Cohomology, Volumes, and Theta Characteristics on Metric Graphs

Junaid Hasan

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Farbod Shokrieh, Chair

Gaku Liu

Cynthia Vinzant

Program Authorized to Offer Degree:

Mathematics

University of Washington

Abstract

Cohomology, Volumes, and Theta Characteristics on Metric Graphs

Junaid Hasan

Chair of the Supervisory Committee:
Professor Farbod Shokrieh
Department of Mathematics

This thesis develops tropical and combinatorial analogues of several classical constructions in algebraic geometry, with metric graphs as the central objects of study.

In the first part, we establish a tropical de Rham theorem for rational polyhedral spaces: the natural map from Lagerberg superform cohomology to the singular tropical cohomology of Itenberg–Katzarkov–Mikhalkin–Zharkov is an isomorphism of bigraded rings, compatible with Poincaré duality on both sides. This extends prior work of Jell–Shaw–Smacka, who proved the isomorphism as vector spaces, by equipping singular tropical cohomology with an explicit ring structure via cup and cap products.

In the second part, we apply this machinery to principally polarized tropical abelian varieties. A polarization $\phi\colon N \to \Lambda^\vee$ determines a canonical constant $(1,1)$-superform ν_ϕ whose cohomology class is the tropical analogue of the first Chern class of a polarization line bundle. For a tropical Jacobian $\mathrm{Jac}(\Gamma)$, the pullback of ν_ϕ via the Abel–Jacobi map recovers the Zhang canonical measure on Γ. We also prove the tropical Poincaré formula $[\widetilde{W}_d] = [\Theta]^{g-d}/(g-d)!$ in $H_{S,d,d}(\mathrm{Jac}(\Gamma))$ and compute the volumes $\int_{\widetilde{W}_d} \frac{1}{d!}\nu_\phi^{\wedge d} = \binom{g}{d}$.

In the third part, we study the Voronoi polytope $\mathrm{Vor}(G) \subset H_1(G,\mathbb{R})$ of a metric graph and its zonotopal subdivision by spanning trees. We solve the cell membership problem

for this subdivision by introducing (f,q)-optimal spanning trees, and apply the solution to produce, for each 2-torsion class $\gamma \in H_1(G,\mathbb{Z}/2\mathbb{Z})$, a break divisor $B(f_\gamma)$ whose class $[B(f_\gamma)-q]$ is Zharkov's tropical theta characteristic $\mathcal{K}_\gamma \in \mathrm{Pic}^{g-1}(\Gamma)$.

In the fourth part, we characterize break divisors as the unique minimizers of a quadratic energy functional $E_F(D) = [D-\tilde{\mu}_F]^T L^\dagger [D-\tilde{\mu}_F]$ defined via the Moore–Penrose pseudoinverse of the graph Laplacian, under a positive resistance curvature hypothesis.

Finally, we present joint work on integral aspects of Fourier duality for abelian varieties. Using Pappas's integral Grothendieck–Riemann–Roch theorem, we construct integral lifts F, F^t of the classical Fourier–Beauville transforms on Chow groups and prove that Beauville's decomposition, inversion formula, and $\mathfrak{sl}_2$-action all hold over $\Lambda = \mathbb{Z}[1/(2g+d+1)!]$.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Farbod Shokrieh, for his guidance, generosity, and mathematical insight throughout my doctoral studies. I thank my committee members Gaku Liu and Cynthia Vinzant for their time and helpful feedback.

Chapters 2 and 3 are based on joint work in progress with Caelan Ritter and Farbod Shokrieh. Chapter 4 is based on joint work in progress with Amzi Jeffs and Farbod Shokrieh. Chapter 5 is based on joint work in progress with Matt Baker, Karel Devriendt and Farbod Shokrieh. Chapter 6 is based on joint work with Hazem Hassan, Milton Lin, Marcella Manivel, Lily McBeath, and Ben Moonen, which arose at the Arizona Winter School 2024.

I am grateful to all of these collaborators for many stimulating conversations.

