ can be identified with map $(g, h) : X \rightarrow \mathbf{TP}^1$. This alternative description shows that $\text{Mod}(\mathcal{X}, g, h)$ can be regarded as a tropical blow-up of X at $V(g) \cap V(h)$.
- We have met the construction of $\text{Mod}(\mathcal{X}, g, h)$ before in our discussion of floor decompositions. Let us recall the notations from Proposition 2.6. Let $V(f)$ be a floor-decomposed hypersurface in \mathbf{R}^n , $f = \sum_{i=0}^m f_i x_n^i$, $0 \neq f_i \in \mathbf{T}[x_1, \dots, x_{n-1}]$. Let F_1, \dots, F_m and E_0, \dots, E_m denote the corresponding floors and elevators, respectively. Then for all $m = 1, \dots, m$ we have

$$E_{i-1}, F_i, E_i \subset \text{Mod}(\mathbf{R}^n, f_{i-1}, f_i)$$

by Proposition 2.6.2. In other words, $\text{Mod}(\mathbf{R}^n, f_{i-1}, f_i)$ is equal to the floor F_i with elevators E_{i-1} and E_i attached and extended upwards and downwards to infinity. The operation of “gluing together” the various modifications $\text{Mod}(\mathbf{R}^n, f_{i-1}, f_i)$ to obtain $V(f)$ is analogous to the symplectic sum operation in symplectic geometry.

Exercise 4.5.7

Show that both $\text{Mod}(X, F)$ and $\text{Mod}(X, g, h)$ are independent of the chosen (sufficiently fine) polyhedral structure \mathcal{X} .

4.6 The projection formula

The goal of this section is to establish a simple type of projection formula which relates divisors of pull-backs of piecewise integer functions with push

forwards of tropical cycles. We start by introducing the necessary terminology.

4.6.1 Integer affine maps

Recall that an integer affine map $\Phi : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a map of the form $x \mapsto Ax + b$ where $A \in \text{Mat}(m \times n, \mathbf{Z})$ and $b \in \mathbf{R}^m$. Let $X \subset \mathbf{R}^n$ and $Y \subset \mathbf{R}^m$ be subsets. A map $\varphi : X \rightarrow Y$ is *integer affine* if it is the restriction of an integer affine map $\Phi : \mathbf{R}^n \rightarrow \mathbf{R}^m$.

Lemma 4.6.1

Let $X \subset \mathbf{R}^n$ and $Y \subset \mathbf{R}^m$ be polyhedral sets and $\varphi : X \rightarrow Y$ be an integer affine map. Then there exist polyhedral structures \mathcal{X} of X and \mathcal{Y} of Y such that $\varphi(\sigma) \in \mathcal{Y}$ for all $\sigma \in \mathcal{X}$.

Proof. We choose polyhedral structures \mathcal{X}' of X and \mathcal{Y}' of Y and $\Phi : \mathbf{R}^n \rightarrow \mathbf{R}^m$ an integer affine map with $\Phi|_X = \varphi$. Since $\varphi(\sigma') \subset \mathbf{R}^m$ is a polyhedron for all $\sigma' \in \mathcal{X}'$, we can write $|Y|$ as union of polyhedra

$$|Y| = \bigcup_{\tau' \in \mathcal{Y}'} \tau' \cup \bigcup_{\sigma' \in \mathcal{X}'} \varphi(\sigma'). \quad (4.8)$$

Applying Proposition 4.1.1 (a) to this union, we obtain a polyhedral subdivision

$$\mathcal{S} = \mathcal{G}_{\kappa_1} \cap \dots \cap \mathcal{G}_{\kappa_l},$$

of \mathbf{R}^m such that the polyhedra in Equation (4.8) are unions of cells in \mathcal{S} . The pull back subdivision of \mathbf{R}^n is

$$\mathcal{T} = \mathcal{G}_{\kappa_1 \circ \Phi} \cap \dots \cap \mathcal{G}_{\kappa_l \circ \Phi}$$

A simple computation shows that $\mathcal{Y} = \mathcal{S} \cap \mathcal{Y}' = \mathcal{S}|_Y$ and $\mathcal{X} = \mathcal{T} \cap \mathcal{X}'$ fulfil the required condition. \square

If \mathcal{X} and \mathcal{Y} satisfy the conditions of the lemma, they are called *sufficiently fine* for φ . We denote by $\varphi_*(\mathcal{X})$ the part of the subcomplex $\{\varphi(\sigma) : \sigma \in \mathcal{X}\} \subset \mathcal{Y}$ of pure dimension $\dim(X)$. Note that the facets of $\varphi_*(\mathcal{X})$ are of the form $\varphi(\sigma)$ for a facet $\sigma \in \mathcal{X}$ for which $\varphi|_\sigma$ is injective.

4.6.2 Push forward of tropical cycles

Definition 4.6.2

Let $\varphi : X \rightarrow Y$ be an integer affine map between tropical cycles and let \mathcal{X} and \mathcal{Y} be sufficiently fine polyhedral structures. The *push forward* of \mathcal{X} along φ is the polyhedral complex $\varphi_*(\mathcal{X})$ whose facets σ' carry the weight

$$\omega_{\varphi_*(\mathcal{X})}(\sigma') = \sum_{\substack{\sigma \in \mathcal{X} \\ \varphi(\sigma) = \sigma'}} \omega_{\mathcal{X}}(\sigma) \cdot [\mathbf{L}_{\mathbf{Z}}(\sigma') : d\varphi|_{\sigma}(\mathbf{L}_{\mathbf{Z}}(\sigma))]. \quad (4.9)$$

The *push forward* $\varphi_*(X)$ is the tropical cycle defined by $\varphi_*(\mathcal{X})$.

Note that since σ' is a facet of $\varphi_*(\mathcal{X})$ and $\varphi(\sigma) = \sigma'$, it follows that $\varphi|_{\sigma}$ is injective and hence $[\mathbf{L}_{\mathbf{Z}}(\sigma') : d\varphi|_{\sigma}(\mathbf{L}_{\mathbf{Z}}(\sigma))] \in \mathbf{Z}$ is finite. To show that $\varphi_*(\mathcal{X})$ defines a tropical cycle, we need to prove the following lemma.

Lemma 4.6.3

The weighted polyhedral complex $\varphi_(\mathcal{X})$ is balanced.*

Proof. By Exercise 4.6.10 it suffices to prove the local case. More precisely, we can restrict to the case where \mathcal{X} is a one-dimensional fan. Let ρ_1, \dots, ρ_l denote the rays of \mathcal{X} with primitive generators v_1, \dots, v_l . We may assume $\varphi(v_i) = 0$ if and only if $i = k + 1, \dots, l$ for some $0 \leq k \leq l$. Then the rays of $\varphi_*(\mathcal{X})$ are spanned by the vectors $\varphi(v_i)$, $i = 1, \dots, k$, and $\varphi(v_i)$ is equal to $[\mathbf{L}_{\mathbf{Z}}(\varphi(\rho_i)) : \varphi(\mathbf{L}_{\mathbf{Z}}(\rho_i))] \cdot w_i$, where w_i is primitive. Hence the balancing condition for $\varphi_*(\mathcal{X})$ turns into

$$\begin{aligned} \sum_{i=1}^k \omega(\varphi(\rho_i))w_i &= \sum_{i=1}^k \omega(\rho_i)\varphi(v_i) \\ &= \sum_{i=1}^l \omega(\rho_i)\varphi(v_i) = \varphi\left(\sum_{i=1}^l \omega(\rho_i)v_i\right) = 0, \end{aligned}$$

which finishes the proof. \square

Remark 4.6.4

Let us make a few remarks.

- Note that the support of a push forward satisfies $|\varphi_*(X)| \subset \varphi(|X|) \subset |Y|$.

- If $\varphi(|X|)$ is a polyhedral set of dimension strictly lower than $\dim(X)$, then $\varphi_*(X) = 0$.
- An integer affine map $\varphi : X \rightarrow Y$ induces a group homomorphism $\varphi_* : Z_k(X) \rightarrow Z_k(Y)$ given by $Z \mapsto \varphi_*(Z) := (\varphi|_Z)_*(Z)$. This follows directly from the linearity of the weight formula 4.9.
- Push forwards are functorial. This is a straightforward computation using the multiplicativity of indices

$$[\mathbf{Z}^n : AB(\mathbf{Z}^n)] = [\mathbf{Z}^n : A(\mathbf{Z}^n)] \cdot [\mathbf{Z}^n : B(\mathbf{Z}^n)]$$

for two (injective) integer linear maps $A, B : \mathbf{R}^n \rightarrow \mathbf{R}^n$ (or, equivalently, $|\det(AB)| = |\det(A)| \cdot |\det(B)|$).

Example 4.6.5

Let us consider the standard line $X = L_1 = V("0+x+y") \subset \mathbf{R}^2$ and the map $\varphi : \mathbf{R}^2 \rightarrow \mathbf{R}, (x, y) \mapsto x+y$. The standard structure on L_1 and the polyhedral structure $\{\mathbf{R}_{\geq 0}, \mathbf{R}_{\leq 0}, \{0\}\}$ are sufficiently fine for $\varphi|_X$. Let ρ_0, ρ_1, ρ_2 be the rays of X spanned by $(1, 1), (-1, 0), (0 - 1)$. The generators are mapped under φ to

$$(-1, 0), (0, -1) \mapsto -1 \qquad (1, 1) \mapsto 2.$$

Therefore the weights of $\varphi_*(X)$ are given by

$$\begin{aligned} \omega(\mathbf{R}_{\leq 0}) &= [\mathbf{Z} : \varphi(L_Z(\rho_1))] + [\mathbf{Z} : \varphi(L_Z(\rho_2))] &= 1 + 1 = 2, \\ \omega(\mathbf{R}_{\geq 0}) &= [\mathbf{Z} : \varphi(L_Z(\rho_0))] &= 2, \end{aligned}$$

and so $\varphi_*(X) = 2[\mathbf{R}]$.

4.6.3 The projection formula

Definition 4.6.6

Let $\varphi : X \rightarrow Y$ be an integer affine map between tropical cycles and let $f : Y \rightarrow \mathbf{R}$ be a piecewise integer affine function on Y . The *pull back* of f along φ is the piecewise integer affine function $\varphi^*f = f \circ \varphi$ on X .

Remark 4.6.7

In order to construct a sufficiently fine polyhedral structure for φ^*f , it suffices to start with a sufficiently fine structure \mathcal{Y}' for f and then apply the construction of Lemma 4.6.1.

We are now ready to state the main the projection formula.

Theorem 4.6.8 (Projection formula)

Let $\varphi : X \rightarrow Y$ be a tropical morphism and let $f : Y \rightarrow \mathbf{R}$ be a piecewise integer affine function on Y . Then the following formula holds:

$$f \cdot \varphi_*(X) = \varphi_*(\varphi^*(f) \cdot X) \in Z_{\dim X - 1}(Y)$$

Proof. We choose polyhedral structures \mathcal{X} and \mathcal{Y} which are sufficiently fine for φ , f and hence also $\varphi^*(f)$ (see Remark 4.6.7). Set $k = \dim(X)$. We have to compare the weights of a cell $\sigma' \in \mathcal{Y}^{(k-1)}$ on both sides. Note that this weights can be computed locally on both sides. Therefore it suffices to treat the case where \mathcal{X} is a one-dimensional fan, $\varphi : X \rightarrow Y$ is integer linear and f is a piecewise linear function on $f_*(\mathcal{X})$. Given a ray ρ , we use the shorthand $v_{\rho/\{0\}} = v_\rho$. Denoting the rays of \mathcal{X} and $\varphi_*(\mathcal{X})$ by ρ and ρ' , respectively, we obtain

$$\begin{aligned} \omega_{\varphi_*(\varphi^*(f) \cdot \mathcal{X})}(\{0\}) &= \omega_{\varphi^*(f) \cdot \mathcal{X}}(\{0\}) \\ &= \sum_{\rho \in \mathcal{X}} \omega_{\mathcal{X}}(\rho) f(\varphi(v_\rho)) \\ &= \sum_{\rho' \in f_*(\mathcal{X})} \sum_{\substack{\rho \in \mathcal{X} \\ \varphi(\rho) = \rho'}} \omega_{\mathcal{X}}(\rho) [L_{\mathbf{Z}}(\rho') : \varphi(L_{\mathbf{Z}}(\rho))] f(v_{\rho'}) \\ &= \sum_{\rho' \in f_*(\mathcal{X})} \omega_{\varphi_*(\mathcal{X})}(\rho') f(v_{\rho'}) \\ &= \omega_{f \cdot \varphi_*(\mathcal{X})}(\{0\}). \end{aligned}$$

□

Exercise 4.6.9

Prove that the definition of the push forward $\varphi_*(X)$ is independent of the (sufficiently fine) polyhedral structures \mathcal{X} and \mathcal{Y} .

Exercise 4.6.10 (Locality of push forwards)

Let $\varphi : X \rightarrow Y$ be an integer affine map between tropical cycles and let \mathcal{X} and \mathcal{Y} be sufficiently fine polyhedral structures.

- (a) For any cell $\tau \in \mathcal{X}$, show that (the linear part of) φ induces a canonical “local” integer affine map $\varphi^\tau : \text{Star}_X(\tau) \rightarrow \text{Star}_Y(\varphi(\tau))$ for which $\text{Star}_{\mathcal{X}}(\tau)$ and $\text{Star}_{\mathcal{Y}}(\varphi(\tau))$ are sufficiently fine.
- (b) Show that for any cell $\tau' \in \mathcal{Y}$ of dimension m , we have

$$\text{Star}_{\varphi_*(\mathcal{X})}(\tau) = \sum_{\substack{\tau \in \mathcal{X}^{(m)} \\ \varphi(\tau) = \tau'}} [\text{L}_Z(\tau') : d\varphi|_\tau(\text{L}_Z(\tau))] \cdot \varphi^\tau_*(\text{Star}_{\mathcal{X}}(\tau)).$$

4.7 Examples and further topics

4.7.1 Tropical Bernstein-Kouchnirenko Theorem

The Bernstein-Kouchnirenko Theorem states that the number of common zeros in $(\mathbf{C}^\times)^n$ of a system of n generic Laurent polynomials in n variables is equal to the mixed volume of the associated Newton polytopes (see [Ber76; Kou76]). The *mixed volume* $\text{MV}(P_1, \dots, P_n)$ can be defined as the unique function evaluated on n -tuples of (lattice) polytopes which is

- (a) symmetric,
- (b) multilinear, i.e. for all $\lambda, \mu \in \mathbf{N}$

$$\text{MV}(\lambda P_1 + \mu P'_1, \dots, P_n) = \lambda \text{MV}(P_1, \dots, P_n) + \mu \text{MV}(P'_1, \dots, P_n),$$

- (c) normalized by $\text{MV}(P, \dots, P) = \text{Vol}(P)$.

In this subsection, we want to prove the following tropical analogue.

Theorem 4.7.1

Let $f_1, \dots, f_n \in \mathbf{T}[x_1^\pm, \dots, x_n^\pm]$ be tropical Laurent polynomials with associated Newton polytopes P_1, \dots, P_n . Then

$$\deg(f_1 \cdots f_n \cdot \mathbf{R}^n) = \text{MV}(P_1, \dots, P_n).$$

We start by showing that the degree on the left hand side only depends on the Newton polytopes of the polynomials, not the polynomials themselves. We put the details of this verification in exercises below. For a lattice polytope P with vertices j_1, \dots, j_k , we set

$$f_P = \sum_{i=1}^k x^{j_i} \mathbf{T}[x_1^\pm, \dots, x_n^\pm].$$

Lemma 4.7.2

The degree of tropical 0-cycle $f_1 \cdots f_n \cdot \mathbf{R}^n$ only depends on the Newton polytopes P_1, \dots, P_n .

Proof. We set $g_i := f_{P_i}$, $i = 1, \dots, n$. Since $\text{NP}(f_i) = P_i = \text{NP}(g_i)$, by Exercise 4.7.4 (a) the functions $f_i - g_i$ are bounded. Applying Exercise 4.7.4 (b) to the 1-cycle $X = f_2 \cdots f_n \cdot \mathbf{R}^n$, we obtain

$$\deg(f_1 \cdots f_n \cdot \mathbf{R}^n) = \deg(g_1 \cdot f_2 \cdots f_n \cdot \mathbf{R}^n) = \deg(f_2 \cdots f_n \cdot g_1 \cdot \mathbf{R}^n),$$

where we used Proposition 4.4.18 in the second equality. Repeating this procedure we arrive at $\deg(f_1 \cdots f_n \cdot \mathbf{R}^n) = \deg(g_1 \cdots g_n \cdot \mathbf{R}^n)$, and the claim follows. \square

Proof of Theorem 4.7.1. By what we have said before, it suffices to prove the following: Given lattice polytopes P_1, \dots, P_n and setting $f_i = f_{P_i}$, the function $\deg(f_1 \cdots f_n \cdot \mathbf{R}^n)$ satisfies the characteristic properties of the mixed volume function on page 136. Symmetry follows from Proposition 4.4.18, multilinearity follows from

$$f_{\lambda P + \mu P'} = \lambda f_P + \mu f_{P'}$$

and the linearity of $f \mapsto f \cdot X$ (see Remark 4.4.15). It remains to prove normalization, which we put as an extra theorem for clarity. \square

Theorem 4.7.3

Let $f \mathbf{T}[x_1^\pm, \dots, x_n^\pm]$ be a tropical Laurent polynomial with associated Newton polytope P . Then $\deg(f \cdot \mathbf{R}^n) = \text{Vol}(P)$.

Proof. To prove the theorem, we proceed as follows. First, recall that any lattice polytope P admits a regular triangulation \mathcal{T} (choose generic lift of

the vertices to \mathbf{R}^{n+1} , see e.g. [GKZ08, Proposition 1.5]). Let g be a tropical polynomial with $\text{SD}(g) = \mathcal{T}$. Using Lemma 4.7.2 again, it suffices to show $\deg(g^n \cdot \mathbf{R}^n)$. But $g^n \cdot \mathbf{R}^n$ consists of a finite number of points dual to the full-dimensional simplices of \mathcal{T} . Since the weight of these points in $g^n \cdot \mathbf{R}^n$ can be computed locally, we reduced the statement to the case of simplices.

Let us assume that $P = \Delta$ is a lattice simplex and $f = f_\Delta$. Let us first consider the case $f = f_1 := "0 + x_1 + \dots + x_n"$, i.e. $\Delta = \Delta_1$ is equal to the minimal standard simplex. Then the claim $\deg(f_1^n \cdot \mathbf{R}^n) = 1$ follows from a direct calculation (or from Examples 4.3.8 and 4.4.20).

Now let us go back to the case of an arbitrary lattice simplex Δ . Note that there exists a unique integer affine map $\psi : (\mathbf{R}^n)^* \rightarrow (\mathbf{R}^n)^*$ such that $\psi(\Delta_1) = \Delta$. Let $\varphi : \mathbf{R}^n \rightarrow \mathbf{R}^n$ denote the dual map such that $\varphi^*(f_1) = f$. Then the volume of Δ can be expressed as $\text{Vol}(\Delta) = [\mathbf{Z}^n : d\psi(\mathbf{Z}^n)] = [\mathbf{Z}^n : d\varphi(\mathbf{Z}^n)]$. Finally, using the projection formula 4.6.8 we get

$$\begin{aligned} \deg(f^n \cdot \mathbf{R}^n) &= \deg(\varphi_*(\varphi^*(f_1)^n \cdot \mathbf{R}^n)) = \deg(f_1^n \cdot \varphi_*(\mathbf{R}^n)) \\ &= \text{Vol}(\Delta) \deg(f_1^n \cdot \mathbf{R}^n) = \text{Vol}(\Delta), \end{aligned}$$

which finishes the proof. □

Exercise 4.7.4

Show that the following statements hold.

- (a) Let $f, g \in \mathbf{T}[x_1^\pm, \dots, x_n^\pm]$ be tropical Laurent polynomials. If $\text{NP}(f) = \text{NP}(g)$, then the function " f/g " = $f - g$ is bounded on \mathbf{R}^n .
- (b) Let $X \subset \mathbf{R}^n$ be a one-dimensional tropical cycle and let $h : X \rightarrow \mathbf{R}$ be a bounded rational function, then $\deg(h \cdot X) = 0$.

4.7.2 Volumes of faces of polytopes

There is an interesting equation satisfied by the volumes of the faces of a (lattice) polytope which can be expressed in terms of the tropical balancing condition.

Theorem 4.7.5

Let $P \in \mathbf{R}^n$ be a lattice polytope and let $\mathcal{N} = \mathcal{N}(P)$ be its normal fan. For a

face F of P , we denote the dual cone in \mathcal{N} by σ_F . Then for any $k = 0, \dots, n$, the polyhedral fan $\mathcal{N}^{(k)}$ with weights

$$\omega(\sigma_F) = \text{Vol}(F)$$

is balanced.

We may generalize this by considering the subdivision $\mathcal{S}(f)$ of \mathbf{R}^n associated to a tropical Laurent polynomial f . Recall from Section 2.4 that we can equip each cell $\sigma \in \mathcal{S}(f)$ with the weight $\omega(\sigma) = \text{Vol}(D_\sigma)$, where $D_\sigma \in \text{SD}(f)$ is the dual cell to σ . In Theorem 2.4.7 we proved that the hypersurface complex $\mathcal{S}(f)^{(n-1)}$ is balanced. We now extend this to arbitrary codimensions.

Theorem 4.7.6

The polyhedral complex $\mathcal{S}(f)^{(k)}$ with weights $\omega(\sigma) = \text{Vol}(D_\sigma)$ is balanced for all $k = 0, \dots, n$.

It is clear that Theorem 4.7.6 follows from Theorem 4.7.5 applied locally. On the other hand, both statements have a nice conceptual explanation in terms of tropical intersection products.

Theorem 4.7.7

Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a tropical Laurent polynomial with associated subdivision $\mathcal{S}(f)$. Then for all $k = 0, \dots, n$, the intersection product $f^{n-k} \cdot \mathcal{S}(f)$ is equal to the polyhedral complex $\mathcal{S}(f)^{(k)}$ with weights $\omega(\sigma) = \text{Vol}(D_\sigma)$.

Proof. Using the local computation of $\omega(\sigma)$ in $f^{n-k} \cdot \mathcal{S}(f)$, the statement follows immediately from Theorem 4.7.3. □

Proof of Theorems 4.7.5 and 4.7.6. We know that $f^{n-k} \cdot \mathcal{S}(f)$ is balanced by Proposition 4.4.7. This implies Theorem 4.7.6. The special case $f = f_p$ proves Theorem 4.7.5. □

4.7.3 Intersecting with the diagonal

An alternative way for constructing stable intersections is given in [AR10]. The approach is based on intersecting with the diagonal $\Delta \subset \mathbf{R}^n \times \mathbf{R}^n$ and only uses the divisor construction while circumventing the moving lemma.

The approach is more formal but sometimes simplifies arguments (for example, for showing the compatibility of the divisor construction with stable intersection) can be generalized to situations where usual stable intersection (i.e. moving cycles around) is not available. We briefly state the main ideas here.

Given tropical $X \subset \mathbf{R}^n$ and $Y \subset \mathbf{R}^m$, the cartesian product $X \times Y \subset \mathbf{R}^{n+m}$ with weights $\omega((x, y)) = \omega(x)\omega(y)$ is a tropical cycle of dimension $\dim(X) + \dim(Y)$. We are particularly interested in the case $n = m$. Let us denote by x_1, \dots, x_n and y_1, \dots, y_n the coordinates of the first and second factor of $\mathbf{R}^n \times \mathbf{R}^n$, respectively.

Definition 4.7.8

Given a tropical cycle $Z \subset \mathbf{R}^n \times \mathbf{R}^n$ of dimension k , the *intersection with the diagonal* is

$$\Delta \cdot Z := \max\{x_1, y_1\} \cdots \max\{x_n, y_n\} \cdot Z \in Z_{k-n}(\mathbf{R}^{2n}).$$

Given two tropical cycles $X, Y \subset \mathbf{R}^n$ of dimension k and l , respectively, the Δ -*intersection product* of X and Y is

$$X.Y := \pi_*(\Delta \cdot (X \times Y)) \in Z_{k+l-n}(\mathbf{R}^n),$$

where $\pi : \mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}^n$ denotes the projection to the first factor.

Note that $\Delta \cdot \mathbf{R}^{2n} = \{(x, x) : x \in \mathbf{R}\}$ with constant weight 1. It is easy to check that the definitions are invariant under (integer invertible) coordinate changes and switching the two factors of \mathbf{R}^n . Moreover, $X.Y$ coincides with stable intersection.

Theorem 4.7.9

For any two tropical cycles $X, Y \subset \mathbf{R}^n$ we have $X.Y = X \cdot Y$.

Proof. It is easy to check that $X.Y$ satisfies the same locality formula as $X \cdot Y$ (see Exercise 4.7.10). It therefore suffices to prove the case $k + l = n$ and X, Y fan cycles. Since by Exercise 4.7.11 $X.Y$ and $X \cdot Y$ agree in the transversal case, it suffices to show $\deg(X.Y) = \deg(X.(Y + \nu))$ for any $\nu = (\nu_1, \dots, \nu_n) \in \mathbf{R}^n$. But note that $X.(Y + \nu)$ can be described as $\pi_*(\Delta_\nu \cdot (X \times Y))$ for the modified diagonal

$$\Delta_\nu \cdot Z := \max\{x_1, y_1 + \nu_1\} \cdots \max\{x_n, y_n + \nu_n\} \cdot Z.$$

Obviously, the functions $\max\{x_i, y_i + v_i\} - \max\{x_i, y_i\}$ are bounded for all $i = 1, \dots, n$. Using Exercise 4.7.4 and the trick in the proof of Lemma 4.7.2 this implies $\deg(\Delta \cdot (X \times Y)) = \deg(\Delta_v \cdot (X \times Y))$, which proves the claim. \square

Exercise 4.7.10

Let $X, Y \subset \mathbf{R}^n$ be two tropical cycles with polyhedral structures \mathcal{X} and \mathcal{Y} , respectively. Then for any cell $\tau \in (\mathcal{X} \cap \mathcal{Y})^{(k+l-n)}$ we have

$$\text{Star}_{\mathcal{X} \cdot \mathcal{Y}}(\tau) = \text{Star}_{\mathcal{X}}(\tau) \cdot \text{Star}_{\mathcal{Y}}(\tau).$$

Exercise 4.7.11

Show the following statements.

- (a) Let $\lambda_1, \dots, \lambda_n : \mathbf{R}^n \rightarrow \mathbf{R}$ be integer linear forms such that $\langle \lambda_1, \dots, \lambda_n \rangle_{\mathbf{R}} = (\mathbf{R}^n)^*$. Then

$$\max\{\lambda_1, 0\} \cdots \max\{\lambda_n, 0\} \cdot \mathbf{R}^n$$

is equal to the origin equipped with weight $[(\mathbf{Z}^n)^* : \langle \lambda_1, \dots, \lambda_n \rangle_{\mathbf{Z}}]$.

- (b) Assume that X and Y are affine subspaces of \mathbf{R}^n (equipped with constant weight 1) such that $L(X) + L(Y) = \mathbf{R}^n$. Then $X \cdot Y$ is equal to the affine space $X \cap Y$ equipped with weight $[\mathbf{Z}^n : L_{\mathbf{Z}}(X) + L_{\mathbf{Z}}(Y)]$.
- (c) Let $X, Y \subset \mathbf{R}^n$ be two tropical cycles intersecting transversally. Then $X \cdot Y = X \cdot Y$.

4.7.4 Irreducible tropical cycles

A tropical k -cycle X is *irreducible* if $Z_k(X) = \mathbf{Z}[X]$. A tropical cycle X of dimension k is *\mathbf{Q} -irreducible* if $Z_k(X) \otimes \mathbf{Q} = \mathbf{Q}[X]$. In this case, the underlying polyhedral set $|X|$ is called *irreducible*. In other words, there is a unique balanced weight function on $|X|$ up to multiplying with a constant. Clearly, if X is \mathbf{Q} -irreducible and $\gcd\{\omega_X(x) : x \in X^{\text{gen}}\} = 1$, then X is irreducible.

A tropical k -cycle X is *(\mathbf{Q} -)irreducible in codimension one* if for one (and then for any) polyhedral structure \mathcal{X} the stars $\text{Star}_{\mathcal{X}}(\tau)$ are (\mathbf{Q} -)irreducible for all $k-1$ -cells $\tau \in \mathcal{X}$. We call X *connected in codimension one* if for one (and then for any) polyhedral structure \mathcal{X} the set $|X| \setminus |\mathcal{X}^{(k-2)}|$ is connected.

Exercise 4.7.12

Show the two “and then for any” parts. Show the following statements. If X is (\mathbf{Q}) -irreducible and connected in codimension one, then X is (\mathbf{Q}) -irreducible.

Exercise 4.7.13

Show that if X is a k -cycle irreducible in codimension one and $f : X \rightarrow \mathbf{R}$ is a piecewise integer affine function, then $|\operatorname{div}(f)| = V(f)$. Let us denote by $\mathcal{A}ff(X) \subset \operatorname{PA}_{\mathbf{Z}}(X)$ the locally affine functions on X . Show that the sequence

$$0 \longrightarrow \mathcal{A}ff(X) \longrightarrow \operatorname{PA}_{\mathbf{Z}}(X) \xrightarrow{\operatorname{div}} Z_{k-1}(X) \longrightarrow 0$$

is exact. Give counterexamples to these statements for X irreducible, but not irreducible in codimension one.

Exercise 4.7.14

Construct an irreducible tropical fan k -cycle X and a fan $(k-1)$ -subcycle D such that there exists no $f \in \operatorname{PA}_{\mathbf{Z}}(X)$ such that $\operatorname{div}(f) = D$. Can you go on to X irreducible in codimension one and all weights are 1?

4.7.5 Toric intersection theory

Predecessors of some the constructions of this chapter can be found in intersection theory for (classical) toric varieties. Let us outline some of these connections. For simplicity, we restrict to compact varieties over the base field \mathbf{C} . For a more details on the presented material and generalizations, we refer the reader to [Rau16; Kat09; Rau09; Kat12] and references therein. For more background on toric (co-)homology theory, please consult [Dan78, Chapter 10], [Ful93, Chapter 5], [CLS11, Chapter 12].

Let Ξ be a complete pointed fan in \mathbf{R}^n . We denote by Ξ_k the subset of k -cones and by $F_k(\Xi)$ the free abelian group generated by Ξ_k , i.e. the group of formal linear combinations of k -cones with integer coefficients. We denote by $Z_k(\Xi)$ the group of balanced weight functions on the k -cones of Ξ . By the discussion in Section 4.1, $Z_k(\Xi)$ is equal to the group of tropical k -cycles supported on $|\Xi^{(k)}|$. Regarding $Z_k(\Xi)$ as a sublattice of $F_k(\Xi)^*$, we can replace the balancing condition by its dual version as follows. For any $(k-1)$ -cell τ , we regard the lattice $L_{\mathbf{Z}}(\tau)^\perp$ of integer linear functions λ that

vanish on $L_{\mathbf{Z}}(\tau)$. Evaluation of λ at primitive generators provides a map

$$\begin{aligned} \text{ev}_{\tau} : L_{\mathbf{Z}}(\tau)^{\perp} &\rightarrow F_k(\Xi), \\ \lambda &\mapsto \sum_{\substack{\sigma \in \Xi_k \\ \tau \subset \sigma}} \lambda(v_{\sigma/\tau})[\sigma]. \end{aligned}$$

Note that the choice of primitive generators does not matter since $\lambda|_{L_{\mathbf{Z}}(\tau)} = 0$. Obviously, a weight function $\omega \in F_k(\Xi)^*$ is balanced around τ if and only if $\omega \in \text{Im}(\text{ev}_{\tau})^{\perp}$. Let $R_k(\Xi) = \otimes_{\tau \in \Xi_{k-1}} \text{Im}(\text{ev}_{\tau})$ the sum of all these sublattices. It follows that

$$Z_k(\Xi) = R_k(\Xi)^{\perp} = \text{Hom}(F_k(\Xi)/R_k(\Xi), \mathbf{Z}). \quad (4.10)$$

Let now $\mathbf{X} = \mathbf{C}X_{\Xi}$ be the associated complex compact toric variety to Ξ . Then the objects from above have the following (mostly straightforward) reinterpretation. We denote the complex algebraic torus by $T = (\mathbf{C}^{\times})^n$.

- Ξ_k : closures of $(n - k)$ -dimensional torus orbits, \mathbf{X}_{σ}
- $F_k(\Xi)$: torus-invariant algebraic $(n - k)$ -cycles, generated by $[\mathbf{X}_{\sigma}]$
- $L_{\mathbf{Z}}(\tau)^{\perp}$: (monomial) rational functions x^{λ} on \mathbf{X} which are regular invertible on $\mathbf{C}U_{\tau}$
- $\text{ev}_{\tau}(\lambda)$: principal divisor of x^{λ} restricted to \mathbf{X}_{τ}
- $R_k(\Xi)$: torus-invariant cycles rationally equivalent to zero

Let $A_k(\mathbf{X})$ denote the k -th Chow group of \mathbf{X} . In our case, it can be completely described in terms of torus-invariant cycles and relations, hence $A_{(n-k)}(\mathbf{X}) = F_k(\Xi)/R_k(\Xi)$. Let $A^k(\mathbf{X})$ denote the “operational” Chow cohomology groups (see [Ful98]). For compact \mathbf{X} , these can be described as $A^k(\mathbf{X}) = \text{Hom}(A_k(\mathbf{X}), \mathbf{Z})$. In view of Equation 4.10, we find $A^{n-k}(\mathbf{X}) = Z_k(\Xi)$. In other words, Chow cohomology classes correspond to tropical fan cycles on $|\Xi^{(k)}|$, and the identification is given by $\omega_{\alpha}(\sigma) = \text{deg}(\alpha \cap [\mathbf{X}_{\sigma}])$ for any $\alpha \in A^{n-k}(\mathbf{X})$. Chow cohomology carries a ring structure given by the cup product $\cup : A^k(\mathbf{X}) \times A^l(\mathbf{X}) \rightarrow A^{k+l}(\mathbf{X})$. It turns out that under the identification $A^{n-k}(\mathbf{X}) = Z_k(\Xi)$, the cup product is replaced by stable intersection of tropical cycles (the “fan displacement rule”, see [FS97]).

We denote by $\mathrm{PL}_{\mathbf{Z}}(\Xi)$ the group of piecewise integer linear functions $\varphi : \mathbf{R}^n \rightarrow \mathbf{R}$ which are linear on the cones of Ξ (i.e., Ξ is sufficiently fine). The integer linear functions form a sublattice $(\mathbf{Z}^n)^* \subset \mathrm{PL}_{\mathbf{Z}}(\Xi)$. Again, there are toric wordings for that.

- $\mathrm{PL}_{\mathbf{Z}}(\Xi)$: torus-invariant Cartier divisors
- $(\mathbf{Z}^n)^*$: principal torus-invariant Cartier divisors

Let $\mathrm{Pic}(\mathbf{X})$ denote the Picard group of \mathbf{X} . Again, in the toric case it suffices to consider torus-invariant Cartier divisors and relations and therefore $\mathrm{Pic}(\mathbf{X}) = \mathrm{PL}_{\mathbf{Z}}(\Xi)/(\mathbf{Z}^n)^*$. There are natural maps $\mathrm{Pic}(\mathbf{X}) \rightarrow A^1(\mathbf{X})$ and more generally $\mathrm{Pic}(\mathbf{X}) \times A^k(\mathbf{X}) \rightarrow A^{k+1}(\mathbf{X})$ given by intersecting with Cartier divisors. It turns out that under the identification $\mathrm{Pic}(\mathbf{X}) = \mathrm{PL}_{\mathbf{Z}}(\Xi)/(\mathbf{Z}^n)^*$, these maps are replaced by the tropical divisor construction $\mathrm{PL}_{\mathbf{Z}}(\Xi) \times Z_{n-k}(\Xi) \rightarrow Z_{n-k-1}(\Xi)$, $(\varphi, X) \mapsto \varphi \cdot X$. Note that in particular $\lambda \cdot X = 0$ for $\lambda \in (\mathbf{Z}^n)^*$.

In codimension one, some interesting special cases occur. For example, the fact that the map $\mathrm{Pic}(\mathbf{X}) \rightarrow A^1(\mathbf{X})$ is an isomorphism can be deduced from Theorem 2.4.10 and Corollary 4.4.12 using the previous identifications. Indeed, a simple adaptation of the proofs shows that the map $\mathrm{PL}_{\mathbf{Z}}(\Xi) \rightarrow Z_{n-1}(\Xi)$, $\varphi \mapsto \mathrm{div}(\varphi)$, is surjective with kernel $(\mathbf{Z}^n)^*$, which implies $\mathrm{Pic}(\mathbf{X}) \cong A^1(\mathbf{X})$. Another interesting map is $\mathrm{Pic}(\mathbf{X}) \rightarrow A_{n-1}(\mathbf{X})$ given by sending a Cartier divisor to its associated Weil divisor. This relates to the map $\mathrm{ev} : \mathrm{PL}_{\mathbf{Z}}(\Xi) \rightarrow F_1(\Xi)$, $\varphi \mapsto \sum_{\rho \in \Xi_1} \varphi(v_\rho)[\rho]$. Clearly, this map is injective. Moreover, if Ξ is unimodular (or at least simplicial), then the map is bijective (bijective after tensoring with \mathbf{Q} , respectively). In this case, the inverse map is given by linear extension of φ on cones given its values on the generators. It follows that $\mathrm{Pic}(\mathbf{X}) \cong A_{n-1}(\mathbf{X})$ and $\mathrm{Pic}(\mathbf{X}) \otimes \mathbf{Q} \rightarrow A_{n-1}(\mathbf{X}) \otimes \mathbf{Q}$ if \mathbf{X} is smooth or \mathbf{Q} -smooth, respectively.

References

Parts of the theory presented here appeared in many papers in the early days of tropical geometry. Stable intersection in the setting of tropical cycles was introduced in [Mik06]. A formal approach to tropical intersection theory covering most of the material presented here was developed in [AR10; Rau16] (see also [Rau09; Kat09]).

5 Projective tropical geometry

5.1 Tropical cycles in toric varieties

5.1.1 The quick and formal definition

Our first objective is to define tropical cycles in tropical toric varieties, in particular, in \mathbf{T}^n and \mathbf{TP}^n . Again, there are several equivalent definitions, depending on our preferred viewpoint. Let us start with the quick and formal definition.

Let Ξ be a fan in \mathbf{R}^n . In this section, we denote its cones by $\vartheta, \rho \in \Xi$, to distinguish them from the cells $\tau, \sigma \in \mathcal{X}$ of polyhedral complexes to appear soon. We denote by $W = W_\Xi$ the tropical toric variety associated to Ξ . Let us recall that to any $\vartheta \in \Xi$ we can associate a torus orbit \mathbf{R}_ϑ as well as its closure $W_\vartheta = \overline{\mathbf{R}_\vartheta} \subset W$. They have dimension $\dim(W_\vartheta) = \dim(\mathbf{R}_\vartheta) = n - \dim(\vartheta)$. In particular, recall that $\mathbf{R}_\vartheta \cong \mathbf{R}^{n - \dim(\vartheta)}$.

Definition 5.1.1

A tropical k -cycle in W_Ξ is a formal sum

$$X = \sum_{\vartheta \in \Xi} X_\vartheta,$$

where each X_ϑ is a tropical k -cycle in \mathbf{R}_ϑ . Thus the group of tropical k -cycles in W_Ξ is

$$Z_k(W_\Xi) := \bigoplus_{\vartheta \in \Xi} Z_k(\mathbf{R}_\vartheta).$$

The support of X is $|X| = \bigcup_{\vartheta \in \Xi} \overline{X_\vartheta}$, where the closure is taken in W .

A k -cycle X is of pure sedentarity ϑ if it consists of a single summand $X = X_\vartheta$ (i.e. $X_\rho = 0$ for all $\vartheta \neq \rho \in \Xi$). Moreover, if $\rho = \{0\}$, X is of sedentarity zero.

Note that this definition, in the special case of $k = n - 1$, agrees with the definition of tropical divisors in Definition 3.4.4.

5.1.2 Polyhedra and their peers in toric varieties

It will be useful later to also extend the language of polyhedral complexes and sets to the toric setting. This will lead to other, equivalent descriptions of tropical k -cycles. A (generalized) polyhedron σ in W is the closure in W of a (usual) polyhedron $\sigma' \subset \mathbf{R}_\vartheta$ for some $\vartheta \in \Xi$. We refer to ϑ as the *sedentarity* of σ (and σ^m). Moreover, we call $\sigma^m := \sigma' = \sigma \cap \mathbf{R}_\vartheta$ the *mobile part* of σ , while $\sigma^s := \sigma \setminus \sigma^m$ is the *sedentary part* of σ . The *dimension* of σ is the dimension of σ^m . Analogously, we define the linear span $L(\sigma)$, recession cone $\text{rc}(\sigma)$, etc., by referring to the corresponding notions for σ^m .

There are two types of *faces* of σ : For any face $\tau' \subset \sigma^m$, we call the closure $\overline{\tau'} \subset \sigma \subset W$ a *mobile face of σ* (or *face of same sedentarity*). On the other hand, for all $\rho \in \Xi$ the intersection $\sigma \cap \mathbf{R}_\rho$ is either empty or a (usual) polyhedron in \mathbf{R}_ρ , see Proposition 3.2.6. If non-empty, we call $\overline{\sigma \cap \mathbf{R}_\rho} \subset \sigma \subset W$ a *sedentary face of σ* (or *face of higher sedentarity*). Note that in general $\sigma \cap W_\rho$ is a disjoint union of sedentary faces. Figure 5.1 an Example of a generalized polyhedron with its sedentary faces.




Figure 5.1: The closures of a standard polyhedron in \mathbf{T}^2 with its new faces

We can directly generalize the definitions of polyhedral complexes, equivalence of (weighted) complexes, polyhedral sets, etc. to the toric case by repeating the definitions word by word, only replacing polyhedra (and its faces) by its generalized versions of polyhedra in W . A polyhedral set $X \subset W$ is of *pure sedentarity* ϑ if $X = \overline{X \cap \mathbf{R}_\vartheta}$. In this case, the mobile and sedentary part of X are $X^m = X \cap \mathbf{R}_\vartheta$ and $X^s = X \setminus X^m$, respectively. A polyhedral complex \mathcal{X} is of pure sedentarity ϑ if $|\mathcal{X}|$ is so. The mobile and sedentary part of \mathcal{X} are $\mathcal{X}^m = \{\sigma \cap \mathbf{R}_\vartheta : \sigma \in \mathcal{X}\}$ and $\mathcal{X}^s = \{\sigma \in \mathcal{X} : \sigma \subset |\mathcal{X}|^s\}$,

respectively.

Given a weighted polyhedral complex \mathcal{X} in W (a polyhedral structure of a weighted polyhedral set X in W , respectively), the *balancing condition in W* requires that for all $\tau \in \mathcal{X}^{(k-1)}$ we have

$$\sum_{\substack{\sigma \supset \tau \\ \text{same sed.}}} \omega(\sigma) v_{\sigma/\tau} = 0 \pmod{L_{\mathbf{Z}}(\tau)}. \quad (5.1)$$

Here, the sum runs through all facets $\sigma \in \mathcal{X}$ containing τ and of *same sedentarity*, in which case the notion of primitive generator $v_{\sigma/\tau} := v_{\sigma^m/\tau^m}$ is well-defined. Note that if τ is of sedentarity ϑ , then the computation takes place completely inside $\mathbf{Z}_{\vartheta} \subset \mathbf{R}_{\vartheta}$. We can now reformulate Definition 5.1.1 in the following equivalent ways (cf. 4.2.2 and 4.2.3).

Definition 5.1.2

A *tropical k -cycle X in W* is the equivalence class of balanced polyhedral complex \mathcal{X} in W of pure dimension k . Equivalently, X is a balanced polyhedral set in W of pure dimension k . If all weights are positive, X is a *tropical subspace in W* . Moreover, X is of pure sedentarity ϑ if $|X|$ is so.

Remark 5.1.3

Having rephrased our definition like this, when writing a cycle X as sum of its pure sedentarity parts $X = \sum_{\vartheta} X_{\vartheta}$, the X_{ϑ} will from now on be *closed cycles in W_{ϑ}* . The mobile parts $X_{\vartheta} \cap \mathbf{R}_{\vartheta}$ that were used in Definition 5.1.1 will be denoted by X_{ϑ}^m .

Example 5.1.4

For each k with $0 \leq k \leq n$, let $L_k^m \subset \mathbf{R}^n$ be the standard tropical k -plane defined in Example 4.2.6. Its closure $L_k \subset \mathbf{TP}^n$ is a sedentarity zero k -cycle in \mathbf{TP}^n and is called the *standard tropical projective k -plane*. Figure 5.2 illustrates the example $L_2 \subset \mathbf{TP}^3$. Note that L_2^s consists of four tropical lines of higher sedentarity, one for each coordinate hyperplane at infinity. More generally, for any boundary stratum $\mathbf{TP}_I^n \cong \mathbf{TP}^{n-|I|}$, we have $L_k \cap \mathbf{TP}_I^n = L_{k-|I|} \subset \mathbf{TP}^{n-|I|}$. On a set-theoretic level, this follows immediately from the recursive structure of the stars of L_k^m , see Exercise 4.2.12. Regarding the weights, we will be more precise in Section .




Figure 5.2: A tropical plane in \mathbf{TP}^3

5.1.3 Push forwards along toric morphisms

Let Ξ and Ξ' be pointed fans in \mathbf{R}^n and \mathbf{R}^m , respectively. Let $\varphi : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be an integer affine map such that for all $\vartheta \in \Xi$ there exists $\vartheta' \in \Xi'$ with $d\varphi(\vartheta) \subset \vartheta'$ (we may assume $d\varphi(\vartheta^\square) \subset \vartheta'^\square$). Then there exists a unique continuous extension to a map between toric varieties $\varphi : W_\Xi \rightarrow W_{\Xi'}$. More concretely, for each $\vartheta \in \Xi$ and $\vartheta' \in \Xi'$ such that $d\varphi(\vartheta^\square) \subset \vartheta'^\square$ we have $\varphi(\mathbf{R}_\vartheta) \subset \mathbf{R}_{\vartheta'}$ and the diagram

$$\begin{array}{ccc} \mathbf{R}^n & \xrightarrow{\varphi} & \mathbf{R}^m \\ \downarrow \pi_\vartheta & & \downarrow \pi_{\vartheta'} \\ \mathbf{R}_\vartheta & \xrightarrow{\varphi} & \mathbf{R}_{\vartheta'} \end{array}$$

commutes. Such a map $\varphi : W_\Xi \rightarrow W_{\Xi'}$ is called an *integer affine map* (or *monomial map*) between toric varieties. If there exists an inverse integer affine map $\psi : W_{\Xi'} \rightarrow W_\Xi$, φ is an *integer affine isomorphism*.

Definition 5.1.5

Let $\varphi : W_\Xi \rightarrow W_{\Xi'}$ be an integer affine map between toric varieties and let $X = \sum_{\vartheta} X_\vartheta$ be a tropical k -cycle in W_Ξ . The *push forward* of X along φ is

the tropical k -cycle in $W_{\Xi'}$ given by

$$\varphi_*(X) := \sum_{\vartheta} \overline{\varphi_*(X_{\vartheta}^m)}.$$

Here, $\varphi_*(X_{\vartheta}^m)$ denotes the push forward of along $\varphi : \mathbf{R}_{\vartheta} \rightarrow \mathbf{R}_{\vartheta'}$, where ϑ and ϑ' are such that $d\varphi(\vartheta^{\square}) \subset \vartheta'^{\square}$.

5.2 Projective degree

5.2.1 Intersecting with the standard projective hyperplane

In this section, we want to define the degree of tropical k -cycle in \mathbf{TP}^n . We will use the following notations. First recall that the cones σ_I of the fan defining \mathbf{TP}^n are labelled by subsets $I \subsetneq \{0, \dots, n\}$. In the following, we will always use I instead of σ_I at the appropriate places. For example, given a tropical cycle $X \subset \mathbf{TP}^n$, we will refer to its pure sedentarity parts by $X_I \subset \mathbf{TP}_I^n$. For any $I \subsetneq \{0, \dots, n\}$, we denote by $H_I \subset \mathbf{TP}_I^n$ the standard projective $(n - |I| - 1)$ -plane (via the canonical identification $\mathbf{TP}_I^n \cong \mathbf{TP}^{n-|I|}$). In the case $I = \emptyset$, we use the shorthand $H := H_{\emptyset} \subset \mathbf{TP}^n$ for the *standard projective hyperplane* in \mathbf{TP}^n . As explained in Example 5.1.4, we have $H \cap \mathbf{TP}_I^n = H_I$ for all I (at least set-theoretically). By slight abuse of notation, we will define the stable intersection of any cycle with H as follows.

Definition 5.2.1

Let $X = \sum_I X_I$ be a tropical k -cycle in \mathbf{TP}^n . The *stable intersection* of H and X is

$$H \cdot X := \sum_I \overline{H_I^m \cdot X_I^m},$$

where each product $H_I^m \cdot X_I^m$ refers to the stable intersection of cycles in \mathbf{R}_I .

In analogy to the classical case, we define the projective degree of a k -cycle as the number of intersection points with k hyperplanes. In contrast to the classical case, where these hyperplanes are typically chosen generically such as to obtain a simple count of (transversal) intersection points, we will only use the fixed standard hyperplane $H \subset \mathbf{TP}^n$. Recall that the degree of

0-cycle A , regarded as a formal sum of points $A = \sum_p a_p \cdot [p]$, is $\deg(A) = \sum_p a_p \in \mathbb{Z}$.

Definition 5.2.2

Let X be a tropical k -cycle in \mathbf{TP}^n . The (projective) degree of X is

$$\deg(X) := \deg(H^k \cdot X).$$

where H denotes the standard projective hyperplane in \mathbf{TP}^n .

Example 5.2.3

From our computations in example 4.3.8 it follows that the standard projective k -planes $L_k \in \mathbf{TP}^n$ have degree 1.

In the following Proposition, we show that a local part of tropical subspace X can only have smaller degree than X itself.

Proposition 5.2.4

Let X be a tropical subspace of \mathbf{R}^n and let \mathcal{X} be a polyhedral structure. Then the followings holds.

- (a) For every $\tau \in \mathcal{X}$, we have $\deg(\overline{\text{Star}_X(\tau)}) \leq \deg(\overline{X})$.
- (b) For every facet $\sigma \in \mathcal{X}$, we have $\omega(\sigma) \leq \deg(\overline{X})$.
- (c) If $\deg(\overline{X}) = 0$, then $X = 0$.

Here, all closures are taken in $\mathbf{TP}^n \supset \mathbf{R}^n$.

Proof. Obviously, (a) \Rightarrow (b) \Rightarrow (c), hence it remains to prove (a). Note that by 4.3.6 the degree is invariant under translations, e.g. $\deg(\overline{X}) = \deg(X \cdot ((H^m)^k + p))$ for some $p \in \tau^\square$. As both factors in this product are subspaces, all points in $X \cdot (H^k + p)$ have positive weight. But by the locality of the intersection product (again 4.3.6), the weight of p in this product is equal to $\deg(\overline{\text{Star}_X(\tau)})$, and hence the claim follows. \square

5.2.2 The degree of projective hypersurfaces

As a more general example, let us compute the degree of projective hypersurfaces.

Proposition 5.2.5

Let F be a homogeneous tropical polynomial of degree d and $V(F) \subset \mathbf{TP}^n$ its associated hypersurface (see Definition 3.4.6). Then $\deg(V(F)) = d$.

Proof. As in Definition 3.4.6, let $H_i = V(x_i)$ denote the i -th coordinate hyperplane. Obviously $\deg(H_i) = 1$. On the other hand, dividing F by x_i decreases the degree of F by 1. Hence we may assume that F is monomialfree.

Let $f = F(x_0 = 0) \in \mathbf{T}[x_1, \dots, x_n]$ be the dehomogenization of F and let $L = L_1 = (H^m)^{n-1} \subset \mathbf{R}^n$ be the standard line. It remains to show $\deg(f \cdot L) = \deg(f)$. Since f is monomialfree, the function will be eventually constant on each ray $\mathbf{R}_{\leq 0}e_i$, $i = 1, \dots, n$, of L . On the other hand, on the ray $\mathbf{R}_{\leq 0}e_i$ the slope of f will ultimately be $d = \deg(f)$. It follows by Exercise 4.4.22 that $\deg(f \cdot L) = d$ as required. \square

5.2.3 Projective tropical lines

Definition 5.2.6

A (projective) tropical line is a tropical subspace of \mathbf{TP}^n of dimension and degree 1.

Let us recall a trivial fact from classical algebraic geometry: Any line in \mathbf{CP}^n is isomorphic to \mathbf{CP}^1 . In tropical geometry, however, there exist (topologically) different projective lines. Of course, the archetype is still \mathbf{TP}^1 . But already the standard line L_1 in \mathbf{TP}^2 is different from \mathbf{TP}^1 . L contains 3 infinite points and one vertex, whereas \mathbf{TP}^1 has 2 infinite points and no vertex at all. Figure 5.3 depicts a tropical line in \mathbf{TP}^3 containing a bounded edge. On the other hand, tropical lines are of course very special 1-cycles. Some properties are summarized in the following Proposition. Note that any 1-cycle carries a canonical coarsest polyhedral structure. In the following, all statements about edges and vertices refer to this polyhedral structure.

Lemma 5.2.7

Let L be a tropical line in \mathbf{TP}^n of sedentarity zero. Then L has the following properties:

- (a) All weights of L are 1.




Figure 5.3: A tropical line in \mathbf{TP}^3

- (b) L is contractible. In other words, it is a rational graph.
- (c) Let v_0, \dots, v_m be the primitive generators of the rays of L^m (pointing towards infinity). Then $m \leq n$ and

$$v_j = - \sum_{i \in I_j} e_i$$

for a suitable subset $I_j \subset \{0, \dots, n\}$. Moreover, $\{0, \dots, n\}$ is the disjoint union of I_0, \dots, I_m .

- (d) The graph L^m has at most $m - 3$ bounded edges and at most $m - 2$ finite vertices. Equality holds when all vertices of L^m are 3-valent.

Proof. Statement (a) follows from Proposition 5.2.4.

For statement (b), assume that L contains a cycle C . We can find a vector $v \in \mathbf{R}^n$ such that the translated hyperplane H and L only intersect in \mathbf{R}^n and such that $H^m + v$ intersects C in a transversal intersection point p . Since the facets of $H + v$ subdivide \mathbf{TP}^n into $n + 1$ connected components, the cycle C must intersect $H + v$ in a second point in \mathbf{R}^n . It follows that $\deg((H^m + v) \cdot L) \geq 2$, a contradiction.

Let us now prove the assertion (c). We start by showing that each v_j is a simple sum of the standard primitive vectors e_0, \dots, e_n . By relabelling the coordinates we may assume $v_j \in (\mathbf{Z}_{\leq 0})^n$. We have to show that all

coordinate entries of v_j are either 0 or -1 . Assume contrarily that the first coordinate entry is lower than -1 . Then we take a hyperplane H transversal to L and such that the facet of H spanned by $-e_2, \dots, -e_n$ and the ray of L with direction v_j intersect. As the lattice index contributing to the local intersection multiplicity is equal to the first coordinate entry of v_j , we get a contradiction again. Thus we showed

$$v_j = - \sum_{i \in I_j} e_i$$

for a suitable subset $I_j \subset \{0, \dots, n\}$. Next, we show that these sets are pairwise disjoint. Assume $i \in I_j \cap I_{j'}$. Then any hyperplane H which we move far towards $-e_i$ will intersect both rays generated by v_j and $v_{j'}$, which is impossible. It remains to show $\{0, \dots, n\} = \bigcup_{j=0}^m I_j$. But note $\sum_{j=0}^m v_j = 0$ by Part (a) and Lemma 4.2.8. Since $\sum_{i \in I} e_i = 0$ if and only if $I = \{0, \dots, n\}$, the claim follows.

Statement (d) follows immediately from the fact established so far that L is a rational graph with $m + 1$ 1-valent vertices (the points of higher sedentarity). \square

A line $L \subset \mathbf{TP}^n$ of sedentarity zero is called *non-degenerate* if L^m it has $n + 1$ rays (with direction vectors $-e_0, \dots, -e_n$). We will come back to the relationship between tropical lines (or more generally, rational curves) and \mathbf{TP}^1 in Section 5.5.

5.3 Modifications in toric varieties

In this section, we want to extend the divisor and modification constructions from Chapter 4 to cycles in toric varieties. Recall that both constructions are closely related. The modification along a function f is the completion of its graph by adding “downward” faces projecting to the divisor of f . Vice versa, the divisor of f can be obtained from the modification by “intersecting with ∞ ”. In the toric case, it is useful to study both constructions simultaneously right away. We start, however, with formalizing the concept of “intersecting with ∞ ”.

5.3.1 Intersections with toric boundary divisors

Let Ξ be a fan in \mathbf{R}^n . We denote the primitive generator of a ray $\rho \in \Xi^{(1)}$ by v_ρ . In order to intersect with a specific divisor $D = W_\rho$ of W_Ξ , we need a partial smoothness or unimodularity assumption with respect to ρ .

Definition 5.3.1

A ray ρ is *Cartier* in Ξ if there exists a function $f : |\Xi| \rightarrow \mathbf{R}$ which is integer linear on each cone of Ξ and such that $f(v_\rho) = -1$ and $f(v_{\rho'})$ for all other rays $\rho' \neq \rho$. A divisor $D = W_\rho$ is *Cartier* in W_Ξ if ρ is Cartier in Ξ .

In such function f exists, it is obviously unique and we denote it by f_ρ or f_D .

Remark 5.3.2

Note that ρ is Cartier in Ξ if and only if the following holds. Any cone $\vartheta \in \Xi$ containing ρ has a unique face ϑ' not containing ρ and $L_{\mathbf{Z}}(\vartheta) = L_{\mathbf{Z}}(\vartheta') + \mathbf{Z}v_\rho$. The terminology is adapted from classical toric geometry, where the condition is equivalent to CW_ρ being a Cartier divisor in CW_Ξ (in the classical sense).

In order to define intersections with toric divisors, we need to introduce the notion of primitive generators for faces of higher sedentarity. Let $\sigma \subset W$ be a polyhedron of dimension k and sedentarity zero and let $\tau \subset \sigma$ be a $(k-1)$ -face of higher sedentarity, say ϑ . By Proposition 3.2.6, it follows that $\text{rc}(\sigma^m) \cap \vartheta^\square \neq \emptyset$ and that $\pi_\vartheta(\sigma^m) = \tau^m$, where $\pi_\vartheta : \mathbf{R}^n \rightarrow \mathbf{R}_\vartheta = \mathbf{R}^n / L(\vartheta)$ is the canonical projection. In particular,

$$\dim(\text{rc}(\sigma^m) \cap \vartheta) = \dim(L(\text{rc}(\sigma^m)) \cap L(\vartheta)) = \dim(L(\sigma^m) \cap L(\vartheta)) = 1,$$

and hence $\text{rc}(\sigma^m) \cap \vartheta$ is a (one-dimensional) ray.

Definition 5.3.3

Let $\sigma \subset W$ be a polyhedron of dimension k and sedentarity zero and let $\tau \subset \sigma$ be a $(k-1)$ -face of higher sedentarity, say ϑ . Let $-v_{\sigma/\tau} \in \mathbf{Z}^n$ be the primitive generator of $\text{rc}(\sigma^m) \cap \vartheta$. Then its negative, $v_{\sigma/\tau}$, is the (*sedentary*) *primitive generator* of σ modulo τ .

If σ is of sedentarity ϑ' , we replace \mathbf{R}^n by $\mathbf{R}_{\vartheta'}$ and Ξ by $\text{Star}_\Xi(\vartheta')$ and obtain the sedentary primitive generator $v_{\sigma/\tau} \in \mathbf{Z}_{\vartheta'} = \mathbf{Z}^n / L_{\mathbf{Z}}(\vartheta')$.

We are now ready to define intersections with toric divisors D .

Definition 5.3.4

Let $X \subset W$ be a k -cycle of sedentarity zero and let \mathcal{X} be a polyhedral structure. Let $D = W_\rho$ be a toric divisor which is Cartier, with associated function $f_D : |\Xi| \rightarrow \mathbf{R}$. The (sedentary) intersection $D \cdot \mathcal{X}$ is the weighted polyhedral complex $\mathcal{X}|_D$ with weights

$$\omega(\tau) = \sum_{\sigma \supset \tau} \omega(\sigma) f_D(v_{\sigma/\tau}).$$

Here, τ of $(k-1)$ -cell of $\mathcal{X}|_D$ and the sum runs through the k -cells $\sigma \in \mathcal{X}$ (hence of sedentarity zero) containing τ . The (sedentary) intersection $D \cdot X$ is the tropical $(k-1)$ -cycle represented by $D \cdot \mathcal{X}$.

Getting routine by now, it remains to show that $D \cdot \mathcal{X}$ is balanced.

Proposition 5.3.5

The weighted polyhedral complex $D \cdot \mathcal{X}$ in W is a balanced.

Proof. Let α be a $(k-2)$ -cell of $D \cdot \mathcal{X}$ of pure sedentarity ϑ . We denote by $\pi = \pi_\vartheta : \mathbf{R}^n \rightarrow \mathbf{R}_\vartheta$ the sedentary projection. Recall that the balancing condition around α only involves the $(k-1)$ -cells $\beta \supset \alpha$ of $D \cdot \mathcal{X}$ of same sedentarity ϑ . Throughout this proof, we use β to denote such cells and use σ to denote k -cells in \mathcal{X} (hence of sedentarity zero) that contain α .

For any flag $\alpha \subset \beta \subset \sigma$, recall that $\pi(\sigma^m) = \beta^m$ and α^m is a (usual) face of β^m . It follows that $\sigma^m \cap \pi^{-1}(\alpha^m)$ is a face of σ^m ; we denote its closure by τ . Note that $\text{rc}(\tau^m) \cap \vartheta = \text{rc}(\sigma^m) \cap \vartheta$. Thus τ is $(k-1)$ -cell of \mathcal{X} of sedentarity zero such that $\tau \cap \mathbf{R}_\vartheta = \alpha^m$ and such that the sedentary primitive generators $v_{\tau/\alpha} = v_{\sigma/\beta}$ coincide. Moreover, given a (mobile) primitive generator $v_{\sigma/\tau}$, the vector $v_{\beta/\alpha} = \pi(v_{\sigma/\tau})$ is a primitive generator for β modulo α . This follows directly from $L_Z(\beta^m) = \pi(L_Z(\sigma^m))$ and $L_Z(\alpha^m) = \pi(L_Z(\tau^m))$.

We add the convention that τ always denotes a $(k-1)$ -cells of \mathcal{X} of sedentarity zero containing α . Given a facet $\sigma \supset \alpha$ two cases can occur. Either $\sigma \cap \mathbf{R}_\vartheta$ is of dimension $k-1$. Then the previous considerations apply. Otherwise, $\sigma \cap \mathbf{R}_\vartheta = \alpha^m$ is of dimension $k-2$. In this case, $\pi(L(\sigma)) = L(\alpha)$ and hence $\pi(v_{\sigma/\tau}) \in L(\alpha)$ for for any $\tau \subset \sigma$. We are now ready to combine our considerations in the following calculation, keeping in mind the various

conventions for β , σ and τ . We get

$$\begin{aligned} \sum_{\beta} \omega_{D \cdot \mathcal{X}}(\beta) v_{\beta/\alpha} &= \sum_{\sigma \supset \beta} \omega_{\mathcal{X}}(\sigma) f_D(v_{\sigma/\beta}) v_{\beta/\alpha} \\ &= \sum_{\tau} f_D(v_{\tau/\alpha}) \pi \left(\sum_{\sigma \supset \tau} \omega_{\mathcal{X}}(\sigma) v_{\sigma/\tau} \right) \pmod{L(\alpha)}. \end{aligned}$$

Since for all τ the argument of π in the last expression is in $L(\tau)$ (balancing around τ), it follows that the total sum is contained in $\pi(L(\tau)) = L(\alpha)$, which proves the claim. \square

Remark 5.3.6

Clearly, the definition of $D \cdot X$ can be extended to cycles X of pure sedentarity ϑ as long as $\rho \not\subset \vartheta$. More precisely, if $\vartheta + \rho$ is a cone of Ξ , we consider X as a sedentarity zero cycle in $W' = W_{\text{Star}_{\Xi}(\vartheta)}$ and let D' be the toric divisor in W' associated to the ray $(\vartheta + \rho)/\vartheta$. We set $D \cdot X := D' \cdot X$, using the previously defined construction in W' . Otherwise, we set $D \cdot X = 0$. Note, however, that in the case $\rho \subset \vartheta$ the intersection $D \cdot X$ is not well-defined (at least, on the cycle level).

Exercise 5.3.7

Prove that the multiplicities introduced in Exercise 2.5.5 correspond to the intersection of a curve $C \subset \mathbf{TP}^2$ of sedentarity zero with the toric divisor L_{∞} corresponding to the ray $\mathbf{R}_{\geq 0}(1, 1)$.

5.3.2 Modifications and divisors

Let us first introduce the class of functions that we want to consider. In the following, $W = W_{\Xi}$ is a fixed tropical toric variety.

Definition 5.3.8

Let $X = \sum_{\vartheta \in \Xi} X_{\vartheta}$ be tropical k -cycle in W . A *piecewise integer affine function* on X is a collection of piecewise integer affine functions $f_{\vartheta} : X_{\vartheta}^m \rightarrow \mathbf{R}$ for each $\vartheta \in \Xi$. In particular, the group of piecewise integer linear functions on X is

$$\text{PA}_{\mathbf{Z}}(X) = \bigoplus_{\vartheta \in \Xi} \text{PA}_{\mathbf{Z}}(X_{\vartheta}^m).$$

If X is of pure sedentarity, f is just a piecewise integer affine function $f : X^m \rightarrow \mathbf{R}$.

For general cycles $X = \sum_{\vartheta} X_{\vartheta}$ we set $X^m = \bigsqcup_{\vartheta} X_{\vartheta}^m \subset W$. Then a piecewise integer affine function is determined by the induced function $f : X^m \rightarrow \mathbf{R}$, which we usually identify with f . One should keep in mind, however, that this function can be rather peculiar. For example, in the case of mixed sedentarity it will in general not be continuous.

Definition 5.3.9

Let $X = \sum_{\vartheta \in \Xi} X_{\vartheta}$ be tropical k -cycle in W and $f : X^m \rightarrow \mathbf{R}$ a piecewise integer affine function on X . The *modification of X along f* is the k -cycle in $W \times \mathbf{TP}^1$ given by

$$\text{Mod}(X, f) = \sum_{\vartheta} \overline{\text{Mod}(X_{\vartheta}^m, f_{\vartheta})}.$$

On a formal level, we could think of $\text{Mod}(X, f)$ as a cycle in $W \times \mathbf{R}$ as well, staying closer to the definition in Section 4.5. However, we already gave some motivation for compactifying the second factor in Remark 4.5.4, and will find more reasons shortly. Note that piecewise integer affine functions on X may have points of indeterminacy at higher sedentarity even after extending the range of values from \mathbf{R} to \mathbf{TP}^1 . Here is a typical example.

Example 5.3.10

Consider the Laurent monomial $f = "x_2/x_1" = x_2 - x_1 : \mathbf{R}^2 \rightarrow \mathbf{R}$. We consider it here as a element in $\text{PA}_2(\mathbf{T}^2)$. Note that f obviously be extended to continuous function on $\mathbf{T}^2 \setminus \{(-\infty, -\infty)\} \rightarrow \mathbf{TP}^1$, the value at $(-\infty, -\infty)$ is not well-defined. Let us compute $\text{Mod}(\mathbf{T}^2, f)$. Note that $f|_{\mathbf{R}^2}$ is just a linear function. Taking the closure of its graph $\{x \in \mathbf{R}^3 : x_2 = x_1 + x_3\}$ in $\mathbf{T}^2 \times \mathbf{TP}^1$, we see easily that the extra added points form the union of three coordiante lines

$$\begin{aligned} \text{Mod}(\mathbf{T}^2, f) \setminus \mathbf{R}^3 &= \mathbf{T} \times \{-\infty\} \times \{-\infty\} \\ &\cup \{-\infty\} \times \mathbf{T} \times \{+\infty\} \\ &\cup \{-\infty\} \times \{-\infty\} \times \mathbf{TP}^1 \end{aligned}$$

(see Figure 5.4). While the first two lines are due to the extension of f to $\mathbf{T}^2 \setminus \{(-\infty, -\infty)\} \rightarrow \mathbf{TP}^1$, the last line $\{-\infty\} \times \{-\infty\} \times \mathbf{TP}^1$ reflects the




Figure 5.4: The modification $\text{Mod}(\mathbf{T}^2, x_1 - x_2)$ with its 3 boundary lines

fact that f is not well-defined at $(-\infty, -\infty)$. In fact, $\text{Mod}(\mathbf{T}^2, f)$ can be understood as the *blow up* of \mathbf{T}^2 at $(-\infty, -\infty)$.

We denote by $H_- = W \times \{-\infty\}$ and $H_+ = W \times \{+\infty\}$ the two canonical toric divisors in $W \times \mathbf{TP}^1$. Clearly, the projection $\pi : W \times \mathbf{TP}^1 \rightarrow W$ induces an isomorphism of cycle groups $\pi_* : Z_k(H_\pm) \rightarrow Z_k(W)$. By definition, the modification $\text{Mod}(X, f)$ is a tropical cycle in $W \times \mathbf{TP}^1$. It has parts of pure sedentarity $\vartheta \times \{0\}$, where ϑ is a cone in Ξ . Hence by Remark 5.3.6, the sedentary intersections $H_- \cdot \text{Mod}(X, f)$ and $H_+ \cdot \text{Mod}(X, f)$ are well-defined.

Definition 5.3.11

Let $X = \sum_{\vartheta \in \Xi} X_\vartheta$ be tropical k -cycle in W and $f : X^m \rightarrow \mathbf{R}$ a piecewise integer affine function on X . The *divisor of f* , denoted by $\text{div}(f)$ or $f \cdot X$, is the tropical $k - 1$ -cycle

$$\text{div}(f) = \pi_*(H_- \cdot \text{Mod}(X, f) - H_+ \cdot \text{Mod}(X, f)) \in Z_{k-1}(W).$$

Note that this time the balancing condition for $\text{div}(f)$ follows from Proposition 5.3.5.

Remark 5.3.12

Let us assume that X is a tropical subspace. Then the domains of $\text{div}(f)$ of positive and negative weight can be regarded as “zeros” and “poles” of f , respectively. Note that poles can arise in two ways: From $H_+ \cdot \text{Mod}(X, f)$,

obviously, but also from facets in $\text{Mod}(X, f)$ pointing downwards but carrying negative weights. As discussed in Remark 4.5.4, in \mathbf{R}^n -case only the poles of the second type occur and it is therefore sufficient to intersect with H_- . Again, for general piecewise integer affine functions there seems to be no canonical way to repair this asymmetry between zeros and poles.

For rational functions, on the other hand, we can of course adapt Definition 4.5.5 as follows. For simplicity, let us assume that X is a tropical subspace of W of sedentarity zero. Let $f = "g/h" : X^m \rightarrow \mathbf{R}$ be a rational function given as the quotient of two polynomial functions g, h . Then *modification of the quotient "g/h"* is given by

$$\text{Mod}(X, g, h) := \overline{\text{Mod}(\mathbf{X}^m, g, h)} \subset W \times \mathbf{TP}^1.$$

It is easy to check that $\text{div}(f) = \text{div}(g) - \text{div}(h)$, $\text{div}(g) = \pi_*(H_- \cdot \text{Mod}(X, g, h))$ and $\text{div}(h) = \pi_*(H_+ \cdot \text{Mod}(X, g, h))$. Hence this time the zeros and poles of f come as two separate cycles with positive weights. We will consider another generalization of this in Section 5.3.4.

We would like to give an explicit formula for the weights of $\text{div}(f)$ avoiding the detour via $\text{Mod}(X, f)$. For the sedentarity zero part we can obviously still use the weight formula from Equation (4.5). Actually, something similar is true at the boundary. Actually, the formula can be easily adapted to cover faces of higher sedentarity as well. It enough to consider the sedentarity zero case.

Let $X \subset W$ be a tropical k -cycle of sedentarity zero and let $f : X^m \rightarrow \mathbf{R}$ be a piecewise integer affine function. A polyhedral structure \mathcal{X} is called *sufficiently fine* for f if \mathcal{X}^m is sufficiently fine.

Corollary 5.3.13

Let $X \subset W$ be a tropical k -cycle of sedentarity zero and let $f : X^m \rightarrow \mathbf{R}$ be a piecewise integer affine function. Let \mathcal{X} be a sufficiently fine polyhedral structure and $\tau \in \mathcal{X}$ a $(k - 1)$ -cell of higher sedentarity. Then the weight of τ in $\text{div}(f)$ is equal to

$$\omega(\tau) = \sum_{i=1}^k \omega(\sigma_i) df|_{\sigma_i^m}(v_{\sigma_i/\tau}), \tag{5.2}$$

where $\sigma_1, \dots, \sigma_k$ denote the facets of \mathcal{X} containing τ .