This work was supported in part by NSF CAREER grant DMS-2044564.

Chapter 1
INTRODUCTION

1.1 Overview

This thesis explores the geometry, topology, and combinatorics of tropical varieties and metric graphs. In particular, we develop and apply analogues of classical algebraic tools—such as de Rham cohomology, Deligne pairings, theta characteristics, and Fourier–Mukai transforms—in the tropical and combinatorial settings.

In Chapter 2, we establish a tropical analogue of the de Rham theorem. Following the introduction of superforms by Lagerberg [Lag12] and their extension to rational polyhedral spaces [Gub16, JRS18], and the singular tropical cohomology introduced by Itenberg–Katzarkov–Mikhalkin–Zharkov [IKMZ19], we construct a natural isomorphism of rings between the Dolbeault superform cohomology and the singular tropical cohomology of a rational polyhedral space.

In Chapter 3, we study principally polarized tropical abelian varieties. Drawing motivation from the canonical Arakelov $(1,1)$-form on a complex Jacobian, we construct a canonical $(1,1)$-superform on a principally polarized tropical abelian variety. We show that when this variety is the Jacobian of a metric graph, the restriction of this form along the Abel–Jacobi map recovers the canonical Zhang measure on the graph.

In Chapter 4, we study break divisors and tropical theta characteristics on metric graphs. The central objects are the Voronoi polytope $\mathrm{Vor}(G) \subset H_1(G, \mathbb{R})$ and its zonotopal subdivision indexed by spanning trees of G. We show that cell membership in this subdivision is equivalent to feasibility of the Bellman–Ford shortest-path relaxation, and use the resulting algorithm to produce, for each 2-torsion class $\gamma \in H_1(G, \mathbb{Z}/2\mathbb{Z})$, a break divisor $B(f_\gamma)$ whose class $[B(f_\gamma) - q]$ is Zharkov's tropical theta characteristic $\mathcal{K}_\gamma \in \mathrm{Pic}^{g-1}(\Gamma)$.

In Chapter 5, we reexamine break divisors—the unique effective representatives of degree-g divisor classes on a metric graph of genus g [MZ08, ABKS14]. We characterize these break divisors analytically as the unique minimizers of a natural energy functional defined using the Moore–Penrose pseudoinverse of the graph Laplacian, tying together discrete divisors with continuous potential theory.

Finally, Chapter 6 presents joint work on integral aspects of Fourier duality for abelian varieties. Using Pappas's integral Grothendieck–Riemann–Roch theorem, we construct integral lifts F, F^t of the classical Fourier–Beauville transforms on Chow groups of abelian schemes, prove an integral inversion formula, and show that Beauville's decomposition and $\mathfrak{sl}_2$-action hold over $\Lambda = \mathbb{Z}[1/(2g+d+1)!]$.

Chapter 2

TROPICAL DE RHAM THEORY

2.1 Introduction

Tropical varieties arise as degenerations of families of algebraic varieties; while tropicalization discards much information, it retains certain enumerative data of the original family. Alternatively, tropical varieties may be understood as skeleta of Berkovich spaces.

Two cohomology theories attach naturally to a rational polyhedral space X. The first, *superform cohomology* $H_D^{p,q}(X)$, originates with Lagerberg's theory of superforms on $\mathbb{R}^n$ [Lag12] and was extended to Berkovich spaces by Chambert-Loir–Ducros [CLD12] and to rational polyhedral spaces by Gubler [Gub16] and Jell [JRS18]. The complex of superforms is a bigraded bidifferential algebra, so $H_D^{\bullet,\bullet}(X)$ inherits a natural ring structure via the wedge product, together with a Poincaré duality map constructed by Jell–Rau–Shaw [JRS18].

The second theory, *singular tropical cohomology* $H_S^{p,q}(X)$, was introduced by Itenberg–Katzarkov–Mikhalkin–Zharkov [IKMZ19] and has been studied from a sheaf-theoretic perspective by Gross–Shokrieh [GS23a], who use the derived category to define cup products, cap products, and Verdier duality. Despite these developments, no explicit ring structure or Poincaré duality map on $H_S^{\bullet,\bullet}(X)$ appeared in the literature prior to this work.