Recall that if τ is of sedentarity zero and the primitive generators around τ are chosen such that $\sum_{i=1}^k \omega(\sigma_i)v_{\sigma_i/\tau} = 0$, then the formula for $\omega(\tau)$ is identical to the one presented here (see Remark 4.4.5).

Proof. Each pair $\tau \subset \sigma_i$ corresponds to a pair $\tilde{\tau} \subset \tilde{\sigma}_i$ of faces in $\text{Mod}(X, f)$, and primitive generators are given by

$$v_{\tilde{\sigma}/\tilde{\tau}} = (v_{\sigma/\tau}, df|_{\sigma_i^m}(v_{\sigma_i/\tau})) \in \mathbf{Z}^n \times \mathbf{Z}.$$

We set $s_i = df|_{\sigma_i^m}(v_{\sigma_i/\tau}) \in \mathbf{Z}$. Let ϑ be the sedentarity of τ . Depending on the cases $s_i > 0$, $s_i = 0$, or $s_i < 0$, the sedentarity of $\tilde{\tau}$ is $\vartheta \times \mathbf{R}_{\leq 0}e_{n+1}$, $\vartheta \times \{0\}$ or $\vartheta \times \mathbf{R}_{\geq 0}e_{n+1}$, respectively. In the middle case, the pair $\tilde{\tau} \subset \tilde{\sigma}_i$ neither contributes to Equation (5.2) nor to either of the intersections $H_{\pm} \cdot \text{Mod}(X, f)$. If $s_i > 0$, the pair $\tilde{\tau} \subset \tilde{\sigma}_i$ contributes to $H_- \cdot \text{mod}(X, f)$ with weight $\omega(\tilde{\sigma}_i)f_{H_-}(v_{\tilde{\sigma}/\tilde{\tau}})$. But note that $\omega(\tilde{\sigma}_i) = \omega(\sigma_i)$ and $f_{H_-}(v_{\tilde{\sigma}/\tilde{\tau}}) = s_i$ (see Example). So the contribution of $\tilde{\tau} \subset \tilde{\sigma}$ equals the contribution of $\tau \subset \sigma_i$ to Equation (5.2). The case $s_i < 0$ works symmetrically. \square

Example 5.3.14

H_{\pm} Cartier in $W \times \mathbf{TP}^1$.

Example 5.3.15

The Laurent monomial $f = "x_2/x_1" = x_2 - x_1$ from Example 5.3.10 has divisor $\text{div}(f) = V(x_2) - V(x_1) \subset \mathbf{T}^2$.

Example 5.3.16

If D is Cartier then $D = \text{div}(f_D)$.

Explain subdivision of W induced by $\Xi??$

5.3.3 Line bundles of tropical toric varieties

In the previous section, we considered modifications in $W \times \mathbf{TP}^1$, the trivial \mathbf{TP}^1 -bundle over W , provided we are given a piecewise integer affine function on X . Starting from homogeneous polynomials F , we can similarly construct a modification $\text{Mod}(X, F)$ which this time sits naturally in a *non-trivial* line bundle over W . In this subsection we introduce these objects in the generality needed here, postponing a general discussion of tropical line bundles to .

Let Ξ be a pointed fan in \mathbf{R}^n and let $f : |\Xi| \rightarrow \mathbf{R}$ be a function which is integer linear on every cone of Ξ . In other words, f is piecewise integer linear and Ξ is sufficiently fine. For any $\vartheta \in \Xi$, we denote by $\tilde{\vartheta} = \Gamma(\vartheta, f)$ the cone in \mathbf{R}^{n+1} equal to the graph of $f|_{\vartheta}$. Furthermore, we set

$$\vartheta_{\leq} := \tilde{\vartheta} + \mathbf{R}_{\leq 0}e_{n+1}, \quad \vartheta_{\geq} := \tilde{\vartheta} + \mathbf{R}_{\geq 0}e_{n+1},$$

and consider the fans in \mathbf{R}^{n+1}

$$\Xi_{\leq}(f) := \{\tilde{\vartheta}, \vartheta_{\leq} : \vartheta \in \Xi\}, \tag{5.3}$$

$$\Xi(f) := \{\tilde{\vartheta}, \vartheta_{\leq}, \vartheta_{\geq} : \vartheta \in \Xi\}. \tag{5.4}$$

Exercise 5.3.17

Show that $\Xi_{\leq}(f)$ and $\Xi(f)$ are fans.

Definition 5.3.18

The *tropical line bundle* $\mathcal{O}_W(f)$ associated to f on $W = W_{\Xi}$ is the toric variety associated to $\Xi_{\leq}(f)$ together with the canonical projection $\pi : \mathcal{O}_W(f) \rightarrow W$.

The *tropical projective line bundle* (or \mathbf{TP}^1 -bundle) $\mathcal{P}_W(f)$ associated to f on $W = W_{\Xi}$ is the toric variety associated to $\Xi(f)$ together with the canonical projection $\pi : \mathcal{P}_W(f) \rightarrow W$.

Exercise 5.3.19

Show that a fiber of $\pi : \mathcal{O}_W(f) \rightarrow W$ can be identified with \mathbf{T} . Show that a fiber of $\pi : \mathcal{P}_W(f) \rightarrow W$ can be identified with \mathbf{TP}^1 .

Example 5.3.20

If $f \equiv 0$ is constant zero, we have $\mathcal{O}_W(f) = W \times \mathbf{T}$ and $\mathcal{P}_W(f) = W \times \mathbf{TP}^1$.

Example 5.3.21

For $W = \mathbf{TP}^n$, $f = "0 + x_1 + \dots + x_n"$ and $d \in \mathbf{Z}$, we obtain the tropical line bundles $\mathcal{O}(d) := \mathcal{O}_{\mathbf{TP}^n}(d \cdot f)$ (sometimes called the *twisting bundles*) their projective counterparts and $\mathcal{P}(d) := \mathcal{P}_{\mathbf{TP}^n}(d \cdot f)$. Let us describe the gluing maps for $\mathcal{O}(d)$ explicitly. We denote by $U_i = \{x \in \mathbf{TP}^n : x_i \neq -\infty\}$ the standard affine charts of \mathbf{TP}^n . On the overlap $U_i \cap U_j$ the function $a_{ij}(x) := d(x_j - x_i)$ is a well-defined \mathbf{Z} -invertible integer linear function (in particular, it extends to the sedentary points of $U_i \cap U_j$). The result of gluing the affine

charts $U_i \times \mathbf{T}$ along the maps

$$\begin{aligned} \phi_{ij} : (U_i \cap U_j) \times \mathbf{T} &\rightarrow (U_i \cap U_j) \times \mathbf{T} \\ (x, t) &\mapsto (x, t + a_{ij}(x)) \end{aligned}$$

can be canonically identified with $\mathcal{O}(d)$. The same argument applies to $\mathcal{P}(d)$.

Remark 5.3.22

Let us assume that Ξ is complete and $f : |\Xi| \rightarrow \mathbf{R}$ is strongly convex, i.e. f is strictly convex on any line in \mathbf{R}^n . Then the convex hull $\text{Conv}(\Gamma(f))$ of the graph of f is a pointed cone in \mathbf{R}^{n+1} . Adding this cone to $\Xi_{\leq}(f)$, we obtain a complete fan $\overline{\Xi}_{\leq}(f)$ of \mathbf{R}^{n+1} . The associated tropical toric variety is denoted by $\overline{\mathcal{O}}_W(f)$.

Example 5.3.23

For $W = \mathbf{TP}^n$, the variety $\overline{\mathcal{O}}_W(d)$ is called a *tropical weighted projective space* $\mathbf{TP}(1, \dots, 1, d)$.

Example 5.3.24

Let $\lambda : \mathbf{R}^n \rightarrow \mathbf{R}$ be an integer linear function. Then the fans $\Xi_{\leq}(f)$ and $\Xi_{\leq}(f + \lambda)$ are canonically isomorphic under the action of $\text{GL}(n+1, \mathbf{Z})$. The same is true for $\Xi(f)$ and $\Xi(f + \lambda)$. Hence we can canonically identify $\mathcal{O}_W(f) \cong \mathcal{O}_W(f + \lambda)$ and $\mathcal{P}_W(f) \cong \mathcal{P}_W(f + \lambda)$.

Remark 5.3.25

Recall the toric degree map $\delta : \mathbf{Z}^l \rightarrow B$ from Remark 3.4.8. Fixing a toric degree $d \in B$, we consider the rational polyhedron $P(d) := (\delta^{-1}(d) \otimes \mathbf{R}) \cap \mathbf{R}_{\geq 0}^l$. Up to translation P_d determines a polyhedron in \mathbf{R}^n via the map $Z^n \rightarrow \mathbf{Z}^l$. On the dual side, the support function f_d of P_d (i.e. the Legendre dual of the characteristic function of P_d) is a function on $|\Xi|$ which is linear on the cones of Ξ . We call d *Cartier* if $f(d)$ is integer linear (if Ξ is complete, this is equivalent to P_d being a lattice polytope). Under this assumption, we get associated line bundles $\mathcal{O}_W(d) := \mathcal{O}_W(f_d)$ and $\mathcal{P}_W(d) := \mathcal{P}_W(f_d)$. Moreover, if Ξ is complete and P_d is an n -dimensional polytope, then f_d is strongly convex and we set $\overline{\mathcal{O}}_W(d) := \overline{\mathcal{O}}_W(f_d)$.

5.3.4 Homogeneous modifications

Let $F \in \mathbf{T}[x_0, \dots, x_n]$ be a homogeneous polynomial of degree d . Recall that F does not define a function on \mathbf{TP}^n . On the standard affine chart $U_i = \{x \in \mathbf{TP}^n : x_i \neq -\infty\}$, we can consider the dehomogenization f_i of F . In homogeneous coordinates, f_i is given by the quotient $F_i := "F/x_i^d" = F - dx_i$. On the overlap $U_i \cap U_j$, we have

$$F_i(x) - F_j(x) = d(x_j - x_i) = a_{ij}(x), \quad (5.5)$$

where a_{ij} denotes the transition functions from Example 5.3.21. Let now $X \subset \mathbf{TP}^n$ be a tropical k -cycle of sedentarity zero. We set $X_i = X \cap U_i$. On each chart, we may consider the modification $M_i := \text{Mod}(X_i, f_i|_{X_i})$. Since f_i is a polynomial, M_i is contained in $U_i \times \mathbf{T}$ (versus $U_i \times \mathbf{TP}^1$). Moreover, it follows from Equation (5.5) that restricted to $(U_i \cap U_j) \times \mathbf{T}$, M_i and M_j are identified via ϕ_{ij} . In more precise terms,

$$\phi_{ij*}(M_j|_{(U_i \cap U_j) \times \mathbf{T}}) = M_i|_{(U_i \cap U_j) \times \mathbf{T}}.$$

It follows that the various modifications can be glued to a tropical k -cycle $\text{Mod}(X, F)$ in $\mathcal{O}(d)$.

Definition 5.3.26

Let $F \in \mathbf{T}[x_0, \dots, x_n]$ be a homogeneous polynomial of degree d and $X \subset \mathbf{TP}^n$ a tropical k -cycle of sedentarity zero. We denote by $X_i = X \cap U_i$ the affine pieces of X . The *modification* $\text{Mod}(X, F)$ of X along F is the tropical k -cycle in $\mathcal{O}(d)$ such that for any $i = 0, \dots, n$ we have

$$\text{Mod}(X, F) \cap (U_i \times \mathbf{T}) = \text{Mod}(X_i, f_i).$$

The previous discussion shows that $\text{Mod}(X, F)$ exists and is unique. We can immediately extend the construction to a quotient of two homogeneous polynomials.

Definition 5.3.27

Let $G, H \in \mathbf{T}[x_0, \dots, x_n]$ be homogeneous polynomials and set $d = \deg(G) - \deg(H)$. Let $X \subset \mathbf{TP}^n$ a tropical k -cycle of sedentarity zero. We denote by $X_i = X \cap U_i$ the affine pieces of X . The *modification* $\text{Mod}(X, G, H)$ of X

along the quotient “ G/H ” is the tropical k -cycle in $\mathcal{P}(d)$ such that for any $i = 0, \dots, n$ we have

$$\text{Mod}(X, F) \cap (U_i \times \mathbf{TP}^1) = \text{Mod}(X_i, g_i, h_i).$$

Example 5.3.28

Let us look at two examples of modifications of \mathbf{TP}^2 along *linear* polynomials. Let us first consider the standard line given by $f = “x_0 + x_1 + x_2”$. Then $\text{Mod}(\mathbf{TP}^2, f) = V(“x_0 + x_1 + x_2 + x_3”)$ is just the standard plane in \mathbf{TP}^3 (cf. Figure 5.4). Let us now consider the degenerated polynomial $g = “x_1 + x_2”$, hence $(0 : -\infty : -\infty) \in V(g)$. Correspondingly, the modification $\text{Mod}(\mathbf{TP}^2, g)$ is still a plane in \mathbf{TP}^3 , but degenerated as well. In particular, it contains the torus fixed points $(0 : -\infty : -\infty : -\infty)$. Both examples are illustrated in Figure 5.5.




Figure 5.5: Modifying \mathbf{TP}^2 along two different lines

Remark 5.3.29

Clearly, both construction can be extended to arbitrary toric varieties W as follows. Let $d \in B$ be a toric degree which is Cartier (see Remark 5.3.25). Recall that $\mathcal{O}_W(d)$ and $\mathcal{P}_W(d)$ denote the corresponding (projective) line bundles. Let F be a homogeneous polynomial of degree d (see Remark 3.4.8). There exist well-defined dehomogenizations $f_i \in \mathbf{T}[y_1, \dots, y_n]$

(the translation part is fixed by requiring that if F has full support P_d , then all f_i are monomialfree polynomials), and corresponding shifts $F_i \in \mathbf{T}[x_1^\pm, \dots, x_l^\pm]$. It is easy to check that on overlaps $U_\vartheta \cap U_{\vartheta'}$, the transition functions for the F_i coincide with the transition functions for $\mathcal{O}_W(d)$. Hence we can again define $\text{Mod}(X, F) \subset \mathcal{O}_W(d)$ as the tropical k -cycle such that $\text{Mod}(X, F) \cap (U_\vartheta \times \mathbf{T}) = \text{Mod}(X \cap U_\vartheta, f_i)$ for all $\vartheta \in \Xi$. In exactly the same way we define $\text{Mod}(X, G, H) \subset \mathcal{P}_W(d)$ for two homogeneous polynomials G, H such that $d = \deg(G) - \deg(H)$ is Cartier.

Definition 5.3.30

Two tropical k -cycles $X \subset W, Y \subset W'$ are *isomorphic*, denoted $X \cong Y$, if there exist open toric subvarieties $U \subset W$ and $U' \subset W'$ such that $X \subset U$ and $Y \subset U'$ and an integer affine isomorphism $\varphi : U \rightarrow U'$ such that $\varphi_*(X) = Y$.

We say X and Y are (*modification*) *equivalent*, denoted by $X \equiv Y$, if there exists a chain of tropical k -cycles $X = X_0, X_1, \dots, X_l = Y$, $X_i \subset W_i$ such for each consecutive pair X_i, X_{i+1} either $X_i \cong \text{Mod}(X_{i+1}, F_i)$ or $X_{i+1} \cong \text{Mod}(X_i, F_i)$ for suitable homogeneous polynomials F_i on W_i or W_{i+1} , respectively.

Exercise 5.3.31

Show that the tropical cycle $D \cdot X$ does not depend on the chosen polyhedral structure \mathcal{X} .

Exercise 5.3.32

Show that $\text{Mod}(W, F) = V("y + F(x)")$. Show that $\text{Mod}(X, F) = "y + F(x)" \cdot \pi^{-1}(X)$ with $\pi : \mathcal{O}_W(d) \rightarrow W$. Show that $\text{div}(F) = V(F)$.

Remark 5.3.33

Cycles of higher sedentarity

5.3.5 Linear modifications

Theorem 5.3.34

$$\deg(\text{Mod}(X, F)) = \deg(X), \deg(\text{div}(f)) = \deg$$

5.4 Configurations of hyperplanes in projective space

A difference between classical and tropical varieties is that tropical varieties always are equipped with a divisor of special points, the points of higher sedentarity. The main meaning of modifications and contractions is to add or remove (sedentarity zero) hypersurfaces to this divisor if necessary.

As an example, let us regard tuples of three distinct points in \mathbf{TP}^1 . For the classical projective line \mathbf{CP}^1 , all such tuples are equivalent, as we can always find an automorphism of \mathbf{CP}^1 which maps the tuple to $0, 1, \infty$.

In contrast, tropical \mathbf{TP}^1 has two distinguished points $\pm\infty$. They are topologically different from all finite points and therefore the classical statements can obviously not be translated directly. Indeed, as we will see later, the automorphisms of \mathbf{TP}^1 are translations of the finite part by a real constant (which keeps the infinite points fixed) and reflections at any real number (which exchanges the infinite points). We see that automorphisms do not change the sedentarity of points. Instead, we are forced to use modifications here. Let us study this in more details.

First, we may fix 3 points in \mathbf{TP}^1 as reference points (like $0, 1, \infty$ for \mathbf{CP}^1). A natural choice is $-\infty, 0, \infty$. After what we just said above, it is even more natural to modify \mathbf{TP}^1 along 0 to get the line $L = V("x_0 + x_1 + x_2")$ in \mathbf{TP}^2 . Now all our three reference points p_0, p_1, p_2 are infinite and are symmetrically given by $x_0 = -\infty, x_1 = -\infty$ resp. $x_2 = -\infty$. Now let $q_0, q_1, q_2 \in L$ be any other configuration of three distinct points. Then we can formulate the following statement.

Lemma 5.4.1

There exists a series of modifications and contractions of L to an isomorphic line L' which transforms the points q_0, q_1, q_2 to the infinite points p_0, p_1, p_2 .

Proof. We modify along the points q_0, q_1, q_2 , obtaining a line in \mathbf{TP}^5 with coordinates $x_0, x_1, x_2, y_0, y_1, y_2$. We now contract three times, namely we forget the original coordinates x_0, x_1, x_2 of \mathbf{TP}^2 . We end up with a line L' in \mathbf{TP}^2 whose infinite points q'_0, q'_1, q'_2 correspond to q_0, q_1, q_2 . L' is non-degenerated as the points q_0, q_1, q_2 are distinct, hence also q'_0, q'_1, q'_2 . Hence




Figure 5.6: Arrangements of three points in \mathbf{TP}^1 via modifications

L and L' are isomorphic and the identification $p_i = q'_i$ satisfies the required properties. \square

Basically the same is true for higher dimensions. Again as a reference we may fix the standard hyperplane $H \in \mathbf{TP}^{n+1}$ with its $n+2$ planes P_0, \dots, P_{n+1} at infinity. Let Q_0, \dots, Q_{n+1} be another collection of hyperplanes in H . We assume that this collection is generic, i.e. the intersection of any choice of $n+1$ of these hyperplanes is empty. Then the following is true.

Lemma 5.4.2

There exists a sequence of modifications and contractions of H to an isomorphic hyperplane H' which transforms the chosen hyperplanes Q_0, \dots, Q_{n+1} to the infinite hyperplanes P_0, \dots, P_{n+1} .

Under the assumption that each Q_i is obtained as the stable intersection $H \cdot H_i$ of H with another hyperplane H_i of \mathbf{TP}^{n+1} , the proof of this statement is completely analogous to the one-dimensional case. The assumption is not necessarily satisfied, but these issues will be addressed more thoroughly in the following chapters.

Instead of hyperplanes, we may also consider points in \mathbf{TP}^n . Let us consider the case $n = 2$. Let $p_0, p_1 \in \mathbf{TP}^2$ be two distinct points. Then there exists a line $L \subset \mathbf{TP}^2$ containing p_0 and p_1 . The modification along L gives a hyperplane $H \subset \mathbf{TP}^3$, and we identify L with the line $\bar{L} = H \cap \{x_3 = -\infty\}$ of sedentarity (at least) 1. Consistently we identify p_0 and p_1 with the points $\bar{p}_0, \bar{p}_1 \in \bar{L}$ which are mapped to p_0 resp. p_1 by the contraction map. For any point $p \in \mathbf{TP}^2 \setminus L$, there is a unique lift to H , denoted by \bar{p} . Note that H can be contracted in any of the 4 standard directions $-e_0, \dots, -e_3$ to projective

space \mathbf{TP}^2 . For a point configuration $p_0, p_1, p_2, \dots, p_n$ and a chosen contraction of H , the images of $\bar{p}_0, \bar{p}_1, \bar{p}_2, \dots, \bar{p}_n$ give a new point configuration in \mathbf{TP}^2 . Two point configurations are called *projectively equivalent* if they can be connected by a series of the above construction and automorphisms of \mathbf{TP}^2 (generated by translations and permutations of the coordinates).

Lemma 5.4.3

Let p_0, \dots, p_3 be 4 generic points in \mathbf{TP}^2 (i.e. no three of them are contained in a line). Then this point configuration is tropically equivalent to the configuration $(-\infty : 0 : 0), (0 : -\infty : 0), (0 : 0 : -\infty), (0 : 0 : 0)$.

Proof. Let $s = \text{sed}(p_0) + \text{sed}(p_1) + \text{sed}(p_2)$ be the sum of the sedentarities of the first three points. We are done if $s = 6$, because this implies $\{p_0, p_1, p_2\} = \{(-\infty : 0 : 0), (0 : -\infty : 0), (0 : 0 : -\infty)\}$ and, as the p_i are generic, p_3 must be a finite point in \mathbf{R}^2 . Thus, we can use an automorphism of \mathbf{TP}^2 to reorder the first 3 points correctly and to move p_3 to $(0, 0, 0)$.

It remains to show that, if $s < 6$, there is always a choice of modification-contraction as above such that the obtained tropically equivalent point configuration is still generic and has sedentarity $s' > s$. Then the claim follows. So let us assume $s < 6$. Let L_0, L_1, L_2 be the three lines passing through $\{p_1, p_2\}, \{p_0, p_2\}$ resp. $\{p_0, p_1\}$. Note that $s < 6$ implies

$$\{L_0, L_1, L_2\} \neq \{V(x_0), V(x_1), V(x_2)\}.$$

In other words, one of the lines, say L_0 , must be of sedentarity zero. Moreover, there exists at least one coordinate lines, say $V(x_0)$, which contains at most one of the points p_0, p_1, p_2 . We modify along L_0 and then contract $V(x_0)$ (to be precise, we contract by forgetting the coordinate x_0). Let $H = \text{Mod}(\mathbf{TP}^2, L_0)$ be the modified plane in \mathbf{TP}^3 . By construction we have $\text{sed}(\bar{p}_0) = \text{sed}(p_0) + 1$, $\text{sed}(\bar{p}_1) = \text{sed}(p_1) + 1$ and $\text{sed}(\bar{p}_2) = \text{sed}(p_2)$. When contracting (let us call the contraction π), the sedentarity of a point drops (by one) if and only if it is contained in the contracted line at infinity $\overline{V(x_0)} = \text{Mod}(V(x_0), L_0 \cap V(x_0))$. By our choice, $\overline{V(x_0)}$ contains at most one of the points $\bar{p}_0, \bar{p}_1, \bar{p}_2$. It follows

$$\text{sed}(\pi(\bar{p}_0)) + \text{sed}(\pi(\bar{p}_1)) + \text{sed}(\pi(\bar{p}_2)) \geq s + 2 - 1 > s,$$

so indeed, we increased the sedentarity of the new point configuration.




Figure 5.7: Two examples of the modification-contraction process

It remains to check that $\pi(\bar{p}_0), \dots, \pi(\bar{p}_3)$ are still generic. By construction, $\pi(\bar{p}_0), \pi(\bar{p}_1), \pi(\bar{p}_2) \notin \pi(\bar{L}_\infty)$. Let M be a line passing through at least 3 of the points $\pi(\bar{p}_0), \dots, \pi(\bar{p}_3)$. For any $x \in M \cap \pi(\bar{L}_\infty)$, we can choose the “height” of $\text{Mod}(M, x)$ such that $\text{Mod}(M, x) \subset H$. We choose $x = \pi(\bar{p}_3)$ if $\pi(\bar{p}_3) \in M$, or any other $x \in M \cap \pi(\bar{L}_\infty)$ otherwise. Then $\text{Mod}(M, x)$ is a line in H which contains at least 3 of the points $\bar{p}_0, \dots, \bar{p}_3$. Contracting along the original modification gives a line containing at least 3 of the points p_0, \dots, p_3 . This contradicts to the assumption that the p_0, \dots, p_3 are generic. \square

Similar arguments also work in higher dimensions.

5.5 Equivalence of rational curves

Let us anticipate the definition of an abstract tropical variety in its easiest case, namely for (rational) tropical curves. An *abstract tropical smooth curve* is a tuple $C = (\Gamma, d)$, where Γ is a graph (i.e. a topological space homeomorphic to a one-dimensional simplicial complex) and d is a complete inner metric on the “finite part” $\Gamma^\circ := \Gamma \setminus \partial\Gamma$. Here $\partial\Gamma$ denotes the set of 1-valent vertices of Γ (which is clearly independent of the chosen simplicial struc-

ture). It follows that, for a chosen simplicial structure, we call the edges containing an 1-vertex *ends* of C . Other edges are called *bounded edges*. A bounded edge is homeomorphic to an interval $[0, l] \subset \mathbf{R}, l > 0$, and we call l the *length* of the edge. An end is homeomorphic to $[-\infty, 0]$, where $-\infty$ is identified with the 1-valent vertex. Therefore, C is completely described its combinatorial graph Γ and by a positive real length $l(E)$ for any bounded edge E . Finally, C is called *rational* if Γ is a tree.

The reader might be surprised that the definition of an abstract curve does not require any balancing condition. In fact, as we will later see, given the valence of a point, there is only one smooth local model in dimension 1. For a point of valence $n + 1$, this local model is given by the standard line in \mathbf{TP}^n (i.e. the 1-dimensional fan with $n + 1$ rays pointing to the standard directions $-e_0, \dots, -e_n$) or \mathbf{T} , if $n = 0$. Hence the tropical structure of a smooth curve is completely determined by its graph and the balancing condition at each (finite) point is somewhat hidden by the uniqueness of local models.

An *isomorphism* of two abstract tropical curves C, D is a continuous map $C \rightarrow D$ which restricts to an isometry $C^\circ \rightarrow D^\circ$.

Lemma 5.5.1

Any smooth rational tropical curve C is isomorphic to a sequence of modification of \mathbf{TP}^1 along single points.

Proof. Assume that $C = (\Gamma, d)$ has more than two 1-valent vertices. Then we can contract one of them, i.e. we just remove the vertex and the interior of the adjacent edge E to obtain a new abstract curve $C' = C \setminus E$. Let $x \in C'$ be the point where E was attached to C' . When we modify C' along x , we just glue an interval $[-\infty, 0]$ to C' (identifying 0 to x). Hence $\text{Mod}(C', x)$ is isomorphic to C . Now we repeat this process until we end up with a curve with only two ends. Clearly its graph must be linear and hence the curve is isomorphic to \mathbf{TP}^1 , which proves the claim. \square

5.6 Equivalence of linear spaces

In classical algebraic geometry, a linear space of dimension n (embedded in some big \mathbf{CP}^N) is always isomorphic to \mathbf{CP}^n . In tropical geometry, linear spaces of a fixed dimension n , i.e. positive tropical n -cycles of degree 1

in some \mathbf{TP}^N), can look quite differently. This is already happens for lines where, aside from the most natural "model" \mathbf{TP}^1 , we may consider lines in \mathbf{TP}^N with a bigger number of infinite points and more complicated (though rational) graphs. Again, tropical equivalence generated by tropical modifications is the right notion to identify all these different models/embeddings of tropical lines and linear spaces.

Lemma 5.6.1

Let $L \subseteq \mathbf{TP}^N$ be a tropical linear space of dimension n . Then L is a multiple modification of \mathbf{TP}^n . In particular, all linear spaces of dimension n are tropically equivalent to each other.

Proof. We choose a torus fixed point $p \in \mathbf{TP}^N$ such that $p \notin L$ and project along p . More precisely, we consider the projection map $\pi : \mathbf{TP}^N \setminus \{p\} \rightarrow \mathbf{TP}^{N-1}$ which in each affine chart not containing p is just given by $\mathbf{T}^N \rightarrow \mathbf{T}^{N-1}$ forgetting one coordinate (say, after reordering, the last one). Let $L' = \pi(L)$ be the image of L , For each point $x \in L'$, the preimage $\pi^{-1}(x) \cap L$ is either a single point or an interval of the form $[-\infty, a]$. In any other case, (e.g. an interval) one can the existence of a nearby fiber which contains at least two isolated points which contradicts the assumption that L is of degree 1. It follows that L , equipped with trivial weights, is a linear space again. Moreover, a generic fiber $\pi^{-1}(x) \cap L$ consists of only one point, say \tilde{x} , and we may define, in each affine chart, a piecewise linear function f on L' whose value at a generic x is given by the last coordinate of \tilde{x} . Let us now modify L' along f (in each chart). By construction, the graph of f is contained in L and ... Therefore L is in fact equal to $\text{Mod}(L', f)$. If $L' \subsetneq \mathbf{TP}^{N-1}$, we continue until we finally end up with \mathbf{TP}^n . □

6 Tropical cycles and the Chow group

6.1 Tropical cycles

We defined tropical cycles in toric varieties in section 5.1. The main idea was that tropical cycles should be polyhedral sets (with rational slopes) whose generic points carry multiplicities such that around each codimension one cell the balancing condition holds. As all these requirements are of a local nature, it is easy to extend the definition to arbitrary tropical varieties.

Let V be a tropical variety. A subset $X \subseteq V$ is called a *polyhedral set* if it is (finite closed) polyhedral in any chart. Recall that in charts with points of higher sedentarity, we define a polyhedron to be the closure of a usual polyhedron in \mathbf{R}^n . Let $x \in X$ be a point. We define the *speciality* of x , denoted by $\text{spec}_X(x)$ or just $\text{spec}(x)$, to be the minimal codimension of the polyhedra P such that $x \in P^\square \subseteq X$. Points with $\text{spec}_X(x) = 0$ are called generic points. Let X be of pure dimension m . Then the closure of all points of speciality $m - k$ is called the *k-skeleton* of X , denoted by $X^{(k)}$.

$$X^{(k)} := \overline{\{x \in X : \text{spec}(x) = m - k\}} = \{x \in X : \text{spec}(x) \geq m - k\}.$$

A polyhedral set is called *weighted* if it is equipped with a locally constant function $\text{mult} : X^{\text{gen}} \rightarrow \mathbf{Z} \setminus \{0\}$, where X^{gen} denotes the set of generic points of X . If $P \subseteq X$ is a polyhedron of maximal dimension m , then we also write $\text{mult}(P)$ for the number $\text{mult}(x)$ for any $x \in P^\square$.

Definition 6.1.1

Let V be a tropical variety. A *tropical k-cycle* X of V is a weighted polyhedral set $X \subseteq V$ of pure dimension k such that for any chart $U \in V$, any polyhedral

structure of $X \cap U$ and any codimension one cell $\tau \subset X \cap U$ the balancing condition

$$\sum_{\substack{\tau \subset \sigma \text{ facet} \\ \text{sed}(\tau) = \text{sed}(\sigma)}} \text{mult}(\sigma) v_{\sigma/\tau} = 0 \pmod{\mathbf{R}\tau}$$

is satisfied.

X is called of *pure sedentarity* if all its generic points have the same sedentarity (in V).

X is called *effective* if all its weights are positive, i.e. $\text{mult} : X^{\text{gen}} \rightarrow \mathbf{N}$.

Remark 6.1.2

Note that an effective tropical cycle $X \subseteq V$ is a tropical space itself by restricting the charts of V to X . In other words, in our terminology effective tropical cycles are just the closed tropical subvarieties of V . In particular, V satisfies the requirement of a cycle itself. Hence V is the fundamental cycle of itself.

Given two tropical cycles X_1 and X_2 , we can form the sum $X_1 + X_2$. We just take the union $X_1 \cup X_2$ (which is again a polyhedral set) and add weights. This means for a generic point x of $X_1 \cup X_2$ we set $\text{mult}_{X_1 + X_2}(x) = \text{mult}_{X_1}(x) + \text{mult}_{X_2}(x)$ (where $\text{mult}_X(x) = 0$ if $x \notin X$). If this sum turns out to be zero, we just remove the point. Thus, in general, $X_1 + X_2$ is only supported on a subset of $X_1 \cup X_2$. It is straightforward to check that $X_1 + X_2$ still satisfies the balancing condition. So the set of all cycles in V , denoted by $Z_*(V) = \bigoplus_k Z_k(V)$, forms a group under addition with neutral element the empty cycle $0 := \emptyset$.

One further remark: As above, we will always ignore points of weight zero. This is to say, whenever a construction (like summing two cycles) produces zero weights, we just discard these points.

6.2 Push-forwards of tropical cycles

Let V and W be tropical varieties and let $f : X \rightarrow Y$ be a tropical morphism. Given a tropical cycle $X \in Z_k(V)$, we define its push-forward $f_*(X) \in Z_k(W)$ as follows. Let $y \in f(X)$ be a point such that $X_y := f^{-1}(y) \cap X$ is isolated and generic in X (in particular, y is generic in $f(X)$ and X_y is finite). Fix

$x \in X_y$ and choose charts around x and y . Then f induces a map of lattices $df_x^Z : T_x^Z X \rightarrow T_y^Z f(X)$, and we define

$$\omega_{f_*(X)}(y) := \sum_{x \in X_y} [T_y^Z f(X) : \text{Im}(df_x^Z)] \cdot \omega_X(x). \quad (6.1)$$

Proposition 6.2.1

In the situation above, there is a unique k -cycle supported on $f(X)$ whose weight function agrees with (6.1) for sufficiently generic points. This cycle is called the push-forward of X , denoted by $f_(X) \in \mathbf{Z}_k(W)$.*

Proof. $f(X)$ is a polyhedral set in W . If $\dim(f(X)) < k$, then $f_*(X) = 0$. If $\dim(f(X)) = k$, let Y denote its k -dimensional part. Then equation (6.1) defines a locally constant weight function on a open polyhedral dense subset $U \in Y$. We have to show that this weight function satisfies the balancing condition.

Choose $y \in Y$ and let S denote the part of $\text{Star}_Y(y)$ given by points of the sedentarity (in Y). We have to show that S with weight function (6.1) is balanced. We use the following locality statement. Let $X_y := f^{-1}(y) \cap X$ and let $Z \subseteq X_y$ be the set of vertices of X_y . For each $z \in Z$, let S_z denote the sedentarity part of $\text{Star}_X(z)$ as above. Then, for $y' \in S$, we have

$$\omega_{f_*(X)}(y') = \sum_{z \in Z} \omega_{f_*(S_z)}(y').$$

This follows from the fact that when we let y' converge to y , then the preimage points in $X_{y'}$ have to converge to points in Z . Using this equation, we can assume that X is a fan, and f is integer linear. The balancing condition is a condition for the ridges of $f(X)$. Hence, by applying locality one more time and using the fact that dividing by a lineality space is compatible with the lattice index showing up in (6.1), we can assume that X is one-dimensional. We denote by u_ρ the primitive generator of a ray ρ . For any ray ρ' of Y we have

$$\omega_{f_*(X)}(\rho') = \sum_{\substack{\rho \subseteq X \\ f(\rho) = \rho'}} [T_{\rho'}^Z f(X) : \text{Im}(df_\rho^Z)] \omega_X(\rho).$$

Note that the primitive generators are related by

$$f(u_\rho) = [T_{\rho'}^Z f(X) : \text{Im}(df_\rho^Z)] u_{\rho'}.$$

Hence the balancing condition for Y follows from

$$\begin{aligned}
 \sum_{\rho' \subseteq Y} \omega_{f_*(X)}(\rho') u_{\rho'} &= \sum_{\rho' \subseteq Y} \sum_{\substack{\rho \subseteq X \\ f(\rho) = \rho'}} [T_{\rho'}^Z f(X) : \text{Im}(df_{\rho'}^Z)] \omega_X(\rho) u_{\rho'} \\
 &= \sum_{\rho \subseteq X} \omega_X(\rho) f(u_{\rho}) \\
 &= f \left(\sum_{\rho \subseteq X} \omega_X(\rho) u_{\rho} \right) = f(0) = 0.
 \end{aligned} \tag{6.2}$$

□

6.3 Linear equivalence of cycles

Let V be a tropical space of pure dimension n and consider the variety $V \times \mathbf{TP}^1$. Let Z be a cycle in $V \times \mathbf{TP}^1$ which is the closure of a cycle in $V \times \mathbf{R}$. Then Z can be intersected with $V_{-\infty} := V \times \{-\infty\}$. Namely, in every chart $U \times \mathbf{T}$ we can use Definition 5.3.4. It is easy to check that the results agree on the overlaps and therefore can be glued together to give the cycle $Z_{-\infty} := Z \cdot V_{-\infty}$. As $V \times \{-\infty\} = V$, we think of $Z_{-\infty}$ as a cycle in V (of dimension $\dim(Z) - 1$). In the same way, we can construct $Z_{+\infty} := Z \cdot V_{+\infty}$. This construction suffices to translate the classical definition of linear equivalence to the tropical world.

Definition 6.3.1

Let X_1, X_2 be two k -cycles in the tropical space V . Then X_1 and X_2 are called *linearly equivalent*, denoted by $X_1 \sim X_2$, if there exists a $(k+1)$ -cycle $Z \subseteq V \times \mathbf{TP}^1$ such that

- Z is the closure of a cycle in $V \times \mathbf{R}$,
- $X_1 = Z_{-\infty}$, and
- $X_2 = Z_{+\infty}$.

Lemma 6.3.2

The relation \sim defined in the previous definition is an equivalence relation.

Furthermore, we have

$$X_1 \sim X_2, Y_1 \sim Y_2 \Rightarrow X_1 + Y_1 \sim X_2 + Y_2.$$

Proof. To show $X \sim X$, we take $Z = X \times \mathbf{TP}^1$. The symmetry of \sim follows from the symmetry $\mathbf{TP}^1 \rightarrow \mathbf{TP}^1 : x \mapsto -x$. To show the compatibility with sums, let Z be the cycle in $V \times \mathbf{TP}^1$ showing $X_1 \sim X_2$ (and, analogously, Z' for $Y_1 \sim Y_2$). Then the sums $Z + Z'$ shows $X_1 + Y_1 \sim X_2 + Y_2$. Finally, using the additivity twice, we find

$$X_1 \sim X_2, X_2 \sim X_3 \Rightarrow X_1 + X_2 \sim X_2 + X_3 \Rightarrow X_1 \sim X_3,$$

which finishes the proof. \square

6.4 Algebraic equivalence of cycles

As in classical algebraic geometry, we may replace \mathbf{TP}^1 by any other smooth tropical curve to obtain another, weaker equivalence relation for tropical cycles. Let C be a smooth tropical curve, and let V be a tropical space. For a cycle $Z \subseteq V \times C$, we have to make sense of

$$Z_p = Z \cdot (V \times \{p\})$$

for any $p \in C$. If $\text{sed}(p) = 1$, C looks like \mathbf{T} near p and we can use Definition 5.3.4 again (assuming that Z is the closure of a cycle in $V \times C \setminus \{p\}$). If $\text{sed}(p) = 0$, then C locally near p looks like the line $L \subseteq \mathbf{R}^m$ with single vertex sitting at 0. So we can pull-back the function $\max\{x_1, \dots, x_m, 0\}$ to $V \times U$, where U is a neighbourhood of p . Let us denote this pull-back by φ_p . Then we define $Z_p = \text{div}(\varphi_p|_{Z \cap V \times U})$.

Definition 6.4.1

Let X_1, X_2 be two k -cycles in the tropical space V . Then X_1 and X_2 are called *algebraically equivalent*, denoted by $X_1 \stackrel{\text{alg}}{\sim} X_2$, if there exists a tropical curve C , two points $p_1, p_2 \in C$ and a $(k+1)$ -cycle $Z \subseteq V \times C$ such that

- Z is the closure of a cycle in $V \times C \setminus \{p_1, p_2\}$,
- $X_1 = Z_{p_1}$, and

- $X_2 = Z_{p_2}$.

Lemma 6.4.2

The relation $\overset{\text{alg}}{\sim}$ defined in the previous definition is an equivalence relation. Furthermore, we have

$$X_1 \overset{\text{alg}}{\sim} X_2, Y_1 \overset{\text{alg}}{\sim} Y_2 \Rightarrow X_1 + Y_1 \overset{\text{alg}}{\sim} X_2 + Y_2.$$

Proof. It suffices to show that the cycles X such that $X \overset{\text{alg}}{\sim} 0$ form a subgroup of $Z_*(V)$. So let X_1, X_2 be two tropical cycles which are algebraically equivalent to zero. We want to show that $X_1 - X_2 \overset{\text{alg}}{\sim} 0$. According to our definition, there exists a curve C_1 , two points p, p' and a cycle $Z_1 \subseteq V \times C_1$ such that $(Z_1)_p = X_1$ and $(Z_1)_{p'} = 0$. Analogously, we find $C_2, q, q' \in C_2$ and $Z_2 \subseteq V \times C_2$ such that $(Z_2)_q = X_2$ and $(Z_2)_{q'} = 0$. □

6.5 Linear systems

Let V be a tropical space of pure dimension n . Then a (Weil-)divisor $D \subset V$ is a subcycle of dimension $n - 1$. Therefore, the group of all divisors is $Z_{n-1}(V)$.

Definition 6.5.1

Let V be a tropical variety and let $D \subset V$ be a divisor. Then the set of all effective linearly equivalent divisors

$$L(D) = |D| := \{D' \geq 0 : D' \sim D\}$$

is called the *complete linear system* of D .

6.6 Recession Fans

Exercise 6.6.1

Let $\sigma \subset \mathbf{R}^n$ be a polyhedron. Show that the set of faces of $\text{rc}(\sigma)$ is equal to the set of cones $\text{rc}(\tau)$ for all faces $\tau \subset \sigma$.

Exercise 6.6.2

Let \mathcal{S} be a polyhedral subdivision of \mathbf{R}^n . A *path* of cells in \mathcal{S} is a sequence $\sigma_0, \dots, \sigma_n, \sigma_i \in \mathcal{S}$, such that for any i either $\sigma_i \subset \sigma_{i+1}$ or $\sigma_{i+1} \subset \sigma_i$. Let σ and σ' be cells of \mathcal{S} such that $v \in \text{rc}(\sigma) \cap \text{rc}(\sigma') \neq \emptyset$.

- (a) Show that there is a path of cells connecting σ and σ' such that $v \in \text{rc}(\sigma_i)$ for all i . You might want to reduce to the 2-dimensional case by choosing two generic points in σ and σ' .
- (b) Assume that $v \in \text{rc}(\sigma)^\square$. Conclude that the $\text{rc}(\sigma)$ is equal to the face of $\text{rc}(\sigma_i)$ containing v for all i .
- (c) Assume that $v \in \text{rc}(\sigma)^\square \cap \text{rc}(\sigma')^\square$. Conclude that $\text{rc}(\sigma) = \text{rc}(\sigma')$.
- (d) Prove that $\{\text{rc}(\sigma) : \sigma \in \mathcal{X}\}$ forms a fan.

7 Tropical manifolds

Lemma 7.0.1

Let $U \subset \mathbf{R}^n$ be open and connected and let $f : U \rightarrow \mathbf{R}$ be a smooth function. Then the following conditions are equivalent.

(a) For all $x \in U$ and $v \in \mathbf{Z}^n$, we have $df_x(v) \in \mathbf{Z}$.

(b) f is integer affine, i.e., of the form

$$f(x) = a + jx,$$

with $a \in \mathbf{R}$ and $j \in \mathbf{Z}^n$.

Proof. Let f be a function satisfying the first condition. This means $df_x \in (\mathbf{Z}^n)^* = \mathbf{Z}^n$ for all $x \in U$. As U is connected and \mathbf{Z}^n is discrete, df_x is constant to, say, $j \in \mathbf{Z}^n$. It follows that f is of the form $a + jx$. The other implication is trivial. \square

In the previous chapters all the objects we considered were embedded in some ambient variety, mostly \mathbf{R}^n or one of its toric compactifications, e.g. \mathbf{T}^n or \mathbf{TP}^n . We will now proceed to define what we mean by an *abstract* tropical space. We will do this by following the standard strategy using an atlas of charts and gluing maps. We will proceed step by step, starting with a very general definition allowing rather strangely behaved spaces and ending with the quite restrictive definition of smooth tropical spaces.

7.1 Tropical spaces

To get started, we have to describe the local building blocks we want to use in order to create abstract tropical spaces. We will give the relevant definitions here, noting that many of the notions we considered before are

of purely local nature. For the reader's convenience, we repeat some terminology which has appeared before. Recall that for any index set $I \subset [n]$, we denote by $\mathbf{R}_I \subset \mathbf{T}^n$ the torus orbit of points with coordinates $x_i = \infty$ for all $i \in I$.

Definition 7.1.1 (Open polyhedral sets)

A *polyhedron* in \mathbf{T}^n is the closure of a (usual) polyhedron in one of the torus orbits $\mathbf{R}_I \cong \mathbf{R}^{n-|I|}$. A union of polyhedra in \mathbf{T}^n is called a *polyhedral set in \mathbf{T}^n* . An *open polyhedral set* X is defined to be the open subset of a polyhedral set Y in \mathbf{T}^n .

A point $p \in X$ is called *generic* if there exists an open neighbourhood of p homeomorphic to an open subset in affine space \mathbf{R}^m . The set of generic points is denoted by X^{gen} . The integer m is called the *dimension of X at p* , denoted $\dim_X(p)$. Hence $\dim_X : X^{\text{gen}} \rightarrow \mathbf{N}$ is a locally constant function. Its maximum is called the *dimension of X* , denoted by $\dim(X)$. If \dim_X is globally constant, we say X is of *pure dimension*. Let X be pure-dimensional and $p \in X$. The *codimension of p in X* , denoted by $\text{codim}_X(p)$, is the minimum of the codimension of the cell whose relative interior contains p for any polyhedral structure of X .

As before, we now equip open polyhedral sets with weights.

Definition 7.1.2 (Weighted open polyhedral sets)

Let X be an open polyhedral set. A *weight function* on X is a locally constant function $\omega : X^{\text{gen}} \rightarrow \mathbf{Z}$, i.e. a choice of integer weight for every connected component of X^{gen} . The set X equipped with a weight function is called a *weighted open polyhedral set*. A weighted open polyhedral set is called *effective* if all weights are non-negative.

In the following we will always assume that all weights are non-zero. If zero weights show up (for example when adding to weighted open polyhedral sets), we remove them by taking the closure of $\omega^{-1}(\mathbf{Z} \setminus \{0\})$ in X . This is again an open polyhedral set.

We continue with the local formulation of the balancing condition.

Definition 7.1.3 (Open tropical cycles)

Let X be an open polyhedral set of pure dimension m and let $p \in X \cap \mathbf{R}_I$ be a point of sedentarity I . We define the *star of X at p* , denoted by $\text{Star}_X(p)$, to

the fan in \mathbf{R}_I containing all direction vectors v for which $p + \epsilon v \in X$ holds for small $\epsilon > 0$. The m -dimensional part of $\text{Star}_X(p)$, i.e. the closure of the set $\dim_{\text{Star}_X(p)}^{-1}(m)$, is denoted by $\text{Star}_X(p)^{(m)}$. If non-empty, $\text{Star}_X(p)^{(m)}$ inherits a weight function from the weights around p . We call X a *balanced open polyhedral set* or *open tropical cycle* if $\text{Star}_X(p)^{(m)}$ is empty or balanced (according to 2.4.6 and 4.2.2) for all $p \in X$. In other words, after choosing a (open) polyhedral structure for X , we may check the balancing condition as in Equation (5.1).

Note that the balancing condition at point $p \in \mathbf{R}_I$ only takes into account the full-dimensional parts of X near p contained in \mathbf{R}_I . In other words, balancing is checked in the locus of points of same sedentarity.

Balanced open polyhedral sets are the local building blocks for tropical spaces. Let us now consider the maps we use to glue these blocks together.

Recall from chapter 3 that in tropical arithmetics, monomial maps from \mathbf{R}^n to \mathbf{R}^m correspond simply to linear maps which map \mathbf{Z}^n to \mathbf{Z}^m . We call such maps *\mathbf{Z} -linear maps*. In chapter 2 we discussed that the tropical structure of \mathbf{R}^n manifests through the data of integer tangent vectors $T_x^{\mathbf{Z}}\mathbf{R}^n \subset T_x\mathbf{R}^n$ at each point $x \in \mathbf{R}^n$ (cf. Definition). Obviously, \mathbf{Z} -linear maps preserve this tropical structure. However, as we only fix the lattices of integer tangent vectors and not the lattice $\mathbf{Z}^n \subset \mathbf{R}^n$ itself, we may combine the monomial maps with arbitrary translations with translation vector $v \in \mathbf{R}^m$. Still, the tropical structure of integer tangent vectors is preserved. Such maps are called *affine \mathbf{Z} -linear maps*. Conversely, any smooth map $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ whose differential is \mathbf{Z} -linear at each point is affine \mathbf{Z} -linear, as $\text{Mat}(n, m, \mathbf{Z})$ is discrete (cf. Lemma 7.0.1).

Definition 7.1.4 (Tropical morphisms for open tropical cycles)

A map $F : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is called *affine \mathbf{Z} -linear* if it is the sum of a \mathbf{Z} -linear map $\mathbf{R}^n \rightarrow \mathbf{R}^m$ and a translation by an arbitrary vector $v \in \mathbf{R}^m$.

A map $f : \mathbf{R}^n \rightarrow \mathbf{T}^m$ is called *affine \mathbf{Z} -linear* if it is the composition of an affine \mathbf{Z} -linear map $\mathbf{R}^n \rightarrow \mathbf{R}_I$ and the inclusion map $\mathbf{R}_I \hookrightarrow \mathbf{T}^m$, for a suitable torus orbit \mathbf{R}_I . Note that we also include the constant map $f \equiv (-\infty, \dots, -\infty)$ corresponding to $I = [n]$.

Let $X \subseteq \mathbf{T}^n$ and $Y \subseteq \mathbf{T}^m$ be two open polyhedral sets. A map $f : X \rightarrow Y$ is called a *tropical morphism* if it is continuous and locally affine \mathbf{Z} -linear. This means that for each point $p \in X$ there exists an open neighbourhood

U of p in X such that for all $I \subset [n]$ with $U \cap \mathbf{R}_I \neq \emptyset$, the map $f|_{U \cap \mathbf{R}_I}$ is the restriction of an affine \mathbf{Z} -linear map $\mathbf{R}_I \rightarrow \mathbf{T}^m$.

Assume moreover that X and Y are weighted. Then f is called a *tropical isomorphism* if there exists an inverse tropical morphism $g : Y \rightarrow X$ and, moreover, if

$$\omega(x) = \omega(f(x))$$

for all $x \in X^{\text{gen}}$. Note that the existence of g implies that $f(x)$ is a generic point of Y .

Example 7.1.5

The group of tropical automorphisms of \mathbf{R}^n is the semidirect product

$$\text{Aut}(\mathbf{R}^n) = \text{GL}(n, \mathbf{Z}) \ltimes \mathbf{R}^n,$$

where the first factor represents \mathbf{Z} -invertible linear maps and the second factor parameterizes translations.

Example 7.1.6

The group of tropical automorphisms of \mathbf{T}^n is the semidirect product

$$\text{Aut}(\mathbf{T}^n) = S_n \ltimes \mathbf{R}^n,$$

where the symmetric group represents permutations of the variables and the second factor parameterizes translations.

Note that there are some subtleties hidden in the definition of tropical morphisms when higher sedentarity points are involved. The following example shows some of that behaviour.

Example 7.1.7

We will compare four polyhedral sets in \mathbf{T}^2 . Let $u, v, w \in \mathbf{Z}_+^2$ be integer vectors with all entries positive and assume that v, w are linearly independent. We define $X_1 = \mathbf{T}_{\{1\}} \cup \mathbf{T}_{\{2\}}$ (union of the two coordinate axes at infinity), $X_2 = \mathbf{T}_{\{1\}} \cup \mathbf{T}u$ (one coordinate axis plus a ray), $X_3 = \mathbf{T}u \cup (\mathbf{T}u + p)$ for some $p \notin \mathbf{T}u$ (two parallel rays) and $X_4 = \mathbf{T}v \cup \mathbf{T}w$ (two non-parallel rays). All of these four sets contain the corner point $o = (-\infty, -\infty)$ and we are interested in whether or not two such sets are isomorphic in a neighbourhood of o . It turns out that such isomorphic neighbourhoods exist for X_1 and X_2 ,

but for none of the remaining pairs. Note that some authors in other contexts use different definitions for tropical morphisms which may discard or allow for extra isomorphisms in this example.

We are now ready to introduce the rather general notion of a tropical space, obtained from glueing balanced open polyhedral sets. Later on, we will add more and more additional requirements to these spaces to make them more manageable.

Definition 7.1.8

Let X be a topological space. A *tropical atlas* or *tropical structure* on X is a collection of tuples $(U_i, \psi_i, V_i)_i$ subject to the following constraints:

- $X = \bigcup_i U_i$ is an open covering of X .
- For each i , $V_i \subset \mathbf{T}^N$ (for suitable N) is an effective balanced open polyhedral set.
- For each i , $\psi_i : U_i \rightarrow V_i$ is a homeomorphism.
- For each pair i, j with $U_i \cap U_j \neq \emptyset$, the composition map

$$\psi_i \circ \psi_j^{-1} : \psi_j(U_i \cap U_j) \rightarrow \psi_i(U_i \cap U_j)$$

is a tropical morphism with $\omega(x) = \omega(\psi_i \circ \psi_j^{-1}(x))$ whenever x and $\psi_i \circ \psi_j^{-1}(x)$ are generic. In other words, $\psi_i \circ \psi_j^{-1}$ is a tropical isomorphism with inverse $\psi_j \circ \psi_i^{-1}$.

Two atlases on X are called equivalent if their union also forms an atlas. A *tropical space* is a topological space X together with the choice of an equivalence class of atlases.

If all V_i are of pure dimension n , we say that X is of dimension n .

Let us list some obvious properties of a tropical space X .

- X contains an open dense subset of generic points X^{gen} .
- The weights on the various charts are compatible and therefore glue to give a locally constant weight function on X^{gen} .

- For each point $p \in X$, the *codimension* $\text{codim}_X(p)$ is well-defined as the codimension of $\psi_i(p)$ in V_i for any chart containing p . A point $p \in X$ is generic if and only if $\text{codim}_X(p) = 0$.
- For each point $p \in X$, we can consider $\text{Star}_X(p) \subseteq \mathbf{R}^N$ (for suitable N). If non-empty, its n -dimensional part $\text{Star}_X(p)^{(n)}$ is a balanced fan of dimension n . It is well-defined only up to tropical isomorphisms of fans given by \mathbf{Z} -linear maps in $\text{Mat}(N, N', \mathbf{Z})$.

By definition, tropical spaces can locally be stratified as rational polyhedral fans. Note however that the definition does not imply the existence of *global* stratification (e.g. as a CW complex) into rational polyhedral cells (i.e. cells which, locally in each chart, are rational polyhedral). Here is an example.

Example 7.1.9

Let $\Lambda \subset \mathbf{R}^2$ be the lattice generated by the vectors $(1, \pi)$ and $(1, e)$. Then $X = \mathbf{R}^2/\Lambda$ is a tropical space, with charts given by restrictions of the projection map $\mathbf{R}^2 \rightarrow X$. Note that here we stick to integer affine structure on \mathbf{R}^2 (and hence also on X) given by the standard lattice $\mathbf{Z}^2 \subset \mathbf{R}^2$. It follows that X cannot be stratified using rational polyhedra (it can, however, be triangulated by *non-rational* triangles).

Definition 7.1.10

Let X be a tropical space and let $Y \subset X$ be a subset. Then Y is called a *polyhedral subset of X* if Y is closed in X and if the image set $\psi(Y \cap U)$ under every chart (U, ψ, V) of X is an open polyhedral set in \mathbf{T}^N .

Polyhedral subsets often arise as the sub- resp. superlevel sets of the following class of functions.

Definition 7.1.11

Let X be a tropical space and let $f : X \rightarrow M$ be a function to some set M . Then f is called *semiconstant* if it is constant on the relative interior of every cell of every polyhedral structure of every chart of X . Moreover, let \leq be a total order on M . Then f is called *lower (resp. upper) semiconstant* if for every containment of cells $\tau \subset \sigma$ we have $f(\tau^\square) \leq f(\sigma^\square)$ (resp. $f(\tau^\square) \geq f(\sigma^\square)$).

We have the following straightforward statement.

Proposition 7.1.12

Let $f : X \rightarrow M$ be a lower (resp. upper) semiconstant function on a tropical space X and choose $m \in M$. Then the sublevel (resp. superlevel) set

$$\{p \in X : f(x) \leq m\} \text{ (resp. } \{p \in X : f(x) \geq m\})$$

is a polyhedral subset of X .

A typical example of an upper semiconstant function on X is the codimension function codim_X . It gives rise to a corresponding filtration and stratification of X , which we collect in the following definition.

Definition 7.1.13

For $k \in \mathbb{N}$, the superlevel set

$$X^{(k)} := \{x \in X \mid \text{codim}_X(p) \geq n - k\}$$

is called the k -skeleton of X . By the previous proposition, $X^{(k)}$ is a polyhedral subset of X . A connected component of

$$X^{(k)} \setminus X^{(k-1)} = \{x \in X \mid \text{codim}_X(p) = n - k\}$$

is called a *combinatorial stratum* of X . The collection of these strata is called the *combinatorial stratification* of X .

In most situations, we would like to impose some kind of finite type condition on tropical spaces. This regards the local polyhedral sets V_i (a priori, we might allow polyhedral sets which have an infinite number of cells), the number of charts needed to give an atlas of X and, finally, the completeness of X (open polyhedral sets can just stop at some bounded distance — this is usually undesirable for tropical spaces). We will address all these issues in the following definition.

Definition 7.1.14

A tropical space X is called *of finite type* if it admits a tropical atlas $\{(U_i, \psi_i, V_i)\}_i$ subject to the following conditions.

- (a) The number of charts is finite, i.e. $i \in I$ is taken from a finite index set I .
- (b) Each V_i is the open subset of a finite polyhedral set, i.e. the union of finitely many polyhedra. This was implicit in the general definition, but we want to emphasize it here.
- (c) Each chart can be extended, i.e. there exists a chart (U'_i, ψ'_i, V'_i) such that $\overline{V}_i \subset V'_i$ for all i . The closure is taken in \mathbf{T}^N .

A Hausdorff tropical space of pure dimension and of finite type is called a *tropical variety*.

Example 7.1.15

A (finite) balanced effective polyhedral set in \mathbf{T}^n is a tropical space of finite type. More general, tropical cycles, considered as polyhedral sets, in any toric variety are examples of tropical spaces of finite type. Note however that in order to satisfy the third condition from above, it might be necessary to split a chart in \mathbf{T}^n into several smaller ones. For example, tropical cycles in \mathbf{R}^n are of finite type, but in general we have to use more than one embedding $\mathbf{R}^n \subset \mathbf{T}^n$ to satisfy the third condition.

Example 7.1.16

A simple example of a tropical space which is *not* of finite type is the unit interval $(0, 1) \subset \mathbf{T}$.

Example 7.1.17

Here is a strange example of a tropical space of finite type — we call it ρ space, denoted X_ρ . It is obtained from \mathbf{TP}^1 simply by glueing the points $+\infty$ and 0 . The glued point is denoted $O \in X_\rho$. The result is depicted in Figure 7.1. This quotient space can be made a tropical space by using a neighbourhood of $(-\infty, 0) \in V("x_1(x_2 + 0)") \subset \mathbf{T}^2$ as a chart for O . Note that in a certain sense the point O has sedentarity 0 and 1 at the same time, so for general tropical spaces we need to be careful when we speak about sedentarity.




Figure 7.1: The so-called ρ -space X_ρ

7.2 Regularity at Infinity

As we know from chapter 5, tropical varieties come naturally in classes given by the process of modification. This reflects the fact that, when comparing tropical to classical varieties, they correspond not only to a variety but to a tuple (X, D) where X is a variety and $D \subset X$ is a distinguished divisor. In the tropical world, this divisor is given by points of higher sedentarity and can be changed by modifications. Unlike in the classical world, we cannot completely forget this additional choice of divisor. As a consequence, when we want to define a good class of tropical spaces, we are forced to think about requirements imposed on the divisor at infinity. Probably the most important property to impose is *regularity at infinity*. It should be compared to the normal crossing divisor property when considering a tuple (X, D) in the classical world.

We start by proving the local version of the fact that the locus of points of sedentarity k has codimension at most k .

Proposition 7.2.1

Let $X \subseteq \mathbf{T}^N$ be an open polyhedral set of pure dimension n . Let \mathbf{R}_I be a torus orbit of \mathbf{T}^N such that $X \cap \mathbf{R}_I \neq \emptyset$. Then it holds

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

Proof. Choose a polyhedral structure for X . Then the cells contained in \mathbf{R}_I form a polyhedral structure for $X \cap \mathbf{R}_I$. Choose a cell τ of $X \cap \mathbf{R}_I$ of maximal dimension, i.e. $\dim(\tau) = \dim(X \cap \mathbf{R}_I)$. As X is pure-dimensional, there exists a cell σ of X of dimension n containing τ . The maximality of τ

implies $\sigma \cap \mathbf{R}_I = \tau$. Let \mathbf{R}_J be the torus orbit containing the relative interior of σ (we have $J \subset I$). Then we can apply 3.2.6 to $\sigma \subset \mathbf{T}_J$ and obtain

$$n = \dim(\sigma) = \dim(\tau) + \dim(\mathbf{R}\sigma \cap \langle e_i : i \in I \setminus J \rangle) \leq \dim(X \cap \mathbf{R}_J) + |I|.$$

□

Remark 7.2.2

Note that when replacing \mathbf{R}_I by \mathbf{T}_I , the previous proposition is not true anymore (at least as long as we do not impose the balancing condition or something similar). A simple counterexample is given in Figure 7.2. For the same reason, the proof of the proposition can not proceed inductively, i.e. can not be reduced to the case $|I| = 1$.




Figure 7.2: A 2-dimensional polyhedron with a single vertex at infinity

We now consider the case when the inequality of proposition 7.2.1 is sharp.

Proposition 7.2.3

Let $X \subseteq \mathbf{T}^N$ be an open polyhedral set of pure dimension n . Then the following conditions are equivalent.

- (a) For each torus orbit $\mathbf{R}_I \subset \mathbf{T}^N$ such that $X \cap \mathbf{R}_I \neq \emptyset$, we have

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

- (b) For each stratum $\mathbf{T}_I \subset \mathbf{T}^N$ such that $X \cap \mathbf{T}_I \neq \emptyset$, we have

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

(c) For each point $p \in X$, we have

$$\text{sed}_{\mathbf{T}^N}(p) \leq \text{codim}_X(p).$$

Proof. (a) \Rightarrow (c). Let $p \in X$ be an arbitrary point. Assume $p \in \mathbf{R}_I$. Note that a polyhedron in \mathbf{T}^N has the same sedentarity as any of point in its relative interior, hence $\dim(X \cap \mathbf{R}_I) \geq n - \text{codim}_X(p)$. We get

$$\text{codim}_X(p) \geq n - \dim(X \cap \mathbf{R}_I) = |I| = \text{sed}_{\mathbf{T}^N}(p).$$

(c) \Rightarrow (b). Assume $X \cap \mathbf{T}_I \neq \emptyset$. Choose a polyhedral structure of X and a point $p \in P^\square$, where P is a cell of $X \cap \mathbf{T}_I$ of dimension $\dim(X \cap \mathbf{T}_I)$. It follows $\dim(X \cap \mathbf{T}_I) = n - \text{codim}_X(p)$. But $\text{codim}_X(p) \geq \text{sed}_{\mathbf{T}^N}(p) \geq |I|$ by assumption, and hence $\dim(X \cap \mathbf{T}_I) \leq n - |I|$.

(b) \Rightarrow (a). Assume $X \cap \mathbf{R}_I \neq \emptyset$. We have $X \cap \mathbf{R}_I \subset X \cap \mathbf{T}_I$ and hence

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

But by 7.2.1 we also have $\dim(X \cap \mathbf{R}_I) \geq n - |I|$, and hence equality holds. \square

Definition 7.2.4

Let $X \subseteq \mathbf{T}^N$ be an open polyhedral set of pure dimension n . Then X is called *regular at infinity* if the equivalent conditions of 7.2.3 hold. A tropical space X is called *regular at infinity* if it admits a tropical atlas $\{(U_i, \psi_i, V_i)\}_i$ where all V_i are regular at infinity.

Example 7.2.5

Let us give a few examples.

- (a) The polyhedron depicted in Figure 7.2 is regular at infinity.
- (b) The hyperplane $H := V("x_1 + x_2 + x_3") \subset \mathbf{T}^3$ is not regular at infinity, because at $p = (-\infty, -\infty, -\infty) \in H$ we have $\text{sed}_{\mathbf{T}^3}(p) = 3 > 2 = \text{codim}_H(p)$. Note that the "divisor" of points of higher sedentarity consists of the 3 lines $H \cap \mathbf{T}_i$, which meet at p .
- (c) Consider the line $L = V("x_1 + x_2") \subset \mathbf{T}^2$. As a open polyhedral set, L is not regular at infinity, as $\text{sed}_{\mathbf{T}^2}((-\infty, -\infty)) = 2 > 1 = \text{codim}_L((-\infty, -\infty))$. However, as tropical space, we may instead use the chart $L \cong \mathbf{T}^1$ and hence the tropical space L is regular at infinity.

- (d) Let Y be a tropical toric surface and let $X \subset Y$ be a curve (i.e. an effective 1-cycle) such that $X \cap \mathbf{R}^2 \neq \emptyset$. Then the tropical space X is regular at infinity if and only if it does not contain any of the boundary curves and if for each torus fixed point q there is at most one ray in $X \cap \mathbf{R}^2$ converging to q . In this case, any point $p \in X$ of higher sedentarity admits a chart homeomorphic to a neighbourhood of $-\infty$ in \mathbf{T} .
- (e) Let C be a one-dimensional tropical space. Then C is regular at infinity if and only if can be covered by charts which are isomorphic to an open subset of \mathbf{T} or a 1-cycle in \mathbf{R}^N .
- (f) The ρ -space constructed in 7.1.17 is not regular at infinity.
- (g) The integer-affine manifolds defined in 1.3.2 are examples of tropical spaces which are regular at infinity. Indeed, all charts are open subsets of \mathbf{R}^n and hence trivially regular at infinity.

As suggested by these examples, one property of tropical spaces regular at infinity is that the notion of sedentarity makes sense. Indeed, note that for an open polyhedral set $X \subset \mathbf{T}^N$ regular at infinity, we have $X^{\text{gen}} \subset \mathbf{R}^N$. Given this it follows easily that the sedentarity $\text{sed}_{\mathbf{T}^N}(p)$ does not change under an isomorphism of open polyhedral sets regular at infinity. We can hence make the following definition.

Definition 7.2.6

Let X be a tropical space which is regular at infinity. For $p \in X$ we define the *sedentarity of p in X* to be

$$\text{sed}_X(p) := \text{sed}_{\mathbf{T}^N}(\psi(p)),$$

where (U, ψ, V) is a chart with $p \in U$ and V regular at infinity.

Example 7.2.7

The class of tropical spaces X which are regular at infinity and such that every point satisfies $\text{codim}_X(p) = \text{sed}_X(p) = 0$ coincides with the class of integer-affine manifolds defined in 1.3.2.

The main feature of regularity at infinity is that each point of sedentarity k admits a chart which factors into a product of a sedentarity 0 part and

a \mathbf{T}^k -part. Before we formulate this, let us fix some notations. Let X be a tropical space which is regular at infinity. Note that the sedentarity function sed_X is another example of an upper semiconstant function on X . For $k \in \mathbf{N}$, the superlevel set

$$X^{[k]} := \{p \in X : \text{sed}_X(p) \geq k\}$$

is called the k -sedentarity locus of X . By 7.1.12, $X^{[k]}$ is a polyhedral subset of X . We also use the shorthand $X_\infty := X^{[1]}$.

Theorem 7.2.8

Let X be a tropical space which is regular at infinity. Then each point $p \in X$ of sedentarity k has a chart isomorphic to $U \times V \subseteq X^{[k]} \times \mathbf{T}^k$, where U and V are open neighbourhoods of p resp. $(-\infty, \dots, -\infty)$ in $X^{[k]}$ resp. \mathbf{T}^k .

Proof. As we want to prove a local statement, it we may restrict to the case where X is an open polyhedral set in \mathbf{T}^N which is regular at infinity. Choose I such that $X \cap \mathbf{R}_I \neq \emptyset$ and consider $X' := (X \cap \mathbf{R}_I) \times \mathbf{T}^I \subset \mathbf{T}_I \times \mathbf{T}^I = \mathbf{T}^N$. Regularity at infinity implies that X' is a polyhedral set of dimension $\dim(X)$. It follows from 3.2.6 that the intersection $X \cap X'$ contains an open neighbourhood of $p \in X \cap \mathbf{R}_I$ both in X and X' . As the topology on \mathbf{T}^N agrees with the product topology on $\mathbf{T}_I \times \mathbf{T}^I$, this implies that every $p \in X \cap \mathbf{R}_I$ admits a open neighbourhood in X of the form $U \times V$ where U is a open neighbourhood of p in $X \cap \mathbf{R}_I$ and V is a open neighbourhood of $(-\infty, \dots, -\infty)$ in \mathbf{T}^I . This is what we asked for, as $X \cap \mathbf{R}_I$ is obviously an open subset of $X^{[I]}$. \square

Let $Y \subset \mathbf{R}^N$ be a open polyhedral set. Note that a weight function on the open polyhedral set $Y \times \mathbf{T}^k$ induces a weight function on Y (and vice versa). Moreover, if this weight function satisfies the balancing condition (on $Y \times \mathbf{T}^k$), the induced weight function is also balanced (on Y). Hence, with the help of the previous theorem we can turn the k -skeletons into tropical spaces in their own right.

Definition 7.2.9

Let X be a tropical space which is regular at infinity. Then we define a weight function on the k -skeleton $X^{[k]}$ by defining the weight of a generic point $p \in X^{[k]}$ to the weight of a nearby generic point of X . This turns $X^{[k]}$ into a tropical space of pure dimension $n - k$. We call X_∞ the *divisor at infinity* of X .

Theorem 7.2.8 can be refined as follows. Note that by our convention $\text{Star}_X(p)$ only takes into account points of same sedentarity, i.e. we have $\text{Star}_X(p) = \text{Star}_{X[\text{sed}(p)]}(p) \subset \mathbf{R}^N$. Hence by 7.2.8 every point in X admits a chart isomorphic to an open neighbourhood of $(0, \dots, 0, -\infty, \dots, -\infty) \in \text{Star}_X(p) \times \mathbf{T}^{\text{sed}_X(p)}$. On the other hand, every tropical space which can be covered by such charts is obviously regular at infinity. We reformulate this as follows.

Definition 7.2.10

Let F be a tropical fan cycle in \mathbf{R}^N , $U \in \mathbf{R}^N$ a convex open set containing the origin and $k \in \mathbf{N}$. The open neighbourhood

$$(0, \dots, 0, -\infty, \dots, -\infty) \in V := (F \cap U) \times [-\infty, 0)^k \subset F \times \mathbf{T}^k.$$

is called a *standard neighbourhood*.

Let X be a tropical space and $p \in X$. A chart of p to V is called a *standard chart of p* .

Corollary 7.2.11

A tropical space X is regular at infinity if and only if it can be covered by standard charts.

In view of these considerations, the following generalization suggests itself and will be useful later.

Definition 7.2.12

Let F be a tropical fan cycle in \mathbf{R}^N , $U \in \mathbf{R}^N$ a convex open set containing the origin and $k, M \in \mathbf{N}$ with $k \leq M$. We set S_n^k to be the union of coordinate planes of dimension k in \mathbf{T}^n , i.e.

$$S_n^k := (\mathbf{T}^n)^{[n-k]} = \bigcup_{\substack{I \subset [n] \\ |I|=n-k}} \mathbf{T}_I.$$

The open neighbourhood

$$(0, \dots, 0, -\infty, \dots, -\infty) \in V := (F \cap U) \times (S_n^k \cap [-\infty, 0)^n) \subset F \times S_n^k$$

is called a *normal crossing neighbourhood*.

Let X be a tropical space and $p \in X$. A chart of p to V is called a *normal crossing chart of p* . If X can be covered by normal crossing charts, X is called *normal crossing*.

Note that k is an invariant of all normal crossing charts of a point p , hence $\text{sed}_X(p) = k$ is again well-defined.

Example 7.2.13

Let X be a tropical space which is regular at infinity. Then all its skeletons $X^{[k]}$ are normal crossing tropical spaces. It holds $(X^{[k]})^{[k']} = X^{[k+k']}$.

7.3 Tropical morphisms and structure sheaves

After the preceding discussions, the notion of a tropical morphism between tropical spaces is be clear. We state it here for reference.