Jell–Shaw–Smacka [JSS19] proved that the two theories are abstractly isomorphic as real vector spaces:

$$H_D^{p,q}(X) \cong H_S^{p,q}(X).$$

The main contribution of this chapter is to upgrade this to an *isomorphism of bigraded rings*, realized by an explicit chain-level map $\widetilde{\mathrm{dR}}$ analogous to the classical de Rham map. Along the way, we define cup and cap products on singular tropical (co)homology in analogy with the classical constructions on singular cohomology of topological spaces, and we show that

$\widetilde{\mathrm{dR}}$ is compatible with the Poincaré duality maps on both sides.

The precise statements are as follows.

Theorem 2.1 (de Rham isomorphism, Theorem 2.12)**.** *For any rational polyhedral space X, the de Rham map induces an isomorphism of bigraded rings*

$$\widetilde{\mathrm{dR}}\colon H_D^{\bullet,\bullet}(X) \xrightarrow{\sim} H_S^{\bullet,\bullet}(X).$$

Theorem 2.2 (Poincaré duality compatibility, Theorem 2.15)**.** *The de Rham map $\widetilde{\mathrm{dR}}$ commutes with the Poincaré duality maps on $H_D^{\bullet,\bullet}(X)$ and $H_S^{\bullet,\bullet}(X)$.*

Outline. Section 2.2 introduces rational polyhedral spaces, the sheaf of locally integral affine functions, star neighborhoods, and the face structure. Section 2.3 develops the theory of multi-(co)vectors and superforms, culminating in the superform cochain complex $\mathcal{A}^{p,\bullet}(X)$ and its cohomology $H_D^{p,q}(X)$. Section 2.4 defines singular tropical (co)chains and cohomology, including the sign convention in the boundary map that is essential for compatibility with superforms. Section 2.5 introduces the cup and cap products on singular tropical cohomology; these definitions are new to the literature. Section 2.6 constructs the tropical de Rham map dR as an explicit integration pairing and proves that it is a quasi-isomorphism. Section 2.7 proves the two main theorems.

Notation. We establish the notation used throughout this chapter. We use the assignment symbol $:=$ to define quantities. Calligraphic letters such as $\mathcal{C}$ and $\mathcal{P}$ will denote collections of polyhedra. The symbol $\preceq$ denotes the face relation for polyhedra. We will use multi-indices, typically denoted I, to write differential forms compactly.

2.2 Rational Polyhedral Spaces

We introduce rational polyhedral spaces following [GS23a, §2.1] and [JRS18, §2]. For ease of exposition, we restrict throughout to rational polyhedral spaces *without boundary*, i.e., those whose charts land in $\mathbb{R}^n$ rather than $\mathbb{T}^n = [-\infty, \infty)^n$.

2.2.1 Rational polyhedral sets

A polyhedron in $\mathbb{R}^n$ is *rational* if it can be written as a finite intersection of half-spaces whose normal vectors lie in $\mathbb{Z}^n$. A polyhedral complex $\mathcal{P}$ is *rational* if all of its faces are rational. We write $\mathcal{P}^{(k)}$ for the set of k-dimensional faces of $\mathcal{P}$, and we call a face maximal with respect to inclusion a *facet*. The *support* of $\mathcal{P}$ in $\mathbb{R}^n$ is $|\mathcal{P}| := \bigcup_{\sigma \in \mathcal{P}} \sigma$.

A subset $X \subset \mathbb{R}^n$ is a *polyhedral set* if there exists a polyhedral complex $\mathcal{C}$, called a *face structure* for X, such that $X = |\mathcal{C}|$; it is *rational* if it admits a rational face structure.