Definition 7.3.1

Let X and Y be two tropical spaces. A continuous map $f : X \rightarrow Y$ is called a *tropical morphism* if for each pair of charts (U, ψ, V) for X and (U', ψ', V') for Y the composition map

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

is a tropical morphism of open polyhedral sets in the sense of definition 7.1.4.

Which classes of functions are preserved under tropical morphisms? Let us give some examples. Recall that for an open polyhedral set X a tropical morphism $f : X \rightarrow \mathbf{T}$ is just a locally affine \mathbf{Z} -linear function (see 7.1.4).

Definition 7.3.2

Let X be a tropical space. We define the *sheaf of locally affine \mathbf{Z} -linear functions*, denoted $\mathcal{A}ff_X$, by setting

$$\begin{aligned} \mathcal{A}ff_X(U) &:= \{f : U \rightarrow \mathbf{T} : f \text{ tropical morphism}\} \\ &= \{f : U \rightarrow \mathbf{T} : f \text{ locally affine } \mathbf{Z}\text{-linear}\}. \end{aligned} \tag{7.1}$$

Note that the invertible sections $\mathcal{A}ff_X^*(U)$ of this sheaf are given by tropical morphisms/locally affine \mathbf{Z} -linear functions to \mathbf{R} . The sheaf $\mathcal{A}ff_X$ as well as most of the sheaves we will introduce has very special properties, e.g. it is semiconstant (the stalk map $p \mapsto \mathcal{A}ff_{X,p}$ is semiconstant in the sense of 7.1.11) and hence constructible.

Example 7.3.3

Let us give a few examples.

- (a) Let us start with $X = \mathbf{T}^n$. We obviously have

$$\begin{aligned} \mathcal{A}ff_X(\mathbf{R}^n) \setminus \{-\infty\} &\cong (\mathbf{Z}^n)^\vee \times \mathbf{R} \cong \mathbf{Z}^n \times \mathbf{R}, \\ f &\mapsto (df, f(0)). \end{aligned} \tag{7.2}$$

In order to extend f to R_i , we need $df(-e_i) \leq 0$ (or equivalently $df(e_i) \geq 0$) for all $i \in I$. In particular, the the space of non-zero global sections is $\mathcal{A}ff_X(\mathbf{T}^n) \setminus \{-\infty\} \cong \mathbf{N}^n \times \mathbf{R}$. The same holds for every connected open neighbourhood of $O = (-\infty, \dots, -\infty)$.

- (b) Let $F \subset \mathbf{R}^n$ be a tropical fan cycle of sedentarity zero. Let us assume that F is non-degenerate, i.e. not contained in a proper subspace of \mathbf{R}^n . Let U be a connected open neighbourhood of 0. Then

$$\mathcal{A}ff_F(U) \setminus \{-\infty\} = \mathcal{A}ff_{F,0} \setminus \{-\infty\} = \mathcal{A}ff_{\mathbf{R}^n}(\mathbf{R}^n) \setminus \{-\infty\} \cong \mathbf{Z}^n \times \mathbf{R}.$$

- (c) Let X be a tropical space and let (U, ψ, V) be a standard chart for $p \in X$. Hence V is a connected open subset of $\text{Star}_X(p) \times \mathbf{T}^k$. We may assume $\text{Star}_X(p) \subset \mathbf{R}^n$ is non-degenerate. Then

$$\mathcal{A}ff_{X,p} \setminus \{-\infty\} = \mathcal{A}ff_X(U) \setminus \{-\infty\} \cong \mathbf{Z}^n \times \mathbf{N}^k \times \mathbf{R}.$$

Here, the first factor encodes the differential in \mathbf{R}^n -directions (the directions of “same” sedentarity), the second factor encodes the \mathbf{R}^k -directions (the directions to smaller sedentarity) and the last factor again just parameterizes shifts by a constant. Be careful, however, that in general the $(\mathbf{Z}^n)^\vee$ -coordinates are not determined by the restriction of f to the points of k -sedentarity $X^{[k]} \cap U$. Namely, if one of the \mathbf{N}^k -coordinates is positive, this restriction takes constant value

$-\infty$, but the $(\mathbf{Z}^n)^\vee$ -coordinates are still well-defined and might be non-zero. The sheaf $\mathcal{A}ff_X^*$ of invertible sections can be described, via the above identification, as follows.

$$\mathcal{A}ff_{X,p}^* = \mathcal{A}ff_X^*(U) \cong \mathbf{Z}^n \times \{0\} \times \mathbf{R}$$

- (d) Let X_ρ be the ρ -space constructed in 7.1.17. We denote by O the point obtained from glueing 0 and $+\infty$. Then the stalk at O can be described as

$$\mathcal{A}ff_{X_\rho,O} \setminus \{-\infty\} \cong (\mathbf{Z} \times \mathbf{R}) \sqcup (\mathbf{Z}_{>0} \times \mathbf{R}).$$

The first part describes functions with $f(O) \neq -\infty$ using the slope along the 0-branch and the value at O . The second part describes functions with $f(O) = -\infty$ using the slope along the $+\infty$ -branch and a shift parameter.

These examples suggest to make a few more definitions related to the stalks of $\mathcal{A}ff_X$.

Definition 7.3.4

Let X be a tropical space and $p \in X$ be a point. The *cotangent space lattice of P at p* is defined to be

$$\mathbf{Z}T_p^* := \{f \in \mathcal{A}ff_{X,p}^* \mid f(p) = 0\} (\cong \mathcal{A}ff_{X,p}^* / \mathbf{R}). \quad (7.3)$$

$\mathbf{Z}T_p^*$ forms a group with respect to usual addition (tropical multiplication) of germs. Moreover, $\mathbf{Z}T_p^*$ is free abelian and finitely generated. The *cotangent space of P at p* is defined to be the \mathbf{R} -vectorspace

$$T_p^* := \mathbf{Z}T_p^* \otimes \mathbf{R}. \quad (7.4)$$

By dualizing, we obtain the *tangent space (lattice) $\mathbf{Z}T_p \subset T_p$* .

Given a morphism of tropical spaces $f : X \rightarrow Y$, pulling back functions induces a linear map $(df)_x^* : \mathbf{Z}T_{Y,f(x)}^* \rightarrow \mathbf{Z}T_{X,x}^*$ for all $x \in X$. The dual linear map

$$df_x : \mathbf{Z}T_{X,x} \rightarrow \mathbf{Z}T_{Y,f(x)}$$

is called the *differential of f at x* .

Remark 7.3.5

One nice feature of the tangent space T_p is that it contains $\text{Star}_X(p)$ canonically and non-degenerately. Hence we can remove the choice of representation of $\text{Star}_X(p)$ and regard it as embedded in the (abstract) vector space T_p . If X is regular/normal crossing at infinity, $\text{Star}_X(p) \subset T_p$ is a well-defined tropical fan cycle of sedentarity zero.

For completeness, we add the following complementary definition.

Definition 7.3.6

Let X be a tropical space and let $p \in X$ be a point. We define the *slope semigroup of X at p* to be

$$S_p := \text{Aff}_{X,p} \setminus \{-\infty\} / \text{Aff}_{X,p}^*. \tag{7.5}$$

If X is regular at infinity, by 7.3.3 S_p is isomorphic to $\mathbf{N}^{\text{sed}(p)}$. On a standard neighbourhood of p , the isomorphism is given by the slopes of a germ in the e_i -direction (where e_i is the standard basis vector in $\mathbf{R}^k \subseteq \mathbf{T}^k$). All the notions defined above fit together in the exact sequence

$$0 \rightarrow \mathbf{Z}T_p^* \times \mathbf{R} \rightarrow \text{Aff}_{X,p} \setminus \{-\infty\} \rightarrow S_p \rightarrow 0.$$

7.4 Smooth tropical spaces

The category of tropical spaces of finite type and regular at infinity forms a reasonable class of objects for tropical geometry. However, so far we have not imposed any local condition which might reflect the classical concept of smoothness. In particular, we can not expect to get fully-fledged homology and intersection theory in such a generality. In the past, different authors used different concepts of smoothness and it seems that there is no “best” definition of smoothness in the tropical world which reflects all the features of the classical theory.

There is one definition, however, which seems especially useful. This is due to the fact that, on one hand, the definition is simple and fits naturally in the tropical theory developed so far. On the other hand, while being sufficiently restrictive to provide us with many of the properties we would expect from smooth spaces, it is at the same time interesting and flexible

enough — even locally, where it leads to interesting combinatorics related to matroid theory and hyperplane arrangements. In fact, the definition essentially says that a smooth tropical space should locally look like a tropical plane in \mathbf{TP}^n .

Recall that in section 5.2 we defined the degree of a k -cycle in \mathbf{R}^n (resp. \mathbf{TP}^n) to be $\deg(X) = \deg(X \cdot H^k)$, where the product is given by stable intersection and H denotes the standard hyperplane of \mathbf{R}^n (resp. \mathbf{TP}^n). Let $X \subset \mathbf{R}^n$ be an effective tropical fan with $\deg(X) = 1$. A balanced open polyhedral set obtained as a open subset of $\bar{X} \subset \mathbf{T}^n$ is said to be of *degree 1*.

Definition 7.4.1

A tropical space X is called *smooth* if it can be covered by charts to degree 1 open polyhedral sets. We call X a *tropical manifold* if it is of finite type, regular at infinity and smooth.

Note that smoothness does not in general imply regularity at infinity, as example 7.2.5 (b) shows.

Let Λ be a lattice and $\Lambda \times \mathbf{R}$ the associated vector space. Let $X \subset V$ be a tropical fan of pure dimension k . The choice of a basis $B = (v_1, \dots, v_n)$ of \mathbf{Z}^n gives us a tropical isomorphism $\Phi_B : V \rightarrow \mathbf{R}^n$.

Definition 7.4.2

We define the *degree of X with respect to the basis B* to be

$$\deg_B(X) := \deg(\Phi_B(X)).$$

A tropical fan $X \subset V$ is called *degree 1 fan* if it is effective and if there exists a basis B such that $\deg_B(X) = 1$.

Corollary 7.4.3

Let X be a tropical space regular at infinity. Then X is smooth if and only if for all $p \in X$ the star $\text{Star}_X(p) \subset T_p$ is a degree 1 fan.

We will now have a closer look at the local building blocks of tropical manifolds, degree 1 fans. Our goal is to show that degree 1 fans satisfy some very special properties and, summarizing this, can be described in terms of matroids. This correspondence depends on the chosen basis B .

In the following we fix as ambient space \mathbf{R}^n with its standard basis e_1, \dots, e_n (i.e. we compute the degree of a fan with respect to the standard hyperplane

$H = V("x_1 + \dots + x_n + 0")$). Additionally, we set $e_0 = (-1, \dots, -1)$ and $E = \{0, 1, \dots, n\}$. For any subset $S \subset E$, we define the vector $e_S = \sum_{i \in S} e_i$. In particular, $e_E = 0$. To any chain of subsets $\mathcal{S} = (\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_l \subsetneq E)$, we assign the cone $\sigma_{\mathcal{S}} = \mathbf{R}_{\geq 0}e_{S_1} + \dots + \mathbf{R}_{\geq 0}e_{S_l}$. Here, $l = \dim \sigma_{\mathcal{S}}$ is called the length of \mathcal{S} . The collection of $\sigma_{\mathcal{S}}$ for all possible chains of subsets of E forms a unimodular polyhedral fan covering \mathbf{R}^n . It is called the *fine subdivision of \mathbf{R}^n* and denoted by $\text{FS}(\mathbf{R}^n)$. The first theorem we want to prove is that all degree 1 fans are supported on this subdivision (i.e. can be represented as a union of cones of the form $\sigma_{\mathcal{S}}$).

Theorem 7.4.4

Let $X \subset \mathbf{R}^n$ be a degree 1 fan of dimension k . Then X is supported on $\text{FS}(\mathbf{R}^n)^{(k)}$.

Proof. Note that the subdivision $\text{FS}(\mathbf{R}^n)$ can be obtained from intersecting all hyperplane subdivisions given by the hyperplanes $x_i = x_j$. Hence it suffices to show that the linear span $\mathbf{R}\sigma$ of every facet σ of X is the intersection of $n - k$ such hyperplanes. Note that $\mathbf{R}\sigma$, considered as a tropical cycle, is itself of degree 1 by 5.2.4. Hence the following lemma finishes the proof. \square

Lemma 7.4.5

Let $V \subset \mathbf{R}^n$ be a linear space which, considered as a tropical cycle with weight 1, has degree 1. Then there exists a chain $\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_k \subsetneq E$ such that $V = \langle v_{S_1}, \dots, v_{S_k} \rangle$.

Proof. We want to use induction on k . The case $k = 0$ is trivial. Now let V be of dimension $k + 1$. Choose a hyperplane $x_i = x_j$ which does not contain V . Then the stable intersection $V' := V("x_i + x_j") \cdot V$ is again of degree 1. Indeed, in $H^{k-1} \cdot V("x_i + x_j") \cdot V$ we may always replace $V("x_i + x_j")$ by a translated copy of H such that all points in the intersection are contained in the interior of the facet of H contained in the hyperplane $x_i = x_j$. Hence $\deg(V') = \deg(H^k \cdot V) = 1$. So by induction hypothesis $V' = \langle e_{S_1}, \dots, e_{S_k} \rangle$ for a suitable chain \mathcal{S} . Choose an integer vector v that completes e_{S_1}, \dots, e_{S_k} to a lattice basis of V . Let a_i be projective coordinates for v . Adding a suitable linear combination of $e_{S_1}, \dots, e_{S_k}, e_E$, we may assume that $a_i \geq 0$ for all i and that for each $0 \leq j \leq k$ there exists an index $i \in S_{j+1} \setminus S_j$ such that $a_i = 0$. To simplify notation, we permute the coordinates such that $S_j = \{0, \dots, r_j\}$

with $0 \leq r_1 < \dots < r_k < r_{k+1} = n$ and such that $a_{r_j} = 0$ for all $1 \leq j \leq k+1$. Here is a schematic picture of the matrix obtained from this basis of V .

$$\begin{matrix} e_{S_1} \\ e_{S_2} \\ \vdots \\ e_E \\ v \end{matrix} \begin{pmatrix} 1 & \dots & 1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 1 & \dots & 1 & 1 & \dots & 1 & 0 & \dots & 0 \\ \vdots & & \vdots & \vdots & & \vdots & \ddots & & \\ 1 & \dots & 1 & 1 & \dots & 1 & 1 & \dots & 1 \\ * & \dots & 0 & * & \dots & 0 & * & \dots & 0 \end{pmatrix}$$

Our next claim is that v has only 0 and 1 entries. Namely, consider the entry a_i for $i \notin \{r_1, \dots, r_{k+1}\}$. Let I be the complement of $\{i, r_1, \dots, r_{k+1}\}$. Since $|I| = n - k - 1$, σ_I is a facet of H^{k+1} . Now $\deg(H^{k+1} \cdot V) = 1$ implies (after choosing a suitable translation) that if $\mathbf{R}\sigma_I$ and V intersect transversally, the corresponding lattice index should be 1. In the language of determinants this means that after deleting the columns of M labeled by indices in I the square matrix has determinant 0 or ± 1 . However, after the deletion the only non-zero entry in the v -row is a_i . Moreover, after developing this row, the matrix becomes triangular with 1's on the diagonal. Hence $a_i \in \{0, 1\}$. It remains to show that v has non-zero entries in exactly one of the segments $0 \leq r_1 < \dots < r_k < n$, say between r_j and r_{j+1} . This finishes the proof, as this allows us to replace v by e_S , where $S = S_j \cup \{i : a_i = 1\}$. So let us prove this. Assume by contradiction that we find $r_j < i_1 < r_{j+1} < i_2$ such that $a_{i_1} = a_{i_2} = 1$. We set

$$\begin{aligned} I &= E \setminus \{i_1, r_1, \dots, r_{k+1}\}, \\ J &= E \setminus \{i_1, i_2, r_1, \dots, \widehat{r_{j+1}}, \dots, r_{k+1}\}. \end{aligned}$$

Again, deleting the columns labeled by I resp. J , we see easily that the determinant is non-zero in both cases (note that I is exactly of the form as above). Hence $\mathbf{R}\sigma_I$ and $\mathbf{R}\sigma_J$ both intersect V transversely. We now move H^{k+1} such that the translated vertex p has coordinates satisfying

$$p_i = \begin{cases} 1 & i = i_1, \\ 0 & i \in \{i_2, r_1, \dots, r_{k+1}\}, \\ < 0 & \text{otherwise.} \end{cases}$$

Note that $p_i = a_i$ for all $i \notin I$ and $p_i < a_i$ for all $i \in I$, hence $v - p \in \sigma_I$ or equivalently $v \in p + \sigma_I$. One also checks easily that $e_{S_{j+1}} - e_{S_j} - p$ has non-negative entries with zeros labeled by $\{i_1, i_2, r_1, \dots, \widehat{r_{j+1}}, \dots, r_{k+1}\}$, hence

$e_{S_{j+1}} - e_{S_j} \in p + \sigma_j$. It follows that both v and $e_{S_{j+1}} - e_{S_j}$ occur in $(H^{k+1} + p) \cdot V$ and therefore $\deg(V) \geq 2$, a contradiction. \square

Up to now we know that every degree 1 fan is a union of cones $\sigma_{\mathcal{S}}$ for some chains \mathcal{S} , with trivial weights 1. Which chains \mathcal{S} can occur together? To answer this question, let us first give a criterion describing which translations of H^k intersects a given cone $\sigma_{\mathcal{S}}$.

Lemma 7.4.6

Let $p \in \mathbb{R}^n$ be a point with pairwise distinct projective coordinates p_i . Let O be the induced complete order on E such that $p_i > p_j \Leftrightarrow i > j$. Let $\mathcal{S} = (\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_k \subsetneq E)$ be a chain and set $m_j := \max_O(S_{j+1} \setminus S_j)$, with respect to O . Then the following are equivalent:

- (a) $m_0 > m_1 > \dots > m_k$ forms a decreasing sequence with respect to O .
- (b) $(H^k + p) \cap \sigma_{\mathcal{S}} \neq \emptyset$.

In this case the intersection is transversal and we say O is compatible with \mathcal{S} .

Proof. It follows from the definitions that $(H^k + p) \cap \mathbf{R}\sigma_{\mathcal{S}}$ contains a single transversal intersection point x given by $x_i = p_{m_j}$ for all $i \in S_{j+1} \setminus S_j$. Furthermore, $x \in \sigma_{\mathcal{S}}$ if and only if the m_j form a decreasing sequence. Hence the claim follows. \square

Corollary 7.4.7

Let $X \subset \mathbb{R}^n$ be of degree 1 and let $\sigma_{\mathcal{F}}, \sigma_{\mathcal{G}} \subset X$ be two facets. If \mathcal{F} and \mathcal{G} admit a common compatible order, then $\mathcal{F} = \mathcal{G}$.

The next lemma describes the situation around a codimension 1 face.

Lemma 7.4.8

Let $X \subset \mathbb{R}^n$ be a degree 1 fan of dimension k . Let $\mathcal{S} = (\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_{k-1} \subsetneq E)$ be a chain of length $k - 1$ such that $\sigma_{\mathcal{S}} \subset X$. Let $\sigma_{\mathcal{S}_1}, \dots, \sigma_{\mathcal{S}_i}$ be the facets of X containing $\sigma_{\mathcal{S}}$. Then all chains \mathcal{S}_i are of the form $\emptyset \subsetneq \dots \subsetneq S_j \subsetneq T_i \subsetneq S_{j+1} \subsetneq \dots \subsetneq E$ for a fixed $0 \leq j \leq k - 1$ and the $T_i \setminus S_j$ form a partition of $S_{j+1} \setminus S_j$.

Proof. Each chain \mathcal{S}_i is obtained from \mathcal{S} by inserting some set T_i at some position j_i . The balancing condition around $\sigma_{\mathcal{S}}$ is equivalent to the condition that the $T_i \setminus S_{j_i}$ cover each element in $S_{j_i+1} \setminus S_{j_i}$ the same number of times. It follows that the subcollection of facets with some fixed value j_i satisfies the balancing condition independently. However, it follows from 5.2.4 that degree 1 fans are locally irreducible (i.e. $\text{Star}_X(p) = E_1 + E_2$ for two effective cycles E_1, E_2 implies $E_1 = 0$ or $E_2 = 0$). Hence only one value for j_i can occur, say j . It remains to show that the $T_i \setminus S_j$ cover each element of $S_{j+1} \setminus S_j$ exactly once. If an element is covered twice, the two corresponding chains admit a common order, which contradicts to 7.4.7. \square

After these preliminary works, we can now prove that X can be completely recovered from its rays, i.e. from its intersection with $\text{FS}(\mathbf{R}^n)^{(1)}$.

Theorem 7.4.9

Let $X \subset \mathbf{R}^n$ be a degree 1 fan and let $\mathcal{S} = (\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_l \subsetneq E)$ be a chain. Then $\sigma_{\mathcal{S}} \subset X$ if and only if $e_{S_j} \in X$ for all $1 \leq j \leq l$.

The proof uses a certain reordering construction which we choose to present separately in the following lemma.

Lemma 7.4.10

Let $\mathcal{S} = (\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_k \subsetneq E)$ be a chain such that $\sigma_{\mathcal{S}} \subset X$ is a facet of X . Fix $1 \leq l \leq k$ and let O be a complete order on S_l . Then there exists another chain $\mathcal{S}' = (\emptyset \subsetneq S'_1 \subsetneq \dots \subsetneq S'_k \subsetneq E)$ such that

- (a) $S'_j = S_j$ for all $j \geq l$,
- (b) O is compatible with $\emptyset \subsetneq S'_1 \subsetneq \dots \subsetneq S'_l = S_l$,
- (c) $\sigma_{\mathcal{S}'} \subset X$.

We call \mathcal{S}' the reordering of \mathcal{S} by O .

Proof. Let i be the maximal element in S_l with respect to O . Take $1 \leq j \leq l$ minimal such that $i \in S_j$. By 7.4.8 we can replace S_{j-1} by a (unique) set such $i \in S_{j-1}$. Proceeding in this way, we may assume $i \in S_1$. Now consider the maximal element i' in $S_l \setminus S_1$ and proceed analogously to obtain a new chain with $i' \in S_2$. Continuing in this way, we end up with a chain as described in the statement. \square

Proof of theorem 7.4.9. The “only if” direction is clear. Hence let us assume that \mathcal{S} is a chain of length l with $e_{S_j} \in X$ for all j . We prove $\sigma_{\mathcal{S}} \subset X$ by induction on l . By induction hypothesis we assume $\sigma_{\mathcal{S}} \subset X$ and show for given G with $S_l \subsetneq G$ and $e_G \in X$ the concatenated chain also describes a cone of X . The containments in X guarantee that there are maximal chains \mathcal{F} resp. \mathcal{G} completing \mathcal{S} resp. containing G . Choose a complete order of S_l compatible with \mathcal{S} and complete it to an order O on E such that elements in S_l are greater than elements in $G \setminus S_l$ which in turn are greater than elements in $E \setminus G$. Let \mathcal{F}' resp. \mathcal{G}' be the reorderings of \mathcal{F} resp. \mathcal{G} with respect to O . It follows from the construction that the lower part of \mathcal{F} remains unchanged, hence \mathcal{F}' is still a completion of \mathcal{S} . At the same time \mathcal{G}' still contains G by construction. However, \mathcal{F}' and \mathcal{G}' have common order O , hence by 7.4.7 they have to be equal. Therefore $\mathcal{F}' = \mathcal{G}'$ is a completion of $\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_l \subsetneq G \subsetneq E$, which proves the claim. \square

We will now show that there is a one-to-one correspondence between degree 1 fans of \mathbf{R}^n (with respect to the chosen basis) and loopfree matroids supported on the $E = \{0, \dots, n\}$. Recall that a matroid M can for example be given by a set $\mathfrak{F} \subset 2^E$ whose elements are called flats (or closed sets) of M and which satisfies the following three axioms.

- (A) E is a flat.
- (B) If F and G are flats, then so is $F \cap G$.
- (C) Let F be a flat and let G_1, \dots, G_l be all flats which cover F (i.e. $F \subsetneq G_i$ and $F \subset H \subset G_i$ implies $F = H$ or $F = G_i$). Then the sets $G_i \setminus F$ form a partition of $E \setminus F$.

The matroid M is called loopfree if the empty set is a flat as well. Here is how a degree 1 fan can be turned into a matroid.

Theorem 7.4.11

Let X be a degree 1 fan of dimension k . Let \mathfrak{F} be the set of subsets $S \subset \{0, \dots, n\}$ such that $e_S \in X$. Then \mathfrak{F} satisfies the axioms of the lattice of flats of a matroid. It is called the matroid associated to X and denoted by $M(X)$.

Proof. Axiom (A) follows from $e_E = 0 \in X$.

For Axiom (B) choose any order O of E such that the elements in $F \cap G$ are greater than the rest. Let $\sigma_{\mathcal{F}}$ and $\sigma_{\mathcal{G}}$ be facets of X such that \mathcal{F} contains F and \mathcal{G} contains G . Reorder \mathcal{F} by O and let F' be smallest set in the new chain such that $F \cap G \subset F'$. By construction of O , we have $F' \subset F$. Analogously, reordering \mathcal{G} we get a minimal set $G' \subset G$. The two new chains have common order O and hence are equal, and by minimality $F \cap G \subset F' = G' \subset F \cap G$. Hence $F \cap G = F' = G'$ is a flat.

For axiom (C), first note that by 7.4.9 it is clear that the minimal flats G_i containing F are exactly those appearing behind F in some chain \mathcal{F} such $\sigma_{\mathcal{F}}$ is a facet of X . Let $i \in E \setminus F$. Let \mathcal{F} be any such chain that contains F and let O be an order on E such that the largest elements are in F , then i , then the rest. Reordering \mathcal{F} by O we obtain a chain such that the set G behind F contains i . Hence the G_i cover $E \setminus F$. Assume there was a second such G' containing i . Then reordering a chain containing $F \subset G'$ by O produces the same chain as before, hence $G = G'$. This proves the claim. \square

By theorem 7.4.9, the inverse to $X \mapsto M(X)$ should look as follows.

Definition 7.4.12

Let M be a loopfree matroid on E . The associated *matroid fan* $B(M)$ consists of the collection of cones $\sigma_{\mathcal{F}}$, where $\mathcal{F} = (\emptyset \subsetneq F_1 \subsetneq \dots \subsetneq F_l \subsetneq E)$ is a chain of flats of M . In particular, $B(M)$ is a subfan of the fine subdivision of dimension $\text{rank}(M) - 1$ (which is the maximal length of chains of flats in M).

Theorem 7.4.13

Let M be a loopfree matroid. Then $B(M)$ is a degree 1 fan.

Proof. First note that axiom (C) of the matroid axioms guarantees that the statement of lemma 7.4.8 holds and hence $B(M)$ is balanced. So $B(M)$ forms indeed a tropical fan in \mathbf{R}^n and it remains to compute its degree. Our approach is as follows. Let $H \subset \mathbf{R}^n$ be the standard hyperplane. We will show that the stable intersection $H \cdot B(M)$ is the fan of a loopfree matroid again. Therefore, by induction it suffices to prove the claim for the single loopfree matroid of rank 1 on E . Its only flats are \emptyset and E itself. Therefore $B(M)$ is just the origin which is obviously of degree 1.

So, let us now prove that when $\text{rank}(M) > 1$, $H \cdot B(M)$ is a matroid fan again. Indeed, we will show $H \cdot B(M) = B(M')$ where M' is the matroid

obtained from M by removing all flats of rank $\text{rank}(M) - 1$ (which can be easily checked to be an admissible collection of flats again). To do so, we recall that H is given by the homogeneous linear tropical polynomial $l = \max\{x_0, \dots, x_n\}$, where the x_i are the homogeneous coordinates on \mathbf{TP}^n . Therefore $H \cdot B(M)$ can be computed as the divisor of l on $B(M)$. As l is obviously linear on all cones of the fine subdivision, $H \cdot B(M)$ is supported on the codimension one skeleton of $B(M)$. The weight of a codimension one cone $\sigma_{\mathcal{F}}$ of $B(M)$ in $H \cdot B(M)$ can be computed by

$$\omega(\sigma_{\mathcal{F}}) = \sum_{i=1}^k l(e_{H_i}) - l(e_G) - (k-1)l(e_F),$$

where k is the number of facets containing $\sigma_{\mathcal{F}}$, $F \subset G$ is where the rank gap appears in \mathcal{F} and H_i are the filled-in flats corresponding to a facet. To be precise, in this expression l should be replaced by an arbitrary dehomogenization of l . Alternatively, we may set

$$l(e_S) = \begin{cases} -1 & \text{if } S = E, \\ 0 & \text{otherwise.} \end{cases}$$

In any case, it is easy to compute

$$\omega(\sigma_{\mathcal{F}}) = \begin{cases} 1 & \text{if } G = E, \\ 0 & \text{otherwise.} \end{cases}$$

Thus $H \cdot B(M)$ consists of all the cones $\sigma_{\mathcal{F}}$ where \mathcal{F} is a chain of flats in M not containing a flat of rank $\text{rank}(M) - 1$. But this agrees with the definition of $B(M')$. \square

We can summarize this circle of ideas in the following statement.

Corollary 7.4.14

There is a one-to-one correspondence between degree 1 fans in \mathbf{R}^n and loopfree matroids on E given by

$$\begin{aligned} \{\text{degree 1 fans in } \mathbf{R}^n\} &\longleftrightarrow \{\text{loopfree matroids on } E\}, \\ X &\longmapsto M(X), \\ B(M) &\longleftarrow M. \end{aligned}$$

Proof. Theorems 7.4.11 and 7.4.13 show that both maps are well-defined. Theorem 7.4.9 proves that $B(M(X)) = X$. Finally, $M(B(M)) = M$ follows directly from the definitions. \square

Remark 7.4.15

Note that the map $X \mapsto M(X)$ has a completely analogous counterpart in classical algebraic geometry. An embedding of a linear space L in \mathbf{CP}^n is basically equivalent (up to translating \mathbf{TP}^n by the torus action) to a hyperplane arrangement on L . The hyperplanes are given by the intersection of L with the coordinate hyperplanes of \mathbf{CP}^n . In this situation, it is easy to check that the above definition of $M(L)$ is equivalent to the standard definition of the matroid of the hyperplane arrangement.

In fact, we may consider a slight generalization of our statement.

A degree 1 fan X in \mathbf{TP}^n is the closure of a degree 1 fan in an orbit \mathcal{O}_F , $F \subset E$. We think of $M(X)$, which a priori is a matroid on the ground set $E \setminus F$, as a matroid on E by summing with the rank zero matroid on F (hence F is the minimal flat of $M(X)$).

Correspondingly, let M be a (not necessarily loopfree) matroid M and let $F = \bar{\emptyset}$ be its minimal flat. Then $M \setminus F$ is a loopfree matroid and we define

$$B(M) := B(M \setminus F) \subset \mathcal{O}_F.$$

The philosophy here is that loops in the matroid correspond to fans contained in the corresponding coordinate hyperplane at infinity.

With these generalizations, the one-to-one correspondence can be extended to degree 1 fans in \mathbf{TP}^n and (arbitrary) matroids on E .

$$\{\text{degree 1 fans in } \mathbf{TP}^n\} \longleftrightarrow \{\text{matroids on } E\}$$

Theorem 7.4.16

Let $Y \subset \mathbf{R}^n$ be a matroid fan of dimension k and let $V \subset Y$ be a balanced polyhedral complex of dimension $k-1$ contained in Y . Then there exist tropical polynomials f, g such

$$V = \text{div}("f/g"|_Y).$$

Proof. Use FrancoisRau and Alexander Esterov: We can find $X' \subset \mathbf{R}^n$ of dimension $n-1$ such that $X = Y \cdot X'$ (stable intersection). Use the previous

theorem to find $X' = V("f/g")$ and use compatibility of stable intersection and Cartier divisors. \square

Theorem 7.4.17

In the above two theorems, the function $\varphi = "f/g"$ is unique up to adding an affine linear function.

Proof. In the \mathbf{R}^n -case, it follows from the proof that once we fix f on one connected component, it is uniquely determined everywhere else. For a general matroid fan, we can use the fact that it is a (multiple) modification of \mathbf{R}^k along matroid divisors and show that the property survives such a modification. So let us assume VV is a matroid modification of V and $f \cdot VV = 0$ for some rational function f . We want to show that f is affine linear. We prove first that f is the pull-back of a function g on X . Then the claim follows as the function g has to be affine linear by induction assumption.

We check the claim by proving that f is affine linear on every half-ray $VV \subset \{p\} \times \mathbf{R}$, where p is a generic point of *divisor*. Assume there is a point

We write $f = \delta^*g + h$, where g is a function on X and $h = 0$ away from *divisor* $\times \mathbf{R}$. It follows $\delta_*h \cdot VV \subset \text{divisor}$, and, as *divisor* is irreducible, $\delta_*h \cdot VV = a \cdot \text{divisor}$. This implies $g \cdot V = -a \cdot \text{divisor}$ and therefore $g = -a\varphi$, where φ is the modification function, by uniqueness in X . From this follows $f = -ay$. \square

Example 7.4.18

Example: Smooth curves, stable curves, $M_{0,n}, \overline{M}_{0,n}$

7.5 Rational functions, Cartier divisors and modifications on abstract spaces

As an example of constructions that generalize naturally to abstract tropical spaces, let us consider Cartier divisors and modifications of abstract spaces.

Definition 7.5.1

Let $X \subseteq \mathbf{T}^n$ be an open polyhedral set of pure dimension. A function $f :$

$X \rightarrow \mathbf{T}$ is called a *polynomial function* on X if it is the restriction $f = F|_X$ of a tropical polynomial $F \in \mathbf{T}[x_1, \dots, x_n]$. A function $f : X \rightarrow \mathbf{T}$ is called *piecewise \mathbf{Z} -linear* if f is continuous and if there exists a subdivision \mathcal{X} of X such that f restricts to an affine \mathbf{Z} -linear function (or is constant $-\infty$) on the interior of each cell. Assume that X is of sedentarity zero (i.e. $X = \overline{X \cap \mathbf{R}^n}$). A piecewise \mathbf{Z} -linear function $f : X \cap \mathbf{R}^n \rightarrow \mathbf{T}$ is also called a *rational function* on X , denoted by $f : X \dashrightarrow \mathbf{T}$. A general X admits a decomposition $X = \cup_I X_I$ such that X_I is an open polyhedral set of sedentarity zero in \mathbf{T}_I . A *rational function* $f : X \dashrightarrow \mathbf{T}$ is defined to be a tuple of rational functions $(f_I : X_I \dashrightarrow \mathbf{T})_I$. Since open sets of X are again open polyhedral sets, this obtains the presheaves \mathcal{O}'_X and \mathcal{R}'_X of polynomial resp. rational functions on X . The associated sheaves of locally polynomial resp. rational functions are denoted by \mathcal{O}_X resp. \mathcal{R}_X . Note that all these sheaves are sheaves of \mathbf{T} -algebras.

Remark 7.5.2

Both \mathcal{O}'_X and \mathcal{R}'_X are only presheaves in general (recall that we require subdivisions \mathcal{X} of X to be finite). An example is given by the tropical Laurent series

$$f = \sum_{i=0}^{\infty} -\binom{i}{2} x^i \in \mathcal{O}_{\mathbf{T}}(\mathbf{T}) \tag{7.6}$$

with infinitely many tropical zeros.

Exercise 7.5.3

Let $\varphi : X \rightarrow Y$ be a tropical morphism of open polyhedral sets. Show that local polynomiality resp. rationality is preserved under pull-back along φ . This leads to sheaf homomorphisms

$$\varphi^* : \mathcal{O}_Y \rightarrow \mathcal{O}_X \quad \text{and} \quad \varphi^* : \mathcal{R}_Y \rightarrow \mathcal{R}_X. \tag{7.7}$$

Find an example for the fact that the pull-back of a polynomial function is not necessarily polynomial again (only locally).

The exercise shows that the sheaves \mathcal{O}_X and \mathcal{R}_X can be glued along tropical isomorphisms to give sheaves on tropical spaces.

Definition 7.5.4

Let X be a tropical space with atlas (U_i, ψ_i, V_i) . The unique sheaf \mathcal{O}_X on

X such that $\mathcal{O}_{V_i} \cong \mathcal{O}_X|_{U_i}$ along ψ_i^* is called the *sheaf of locally polynomial functions* on X . Analogously, we define the *sheaf of locally rational functions* \mathcal{R}_X . Finally, a *Cartier divisor* on X is a section of the quotient sheaf $\text{CDiv} = \mathcal{R}_X^* / \mathcal{A}ff^*$. Hence it can be described by a collection of non-vanishing rational functions whose domains cover X and which differ by affine \mathbf{Z} -linear functions on the overlaps.

Let $X \subset \mathbf{T}^N$ be an open tropical cycle of pure dimension n and sedentarity zero and let $f : X \dashrightarrow \mathbf{T}$ be a rational function. We may choose a subdivision \mathcal{X} of X such that f is defined and affine \mathbf{Z} -linear on all cells of sedentarity zero. Let $\tau \in \mathcal{X}$ be a cell of codimension one. We define the *order of vanishing* $\omega_f(\tau)$ of f at τ in analogy to the constructions in by

$$\omega_f(\tau) = \sum_{i=1}^k \omega(\sigma_i) f_{\sigma_i}(v_i), \quad (7.8)$$

where $\sigma_1, \dots, \sigma_k$ denote the facets of \mathcal{X} containing τ and $f_{\sigma_i} = df|_{\sigma_i}$ denotes the \mathbf{Z} -linear part of f restricted to σ_i (note that the sedentarity of a facet is zero by assumption). The meaning of v_i depends on the sedentarity of τ . If $\text{sed}(\tau) = 0$, the vector v_i denotes a primitive generator of σ_i modulo τ such that $\sum_i \omega(\sigma_i)v_i = 0 \in \mathbf{Z}^n$. If $\text{Sed}(\tau) = I \neq \emptyset$, the projection $\pi : \mathbf{R}\sigma_i \rightarrow \mathbf{R}_I$ has a one-dimensional kernel K (since $\dim(\pi(\mathbf{R}\sigma_i)) = \dim(\tau) = \dim(\sigma_i) - 1$) and we define v_i to be the unique primitive integer vector in K with non-negative entries (intuitively, $-v_i$ is the primitive integer vector which points to τ (at infinity)).

Definition 7.5.5

Let $X \subset \mathbf{T}^N$ be an open tropical cycle of pure dimension n and sedentarity zero and let $f : X \dashrightarrow \mathbf{T}$ be a rational function. We define the *divisor (of zeros and poles)* of f , denoted by $\text{div}(f)$, to be the weighted open polyhedral set obtained (for a chosen subdivision \mathcal{X} as above) as the union of the codimension one faces τ of \mathcal{X} with weights $\omega_f(\tau)$.

Proposition 7.5.6

The divisor construction satisfies the following properties.

- (a) *The weighted open polyhedral set $\text{div}(f)$ is balanced, hence gives a open tropical cycle.*

(b) The open tropical cycle $\text{div}(f)$ does not depend on the chosen subdivision \mathcal{X} .

(c) Let $l \in \mathcal{A}ff^*(X)$ be a locally affine \mathbf{Z} -linear function on X . Then

$$\text{div}(f) = \text{div}(f + l). \quad (7.9)$$

The last property implies that we can extend the definition to Cartier divisors.

Definition 7.5.7

Let X be a tropical space and let $\varphi \in \text{CDiv}(X)$ be a Cartier divisor. Let $(U_i)_i$ be a open cover of X such that $\varphi|_{U_i} = [f_i]$ for some rational function $f_i \in \mathcal{R}^*(U_i)$. Then the (Weil) divisor $\text{div}(\varphi)$ of φ is the unique tropical subspace of X such that

$$\text{div}(\varphi)|_{U_i} = \text{div}(f_i). \quad (7.10)$$

As we recall from Chapter 5, the divisor construction

8 Tropical curves

In the previous chapters we presented the general concepts of tropical geometry. Instead of developing the general theory further, in the present chapter we want to bring these concepts into action in the easiest non-trivial concrete class of examples. Unsurprisingly, this class is formed by tropical curves.

8.1 Smooth tropical curves

Definition 8.1.1

A *tropical curve* C is a connected tropical space of dimension one. If the tropical space is smooth, we call C a *smooth tropical curve*. If C is a tropical manifold, we call it a *regular tropical curve*.

Let us unwind the definition of tropical manifold to see what that actually means in the curve case. The classification of local models is very easy. We set $L(n+1) = H^{n-1} \subset \mathbf{R}^n$, the fan in \mathbf{R}^n consisting of the $n+1$ rays in the directions $-e_1, \dots, -e_n, e_1 + \dots + e_n$, for suitable n . In the following, we call $L(n)$ the *n-valent line*.

Exercise 8.1.2

Let $F \subset \mathbf{R}^m$ be a one-dimensional tropical fan of degree 1. Show that F is isomorphic to $L(n)$ for suitable $n \geq 2$. Show that $\bar{F} \in \mathbf{T}^m$ is regular at infinity. You can either proceed directly and use the definition of projective degree, or classify matroids of rank 2.

In particular, we observe that smooth tropical curves are automatically regular at infinity and hence a tropical curve is regular if and only if it is smooth and of finite type.

Let C be a smooth tropical curve. The points of higher sedentarity $p \in C^{[1]}$ are called *infinite points* for short. All other points are called *finite*

points of C . Note that by regularity at infinity the only smooth local model at an infinite point is $-\infty$ in \mathbf{T} . It is convenient to extend our notation by setting $L(1) := \mathbf{T}$ (with "vertex" $-\infty$). Then it follows from the above exercise that at any point p the curve C admits a chart to $L(n)$, for suitable $n \in \mathbf{N}$, sending p to the vertex. The number $n =: \text{val}(p)$ is unique and called the *valence* of p . We have $\text{val}(p) = 1$ if and only if $\text{sed}(p) = 1$.

We will now show that smooth tropical curves are the same thing as metric graphs. The first ingredient to this statement is the uniqueness for the local model for each valence n as discussed above. The second ingredient is the fact that in dimension one an integral affine structure is the same thing as an (inner) metric. Let us be more precise.

Definition 8.1.3

A *metric graph (with open ends)* G is a connected one-dimensional finite CW-complex with some 1-valent vertices removed and equipped with a complete inner metric on $G \setminus \{1\text{-valent vertices}\}$. To avoid some tautologies, we exclude the case of a single edge with both endpoints removed (we may subdivide it by a 2-valent vertex, however). We distinguish three kinds of edges. The edges adjacent to an erased resp. remaining 1-valent vertex are called *open* resp. *closed ends*. All other edges are called *inner edges*.

Two metric graphs G, G' are called *isomorphic* if there exists a homeomorphism which is an isometry after removing 1-valent vertices.

Note that the open ends, closed ends, resp. inner edges e of G are isometric to $(-\infty, 0]$, $[-\infty, 0]$, resp. $[0, l_e]$ (for some $l_e \in \mathbf{R}_{>}$). It follows that given a fixed cell decomposition of G , the choice of a complete inner metric is equivalent to specifying a positive real number $l_e \in \mathbf{R}_{>}$ for each inner edge e of G .

Let $L(n+1)$ be the $n+1$ -valent line. We equip $L(n+1)$ with the unique complete inner metric such that all primitive generators $-e_1, \dots, -e_n, e_1 + \dots + e_n$ have length 1. Let $p \in G$ be the vertex of a metric graph and let U_p be the open neighbourhood consisting of p and the interior of all edges adjacent to p (in the presence of loop edges, we subdivide them into two pieces). Then there exists a unique distance-preserving map $U_p \rightarrow L(\text{val}(p))$ up to permuting the rays of $L(\text{val}(p))$. We include the case $\text{val}(p) = 1$ with map $U_p \cong [-\infty, 0] \hookrightarrow \mathbf{T}$. The collection of these maps forms a tropical atlas and give G the structure of a regular tropical curve, denoted (presently) by

$tc(G)$.

Exercise 8.1.4

Check that the collection of maps distance-preserving maps $U_p \rightarrow L(\text{val}(p))$ forms a tropical atlas which is of finite type and smooth.

Proposition 8.1.5

The map $G \mapsto tc(G)$ given by the above construction induces a bijection between the set of isomorphism classes of metric graphs G and the set of isomorphism classes of regular tropical curves (i.e., smooth curves of finite type).

We leave this as an exercise.

Exercise 8.1.6

Let C be a smooth tropical curve.

- (a) Show that C has the structure of a connected finite CW-complex with some 1-valent vertices removed.
- (b) Show that given the metric on $L(n)$, there is a unique inner metric on $C \setminus C^{[1]}$ such that the charts to $L(n)$ are distance-preserving.
- (c) Show that this metric is complete.
- (d) Let C' be another smooth tropical curve and let $\varphi : C \rightarrow C'$ be a homeomorphism of the underlying topological spaces. Show that φ is a tropical isomorphism if and only if it is distance-preserving (after removing the points at infinity).

Note that along this identification, the ends of the graph correspond to rays of the tropical curve, with sedentarity 1 endpoint in the case of closed ends. The inner edges correspond to bounded line segments. In the following, we will use the same letters and jump back and forth freely between the two points of view.

Definition 8.1.7

We define the *genus* of a smooth tropical curve to be $g(C) = b_1(C) = \dim H_1(C, \mathbf{Z})$, the first Betti number of the underlying graph.

Remark 8.1.8

We may drop the condition of finite type and extend the above one-to-one correspondence to any smooth tropical curve. In this case we should relax the completeness condition in the definition of metric graph to the property that the metric cannot be extended to a metric on G . It corresponds to allowing a finite length parameter $l_e \in (0, \infty]$ for open ends.

8.2 Divisors and linear systems

Definition 8.2.1

Let X be a tropical curve. A *divisor* on X is formal finite linear combination of points $D = \sum_{i=1}^n a_i p_i$, $a_i \in \mathbf{Z}$, $p_i \in X$. We say D is *effective* if all $a_i \geq 0$ (denoted by $D \geq 0$). The *degree* of D is defined by $\deg(D) := \sum_i a_i$. The *support* of D is the set $\text{supp}(D) := \{p_i : a_i \neq 0\} \subset C$. We say D is a divisor at infinity if $\text{supp}(D) \subset X^{[1]}$. Formal addition turns the set of all divisors into an abelian group graded by degree and called the *divisor group* of X , denoted $\text{Div}(X) = \bigoplus_{d \in \mathbf{Z}} \text{Div}^d(X)$.

Let $\text{Rat}(X)$ be the group of rational functions of finite type on X (with classical addition). Recall that to any such function f we can associate a divisor $\text{div}(f) \in \text{Div}(X)$. The coefficient of $\text{div}(f)$ at p is called the *order (of vanishing)* at p and denoted by $\text{ord}_p(f) \in \mathbf{Z}$. Moreover, the map $\text{div} : \text{Rat}(X) \rightarrow \text{Div}(X)$ is group homomorphism.

Definition 8.2.2

Two divisors D, D' on X are called *linearly equivalent*, denoted by $D \sim D'$, if $D - D' = \text{div}(f)$ for some $f \in \text{Rat}(X)$. The corresponding quotient group $\text{Div}(X)/\text{Im}(\text{div})$ is called the *divisor class group* of X . The divisors of the form $\text{div}(f)$, $f \in \text{Rat}(X)$ are called *principal divisors* of X .

Proposition 8.2.3

Let X be a compact tropical curve. Then $\deg(\text{div}(f)) = 0$ holds for all $f \in \text{Rat}(X)$.

Proof. Fix a subdivision of X without loop edges on which f is piecewise linear. As X is compact, each edge has two endpoints. The slope of f on an edge contributes to the coefficients of both endpoints with opposite

signs. Hence, when adding up the coefficients the contribution of each edge cancels and the degree vanishes. \square

Let us also note the following fact.

Proposition 8.2.4

Let X be a compact tropical curve and $f \in \text{Rat}(X)$. Then it holds

$$\text{div}(f) = 0 \implies f \equiv \text{const},$$

i.e., the kernel of div consists of constant functions only.

Proof. Let $M \subset X$ be the locus in C where f attains its minimum. If $M \subsetneq X$, the boundary ∂M is non-empty and we may choose some $p \in \partial M$. At this point, all outgoing slopes are positive and at least along one edge the slope must be strictly positive (otherwise p lies in the interior of M). This implies $\text{ord}_p(f) > 0$, which contradicts $\text{div}(f) = 0$. \square

Definition 8.2.5

Let D be a divisor on a tropical curve X . The set of effective divisors linearly equivalent to D

$$|D| := \{D' \in \text{Div}(X) : D \sim D' \geq 0\}$$

is called the *complete linear system* of D . We also use the notation

$$\Gamma(D) := \{f \in \text{Rat}(X) : \text{div}(f) + D \geq 0\} \cup \{-\infty\}.$$

Note that $\text{Rat}(X) \cup \{-\infty\}$ carries the structure of a semiring by setting

$$\begin{aligned} "f + g"(x) &:= \max(f(x), g(x)), \\ "fg"(x) &:= f(x) + g(x), \end{aligned}$$

for all $x \in X \setminus X^{[1]}$ and extending to X by continuity. Moreover, $\text{Rat}(X)$ carries a \mathbf{T} -algebra structure by identifying \mathbf{T} with the constant functions in $\text{Rat}(X)$. By Proposition 8.2.4, if X is compact, $|D|$ is just the projectivization of $\Gamma(D)$ with respect to this multiplication by \mathbf{T} , i.e.,

$$|D| = \mathbf{P}(\Gamma(D)) := \Gamma(D) \setminus \{-\infty\} / \mathbf{R}.$$

Proposition 8.2.6

Let D be an effective divisor on a tropical curve X . Then the following holds.

(a) The set $\Gamma(D)$ is a \mathbf{T} -subalgebra of $\text{Rat}(X)$.

(b) The set $\Gamma(D)$ is finitely generated as a \mathbf{T} -module.

Proof. Cite LINEAR SYSTEMS ON TROPICAL CURVES, CHRISTIAN HAASE, GREGG MUSIKER, AND JOSEPHINE YU □

8.3 Line bundles on curves

Definition 8.3.1

Let X be a tropical curve. A *line bundle (of finite type)* on X consists of a tropical space L together with a tropical morphism $\pi : L \rightarrow X$ such that there exists a finite collection of tuples $(U_i, \psi_i)_i$, where

- the U_i form an open covering of X ,
- the maps $\psi_i : \pi^{-1}(U_i) \rightarrow U_i \times \mathbf{T}$ are tropical isomorphisms such that the diagram

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\psi_i} & U_i \times \mathbf{T} \\ & \searrow \pi & \downarrow \\ & & U_i \end{array}$$

commutes.

The collection (U_i, ψ_i) is called a *local trivialization* for L . Let $\pi' : L' \rightarrow X$ be a second line bundle on X . An *isomorphism of line bundles* is a tropical isomorphism $\Psi : L \rightarrow L'$ which commutes with projections, i.e., $\pi' \circ \Psi = \pi$. The set of isomorphism classes of line bundles on X is called the *Picard group* of X (we will introduce the group structure in a moment) and denoted by $\text{Pic}(X)$.

Remark 8.3.2

Note that the tropical definition is slightly simpler than in the classical case. In the classical case, in way or the other, the vector space structure on the fibers has to be included in the definition and the trivialization is required to respect this structure. (In practice, this is often done by fixing a trivialization and asking the transition maps to be linear on fibers.) As we

will see in a moment, in tropical geometry we get this for free. In some sense, this boils down to the fact that the classical automorphism group $\text{Aut}(\mathbf{C})$ contains many non- \mathbf{C} -linear maps, whereas a tropical automorphism $\phi \in \text{Aut}(\mathbf{T}) \cong \mathbf{R}$ is automatically \mathbf{T} -linear. Similar remarks apply to the definition of isomorphisms of line bundles.

Remark 8.3.3

A line bundle L is smooth if and only if X is smooth. It is a tropical manifold if and only if X is regular.

Let (U_i, ψ_i) be a *local trivialization* of a line bundle L . Consider the induced automorphisms on the overlaps

$$\tilde{\psi}_{ij} := \psi_j \circ \psi_i^{-1} : (U_i \cap U_j) \times \mathbf{T} \rightarrow (U_i \cap U_j) \times \mathbf{T},$$

which acts trivially in the first component. It follows that $\tilde{\psi}_{ij}$ induces automorphisms on $\{x\} \times \mathbf{T}$ for each $x \in U_i \cap U_j$. Since $\text{Aut}(\mathbf{T}) \cong \mathbf{R}$, this gives rise to functions $\psi_{ij} : U_i \cap U_j \rightarrow \mathbf{R}$ such that $\tilde{\psi}_{ij}(x, y) = (x, \psi_{ij}(x)y)$. Moreover, the functions ψ_{ij} are tropical morphisms (write them as $\psi_{ij}(x) = \pi_2(\tilde{\psi}_{ij}(x, 0))$, where π_2 is the projection to the second component) and hence are sections of $\mathcal{A}ff_X^*(U_i \cap U_j)$. They are called the *transition functions* associated to the trivialization (U_i, ψ_i) . The identity $\tilde{\psi}_{jk} \circ \tilde{\psi}_{ij} = \tilde{\psi}_{ik}$ translates to the *cocycle identity*

$$\psi_{ij} - \psi_{ik} + \psi_{jk} \equiv 0 \quad \text{on } U_i \cap U_j \cap U_k.$$

The other way around, any collection of functions $\psi_{ij} \in \mathcal{A}ff_X^*(U_i \cap U_j)$ satisfying the cocycle identity occur as transition functions of a line bundle, and the following statement summarizes this in terms of the first Čech cohomology group $H^1(X, \mathcal{A}ff_X^*)$.

Proposition 8.3.4

Let L be a line bundle on X . The transition functions of any trivialization of L induce a cocycle class $[(\psi_{ij})_{ij}] \in H^1(X, \mathcal{A}ff_X^*)$. The map

$$\begin{aligned} \text{Pic}(X) &\rightarrow H^1(X, \mathcal{A}ff_X^*), \\ L &\mapsto [(\psi_{ij})_{ij}] \end{aligned}$$

is well-defined and bijective. In particular, it induces a group structure on $\text{Pic}(X)$.

Exercise 8.3.5

Prove the proposition (for example, by copying the proof of the classical case). Recall that $\mathcal{A}ff_X^*$ is a constructible sheaf and hence Čech cohomology group can be computed on a suitably fine open covering of X .

Exercise 8.3.6

Let us fix a subdivision of X without loops and double edges and let U_x denote the open neighbourhood of the vertex x consisting of x and the interior of the adjacent edges. Show that the U_v form an admissible open covering for $\mathcal{A}ff_X^*$ and compute the cohomology groups with respect to this covering.

We denote by \bar{L} the \mathbf{TP}^1 -bundle obtained by extending the gluing maps from $(U_i \cap U_j) \times \mathbf{T}$ to $(U_i \cap U_j) \times \mathbf{TP}^1$.

Definition 8.3.7

Let L be a line bundle on X and $(U_i, \psi_i)_i$ be a local trivialization of L . A map $s : X \rightarrow \bar{L}$ is called a *rational section (of finite type)* of L if $\pi \circ s = \text{id}_X$ and $s_i := \pi_2 \circ \psi_i \circ s|_{U_i} \in \text{Rat}(U_i)$ for all i . We call the s_i the *local parts* of s . We define the *order* of s at $x \in U_i$ by $\text{ord}_x(s) := \text{ord}_x(s_i)$ and call $\text{div}(s) := \sum_x \text{ord}_x(s)x \in \text{Div}(X)$ the *divisor* of s . The \mathbf{T} -semifield of all rational sections of L is denoted by $\text{Rat}(L)$.

Exercise 8.3.8

Check that if s is a rational section for some local trivialization of X , then it is so for any trivialization. Check that $\text{ord}_x(s)$ does not depend on the choices made and hence $\text{div}(s)$ is well-defined.

Remark 8.3.9

Every line bundle L has a canonical trivial section which is constant $-\infty$ in every local trivialization. By abuse of notation, let $\{-\infty\} \subset L$ denote the image of this section. Note that any non-trivial rational section s can be restricted to a map $X \setminus X^{[1]} \rightarrow L \setminus \{-\infty\}$, and this restriction determines s uniquely. By abuse of notation, we will often write $s : X \rightarrow L$ in the following.

Proposition 8.3.10

Let L be a line bundle on X and let ψ_{ij} be the transition functions of a given trivialization $(U_i)_i$. Then the following holds.

- (a) A collection of rational functions $s_i \in \text{Rat}(U_i)$ occurs as the local parts of a rational section of L if and only if $s_i - s_j = \psi_{ij}$.
- (b) Let s, s' be two rational sections of L with local parts s_i resp. s'_i . Then there exists a unique global rational function $f \in \text{Rat}(X)$ such that $f|_{U_i} = s_i - s'_i$. We use the notation $f = "s/s'" = s - s'$.
- (c) There exists a rational section s of L .

Proof. Part (a) is clear. For part (b), set $f_i := s_i - s'_i$. Then part (a) implies $f_i - f_j \equiv 0$ on $U_i \cap U_j$, hence we can glue the rational functions f_i to a global rational function f . We leave (c) as an exercise. \square

Exercise 8.3.11

Show that every line bundle L on a tropical curve admits a rational section. You can proceed as follows:

- (a) Reduce to the case where C has no points of higher sedentarity.
- (b) Show that every line bundle admits a trivialization (U_x, ψ_x) of the form given in Exercise 8.3.6 for a suitable graph structure.
- (c) For each U_x and $c \in \mathbf{R}$, set $s_x := -c \text{ dist}(x, \cdot) : U_x \rightarrow \mathbf{R}$. Show that for sufficiently large c , both functions s_x and $s_y + \psi_{xy}$ dominate the other somewhere on $U_x \cap U_y$ (if non-empty).
- (d) Take an appropriate "maximum" of the s_x to obtain a section of L .

The previous proposition should bring to mind the very closely related concept of Cartier divisors. Recall that a *Cartier divisor (of finite type)* on X is a global section of the (pre-)sheaf quotient $\text{Rat} / \mathcal{A}ff^*$. In other words, it can be described by a finite collection $(U_i, f_i)_i$ where the U_i form an open covering of X and $f_i \in \text{Rat}(U_i)$ subject to the condition $f_i - f_j \in \mathcal{A}ff_x^*(U_i \cap U_j)$ for all i, j . Two such collections are considered equal if the merge of the two still satisfies this condition. Any rational function $f \in \text{Rat}(X)$ gives rise to a Cartier divisor given by the single tuple (X, f) . Such Cartier divisors are called *linearly equivalent to zero*. They form a subgroup of the group

of Cartier divisors $\text{CDiv}(X)$. Two Cartier divisors D, D' are called *linearly equivalent* if their difference is linearly equivalent to zero, denoted by $D \sim D'$.

Let $(U_i, f_i)_i$ be the representative of a Cartier divisor on X . Clearly, the functions $f_{ij} = f_i - f_j \in \mathcal{A}ff_X^*(U_i \cap U_j)$ satisfy the cocycle identity. Hence by 8.3.4 they form the transition functions of a line bundle L and by 8.3.10 the f_i define a rational section $s \in \text{Rat}(L)$. Vice versa, let (L, s) be a tuple of a line bundle L and a section $s \in \text{Rat}(L)$. Then for a given trivialization of the local parts $(U_i, s_i)_i$ of s determine a Cartier divisor (again by 8.3.10). A second tuple (L', s') is called *isomorphic* to (L, s) if there exists an isomorphism between the line bundles which identifies the sections. We can now generalize 8.3.4 in the following way.

Proposition 8.3.12

The map

$$\begin{aligned} \{(L, s) : L \text{ line bundle}, s \in \text{Rat}(L)\} / \text{isom.} &\rightarrow \text{CDiv}(X), \\ (L, s) &\mapsto (U_i, s_i)_i, \end{aligned}$$

is a bijection. In particular, it induces a group structure on the source. The induced map

$$\begin{aligned} \text{Pic}(X) &\rightarrow \text{CDiv}(X) / \sim, \\ L &\mapsto [(U_i, s_i)_i] \text{ for some } s \in \text{Rat}(X), \end{aligned}$$

(which is well-defined by 8.3.10) is a bijection. The induced group structure on $\text{Pic}(X)$ coincides with the group structure obtained from 8.3.4.

Exercise 8.3.13

Check the details of the proof.

In total, we get the following commutative diagram of groups.

$$\begin{array}{ccccc} & & \{(L, s)\} / \text{isom.} & \xrightarrow{\cong} & \text{CDiv}(X) \\ & & \downarrow \pi & & \downarrow \\ H^1(X, \mathcal{A}ff_X^*) & \xleftarrow{\cong} & \text{Pic}(X) & \xrightarrow{\cong} & \text{CDiv}(X) / \sim \end{array}$$

relation between line bundles and locally free sheaves of rank 1
 Let us finally have a look at the map

$$\begin{aligned} \{(L, s)\} / \text{isom.} &\rightarrow \text{Div}(X), \\ (L, s) &\mapsto \text{div}(s), \end{aligned}$$

induced by taking divisors of rational sections of line bundles. It is obviously a group homomorphism, but in general neither injective nor surjective.

However, in the smooth case we have the following statement.

Proposition 8.3.14

Let X be a smooth tropical curve. Then the map

$$\begin{aligned} \{(L, s)\} / \text{isom.} &\rightarrow \text{Div}(X), \\ (L, s) &\mapsto \text{div}(s), \end{aligned}$$

and the induced map

$$\begin{aligned} \text{Pic}(X) &\rightarrow \text{Div}(X) / \sim, \\ L &\mapsto \text{div}(s) \text{ for some } s \in \text{Rat}(X), \end{aligned}$$

are group isomorphisms.

We first prove the following lemma covering the local situation.

Lemma 8.3.15

Let $L = L(n+1)$ be the n -valent line and let $R(L)$ be the set of rational functions $f : L \rightarrow \mathbf{R}$ which are affine linear on each ray of $L(n)$. Then the map $\text{ord}_0 : R(L) \rightarrow \mathbf{Z}$ fits into the following exact sequence.

$$0 \rightarrow \mathcal{A}ff^*(L) \rightarrow R(L) \rightarrow \mathbf{Z} \rightarrow 0$$

Proof. The sequence of the statement, when restricted to the case $f(0) = 0$, is dual to the sequence

$$0 \rightarrow \mathbf{Z} \rightarrow \mathbf{Z}^{n+1} \rightarrow \mathbf{Z}^n \rightarrow 0$$

given by $1 \mapsto (1, \dots, 1)$ and $(\omega_0, \dots, \omega_n) \mapsto \sum_i \omega_i(-e_i)$. □

Proof of 8.3.14. By 8.3.10 it suffices to establish the statement for the first map. We prove injectivity first. Let L be a line bundle and $s \in \text{Rat}(L)$ such that $\text{div}(s) = 0$. Let s_i be the local parts of s with respect to some trivialization of L . By 8.3.15 we get $s_i \in \mathcal{A}ff^*(U_i)$. By 8.3.4 this implies the existence of an isomorphism $L \cong X \times \mathbf{T}$ which identifies s with the constant zero function. This proves injectivity. Let now $D \in \text{Div}(X)$. Choose a subdivision of X without loops whose vertices contain the support of D . Let U_ν be the open set containing the vertex ν and the interior of the adjacent edges. By 8.3.15 there exists a collection of functions $(U_\nu, f_\nu)_\nu$ such that $\text{ord}_\nu(f_\nu)$ equals the coefficient of ν in D and such that each f_ν is affine linear on the interior of edges. Hence, the collection determines a Cartier divisor and by 8.3.12 also a tuple (L, s) with $\text{div}(s) = D$. This proves surjectivity. \square

8.4 Riemann-Roch theorem

Let us now turn to the Riemann-Roch theorem for tropical curves. In this section, all tropical curves are compact and smooth.

Definition 8.4.1

Let X be a compact smooth tropical curve and let $L \subseteq |D|$ be a subset of a complete linear system. We define the *rank of L* to be the maximal integer r such that for all effective divisors D' of degree r there exists a divisor $D \in L$ such that $D - D' \geq 0$. The rank of L is denoted by $\text{rk}(L)$. We set $\text{rk}(L) = -1$ if $L = \emptyset$. We write $\text{rk}(D) = \text{rk}(|D|)$ for short.

Obviously, we have $\text{rk}(D) = -1$ if $\text{deg}(D) \leq 0$. Moreover, if $\text{deg}(D) = 0$, then $\text{rk}(D) = 0$ if and only if $D \sim 0$ — otherwise $\text{rk}(D) = -1$.

Let us make a quick comparison to the classical situation. Given a complete linear system and a given point on the curve which is not a base point, the set of divisors "containing" the point forms a hyperplane. Hence our definition requires that the intersection of L with $\text{rk}(L)$ such hyperplanes is non-empty. It easy to check that in the classical situation we just get back the dimension of L .

Remark 8.4.2

Note that there are other non-equivalent definitions for the rank of a tropical linear system $|D|$, for example

- the minimal number of generators of the \mathbf{T} -module $\Gamma(D)$ minus one,
- the (maximal) dimension of the CW-coplex $|D|$.

To convince yourself that these numbers are in general different, take a hyperelliptic genus 3 curve with two bridges and let D be a hyperelliptic divisor. Then these numbers are 3 resp. 2, whereas $\text{rk}(D) = 1$. The above definition was introduced in and is distinguished in particular by the fact that Riemann-Roch holds.

We will present the Riemann-Roch theorem in the form using Serre duality, i.e. with "error" term $\text{rk}(K - D)$.

Definition 8.4.3

Let X be a smooth tropical curve. The *canonical divisor class* $[K] \in \text{Pic}(X)$ is given by

$$K = \sum_{p \in C} (\text{val}(p) - 2)p.$$

Remark 8.4.4

The motivation behind this definition is as follows. Let \mathcal{X} be a classical nodal curve whose whose irreducible components are \mathbf{P}^1 and whose dual graph is equal to X . Then the dualizing sheaf of \mathcal{X} (which is, in some appropriate sense, the limit of the canonical sheaf of a smoothing of \mathcal{X}) restricted to the irreducible component \mathcal{X}_p , $p \in X$ is isomorphic to $\mathcal{O}_{\mathbf{P}^1}(\text{val}(p) - 2)$. Hence K just reflects the degree of the dualizing sheaf restricted to irreducible components. Note however, that K is not "realizable" in general.

We can now state the theorem.

Theorem 8.4.5

Let X be a compact smooth tropical curve of genus g . Let D be a divisor of degree d . Then the following equality holds.

$$\text{rk}(D) - \text{rk}(K - D) = d - g + 1$$

Before proving the statement, we need a few basic facts concerning the rank of a linear system. As some of these facts can be proven most conveniently using the Abel-Jacobi map, we postpone the proof of them to the next chapter.

Exercise 8.4.6

Show that

$$\mathrm{rk}(D) + \mathrm{rk}(D') \leq \mathrm{rk}(D + D') \leq \mathrm{rk}(D) + \deg(D')$$

for all divisors $D, D' \in \mathrm{Div}(X)$. In particular, $-1 \leq \mathrm{rk}(D) \leq \deg(D)$.

Proposition 8.4.7

Let D be a divisor of degree d on a compact smooth tropical curve X of genus g .

- (a) If $d \geq g$, then $|D| \neq \emptyset$.
- (b) The rank of D is bounded from below by $\mathrm{rk}(D) \geq d - g$.
- (c) If $d \geq 2g - 1$, then $\mathrm{rk}(D) = d - g$.

Proof. Part (a) follows from the Jacobi inversion theorem (cf. Theorem 9.3.12 and Remark 9.3.15 (a)). Part (b) follows immediately from (a). For part (c), assume conversely that $d - \mathrm{rk}(D) < g$. Choose a divisor D' of degree $d - \mathrm{rk}(D)$ with $|D'| = \emptyset$ (such a choice exists by the Abel-Jacobi theorem, cf. (cf. Theorem 9.3.6 and Remark 9.3.15 (c)). We have $\deg(D - D') \geq g$ and hence by part (a) there exists an effective divisor $E \in |D - D'|$. But $\deg(E) = \deg(D - D') = \mathrm{rk}(D)$ and hence $|D'| = |D - E| \neq \emptyset$, a contradiction. \square

In the following, it will be useful to consider some rational functions based on the distance function on X . Let us denote the distance of two points in the metric space $X^{[0]}$ by $\mathrm{dist}(x, y) \in \mathbf{R}$. Let $A \subseteq X^{[0]}$ be a non-empty subset. Then we define $\mathrm{dist}(A, x) := \inf_{a \in A} \mathrm{dist}(a, x)$.

Let us now fix a point $p \in X^{[0]}$. We introduce a total preorder on the set of effective divisors on $X \setminus \{p\}$ as follows. We write such a divisor as $D = p_1 + \dots + p_n$ with $\mathrm{dist}(p, p_1) \leq \dots \leq \mathrm{dist}(p, p_n)$ (where we set $\mathrm{dist}(p, p_i) = \infty$ if $p_i \in X^{[1]}$). The increasing sequence of distances is denoted by $\mathrm{dist}(p, D)$. Given two such divisors D, D' , we declare D lower than D' if $\deg(D) < \deg(D')$ or if $\deg(D) = \deg(D')$ and $\mathrm{dist}(p, D)$ is smaller than $\mathrm{dist}(p, D')$ with respect to the lexicographic order. We call this preorder the p -distance.

Definition 8.4.8

A divisor D on X is called p -reduced if its restriction to $X \setminus \{p\}$ is effective and minimal with respect to the p -order among such restrictions in the equivalence class of D .

Exercise 8.4.9

Show that if D is p -reduced and $|D| \neq \emptyset$, then $D \leq 0$.

Here is another characterization of p -reduced divisors using subgraphs $A \subseteq X$ (i.e., closed subsets obtained as the union of cells for some subdivision of X).

Proposition 8.4.10

A divisor D is p -reduced if and only if it is effective on $X \setminus \{p\}$ and for any subgraph $A \subseteq X$ with $p \notin A$ there exists a point $q \neq p$ such that the coefficient a_q of D at q satisfies

$$a_q + \text{val}_A(q) < \text{val}_X(q). \quad (8.1)$$

In intuitive words: Any A is leaking into X and cannot be sealed off using D .

Proof. The characterization based on (8.1) will be referred to as leaking property in the following. Let us first show that p -reducedness implies the leaking property. First note that a p -reduced divisor does not contain infinite points. This is true since an infinite point is linearly equivalent to any other point on the corresponding leaf edge and hence can be replaced by any finite point, thereby decreasing the p -distance. Hence the leaking property holds for any $A \subseteq X^{[1]}$ and it remains to check the case $A^{[0]} = A \cap X^{[0]} \neq \emptyset$. In this case the function $f_A : X^{[0]} \rightarrow \mathbf{R}, x \mapsto \text{dist}(A^{[0]}, x)$ is a well-defined rational function on X . We set $m_A = \max\{f_A, -\epsilon\}$ where $\epsilon > 0$ is chosen such that the ϵ -ball around A does not contain p nor any vertices outside of A . For all $q \in A^{[0]}$ we have $\text{ord}_{m_A}(q) = \text{ord}_{f_A}(q) = \text{val}_A(q) - \text{val}_X(q)$. More precisely, we get

$$\text{div}(m_A) = \sum_{\text{dist}(A, x) = \epsilon} x + \sum_{q \in A^{[0]}} (\text{val}_A(q) - \text{val}_X(q))q.$$

It follows that if the condition $a_q + \text{val}_A(q) \geq \text{val}_X(q)$ holds for all $q \in A$, then $D' := D + \text{div}(m_A)$ is effective. We can see that D' is strictly smaller with

respect to p -distance by looking at a minimizing path from p to A . For points x on this path we have $\text{dist}(A, x) + \text{dist}(x, p) = \text{dist}(A, p)$. Hence D' contains a (new) point with $\text{dist}(x, p) = \text{dist}(A, p) - \epsilon$, and $(D' - D)|_{B_p(\text{dist}(A, p))} > 0$. Hence D' is strictly smaller than D , a contradiction.

Let us now assume D satisfies the leaking property for all A . First of all, note that D does not contain an infinite point q since $A = \{q\}$ implies $a_q < 1$. Let D' be any other divisor equivalent to D and effective away from p and assume that D' . Let f be rational function with $\text{div}(f) = D' - D$. Let A be the maximality domain of f (note that f is bounded since D does not contain infinite vertices). If $p \notin A$, then by the leaking property there exists a point $q \in A$ with $a_q < \text{val}_x(q) - \text{val}_x(q) \leq -\text{ord}_f(q)$. Hence $D + \text{div}(f) = D'$ is not effective at q , a contradiction. It follows $p \in A$. Now, choose a point of minimal distance to p in $\text{Supp}(\text{div}(f))$ (including the case $q = p$). A shortest path from p to q is disjoint from $\text{Supp}(\text{div}(f))$ except for q . Since $A \cap \text{Supp}(\text{div}(f)) = \partial A$ and $p \in A$, this implies $q \in \partial A$ and hence $\text{ord}_f(q) < 0$. It follows that the coefficient of q in D is greater than in D' showing that D has smaller p -distance. (If $q = p$, we use $\text{deg}(D) = \text{deg}(D')$ and hence the restriction of D to $X \setminus \{p\}$ has lower degree than the restriction of D' .) \square

Proposition 8.4.11

Fix $p \in X^{[0]}$. Then every divisor class $[D] \in \text{Pic}(X)$ has a unique p -reduced representative.