Given polyhedra σ and τ, we write $\sigma \preceq \tau$ when σ is a face of τ. The *relative interior* of σ, denoted $\operatorname{relint}(\sigma)$, is the set of points of σ not contained in any proper face. Given a face structure $\mathcal{C}$ on X and a face $\sigma \in \mathcal{C}$, the *star* of σ is the open subset

$$\operatorname{star}(\sigma) := \bigcup_{\substack{\tau \in \mathcal{C} \\ \sigma \preceq \tau}} \operatorname{relint}(\tau).$$

For each $x \in |\mathcal{C}|$, there is a unique face $\sigma \in \mathcal{C}$ with $x \in \operatorname{relint}(\sigma)$. An open neighborhood of x contained in $\operatorname{star}(\sigma)$ that admits a straight-line deformation retraction to x is called a *star neighborhood* of x. In particular, $\operatorname{star}(\sigma)$ is itself a star neighborhood of every point in $\operatorname{relint}(\sigma)$, and every $x \in |\mathcal{C}|$ admits a neighborhood basis of star neighborhoods.

A function $f \colon X \to \mathbb{R}$ is *integral affine* if it has the form $x \mapsto \langle m, x \rangle + a$ for some $m \in (\mathbb{Z}^n)^\vee$ and $a \in \mathbb{R}$. The set of integral affine functions forms a presheaf of $\underline{\mathbb{Z}}$-modules on X; as [GS23a, Example 2.1] shows, it fails to be a sheaf. We let $\mathcal{A}\!\mathit{ff}_X$ denote its sheafification. Concretely, sections of $\mathcal{A}\!\mathit{ff}_X$ are *locally integral affine functions*, i.e., continuous functions $X \to \mathbb{R}$ that are integral affine on an open neighborhood of each point. On a polyhedral set with a fixed face structure, the restriction of any locally integral affine function to a star neighborhood is in factintegral affine. We also consider the sheaf $\mathcal{A}\!\mathit{ff}_{X,\mathbb{R}} := \mathcal{A}\!\mathit{ff}_X \otimes_{\underline{\mathbb{Z}}} \underline{\mathbb{R}}$ of *locally affine functions*, whose sections over star neighborhoods take the form $x \mapsto \langle m, x \rangle + a$ for $m \in (\mathbb{R}^n)^\vee$.

2.2.2 Rational polyhedral spaces

A *rational polyhedral space* is a paracompact, second-countable, Hausdorff topological space X endowed with a sheaf $\mathscr{Aff}_X$ of continuous $\mathbb{R}$-valued functions that is locally modeled by rational polyhedral sets. More precisely, for every $x \in X$, there exist $n \in \mathbb{N}$, a rational polyhedral set $Y \subset \mathbb{R}^n$, an open neighborhood U of x, and an open continuous injection $\psi \colon U \to Y$ inducing an isomorphism $\mathscr{Aff}_{\psi(U)} \xrightarrow{\sim} \psi_* \mathscr{Aff}_X|_U$ via pullback. We call this data a *chart*.

An equivalent definition involves choosing a maximal atlas of charts whose transition maps are restrictions of integral affine maps $\mathbb{R}^m \to \mathbb{R}^n$; see [JRS18, Definition 2.1]. In what follows, we frequently use this perspective to conflate a chart with its image and work in local coordinates.

A *polyhedron* in X is a closed subset contained in the domain of a single chart ψ such that $\psi(\sigma)$ is a polyhedron in $\mathbb{R}^n$. One defines polyhedral complexes in X accordingly; a *face structure* $\mathcal{C}$ on X is always taken to be rational (see [GS23a, Definition 2.4] for details). A *star neighborhood* of $x \in X$ is the preimage of a star neighborhood in $\mathbb{R}^n$ under a face-structure-preserving chart; a chart that is a star neighborhood of some point is called a *star chart*. It follows that X has a basis of star charts.