Proof. First we show existence. It follows from Proposition 8.4.7 that there exists $m \in \mathbf{Z}$ such that $|D + mp| \neq \emptyset$ (choose $m \geq g - \text{deg}(D)$). Pick the minimal such $m \in \mathbf{Z}$. Since $|D + mp|$ is compact, we can choose $E \in |D + mp|$ with minimal p -distance. For any other $D' = -m'p + E' \sim D$, $E' \geq 0$, we have $m' \leq m$, by minimality of m , and hence $\text{deg}(E') \geq \text{deg}(E)$. This together with the minimality of E implies that $D'' = -mp + E$ is p -reduced.

For uniqueness, let $D \sim D'$ be two p -reduced divisors in the same class. Let f be a rational function such that $\text{div}(f) = D' - D$. The leaking property for D implies that p is contained in the maximality domain of f (cf., the proof of the previous proposition). Reversing the roles of D and D' , the same argument shows that p is contained in the minimality domain of f . Hence f is constant and $D = D'$. \square

We continue by defining an important class of divisors of “intermediate” degree $g - 1$.

Definition 8.4.12

An *acyclic orientation* \mathcal{O} of X is a choice of directions of the edges of a subdivision of X which does not admit oriented cycles. We denote by $\text{val}_+(x)$ the number of outgoing edges at x . The divisor

$$K_+ := \sum_{x \in X} (\text{val}_+(x) - 1)x$$

is called the *moderator* with respect to \mathcal{O} .

Taking the opposite orientation $-\mathcal{O}$ gives rise to the moderator $K_- = \sum_{x \in X} (\text{val}_-(x) - 1)x$ counting the incoming edges of \mathcal{O} .

Exercise 8.4.13

Check $\deg(K_+) = \deg(K_-) = g - 1$ and $K_+ + K_- = K$.

Note that the points with negative coefficient in K_+ are exactly the sinks of \mathcal{O} . Since \mathcal{O} is acyclic, there always exists at least one sink, hence K_+ is never effective. This is even true for the full equivalence class of K_+ .

Lemma 8.4.14

For any moderator K_+ it holds $|K_+| = \emptyset$.

Proof. Let f be a rational function on X with $E = K_+ + \text{div}(f) \geq 0$. First note that f cannot have poles at infinity since the coefficient of infinite points in K_+ is at most 0. Let A be the maximality domain of f . The moderator K_+ is produced by some acyclic orientation \mathcal{O} of X . Choose a sink q of the restriction of \mathcal{O} to A (i.e., a point in A such that all adjacent edges in A are incoming). It follows $\text{val}_+(q) \leq \text{val}_X(q) - \text{val}_A(q) \leq -\text{ord}_f(q)$, which shows that E is not effective at q . \square

The following lemma shows in which sense the divisors K_+ "moderate" between the effective and non-effective divisor classes in $\text{Pic}(X)$.

Lemma 8.4.15

For any divisor $D \in \text{Div}(X)$ exactly one of the following two options holds: Either $|D| \neq \emptyset$ or there exists a moderator K_+ such that $|K_+ - D| \neq \emptyset$.

Proof. By lemma 8.4.14 at most one of the two options holds true. It remains to show that if $|D| = \emptyset$, there exists a moderator K_+ with $|K_+ - D| \neq \emptyset$.

Without loss of generality we may assume that D is p -reduced with respect to some finite point $p \in X$. Fix a subdivision of X without loops such that p and the points in $\text{Supp}(D)$ are vertices. We will define an acyclic orientation on X by ordering the vertices and then orienting the edges from higher to lower vertices. The order of the vertices is given inductively as follows. The lowest vertex is $v_1 = p$. Now assume v_1, \dots, v_k have already been chosen. We set A to be the full subgraph of X with vertices $X_{(0)} \setminus \{v_1, \dots, v_k\}$. By the leaking property there exists a vertex v of A where equation (8.1) holds, and we pick $v_{k+1} = v$. Note that the number of outgoing edges at v in the acyclic orientation obtained that way is exactly $\text{val}_X(v) - \text{val}_A(v)$, hence the construction ensures that $K_+ - D$ is effective away from p . At p the constructed orientation has a sink, so the coefficient in K_+ is -1 . However, since $D = \emptyset$, the coefficient of p in D is strictly negative and hence $K_+ - D$ is effective on all of X , proving the statement. \square

A corollary of this proof is the inverse of lemma 8.4.14.

Corollary 8.4.16

Let D be a divisor of degree $g - 1$. If $|D| = \emptyset$ then D is linearly equivalent to a moderator. Moreover, if D is p -reduced, then D is equal to a moderator. Finally, we have $|D| = \emptyset$ if and only if $|K - D| = \emptyset$.

Proof. By lemma 8.4.15 there exists a moderator with $|K_+ - D| \neq \emptyset$. Since $\deg(K_+ - D) = 0$, this implies $K_+ - D \sim 0$. If D is p -reduced, the construction from 8.4.15 even gives $K_+ - D \geq 0$ and hence $K_+ - D = 0$. The final statement follows from the fact that $K_- = K - K_+$ is also a moderator. \square

The above computations lead to a useful criterion for testing non-emptiness of linear systems.

Corollary 8.4.17

Let D be a divisor of degree $d < g - 1$. If $|D + q| \neq \emptyset$ for all q then $|D| \neq \emptyset$.

Proof. If $D = \emptyset$, there exists a moderator K_+ and effective $E \sim K_+ - D$. Choose $q \in \text{Supp}(E)$, then $|K_+ - D - q| \neq \emptyset$ and therefore $|D + q| = \emptyset$. \square

A divisor D is called *special* if $\text{rk}(D) > \max\{d - g, -1\}$. By 8.4.7 D can only be special if $0 \leq d \leq 2g - 2$. In the range $0 \leq d \leq g - 1$, specialness is equivalent to $|D| \neq \emptyset$.

Lemma 8.4.18

A divisor D is special if and only if $K - D$ is special.

Proof. By symmetry we may assume $0 \leq \deg(D) \leq g - 1$. We set $e := g - 1 - d = \deg(K - D) - g + 1$. If $K - D$ is special, then for all effective divisors of degree e we have $|K - D - E| \neq \emptyset$. This implies $|D + E| \neq \emptyset$ by 8.4.16 (since these linear systems have degree $g - 1$). Applying 8.4.17 several times, we obtain $D \neq \emptyset$. Conversely, assume $K - D$ is non-special. Then there exists effective E of degree e with $|K - D - E| = \emptyset$. Again this implies $|D + E| = \emptyset$ and hence $|D| = \emptyset$. \square

We are now ready to prove Riemann-Roch theorem.

Proof of 8.4.5. By exchanging D and $K - D$ it is enough to prove the inequality

$$\operatorname{rk}(D) \geq \operatorname{rk}(K - D) + d - g + 1 =: e.$$

The case $\operatorname{rk}(K - D) = -1$ is contained in 8.4.7, hence we may assume $\operatorname{rk}(K - D) > -1$. Let E be an effective divisor of degree e . We have

$$\operatorname{rk}(K - D + E) \geq \operatorname{rk}(K - D) = \deg(K - D + E) - g + 1,$$

hence $K - D + E$ is special. It follows that $D - E$ is special by 8.4.18 and in particular $|D - E| \neq \emptyset$. This implies $\operatorname{rk}(D) \geq e$ and hence finishes the proof. \square

9 Tropical jacobians

9.1 (Coarse) differential forms, integration, and the Jacobian

The canonical line bundle plays a very special role in the theory of curves. Its sections are differential forms which have special properties and allow for special constructions which are not present in sections of arbitrary line bundles. These peculiarities of the canonical bundle, which manifest classically in the theory of jacobians, have a tropical counterpart, which is presented here.

In the following, we will always assume C to be a smooth tropical curve (and add finite type or compactness when necessary). Hence C can be represented by a metric graph with open and closed ends (where the latter are closed up by infinite vertices). Recall from Definition 7.3.4 that we can define the cotangent sheaf on C in terms of the exact sequence

$$0 \rightarrow \mathbf{R} \rightarrow \mathbf{R} \otimes_{\mathbf{Z}} \mathcal{A}ff_C^* \rightarrow T_C^* \rightarrow 0,$$

where the first map is just the inclusion of constant functions. Note that the tensor product in the middle term is a product over \mathbf{Z} , where \mathbf{Z} acts on $\mathcal{A}ff_C^*$ by (non-tropical) multiplication. In other words, $\mathbf{R} \otimes_{\mathbf{Z}} \mathcal{A}ff_C^*$ is the sheaf of affine linear functions where we also allow non-integer slopes.

Definition 9.1.1

A *residue form* ω is a global section of the cotangent sheaf T_C^* . The set of all residue forms constitutes a real vector space which we denote by $\Omega(C)$.

Remark 9.1.2

Note that residue forms are (locally constant) *classical* differential forms on the interior of edges. The name reflects the fact that these forms appear

as limits of the residues of classical differential forms when tropicalizing a complex curve to a tropical curve. In a sense that can be made precise, residue forms are therefore only “half-tropicalized” objects and should be distinguished from honest tropical differential forms, which can be defined as sections of the canonical line bundle. While the latter are arguably the more intrinsically tropical objects, residue forms play an important role in some contexts, in particular, in the context of tropical jacobians.

Since the function in $\mathcal{A}ff_C^*$ are locally affine linear, any residue form is locally constant (in particular, constant on the interior of edges). Moreover, it is zero on any closed end (since $T_p^* = 0$ for any infinite point p). For a given edge e with endpoint p , we can evaluate ω on the primitive vector $v_{e/p} \in \mathbf{Z}T_p$. If we fix vertex $p \in C$ and let v_1, \dots, v_k denote the primitive generators of all the adjacent edges, then the balancing condition translates to

$$\sum_{i=1}^k \omega(v_i) = 0.$$

(Note that the statements holds trivially for any infinite endpoint by the previous remark). Hence, a residue form can alternatively described by the following data.

Proposition 9.1.3

A residue form ω on C is uniquely described by a assignment of numbers to each oriented edge

$$\omega : \{\text{oriented edge of } C\} \rightarrow \mathbf{R},$$

such that

- *if e^+ and e^- are the two orientations of an edge, then $\omega(e^+) = -\omega(e^-)$,*
- *for any vertex p , if e_1, \dots, e_k denote all the outgoing edges from p , then*

$$\sum_{i=1}^k \omega(e_i) = 0.$$

Of course, due to the equation in the proposition, the values of ω on oriented edges are still largely redundant. In fact, $\Omega(C)$ is a $g + m$ -dimensional

vector space, where

$$g := \dim H_1(C, \mathbf{R}), \quad m := \max(\#\{\text{open ends}\} - 1, 0).$$

To prove this, we need some terminology. A collection of 2-valent points $c_1, \dots, c_k \in C$ is called a *tree cut* if $C \setminus c_1, \dots, c_k$ is simply connected.

Exercise 9.1.4

Show that if c_1, \dots, c_k is a tree cut, then $k = g$.

Proposition 9.1.5

Let c_1, \dots, c_g be a tree cut of C . Let e_1, \dots, e_g be oriented edges with $c_i \in e_i$, and let $e_{g+1}, \dots, e_{g+m+1}$ be the set of open ends of C (oriented away from the end). Then the map

$$\begin{aligned} \Omega(C) &\rightarrow \mathbf{R}^g \times \mathbf{R}^m, \\ \omega &\mapsto (\omega(e_1), \dots, \omega(e_{g+m})), \end{aligned}$$

is an isomorphism of real vector spaces. In particular, the dimension of $\Omega(C)$ is equal to $g + m$.

Proof. The map is obviously linear. Note that the contraction of a closed infinite end does neither affect $\Omega(C)$ nor the map above. We can hence contract all such ends and obtain a curve without such ends (or \mathbf{T}^1 or \mathbf{TP}^1 , which obviously satisfy the statement).

Let us now assume that C is a tree, i.e., $g = 0$. Let us fix some vector $v \in \mathbf{R}^{g+m}$. If C does not contain an inner edge, then v obviously has a unique preimage residue form by the equations in Proposition 9.1.3. Otherwise, let e be an inner edge and C_1 and C_2 be the two curves obtained by cutting C along some interior point of e . We may assume that C_2 contains the end e_{m+1} . By induction, the statement is true for C_1 and there exists a unique residue form ω_1 on C_1 determined by the values of v (restricted to the ends of C_1). In particular, ω_1 has some value on the cut open end e . If we assign $-\omega_1(e)$ to the opposite end in C_2 , the remaining values of v , again by induction, fix a unique residue form ω_2 on C_2 . Since the two residue forms are compatible on e , they can be combined to a residue form ω on C , and the uniqueness on C_1, C_2 implies the uniqueness on C . This finishes the tree case.

Finally, let us assume $g > 0$. We denote by $C' := C \setminus c_1$ the curve of genus $g - 1$ obtained by removing the first cut point. If we denote by e^+, e^- the two freshly created open ends of C' (we may assume that e^+ and e_1 are oriented equally), we can identify the sets

$$\Omega(C) = \{\omega' \in \Omega(C') : \omega(e^+) = -\omega(e^-)\}. \quad (9.1)$$

By induction on g , we may assume that the map $\Omega(C') \rightarrow \mathbf{R}^{g+m+1}$ is an isomorphism (for the tree cut c_1, \dots, c_{g-1} and ends ordered as $e^+, e^-, e_1, \dots, e_{m+1}$). The linear map in question is then given by the concatenation

$$\Omega(C) \rightarrow \Omega(C') \rightarrow \mathbf{R}^{g+m+1} \rightarrow \mathbf{R}^{g+m},$$

where the last map is the projection which forgets the e^- -coordinate (and moves the e^+ -coordinate to the right place). It is hence still an isomorphism by Equation (9.1). \square

We can restrict the aforementioned exact sequence to the integers,

$$0 \rightarrow \mathbf{R} \rightarrow \mathcal{A}ff_C^* \rightarrow \mathbf{Z}T_C^* \rightarrow 0,$$

and obtain a canonical integer lattice in $\Omega(C)$.

Definition 9.1.6

A *integer residue form* ω is a global section of the integer cotangent sheaf $\mathbf{Z}T_C^*$, or, in the description of Proposition 9.1.3, satisfies $\omega(e) \in \mathbf{Z}$ for any oriented edge e . The set of all integer residue forms, denoted by $\Omega_{\mathbf{Z}}(C)$, constitutes a lattice in $\Omega(C)$ and hence turns $\Omega(C)$ into a tropical vector space in the sense of Section 2.1.

As explained above, residue forms are honest classical differential forms on the interior of edges and therefore can be integrated along (compact) paths. To fix notation, let us explicitly define this integration here.

Definition 9.1.7

Let ω be a residue form on C and let e be an oriented bounded edge. We define the *edge integral* of ω along e to be

$$\int_e \omega := l(e)\omega(e).$$

Note that this can be extended to closed leaves, since then $\omega(e) = 0$. Given an oriented 1-chain $\alpha = e_1 + \dots + e_k$ (i.e., a formal sum of oriented bounded edges), we can extend integration by linearity

$$\int_{\alpha} \omega := \sum_{i=1}^k l(e_i) \omega(e_i).$$

Given the opposite orientation e^+ and e^- of an edge e , we obviously get

$$\int_{e^+ + e^-} \omega = l(e)(\omega(e^+) + \omega(e^-)) = 0.$$

It follows that integration gives a well-defined pairing on homology.

Definition 9.1.8

We define *integration pairing* between $\Omega(C)$ and $H_1(C, \mathbf{R})$ by

$$\int : \Omega(C) \times H_1(C, \mathbf{R}) \rightarrow \mathbf{R},$$

$$(\omega, [\alpha]) \mapsto \int_{\alpha} \omega.$$

Proposition 9.1.9

Let C be a compact smooth tropical curve. Then the integration pairing \int is non-degenerate.

Proof. We have to construct for any non-zero residue form ω a suitable cycle α such that $\int_{\alpha} \omega \neq 0$ — and vice versa. Let us start with a non-zero ω . Assume that $\int_{\alpha} \omega = 0$ for all cycles α of C . This would imply that ω can be integrated to a affine linear function

$$f(p) := \int_{p_0}^p \omega \in \mathcal{A}ff(C).$$

By Proposition 8.2.4 we know that f must be a constant function, which implies $\omega = df = 0$.

Vice versa, let α be a cycle determining a non-zero homology class $[\alpha] \in H_1(C, \mathbf{R})$. Let $\omega_{\alpha}(e)$ be the coefficient of the oriented edge e in α (when

writing α as a linear combination of oriented edges using only the given orientation of e). The fact that α is closed translates directly to the zero divergence condition in Proposition 9.1.3, hence this way we obtain a residue form ω_α on C . Integration gives

$$\int_\alpha \omega_\alpha = \sum_e \omega_\alpha(e)^2 l(e) > 0,$$

where the sum runs through the edges of C (with arbitrary orientation). This proves the claim. which proves the claim. \square

Let $\Omega(C)^*$ be the dual vector space of $\Omega(C)$. Equipped with the dual lattice $\Omega_{\mathbf{Z}}(C)^*$, this is again a tropical vector space. By Proposition 9.1.9, the isomorphism $\int : H_1(C, \mathbf{R}) \rightarrow \Omega(C)^*$ allows us to regard $H_1(C, \mathbf{Z})$ as *another* lattice in $\Omega(C)^*$. The quotient of $\Omega(C)^*$ by this second lattice carries a canonical structure as smooth tropical manifold and as group (see Example 7.1.9 and the next section).

Definition 9.1.10

Let C be a compact smooth tropical curve. We define the *Jacobian* of C as the quotient of the tropical vector space $\Omega(C)^*$ by the lattice $H_1(C, \mathbf{Z})$,

$$\text{Jac}(C) := \Omega(C)^*/H_1(C, \mathbf{Z}).$$

Before we proceed to analyze the specific properties of jacobians, it is worthwhile to make a quick stop and have a look at the general features of this construction.

9.2 Abelian varieties, polarization, and theta functions

Analogous to the classical case, the jacobian of a tropical curve is a special example of more general objects called tropical abelian varieties and tropical tori.

The ingredients needed to define a tropical torus are

- a real vector space V ,

- two lattices $L, \Lambda \subset V$.

However, the lattices play different roles in the construction. Namely, we use L to determine the *tropical* structure of V (i.e., the integral tangent vectors), whereas we use Λ to quotient by (i.e., to fix the *metric* structure). Here is the exact definition.

Definition 9.2.1

A *tropical torus* A is the quotient of a tropical vector space $V = \mathbf{R} \otimes L$ by a lattice $\Lambda \subset V$,

$$A = (\mathbf{R} \otimes L) / \Lambda.$$

A tropical torus is compact regular tropical variety in a canonical way.

Exercise 9.2.2

Show that the restrictions of the quotient map $V \rightarrow A$ to sufficiently small domains provide a regular tropical atlas for A .

Example 9.2.3

The jacobian of a compact smooth tropical curve of genus g is tropical torus of dimension g , with $V = \Omega(C)^*$, $L = \Omega_{\mathbf{Z}}(C)^*$ and $\Lambda = H_1(C, \mathbf{Z})$ (embedded via integration).

Remark 9.2.4

Let A be tropical variety homeomorphic to a topological torus $(S^1)^n$ (this implies that A is regular), and let $p \in A$ be a fixed point. Parallel transport of integer tangent vectors produces a monodromy action

$$\gamma : H_1(A, \mathbf{Z}) \rightarrow \mathrm{GL}(\mathbf{Z}T_{A,p}).$$

Remark 9.2.5

tori up to isom are equal to (V, L, Λ) up to isom

Tropical tori often come equipped with additional structure. In the following, we use the identification

$$H^1(A, \mathbf{Z}T_A^*) = \Lambda^* \otimes L^*$$

and think of an element in this space either as a bilinear form on V or as a linear map $P : \Lambda \rightarrow L^*$. We also denote the space of symmetric bilinear forms on V by $\mathrm{Sym}^2 V^* \subset V^* \otimes V^*$.

Definition 9.2.6

A polarization P of the tropical torus A is an element

$$P \in (\Lambda^* \otimes L^*) \cap \text{Sym}^2 V^*,$$

such that the corresponding symmetric bilinear form is positive definite. In other words, P is a scalar product on V such that

$$P(v, w) \in \mathbf{Z} \text{ for all } v \in \Lambda, w \in L.$$

In this case, the pair (A, P) is called a *tropical abelian variety*.

A polarization P is called *principal* if the linear map $P : \Lambda \rightarrow L^*$ is an isomorphism of lattices. In terms of lattice bases v_1, \dots, v_n of Λ and w_1, \dots, w_n of L , this is equivalent to the integer matrix $P(v_i, w_j)$ to be invertible over \mathbf{Z} . A tropical torus together with a principal polarization is called *principally polarized*.

Remark 9.2.7

Given a principal polarization P , we can dualize the representation of A and obtain a canonical isomorphism

$$A \cong (\mathbf{R} \otimes \Lambda^*) / L^*. \tag{9.2}$$

Note that the roles of Λ and L get exchanged.

To any polarization of a tropical torus, we can associate a theta function and a theta divisor.

Definition 9.2.8

Let A be a tropical abelian variety with polarization P . We define the *theta function* $\vartheta : V \rightarrow \mathbf{R}$ by

$$\vartheta(x) := \max_{\lambda \in \Lambda} \{P(\lambda, x) - \frac{1}{2}P(\lambda, \lambda)\}. \tag{9.3}$$

Note that the maximum always exists since P is positive definite (and hence $-P(\lambda, \lambda)$ is a negative definite quadratic functional).

The theta function satisfies $\vartheta(x) = \vartheta(-x)$ and is quasi-periodic in Λ in the following sense.

Lemma 9.2.9

For any $x \in V$ and $\mu \in \Lambda$ the theta function satisfies

$$\vartheta(x + \mu) = \vartheta(x) + P(\mu, x) + \frac{1}{2}P(\mu, \mu). \quad (9.4)$$

In particular, for fixed μ the difference function $\vartheta(x + \mu) - \vartheta(x)$ is affine \mathbf{Z} -linear in x .

Proof. By shifting the λ to $\lambda + \mu$ we can evaluate the theta function as

$$\vartheta(x + \mu) := \max_{\lambda \in \Lambda} \{P(\lambda + \mu, x + \mu) - \frac{1}{2}P(\lambda + \mu, \lambda + \mu)\}. \quad (9.5)$$

Multiplying out this type of term we get

$$\begin{aligned} & P(\lambda + \mu, x + \mu) - \frac{1}{2}P(\lambda + \mu, \lambda + \mu) \\ &= P(\lambda, x) + P(\mu, x) + P(\lambda, \mu) + P(\mu, \mu) - \frac{1}{2}(P(\lambda, \lambda) + 2P(\lambda, \mu) + P(\mu, \mu)) \\ &= P(\lambda, x) - \frac{1}{2}P(\lambda, \lambda) + P(\mu, x) - \frac{1}{2}P(\mu, \mu). \end{aligned} \quad (9.6)$$

Since the first two terms in the last expression coincide with the terms in Equation (9.3), the first statement follows. The second statement holds since $P \in \Lambda^* \otimes L^*$. Hence for fixed $\mu \in \Lambda$, we have $P(\mu, \cdot) \in L^*$ and $P(\mu, x) - \frac{1}{2}P(\mu, \mu)$ is indeed affine \mathbf{Z} -linear on $\mathbf{R} \otimes L$. \square

The quasi-periodicity of ϑ has some interesting consequences. In particular, it follows that $\text{div}(\vartheta) \subset V$ is a Λ -periodic divisor and hence induces a well-defined divisor $\Theta \subset A$. In fact, ϑ defines a Cartier divisor on A (hence a line bundle), and Θ is the Weil divisor associated to this Cartier divisor.

Definition 9.2.10

Given a tropical abelian variety A , we call

$$\Theta = \text{div}(\vartheta)/\Lambda \subset A \quad (9.7)$$

the *Theta divisor* of A .

Remark 9.2.11

The definition depends up to translation on the fixed group structure of A (alternatively, on a fixed marked point $0 \in A$).

9.3 The Abel-Jacobi theorem

Recall that given a compact smooth tropical curve C , we define the jacobian of C as the tropical torus

$$\text{Jac}(C) := \Omega(C)^*/H_1(C, \mathbf{Z}). \quad (9.8)$$

Also recall that this notation implicitly uses the isomorphism

$$\int : H_1(C, \mathbf{R}) \cong \Omega(C)^*. \quad (9.9)$$

We continue to use this identification in the following definition of a canonical polarization of $\text{Jac}(C)$. In the following, we fix arbitrary orientations on the edges of C (for some graph structure) and think of $H_1(C, \mathbf{R})$ as a subspace of the space of inner 1-chains (i.e., the real vector space formally generated by the inner edges of C). As before, the length of an inner edge is denoted by $l(e)$.

Definition 9.3.1

We define the bilinear form P on the space of inner 1-chains as bilinear extension of

$$P(e, f) = \begin{cases} l(e) & \text{if } e = f, \\ 0 & \text{otherwise,} \end{cases} \quad (9.10)$$

where e, f are two inner edges of C .

Proposition 9.3.2

The bilinear form P on $H_1(C, \mathbf{R})$ is symmetric, positive definite and lies in

$$P \in H_1(C, \mathbf{Z})^* \otimes \int^{-1} \Omega_{\mathbf{Z}}(C). \quad (9.11)$$

Moreover, P is principal and hence turns $\text{Jac}(C)$ into a principally polarized abelian variety.

Proof. Let $\mathbf{R}\{\text{inner edges}\}$ denote the space of inner 1-chains. By definition, P has diagonal form with respect to the standard basis of $\mathbf{R}\{\text{inner edges}\}$

and the diagonal entries $l(e)$ are strictly positive. Hence P is symmetric and positive definite on $\mathbf{R}\{\text{inner edges}\}$, thus also on $H_1(C, \mathbf{R})$.

For the integrality statement, note we can think of $\Omega(C)$ as a subspace of $\mathbf{R}\{\text{inner edges}\}$ by Proposition 9.1.5. In fact, this is exactly the same subspace as $H_1(C, \mathbf{R}) \subset \mathbf{R}\{\text{inner edges}\}$, since both spaces are given by the condition that the total in and out flow at each vertex is zero. The induced identification

$$H_1(C, \mathbf{R}) = \Omega(C) \tag{9.12}$$

is denoted by $\omega(\alpha)$ for $\alpha \in H_1(C, \mathbf{R})$ in the following. Note that identifies $H_1(C, \mathbf{Z})$ with $\Omega_{\mathbf{Z}}(C)$. Unwrapping the definition, it turns out that P is nothing else but the bilinear form associated to the isomorphism $H_1(C, \mathbf{R}) \cong H_1(C, \mathbf{R})^*$ obtained by concatenating the above identification (9.12) with integration. In formulas, this just means

$$P(\alpha, \beta) = \int_{\alpha} \omega(\beta). \tag{9.13}$$

The isomorphism $H_1(C, \mathbf{R}) \cong H_1(C, \mathbf{R})^*$ descends to an isomorphism of lattices

$$H_1(C, \mathbf{Z}) \cong \int^{-1} \Omega_{\mathbf{Z}}(C), \tag{9.14}$$

which shows that $P \in H_1(C, \mathbf{Z})^* \otimes \int^{-1} \Omega_{\mathbf{Z}}(C)$ is indeed a principal polarization. \square

Remark 9.3.3

Using the canonical isomorphisms

$$H^1(C, \mathbf{Z}) \cong H_1(C, \mathbf{Z})^*, \quad H^1(C, \mathbf{R}) \cong \mathbf{R} \otimes H^1(C, \mathbf{Z}), \tag{9.15}$$

and the principal polarization P , by Remark 9.2.7 we obtain a second description of the jacobian as

$$\text{Jac}(C) \cong (\mathbf{R} \otimes H^1(C, \mathbf{Z})) / \Omega_{\mathbf{Z}}(C), \tag{9.16}$$

where the embedding $\Omega_{\mathbf{Z}}(C) \hookrightarrow H^1(C, \mathbf{R})$ is given by integration.

Recall from Proposition 8.3.14 that on a smooth curve C , line bundles are in one-to-one correspondence to rational equivalence classes of divisors

$$\text{Pic}(C) \cong \text{Div}(C) / \sim . \quad (9.17)$$

Moreover, by Proposition 8.2.3 we know that two rationally equivalent divisors on a compact curve C have the same degree. This gives rise to a well-defined well-defined surjective group homomorphism

$$\text{deg} : \text{Pic}(C) \cong \text{Div}(C) / \sim \rightarrow \mathbf{Z}. \quad (9.18)$$

The fibers of the degree map split the Picard group into isomorphic pieces

$$\text{Pic}(C) = \sum_{d \in \mathbf{Z}} \text{Pic}^d(C) \quad (9.19)$$

parameterizing line bundles of a given degree d . The Abel-Jacobi theorem states that $\text{Pic}^0(C)$ (and hence any $\text{Pic}^d(C)$) is isomorphic to $\text{Jac}(C)$. Let us proceed to prove this.

As before, we assume in the following that C is a compact smooth curve. We fix a base point $p_0 \in C$. Let $p \in C$ be another point, and choose some path γ in C from p_0 to p . By integration, \int_{γ} defines an element in $\Omega(C)^*$. Choosing another γ' from p_0 to p results in another linear functional $\int_{\gamma'}$, but since $\gamma - \gamma' \in H_1(X, \mathbf{Z})$ forms a closed cycle, the difference

$$\int_{\gamma} - \int_{\gamma'} \in H_1(X, \mathbf{Z}), \quad (9.20)$$

and both functionals determine the same element in $\text{Jac}(C)$, which we denote by

$$\int_{p_0}^p \in \text{Jac}(C). \quad (9.21)$$

Definition 9.3.4

Given a divisor $D = \sum a_i p_i \in \text{Div}^d(C)$ of degree d , we define

$$\mu(D) := \int^D := \sum a_i \int_{p_0}^{p_i} \in \text{Jac}(C). \quad (9.22)$$

This defines a group homomorphism

$$\mu : \text{Div}^d(C) \rightarrow \text{Jac}(C), \quad (9.23)$$

which is called the *tropical Abel-Jacobi map*.

Remark 9.3.5

Note that μ depends on the choice of p_0 for $d \neq 0$. Indeed, choosing another base point p'_0 results in an additive shift of the map by $d \int_{p_0}^{p'_0}$. For $d = 0$ however, the dependence on p_0 vanishes and we obtain a canonically defined map $\text{Pic}^0(C) \rightarrow \text{Jac}(C)$.

Theorem 9.3.6 (Tropical Abel-Jacobi theorem)

The Abel-Jacobi map $\mu : \text{Div}^d(C) \rightarrow \text{Jac}(C)$ is constant on rational equivalence classes,

$$D \sim D' \implies \mu(D) = \mu(D'). \quad (9.24)$$

Moreover, the induced map $\nu : \text{Pic}^d(C) \rightarrow \text{Jac}(C)$ is a bijection.

Recall that $\vartheta : \text{Div}^d(C) \rightarrow \text{Pic}^d(C)$ is the map which associates to the divisor D the line bundle $\vartheta(D)$. The theorem can be summarized in the following commutative diagram.

$$\begin{array}{ccc} \text{Div}^d(C) & & \\ \vartheta \downarrow & \searrow \mu & \\ \text{Pic}^d(C) & \xrightarrow[\nu]{\cong} & \text{Jac}(C) \end{array} \quad (9.25)$$

Similar to the classical situation, the statement naturally splits into two parts called Abel’s theorem (“ ν is well-defined and injective”) and Jacobi inversion theorem (“ ν is surjective”). We start with tropical Abel.

Theorem 9.3.7 (Tropical Abel’s theorem)

Two divisors D, D' are rationally equivalent if and only if $\mu(D) = \mu(D')$.

Proof. To prove the first direction, we show that if D is a divisor rationally equivalent to zero, then $\mu(D) = 0$. Let f be a rational function such that $\text{div}(f) = D$. Choose a graph structure for C which contains the support of

D , and fix an orientation of the edges. Then the slope of f on each inner edge defines an (oriented) 1-chain which we denote by

$$df \in \mathbf{R}\{\text{inner edges}\}. \quad (9.26)$$

By definition, $\mu(D)$ is just the class in $\text{Jac}(C)$ of the linear functional on $\Omega(C)$ given by integration over this 1-chain. Using again the identification $\Omega(C) = H_1(C, \mathbf{R}) \subset \mathbf{R}\{\text{inner edges}\}$, we may now exchange the roles of integration chain and integrand. Namely, for any residue form ω let α be the corresponding 1-cycle, we get

$$\int_{df} \omega = P(df, \alpha) = \int_{\alpha} df = 0, \quad (9.27)$$

(by slight abuse of notation, since df is not balanced at all vertices, hence not a residue form). This shows $\mu(D) = 0$, which proves the first claim.

We proceed to show the second direction. We have to show that any $D \in \text{Div}^0(C)$ with $\mu(D) = 0$ is rationally equivalent to zero. Let γ an oriented 1-chain with boundary D . By definition, $\mu(D)$ is described by integration over γ . Since $\mu(D) = 0$, we can in fact choose γ such that

$$\int_{\gamma} = 0 \in \Omega(C)^*. \quad (9.28)$$

Again, we can essentially turn around domain and integrand and get

$$\int_{\alpha} \gamma = 0 \quad (9.29)$$

for all $\alpha \in H_1(C, \mathbf{R})$. It follows that we can “integrate” γ and obtain a rational function $f := \int \gamma$ whose divisor is equal to the boundary of γ , i.e., $\text{div}(f) = D$. \square

The surjectivity of ν resp. μ can be solved in a rather explicit way, and it is worthwhile to look at this in some detail. The key construction is to pull back translated theta divisors from $\text{Jac}(C)$ to C . To do so, we need a map from C to $\text{Jac}(C)$. In fact, we can use the Abel-Jacobi map μ via the inclusion $C \subset \text{Div}^1(C)$ which associates to each point of C the divisor consisting of only this point.

Lemma 9.3.8

The Abel-Jacobi map

$$\mu|_C : C \rightarrow \text{Jac}(C). \quad (9.30)$$

is a tropical morphism.

Proof. The map is constant in a neighbourhood of the infinite points of C (in fact, it contracts any end to point) since two points on the same end are rationally equivalent. It remains to check the claim for a finite point $p \in C$. Let us fix a chart containing p to the $n + 1$ -valent standard tropical line $L(n + 1) \subset \mathbf{R}^n$ with $n + 1 = \text{val}(p)$. We may assume that the base point p_0 in the definition of μ is equal to p . Let e_0, \dots, e_n denote the oriented edges of C corresponding to the rays of $L(n + 1)$, and let $\varphi_i \in \Omega(C)^*$ denote the evaluations

$$\varphi_i(\omega) := \omega(e_i). \quad (9.31)$$

It holds $\varphi_0 + \dots + \varphi_n = 0$. Moreover, for any point p' on the i -th ray of $L(n + 1)$, we find $\mu(p') = l\varphi_i$ where l is the lattice distance of p' from p . It follows that μ is locally given by the linear map

$$\begin{aligned} \mathbf{R}^n &\rightarrow \Omega(C)^*, \\ (x_1, \dots, x_n) &\mapsto x_1\varphi_1 + \dots + x_n\varphi_n. \end{aligned} \quad (9.32)$$

This shows that μ is a tropical morphism. \square

Let $f : X \rightarrow Y$ be a tropical morphism and let $s = \{(U_i, s_i)\}$ be a Cartier divisor on Y . We define the *pull back* of s along f to be the Cartier divisor

$$f^*s := \{(f^{-1}(U_i), s_i \circ f)\}. \quad (9.33)$$

Exercise 9.3.9

Check that f^*s is a well-defined Cartier divisor on X .

Recall that the polarization P on $\text{Jac}(C)$ provides us with a theta function ϑ which we can either regard as a quasi-periodic function on V or a Cartier divisor on $\text{Jac}(C)$. For any $y \in \text{Jac}(C)$, we the shift function

$$\begin{aligned} \psi_y : \text{Jac}(C) &\rightarrow \text{Jac}(C), \\ x &\mapsto x + y, \end{aligned} \quad (9.34)$$

is a tropical isomorphism. We denote the shifted theta function by

$$\vartheta_y := \psi_{-y}^* \vartheta. \quad (9.35)$$

Definition 9.3.10

The *Jacobi inverse* of $y \in \text{Jac}(C)$ is defined to be the divisor

$$D(y) := \text{div}(\mu|_C^* \vartheta_y) \in \text{Div}(C). \quad (9.36)$$

Note that the induced *Jacobi inversion map*

$$\text{Jac}(C) \rightarrow \text{Div}(C) \quad (9.37)$$

still depends on the choice of a base point p_0 .

Proposition 9.3.11

For any $y \in \text{Jac}(C)$ the divisor $D(y)$ is an effective divisor of degree g .