2.3 Multi-(co)vectors and Superforms

2.3.1 Multi-(co)vectors

Fix a polyhedron σ in $\mathbb{R}^n$, and let $\mathbb{L}(\sigma)$ denote the tangent space to σ at any point of $\operatorname{relint}(\sigma)$, regarded as a subspace of $\mathbb{R}^n$. Following [JRS18, Definition 2.3], we define

$$\mathbf{F}_p(\sigma) := \text{subspace of } \textstyle\bigwedge^p \mathbb{R}^n \text{ generated by } \bigwedge^p \mathbb{L}(\tau) \text{ for each facet } \tau \succeq \sigma.$$

Its dual space $\mathbf{F}^p(\sigma) := \operatorname{Hom}_{\mathbb{R}}(\mathbf{F}_p(\sigma), \mathbb{R})$ is naturally a quotient of $\bigwedge^p (\mathbb{R}^n)^\vee$, and the wedge product on $\bigoplus_p \mathbf{F}^p(\sigma)$ inherited via this quotient is well-defined.

We write $v \otimes \alpha \mapsto \langle v, \alpha \rangle$ for the evaluation pairing $\mathbf{F}_p(\sigma) \otimes \mathbf{F}^p(\sigma) \to \mathbb{R}$.

By interior multiplication, one defines a contraction

$$\mathbf{F}_{p+p'}(\sigma) \otimes \mathbf{F}^{p}(\sigma) \longrightarrow \mathbf{F}_{p'}(\sigma), \qquad v \otimes \alpha \longmapsto v \frown \alpha,$$

characterized by the adjointness property

$$\langle v \frown \alpha,\, \beta \rangle = \langle v,\, \alpha \wedge \beta \rangle \qquad \text{for all } \beta \in \mathbf{F}^{p'}(\sigma). \tag{2.1}$$

If $\sigma \preceq \tau$, there is a natural inclusion $\mu_{\sigma,\tau}\colon \mathbf{F}_p(\tau) \to \mathbf{F}_p(\sigma)$ and a dual restriction map $\mu^*_{\sigma,\tau}\colon \mathbf{F}^p(\sigma) \to \mathbf{F}^p(\tau)$. These satisfy the transitivity condition: for $\sigma \preceq \tau \preceq \upsilon$,

$$\mu_{\sigma,\tau} \circ \mu_{\tau,\upsilon} = \mu_{\sigma,\upsilon}. \tag{2.2}$$

If X is a rational polyhedral space with face structure $\mathcal{C}$, the constructions above are independent of the choice of chart for σ, and we call elements of $\bigoplus_{\sigma\in\mathcal{C}} \mathbf{F}_p(\sigma)$ (respectively $\bigoplus_{\sigma\in\mathcal{C}} \mathbf{F}^p(\sigma)$) *multi-(co)vectors.*

2.3.2 Tropical forms

Let X be a rational polyhedral space with face structure $\mathcal{C}$. The inclusion of locally constant functions into locally affine functions gives a natural injection $\underline{\mathbb{R}} \hookrightarrow \mathcal{A}\!f\!f_{X,\mathbb{R}}$. The *cotangent sheaf* is the quotient

$$\Omega^1_X := \mathcal{A}\!f\!f_{X,\mathbb{R}}/\underline{\mathbb{R}};$$

see [GS23a, §2.2], where Ω^1_X denotes the integral cotangent sheaf $\mathcal{A}\!f\!f_X/\underline{\mathbb{Z}}$ instead. Given a section $f \in \mathcal{A}\!f\!f_{X,\mathbb{R}}(U)$, we denote its image in $\Omega^1_X(U)$ by df. Over a star chart with coordinates $x_1, \dots, x_n$, every locally affine function is affine, so $df = \sum_i m_i\, dx_i$, and $dx_1, \dots, dx_n$ generate $\Omega^1_X(U)$. Relations among these generators are given by the affine constraints on U: if U is contained in a hyperplane $\sum_i m_i x_i = a$, then $\sum_i m_i\, dx_i = 0$ in $\Omega^1_X(U)$. In particular, $\Omega^1_X(U)$ is a real vector space of dimension equal to the affine span of U.

The stalk of Ω^1_X at any point $x \in \operatorname{relint}(\sigma)$ is independent of the star neighborhood chosen; we abbreviate $\Omega^1_X(\sigma) := \Omega^1_X(\operatorname{star}(\sigma))$. If $\sigma \preceq \tau$, restriction gives a canonical surjection $\Omega^1_X(\sigma) \twoheadrightarrow \Omega^1_X(\tau)$, so Ω^1_X is a constructible sheaf with respect to $\mathcal{C}$.

We define the p-th exterior power $\bigwedge^p \Omega^1_X$ as the sheafification of the presheaf $U \mapsto \bigwedge^p \Omega^1_X(U)$ on star charts. As [GS23a, Example 2.9] shows, this sheaf can be nonzero for $p > \dim X$. To remedy this,let $X^\circ := \bigcup_{\tau \text{ facet}} \operatorname{relint}(\tau)$ with inclusion $\iota\colon X^\circ \hookrightarrow X$, and define the sheaf of *tropical p-forms* as the image

$$\Omega^p_X := \operatorname{im}\left(\textstyle\bigwedge^p \Omega^1_X \longrightarrow \iota_*\left(\textstyle\bigwedge^p \Omega^1_X|_{X^\circ}\right)\right), \qquad \Omega^0_X := \underline{\mathbb{R}}.$$

This definition is equivalent to [GS23a, Definition 2.7], and one sees that Ω^p_X is constructible.

Equivalently, $\Omega^p_X \cong \bigwedge^p \Omega^1_X / K$, where K is the kernel of the morphism above. For a face $\sigma \in \mathcal{C}$, a form $\omega \in \bigwedge^p \Omega^1_X(\sigma)$ lies in $K(\sigma)$ if and only if $\langle v, \omega \rangle = 0$ for all $v \in \mathbf{F}_p(\sigma)$. In particular, the contraction descends to a well-defined pairing $\mathbf{F}_p(\sigma) \otimes \Omega^p_X(\sigma) \to \mathbb{R}$.

2.3.3 Superforms and their cohomology

Following the notation of [JSS19], a (p,q)-*superform* on an open set $U \subset \mathbb{R}^n$ is an element of

$$\mathcal{A}^{p,q}(U) := \mathcal{A}^p(U) \otimes_{C^\infty(U)} \mathcal{A}^q(U) \cong C^\infty(U) \otimes_{\mathbb{R}} \textstyle\bigwedge^p (\mathbb{Z}^n)^\vee \otimes_{\mathbb{R}} \textstyle\bigwedge^q (\mathbb{Z}^n)^\vee,$$

where $\mathcal{A}^p(U)$ denotes smooth real p-forms on U. Writing $x_1, \dots, x_n$ for standard coordinates and using multi-index notation $dx_I := dx_{i_1} \wedge \cdots \wedge dx_{i_p}$ for $I = \{i_1 < \cdots < i_p\}$, a superform may be written as

$$\omega = \sum_{|I|=p,\, |J|=q} \omega_{IJ}\, d'x_I \wedge d''x_J, \qquad \omega_{IJ} \in C^\infty(U),$$

where $d'x_I \wedge d''x_J := dx_I \otimes dx_J$. The space $\mathcal{A}^{p,q}(U)$ is generated by primitive forms $h\, d'x_I \wedge d''x_J$.

There are exterior derivatives $d'\colon \mathcal{A}^{p,q} \to \mathcal{A}^{p+1,q}$ and $d''\colon \mathcal{A}^{p,q} \to \mathcal{A}^{p,q+1}$. On primitive forms, d' acts as the usual exterior derivative in the first component; d'' is defined by

$$d''(h\, d'x_I \wedge d''x_J) := (-1)^p \sum_{j=1}^n \frac{\partial h}{\partial x_j}\, d'x_I \wedge d''x_j \wedge d''x_J. \tag{2.3}$$