Proof. The effectiveness follows from the fact that ϑ is locally a tropical polynomial, i.e., a maximum of finitely many terms. To compute the degree, let us first make the construction of $D(y)$ more explicit. Let $c_1, \dots, c_g \in C$ be a tree cut disjoint from $D(y)$ and let e_1, \dots, e_g be oriented edges with $c_i \in e_i$. These choices in fact provide us with explicit bases for all vector spaces and lattices involved. Namely we get lattice bases

$$\begin{aligned} \omega_1, \dots, \omega_g &\in \Omega_{\mathbf{Z}}(C), \\ \varphi_1, \dots, \varphi_g &\in \Omega_{\mathbf{Z}}(C)^*, \\ \alpha_1, \dots, \alpha_g &\in H_1(C, \mathbf{Z}), \end{aligned}$$

where

- the ω_i correspond to the standard basis of \mathbf{R}^g via Proposition 9.1.5, i.e., $\omega_i(e_j) = \delta_{ij}$,
- the φ_i form the dual basis to the ω_i , i.e., $\varphi_i(\omega) = \omega(e_i)$,
- α_i is the unique simple loop in $\tilde{C} \cup \{c_i\}$ which traverses e_i positively oriented (where $\tilde{C} = C \setminus \{c_1, \dots, c_g\}$).

The first and third basis actually coincide via the identification (9.12) and we get

$$P(\alpha_i, \varphi_i) = \delta_{ij}. \quad (9.38)$$

Since \tilde{C} is a tree, we can lift the map $\mu|_{\tilde{C}} : \tilde{C} \rightarrow \text{Jac}(C)$ to a map $\tilde{\mu} : \tilde{C} \rightarrow \Omega(C)^*$ and can compute $D(y)$ as the divisor of the rational function

$\tilde{\mu}^* \varphi_y$. Hence the degree of $D(y)$ is equal to sum of the outgoing slopes of φ_y on the open ends of \tilde{C} . Indeed, let us denote by z_i^+, z_i^- the ends of \tilde{C} coming from the cut point c_i . Here, we use the convention that both ends are oriented outwards and the orientation of z_i^+ is compatible with e_i . Denoting the slope of $\tilde{\mu}^* \varphi_y$ on z_i^\pm by s_i^\pm , a straightforward generalization of Proposition 8.2.3 shows

$$\deg(D(y)) = \sum s_i^+ + s_i^-. \quad (9.39)$$

We now use the following two facts:

- (a) The image under $\tilde{\mu}$ of the primitive generator of z_i^\pm is exactly $\pm \varphi_i \in \Omega_{\mathbb{Z}}(C)^*$. This follows from the simplest case of the calculation in the proof of Lemma 9.3.8 ($n = 1$).
- (b) The difference of the image points of the ends of z_i^+ minus z_i^- in $\Omega(C)^*$ is equal to α_i (since the difference of integration from some base point to the endpoint of z_i^+ minus integration to the endpoint of z_i^- is the same as integration over α_i).

We can now use the quasi-periodicity of ϑ_y and compute the sum of slopes $s_i^+ + s_i^-$ simultaneously by

$$\begin{aligned} s_i^+ + s_i^- &= d(\vartheta_y(x + \alpha_i) - \vartheta_y(x))(\varphi_i) && \text{by (a), (b)} \\ &= P(\alpha_i, \varphi_i) && \text{by 9.2.9} \\ &= 1 && \text{by (9.38)} \end{aligned} \quad (9.40)$$

Since there are exactly g such pairs of ends, we obtain

$$\deg(D(y)) = \sum s_i^+ + s_i^- = g, \quad (9.41)$$

as promised. □

With just a little more effort, we are now ready to prove that the Jacobi inversion map is essentially inverse to the Abel map ν , finishing the proof of Theorem 9.3.6. This is in fact not quite true — we additionally have to shift ν by a constant.

Theorem 9.3.12 (Jacobi inversion theorem)

There exists an element $\kappa \in \text{Jac}(C)$ called the Jacobi inversion constant such that for all $y \in \text{Jac}(C)$ we have

$$\mu(D(y)) + \kappa = y. \tag{9.42}$$

In particular, the map $\mu : \text{Sym}^g(C) \rightarrow \text{Jac}(C)$ and hence all maps $\mu : \text{Div}^d(C) \rightarrow \text{Jac}(C)$ are surjective.

Remark 9.3.13

Note that κ is constant with respect to varying y — it might however depend the base point p_0 . An investigation of this dependency and the meaning of κ will be part of Section 9.4.

In the previous proof, we used the fact that the degree of the divisor of a rational function can be computed in terms of the behavior of the function on the ends of the curve. In fact, we can extract even a little more information from this boundary behavior. This is the extra ingredient we need to prove the Jacobi inversion theorem.

Lemma 9.3.14

Let $C \subset \mathbf{R}^n$ be a bounded tropical curve (not necessarily connected nor smooth). We denote the boundary point of C by

$$\partial C = \{b_1, \dots, b_k\}. \tag{9.43}$$

Let v_1, \dots, v_k be the primitive generators of corresponding open ends of C , oriented outwards (i.e. towards the b_i), and let w_1, \dots, w_k be the weights of the ends. Let f be a rational function of C and let $\text{div}(f) = \sum a_i p_i$ be its divisor. We extend f to the points b_i by continuity. Then the equation

$$\sum a_i p_i = \sum_{j=1}^k w_j (df(v_j) b_j - f(b_j) v_j) \tag{9.44}$$

holds. Here, the sums on both sides refer to sums of vectors in \mathbf{R}^n .

Proof. Both sides of the formula behave linearly with respect to the connected components of C . On the other hand, when cutting the curve into

more pieces by removing a two-valent vertex not contained in the support of $\text{div}(f)$, both sides remain unchanged. It therefore suffices to prove the case $\text{div}(f) = ap$ and p is the only vertex of C . In this case, we may write $b_j = p + l_j v_j$ and then get $f(b_j) = f(p) + l_j df(v_j)$. A straightforward computation then gives

$$\begin{aligned} \sum_{j=1}^k w_j (df(v_j)b_j - f(b_j)v_j) &= \sum_{j=1}^k (w_j df(v_j)p + w_j df(v_j)l_j v_j \\ &\quad - w_j f(p)v_j - w_j l_j df(v_j)v_j) \quad (9.45) \\ &= \sum_{j=1}^k (w_j df(v_j))p = ap. \end{aligned}$$

Here in the second step, the second and fourth term cancel while the third term is zero by the balancing condition. \square

Proof of Theorem 9.3.12. We use exactly the same setup and notation as in the proof of Proposition 9.3.11, i.e. we fix a tree cut of C and consider the lifted map $\tilde{\mu} : \tilde{C} \rightarrow \Omega(C)^*$. We pick $y \in \Omega(C)^*$ from some local chart of $\text{Jac}(C)$ such that the Jacobi inverse $D(y)$ has support disjoint from the tree cut (by choosing different tree cuts, we can cover $\Omega(C)^*$ completely). Since μ is group homomorphism, the image of $D(y) = \sum a_i p_i$ under μ can be computed as

$$\mu(D(y)) = [\sum a_i \tilde{\mu}(p_i)]. \quad (9.46)$$

Here, the sum on the right hand side is a standard sum of vectors in $\Omega(C)^*$, and the square brackets denote its equivalence class in $\text{Jac}(C)$. By projection formula, we have

$$\sum a_i \tilde{\mu}(p_i) = \text{div}(\tilde{\vartheta}_y), \quad (9.47)$$

where $\tilde{\vartheta}_y = \vartheta_y|_{\tilde{\mu}(\tilde{C})}$ is the restriction of the shifted theta function to the image of $\tilde{\mu}$. Hence we can apply Lemma 9.3.14 and obtain

$$\sum a_i \tilde{\mu}(p_i) = \sum_{j=1}^g (s_j^+ b_j^+ - \vartheta_y(b_j^+) \varphi_j) + \sum_{j=1}^g (s_j^- b_j^- + \vartheta_y(b_j^-) \varphi_j)$$

and collecting all constant terms with respect to y on the left, we get

$$\begin{aligned}
 &= \kappa_1 - \sum_{j=1}^g (\vartheta(b_j^- + \alpha_j - y) - \vartheta(b_j^- - y)) \varphi_j \\
 &= \kappa_2 - P(\alpha_j, b_j^- - y) \varphi_j && \text{by 9.2.9} \\
 &= \kappa_3 + P(\alpha_j, y) \varphi_j \\
 &= \kappa_3 + y && \text{by (9.38)}
 \end{aligned}$$

This proves the claim. \square

Remark 9.3.15

The Jacobi inversion theorem and the Abel-Jacobi theorem have some nice corollaries (e.g. used in the proof of Proposition 8.4.7).

- (a) Let $D \in \text{Div}^g(C)$ be an arbitrary divisor of degree g . Then

$$D' := D(\mu(D) + \kappa) > 0 \tag{9.48}$$

is an effective divisor. Jacobi inversion implies $\mu(D) = \mu(D')$, hence Abel's theorem implies $D \sim D'$. It follows that the linear system $|D|$ is non-empty whenever $\deg(D) = g$ (and hence also $\deg(D) \geq g$).

- (b) The previous statement can be generalized to divisors D of arbitrary degree. In the same way as before, we can then conclude

$$D \sim D(\mu(D) + \kappa) + (d - g)p_0. \tag{9.49}$$

- (c) By Lemma 9.3.8, the map

$$\mu|_C^d : C \times \cdots \times C \rightarrow \text{Div}^d(C) \rightarrow \text{Jac}(C) \tag{9.50}$$

is a tropical morphism. By Abel-Jacobi, its image $\mu(\text{Sym}^d(C))$ corresponds to non-empty linear systems of degree d (among all divisor classes of degree d). We have

$$\dim(C^d) = d, \quad \dim(\text{Jac}(C)) = g, \tag{9.51}$$

which implies that $\mu|_C^d$ cannot be surjective for $d < g$. Hence for all $d < g$, there exists a divisor D of degree d with $|D| = \emptyset$.

- (a) Show that $H^1(C, \mathcal{A}ff_C^*) \cong \mathbf{Z}$. Show that a unique generator is given by the 1-cochain C_e which maps the oriented edge e to the covector having slope 1 in the direction of e , and evaluates to zero at all other edges.
- (b) Show that with the above identification, the map $\text{Pic}(C) \rightarrow H^1(C, \mathbf{Z}T_C^*)$ in (9.55) is equal to the degree map $\text{deg} : \text{Pic}(C) \rightarrow \mathbf{Z}$.
- (c) Show that the map $\Omega_{\mathbf{Z}}(C) \rightarrow H^1(C, \mathbf{R})$ in (9.55) is equal to the embedding given by integration.
- (d) Let $\psi : H^1(C, \mathbf{R}) \rightarrow \text{Pic}(C)$ denote the map in the middle of (9.55). Given $D \in \text{Div}^0(C)$, let α be a 1-chain in C with boundary $\partial\alpha = D$. Let $P(\alpha, \cdot) \in H_1(C, \mathbf{R})^* \cong H^1(C, \mathbf{R})$ be the linear functional. Show that

$$\psi(P(\alpha, \cdot)) = \mathcal{O}(D). \quad (9.57)$$

These statements imply the Abel-Jacobi theorem $\text{Jac}(C) \cong \text{Pic}^0(C)$.

9.4 Riemann's theorem

Recall that the appearance of the Jacobi inversion constant $\kappa \in \text{Jac}(C)$ in Theorem 9.3.12 (in general, κ depends on the choice of base point p_0). Riemann's theorem gives a more explicit reinterpretation of this constant. As before, we denote by $\text{Sym}^d(C) \subset \text{Div}^d(C)$ the subset of all effective divisors in $\text{Div}^d(C)$. We set

$$W_d := \mu(\text{Sym}^d(C)) \subset \text{Jac}(C). \quad (9.58)$$

We regard the W_d only as sets here, even though they carry the structure of closed tropical subspaces. We already know $W_g = \text{Jac}(C)$ by the Jacobi inversion theorem. Riemann's theorem identifies W_{g-1} .

Theorem 9.4.1 (Tropical Riemann's theorem)

Let C be a compact smooth tropical curve (with some fixed base point p_0). Then

$$W_{g-1} + \kappa = \Theta, \quad (9.59)$$

where $\Theta \subset \text{Jac}(C)$ is the Theta divisor and κ is the Jacobi inversion constant.

The proof of this statement requires a couple of constructions which are of minor importance for the rest of this text. We will hence move some of the details to the exercises. Our first aim will be to construct special fundamental domains of the quotient map $V \rightarrow A$ for a principally polarized abelian variety A based on the theta divisor.

Exercise 9.4.2

Let A be a principally polarized abelian variety. Let $\tilde{\Theta} \subset V$ be the lift of the Theta divisor to V . Let Δ° be a connected component of $V \setminus \tilde{\Theta}$.

(a) Show that the restriction of the quotient map to

$$\Delta^\circ \rightarrow \text{Jac}(C) \tag{9.60}$$

is injective and has dense image $A \setminus \Theta$.

(b) Show that for generic vectors $w \in V$, the set

$$\Delta_w := \{x \in V : x - \epsilon w \in \Delta^\circ \text{ for sufficiently small } \epsilon\} \tag{9.61}$$

is a fundamental domain for $V \rightarrow A$ consisting of unions of relative interiors of the faces of the closure of Δ° . Moreover, this union is closed with respect to inclusions, i.e. for two faces τ, σ we have

$$\tau \subset \sigma, \tau^\square \subset \Delta_w \Rightarrow \sigma^\square \subset \Delta_w. \tag{9.62}$$

Next, it will be helpful to slightly extend our notion of tree cut. Let Γ be a compact metric graph (where we allow finite closed ends in this section). A *tree map* of Γ is a tuple (Γ', φ) where

- Γ' is a metric tree (with possibly open and closed finite ends),
- $\varphi : \Gamma' \rightarrow \Gamma$ is a bijective isometry.

This generalizes the notion of tree cut in two ways. First, we now also allow cutting at higher-valent vertices. Second, we cut infinitesimally near the cut point, such that e.g. for a two-valent cut point the new ends are open and closed, one each.

Let $c : \tilde{C} \rightarrow C$ be the universal cover of C . We can identify \tilde{C} with the set of all homotopy classes of paths with fixed endpoints, starting in p_0 .

Integration along the path then provides us with a map $\tilde{\mu} : \tilde{C} \rightarrow \Omega(C)^*$, a canonical lift of μ . For $w \in \Omega(C)^*$, let T_w be a connected component of $\tilde{\mu}^{-1}(\Delta_w)$. Let

$$c_w := c|_{T_w} : T_w \rightarrow C \tag{9.63}$$

denote the restriction of the covering map. (The choices of the open cell Δ° and the connected component T_w are not essential in the following, since different choices are isomorphic up to deck transformations of $\tilde{C} \rightarrow C$). The following diagram summarizes the setup.

$$\begin{array}{ccccc}
 & & c_w & & \\
 & & \curvearrowright & & \\
 T_w & \hookrightarrow & \tilde{C} & \xrightarrow{c} & C \\
 \downarrow & & \downarrow \tilde{\mu} & & \downarrow \mu \\
 \Delta_w & \hookrightarrow & \Omega(C)^* & \twoheadrightarrow & \text{Jac}(C)
 \end{array} \tag{9.64}$$

Lemma 9.4.3

For generic $w \in \Omega(C)^*$, the map $c_w : T_w \rightarrow C$ is a tree map.

Proof. By definition, T_w is a tree and c_w is an isometry. Since Δ_w is a fundamental domain for generic w , surjectivity follows as well. It remains to check injectivity. Indeed, let $p \neq p' \in T_w$ with $c_w(p) = c_w(p')$. The path connecting p to p' in T_w projects down to a closed loop $\gamma \in \pi_1(C)$. Since γ is non-zero, it contains a segment which is a simple loop, and hence by relabeling p and p' accordingly, we can assume that γ is a simple loop. In particular, we can assume that its homology class $[\gamma] \in H_1(C, \mathbf{Z})$ is non-zero. This implies $\tilde{\mu}(p') = \tilde{\mu}(p) + [\gamma] \neq \tilde{\mu}(p)$. But again, we may assume that Δ_w is a fundamental domain and therefore $\tilde{\mu}(p) = \tilde{\mu}(p')$, which is a contradiction. \square

Remark 9.4.4

Recall that for any $y \in \text{Jac}(C)$ (or, by abuse of notation, $y \in \Omega(C)^*$) we denote the shifted theta function resp. divisor by ϑ_y resp. Θ_y (and $\tilde{\Theta}_y$). Note that the statement and proof of the previous lemma (and exercises) continue to hold without any changes for Δ° being an open cell of $\tilde{\Theta}_y$ (for any y). Hence, for any $y \in \text{Jac}(C)$ and generic $w \in \Omega(C)^*$ we obtain a

associated tree map denoted by

$$c_w(y) : T_w(y) \rightarrow C. \quad (9.65)$$

We now want to use the construction of tree maps in order to compute the Jacobi inverse $D(y)$ more explicitly. Given a tree map $\varphi : \Gamma' \rightarrow \Gamma$ and a point $q \in \Gamma$, we can define the *cut divisor* by

$$E_\varphi := \sum_{p \in \Gamma'} (\text{val}_\Gamma(\varphi(p)) - \text{val}_{\Gamma'}(p)) \varphi(p) \text{ Div}(\Gamma). \quad (9.66)$$

Note that $\text{val}_\Gamma(\varphi(p)) - \text{val}_{\Gamma'}(p)$ can be interpreted as the number of open ends e in Γ' for which $\varphi(p)$ provides the “missing” endpoint of $\varphi(e)$. For the special tree maps $c_w(y) : T_w(y) \rightarrow C$, we use the notation $E_w(y) := E_{c_w(y)}$.

Proposition 9.4.5

For any $y \in \text{Jac}(C)$ and generic $w \in \Omega(C)^*$, we have

$$D(y) = E_w(y). \quad (9.67)$$

The proof uses reduction to the case of transverse intersection of W_1 and $\Theta(y)$. More precisely, we say that $\mu : C \rightarrow \text{Jac}(C)$ is *transverse* to Θ_y if $\mu^{-1}(\Theta_y)$ consists of finitely many 2-valent vertices of C .

Exercise 9.4.6

Show that for generic y , μ and Θ_y intersect transversely.

Proof of Proposition 9.4.5. First assume that μ and Θ_y are transverse. In this case, $c_w(y)$ in fact corresponds to the ordinary tree cut along the points in $\mu^{-1}(\Theta_y)$ (each point $p \in \mu^{-1}(\Theta_y)$ corresponds bijectively to a closed finite end of $T_w(y)$) and hence

$$E_w(y) = \sum_{p \in \mu^{-1}(\Theta_y)} p. \quad (9.68)$$

(The choice of w only affects which of two associated ends is closed). Let us compute $D(y)$. We first show $\text{supp}(D(y)) = \mu^{-1}(\Theta_y)$. The inclusion $\text{supp}(D(y)) \subset \mu^{-1}(\Theta_y)$ is clear by definition. For the inverse, if $p \in \mu^{-1}(\Theta_y)$ is not in $\text{supp}(D(y))$, this implies that $\vartheta_y \circ \mu$ is affine linear in a neighbourhood of p and hence the whole neighbourhood lies in $\mu^{-1}(\Theta_y)$. This

is a contradiction to transversality, hence the equality of set follows. But $\mu^{-1}(\Theta_y)$ consists of g points by the previous argument, and by Proposition 9.3.11 we know that $D(y)$ is effective and of degree g , hence each point must occur with multiplicity one and the claim follows.

Let us now prove the general case. We may choose generic w such that μ and $\Theta_{y-\epsilon w}$ intersect transversely for all sufficiently small $\epsilon > 0$. By continuity of the divisor and pull-back constructions, it holds

$$\lim_{\epsilon \rightarrow 0} D(y - \epsilon w) = D(y). \quad (9.69)$$

It remains to prove continuity for $E_w(y)$. In terms tree maps, one can check that $T_w(y - \epsilon w)$ is obtained from $T_w(y)$ by shortening all open ends a certain bit, while growing a corresponding number of small closed segments at the “opposite” point $p \in T_w(y)$. In fact, we argued before that this number is just $\text{val}_C(\varphi(p)) - \text{val}_{T_w(y)}(p)$. Since all points in $E_w(y - \epsilon w)$ appear with multiplicity one, this shows

$$\lim_{\epsilon \rightarrow 0} E_w(y - \epsilon w) = E_w(y). \quad (9.70)$$

This proves the claim. □

Exercise 9.4.7

Give the details of the proof of Equation (9.69).

The last important ingredient we need in order to prove Riemann’s theorem is a link to the notion of *moderator* that appeared in Section 8.4. Given a tree map $\varphi : \Gamma' \rightarrow \Gamma$ and a point $q \in \Gamma$, we orient all edges in Γ' towards the unique lift of q and using φ get an induced orientation \mathcal{O}_φ on γ (for a sufficiently fine graph structure). Repeating Definition 8.4.12, we define a divisor

$$K_+ := \sum_{x \in \Gamma} (\text{val}_+(x) - 1)x \quad (9.71)$$

where $\text{val}_+(x)$ denotes the number of outgoing edges at x . (In 8.4.12, we only considered the case when \mathcal{O} is acyclic).

Exercise 9.4.8

Show that $E_\varphi = K_+ + q$.

Lemma 9.4.9

For any $y \in \text{Jac}(C)$, generic $w \in \Omega(C)^*$, and $q \in \text{Jac}(C)$ such that $\mu(q) \notin \Theta_y$, the orientation induced by $c_w(\lambda)$ and q is acyclic.

Proof. Let $\tilde{q} \in T_w(y)$ be the unique lift of q and let $q' := \widetilde{m}u(\tilde{q}) \in \Delta_w$ be its image under the lifted Abel-Jacobi map. Since $\mu(q) \notin \Theta_y$, the point q' sits in fact in the interior, i.e. $q' \in \Delta^\circ$. Note that the orientation on $T_w(y)$ is obviously acyclic, since all edges are oriented towards \tilde{q} . Hence an oriented cycle in C must lift to an oriented path γ_1 in $T_w(y)$, with “open” starting point p_1 and “closed” endpoint p_2 . Let γ_2 the unique (oriented) path from p_2 to \tilde{q} . Let us now change w to a generic vector u close to $-w$. It follows that the new fundamental domain Δ_u contains p_1 , but not p_2 . We choose as $T_u(y)$ the unique connected component of $\widetilde{m}u^{-1}(\Delta_u)$ which contains \tilde{q} . Let γ_3 be the unique path in $T_u(y)$ connecting \tilde{q} to the unique point in the same fiber as p_1 , say p_3 . We denote by γ'_i projected paths in C . Since $\tilde{\mu}(p_1) = \tilde{\mu}(p_3)$, the homology class of $\gamma'_3 \circ \gamma'_2 \circ \gamma'_1$ is zero, and hence

$$[\gamma'_3 \circ \gamma'_2] = -[\gamma'_1] \in H_1(C, \mathbf{Z}). \tag{9.72}$$

We now check on the chain level that this is impossible. Let e be an edge used by γ'_1 such $\tilde{\mu}$ maps its interior to $\Delta_w \cap \Delta_u$, but its endpoint to $\Delta_w \setminus \Delta_u$. Since γ'_1 is an oriented and hence simple path, this edge occurs with non-zero coefficient in the chain representation of $[\gamma'_1]$. On the other hand, this edge can neither be used by γ'_2 (since $\gamma'_2 \circ \gamma'_1$ is still oriented and hence simple) nor γ'_3 (since e corresponds to an open end in $T_v(y)$). Hence the chain representations of $[\gamma'_3 \circ \gamma'_2]$ and $[\gamma'_1]$ are different and the claim follows. \square

We will now formulate the main statement of the section, of which Riemann’s theorem is an easy corollary. Here, the *support* of a linear system is defined to be

$$\text{supp } |D| := \bigcup_{D' \in |D|} \text{supp } D'. \tag{9.73}$$

Theorem 9.4.10

Let D be a divisor of degree g . Set $y = \mu(D) \in \text{Jac}(C)$. Then

$$\text{supp } |D| = \mu|_C^{-1}(\Theta_{y+\kappa}). \tag{9.74}$$

Let us start the proof with the following lemma. Let C be a smooth tropical curve. Let $\Gamma \subset C$ by a proper subgraph. By abuse of notation, we denote by $\partial\Gamma$ the divisor containing each boundary point of Γ with multiplicity 1.

Lemma 9.4.11

Let $D = \sum a_p p$ be a divisor on C . Let $\varphi : \Gamma' \rightarrow \Gamma$ be a tree map such that $D - \partial\Gamma - E_\varphi$ is effective. Then

$$\Gamma \subset \text{supp } |D|. \tag{9.75}$$

Proof. We proceed by induction. We must distinguish two cases. First, let us assume there exists a vertex p in $\partial\Gamma$ such

$$\text{val}_\Gamma(p) > a_p \geq E_\varphi + 1. \tag{9.76}$$

This implies that the lift p' of p in Γ' is connected to at least two edges. Hence p' is not an end in Γ' and $\Gamma' \setminus \{p'\}$ is disconnected. Restricting φ to the closures of the connected components (and Γ to the corresponding image), the assumptions of the lemma still hold. Hence we can apply the induction hypothesis and the claim follows. The case occurs when all boundary points p satisfy $a_p \geq \text{val}_\Gamma(p)$. In this case, we can deform D to an rationally equivalent divisor by moving one point along all the edges adjacent to all boundary points. The corresponding rational function is of the form

$$f(x) = \max\{-\text{dist}(\overline{C \setminus \Gamma}, x), b\} \tag{9.77}$$

for some (not too) negative constant b . Let Γ_b be the subgraph of Γ where $f(x) = b$. If we choose $-b$ equal to the minimal distance of an interior vertex of Γ to the boundary, Γ_b is a “smaller” graph than Γ and $D + \text{div}(f)$ still satisfies the assumption of the lemma (just use the shrunked tree map). Again, by induction the claim follows. \square

Proof of Theorem 9.4.10. We set $y := \mu(D)$. Let us start with the inclusion $\text{supp } |D| \subset \mu|_C^{-1}(\Theta_{y+\kappa})$. Choose $q \notin \mu|_C^{-1}(\Theta_{y+\kappa})$. We have to show $q \notin \text{supp } |D|$ or equivalently $|D - q| = \emptyset$. By the Jacobi inversion theorem, we have $D \sim D(y + \kappa)$. Let K_+ the divisor associated to a tree map $c_w(y + \kappa)$ and the point $q \in C$. By Proposition 9.4.5 and Exercise 9.4.8 we conclude

$$|D - q| = |D(y + \kappa) - q| = |E_w(y + \kappa) - q| = |K_+|. \tag{9.78}$$

Since $\mu(q) \notin \Theta_{y+\kappa}$, we can apply Lemma 9.4.9 which says that the corresponding orientation on C is acyclic and hence K_+ is a moderator in the sense of Definition 8.4.12. Then Lemma 8.4.14 implies $|K_+| = \emptyset$, and we are done.

Let us now prove the second inclusion $\text{supp } |D| \supset \mu|_C^{-1}(\Theta_{y+\kappa})$. Let $q \in \mu|_C^{-1}(\Theta_{y+\kappa})$. Choose a tree map $c := c_w(y + \kappa)$. Since $\mu(q) \in \Theta_{y+\kappa}$, the corresponding image point in Δ_w sits in the interior of some boundary face of Δ_w , say σ . Let Γ' be the connected component of $c^{-1}(\sigma)$ containing q . Let $\Gamma = c(\Gamma') \subset C$. Obviously, $c|_{\Gamma'} : \Gamma' \rightarrow \Gamma$ is again a tree map. If Γ is not yet closed in C , we can add boundary points to Γ and Γ' (just compactify one of the open ends mapping to the open end on Γ), ending up with a compactified tree map $\varphi : \Gamma' \rightarrow \Gamma$. By Lemma 9.4.11 we can finish the proof by showing that $E_c - \partial\Gamma - E_\varphi$ is effective. Since φ is (essentially) just a restriction of c to a smaller domain, $E_c - E_\varphi$ is obviously effective. For a boundary point $p \in \partial\Gamma$, let x be the image point in σ . By the balancing condition, there must be edges of $\tilde{\mu}(\tilde{C})$ adjacent to x but outside of Δ_w . Such edges provide the necessary extra contribution of p in E_c . \square

Proof of Riemann's Theorem 9.4.1. Let D be a divisor of degree $g - 1$. Set $y = \mu(D) = \mu(D + p_0)$. The sequence of equivalences (with Theorem 9.4.10 at third place)

$$y \in W_{g-1} \Leftrightarrow |D| \neq \emptyset \Leftrightarrow p_0 \in \text{supp } |D + p_0| \Leftrightarrow 0 \in \Theta_{y+\kappa} \Leftrightarrow y + \kappa \in \Theta \quad (9.79)$$

proves the claim. In the last step, we use $\Theta_y = \Theta_{-y}$ since $\vartheta(y) = \vartheta(-y)$. \square

Bibliography

- [AR10] Lars Allermann and Johannes Rau. *First steps in tropical intersection theory*. Math. Z. 264.3 (2010), pp. 633–670. arXiv: 0709.3705.
- [Ber76] David N. Bernstein. *The number of roots of a system of equations*. Functional Analysis and its Applications 9 (1976), pp. 183–185.
- [BM07] Erwan Brugallé and Grigory Mikhalkin. *Enumeration of curves via floor diagrams*. Comptes Rendus. Mathématique. Académie des Sciences, Paris 345.6 (2007), pp. 329–334.
- [BM09] Erwan Brugallé and Grigory Mikhalkin. *Floor decompositions of tropical curves: the planar case*. Proceedings of the 15th Gökova geometry-topology conference, Gökova, Turkey, May 26–31, 2008. Cambridge, MA: International Press, 2009, pp. 64–90. arXiv: 0812.3354.
- [BS11] Erwan Brugallé and Kristin M. Shaw. *Obstructions to approximating tropical curves in surfaces via intersection theory*. Preprint (2011). arXiv: 1110.0533.
- [BG09] Winfried Bruns and Joseph Gubeladze. *Polytopes, rings, and K-theory*. New York, NY: Springer, 2009, pp. xiii + 461. DOI: 10.1007/b105283.
- [Cho92] Christian Choffrut. *Rational relations and rational series*. Theor. Comput. Sci. 98.1 (1992), pp. 5–13.
- [CLS11] David A. Cox, John B. Little, and Henry K. Schenck. *Toric varieties*. Vol. 124. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2011, pp. xxiv+841.
- [Dan78] V. I. Danilov. *Geometry of toric varieties*. Russ. Math. Surv. 33.2 (1978), pp. 97–154.

- [DM69] Pierre Deligne and D. Mumford. *The irreducibility of the space of curves of a given genus*. Publications Mathématiques 36 (1969), pp. 75–109.
- [EKL06] Manfred Einsiedler, Mikhail Kapranov, and Douglas Lind. *Non-archimedean amoebas and tropical varieties*. J. Reine Angew. Math. 601 (2006), pp. 139–157. arXiv: math/0408311.
- [Ewa96] Günter Ewald. *Combinatorial convexity and algebraic geometry*. Vol. 168. New York, NY: Springer, 1996, pp. xiv + 372.
- [EI06] Günter Ewald and Masa-Nori Ishida. *Completion of real fans and Zariski-Riemann spaces*. Tohoku Math. J. (2) 58.2 (2006), pp. 189–218.
- [Ful93] William Fulton. *Introduction to toric varieties. The 1989 William H. Roever lectures in geometry*. Princeton, NJ: Princeton University Press, 1993.
- [Ful98] William Fulton. *Intersection theory*. 2nd ed. Berlin: Springer, 1998.
- [FS97] William Fulton and Bernd Sturmfels. *Intersection theory on toric varieties*. Topology 36.2 (1997), pp. 335–353. arXiv: alg-geom/9403002.
- [GKZ08] I. M. Gel'fand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, resultants, and multidimensional determinants*. Reprint of the 1994 edition. Boston, MA: Birkhäuser, 2008, pp. x + 523.
- [IM12] I. Itenberg and G. Mikhalkin. *Geometry in the tropical limit*. Mathematische Semesterberichte 59.1 (2012), pp. 57–73.
- [Ite93] Ilia Itenberg. *Contre-exemples à la conjecture de Ragsdale*. French. Comptes Rendus de l'Académie des Sciences. Série I 317.3 (1993), pp. 277–282.
- [IV96] Ilia Itenberg and Oleg Viro. *Patchworking algebraic curves disproves the Ragsdale conjecture*. The Mathematical Intelligencer 18.4 (1996), pp. 19–28.
- [Kat09] Eric Katz. *A tropical toolkit*. Expo. Math. 27.1 (2009), pp. 1–36. arXiv: math/0610878.
- [Kat12] Eric Katz. *Tropical intersection theory from toric varieties*. Collect. Math. 63.1 (2012), pp. 29–44. arXiv: 0907.2488.

- [KT01] Allen Knutson and Terence Tao. *Honeycombs and sums of Hermitian matrices*. Notices of the American Mathematical Society 48.2 (2001), pp. 175–186.
- [Kou76] Anatoli G. Kouchnirenko. *Polyèdres de Newton et nombres de Milnor*. French. Invent. Math. 32 (1976), pp. 1–31.
- [Mik01] Grigory Mikhalkin. *Amoebas of algebraic varieties*. Preprint (2001). arXiv: math/0108225.
- [Mik04] Grigory Mikhalkin. *Decomposition into pairs-of-pants for complex algebraic hypersurfaces*. Topology 43.5 (2004), pp. 1035–1065. arXiv: math/0205011.
- [Mik05] Grigory Mikhalkin. *Enumerative tropical algebraic geometry in \mathbb{R}^2* . J. Am. Math. Soc. 18.2 (2005), pp. 313–377. arXiv: math/0312530.
- [Mik06] Grigory Mikhalkin. *Tropical geometry and its applications*. Proceedings of the international congress of mathematicians (ICM), Madrid, Spain, August 22–30, 2006. Volume II: Invited lectures. Zürich: European Mathematical Society (EMS), 2006, pp. 827–852. arXiv: math/0601041.
- [Pas08] Mikael Passare. *How to compute $\sum 1/n^2$ by solving triangles*. American Mathematical Monthly 115.8 (2008), pp. 745–752.
- [Pin98] Jean-Eric Pin. *Tropical semirings. Idempotency*. Based on a workshop, Bristol, UK, October 3–7, 1994. Cambridge: Cambridge University Press, 1998, pp. 50–69.
- [Rag06] V. Ragsdale. *On the arrangement of the real branches of plane algebraic curves*. American Journal of Mathematics 28 (1906), pp. 377–404.
- [Rau09] Johannes Rau. *Tropical intersection theory and gravitational descendants*. PhD-Thesis. Technische Universität Kaiserslautern, 2009. URL: <http://kluedo.ub.uni-kl.de/volltexte/2009/2370/>.
- [Rau16] Johannes Rau. *Intersections on tropical moduli spaces*. Rocky Mt. J. Math. 46.2 (2016), pp. 581–662. arXiv: 0812.3678.

- [RST05] Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald. *First steps in tropical geometry*. Idempotent mathematics and mathematical physics, Proceedings of the international workshop, Vienna, Austria, February 3–10, 2003. Providence, RI: American Mathematical Society, 2005, pp. 289–317.
- [Sim88] Imre Simon. *Recognizable sets with multiplicities in the tropical semiring*. Mathematical Foundations of Computer Science. Springer Berlin Heidelberg, 1988, pp. 107–120.
- [SS04] David Speyer and Bernd Sturmfels. *The tropical Grassmannian*. Adv. Geom. 4.3 (2004), pp. 389–411. arXiv: math/0304218.
- [Stu02] Bernd Sturmfels. *Solving systems of polynomial equations*. CBMS Regional Conference Series 97. Providence, RI: American Mathematical Society, 2002.
- [Vig09] Magnus Dehli Vigeland. *The group law on a tropical elliptic curve*. Mathematica Scandinavica 104.2 (2009), pp. 188–204.
- [Vir80] O. Ja. Viro. *Curves of degree 7, curves of degree 8, and the Ragsdale conjecture*. Soviet Mathematics. Doklady 22 (1980), pp. 566–570.
- [Vir79] O. Ya. Viro. *Construction of M-surfaces*. Russian. Funktsional’nyi Analiz i Ego Prilozheniya 13.3 (1979), pp. 71–72.
- [Vir06] Oleg Viro. *Patchworking real algebraic varieties*. Preprint (2006). arXiv: math/0611382.
- [Vir10] Oleg Viro. *Hyperfields for Tropical Geometry I: Hyperfields and dequantization*. Preprint (2010). arXiv: 1006.3034.
- [Zie95] Günter M. Ziegler. *Lectures on polytopes*. Vol. 152. Berlin: Springer-Verlag, 1995, pp. ix + 370.