The sign $(-1)^p$ ensures that $d'' \circ d'' = 0$ and that the Leibniz formula

$$d''(\alpha \wedge \beta) = d''\alpha \wedge \beta + (-1)^{p+q} \alpha \wedge d''\beta \tag{2.4}$$

holds for $\alpha \in \mathcal{A}^{p,q}$ and $\beta \in \mathcal{A}^{p',q'}$, where the wedge product is

$$\left(h_1\, d'x_{I_1} \wedge d''x_{J_1}\right) \wedge \left(h_2\, d'x_{I_2} \wedge d''x_{J_2}\right) := (-1)^{qp'} h_1 h_2\, d'x_{I_1} \wedge d'x_{I_2} \wedge d''x_{J_1} \wedge d''x_{J_2}. \tag{2.5}$$

The pullback of a superform along an integral affine map ψ is defined by pulling back each component as a differential form:

$$\psi^*(h\, d'x_I \wedge d''x_J) := (h \circ \psi)\, \psi^* d'x_I \wedge \psi^* d''x_J. \tag{2.6}$$

These constructions extend to rational polyhedral spaces via charts. Given a polyhedral space X with face structure $\mathcal{C}$, one defines the *superform cochain complex* $\mathcal{A}^{p,\bullet}(X)$ by gluing the local complexes $\mathcal{A}^{p,\bullet}(U)$ over star charts via the transition maps [JSS19, §2]. The q-th cohomology of the complex $\mathcal{A}^{p,\bullet}(X)$ is the *superform cohomology*

$$H_D^{p,q}(X) := H^q(\mathcal{A}^{p,\bullet}(X),\, d'')\,.$$

The wedge product (2.5) makes $H_D^{\bullet,\bullet}(X)$ into a bigraded ring.

Integration of a superform $\omega \in \mathcal{A}^{n,n}(U)$ over a chart $U \subset \mathbb{R}^n$ is defined by

$$\int_U h\, d'x_{[n]} \wedge d''x_{[n]} := (-1)^{\binom{n}{2}} \int_U h\, dx_1 \cdots dx_n, \tag{2.7}$$

where $[n] = \{1, \dots, n\}$; see [Lag12, Definition 2.5] and [JSS19, Definition 2.7].

2.4 Singular Tropical Cohomology

We recall the definition of singular tropical cohomology following [GS23a, §4.7] and [JSS19, §3.1]. Let X be a rational polyhedral space with face structure $\mathcal{C}$.

The standard q-simplex is

$$\Delta^q := \left\{(t_0, \dots, t_q) \in \mathbb{R}^{q+1} \;\middle|\; \textstyle\sum_{i=0}^q t_i = 1,\ t_i \geq 0\right\},$$

with face maps $\delta_i^q \colon \Delta^{q-1} \to \Delta^q$ sending $(t_0, \dots, t_{q-1}) \mapsto (t_0, \dots, t_{i-1}, 0, t_i, \dots, t_{q-1})$ for $i \in \{0, \dots, q\}$. We write e_i for the i-th standard basis vector in $\mathbb{R}^{q+1}$, and given a singular simplex $f \colon \Delta^q \to X$ we abbreviate $f_i := f \circ \delta_i^q$.

A singular q-simplex $f\colon \Delta^q \to X$ is $\mathcal{C}$*-admissible* if it maps the relative interior of each face of Δ^q into the relative interior of some face of $\mathcal{C}$. We write σ_f for the unique face of $\mathcal{C}$ satisfying $f(\mathrm{relint}(\Delta^q)) \subset \mathrm{relint}(\sigma_f)$.

Definition 2.3. The *singular tropical* (p, q)*-chains* on X (with coefficients in $\mathbb{R}$) are

$$C^S_{p,q}(X) := \bigoplus_{\substack{f\colon \Delta^q \to X \\ \mathcal{C}\text{-admissible}}} \mathbf{F}_p(\sigma_f),$$

whose primitive elements are written $v \otimes f$ for $v \in \mathbf{F}_p(\sigma_f)$.

Since f is $\mathcal{C}$-admissible and continuous,$\sigma_{f_i} \preceq \sigma_f$ for each i. The inclusion $\mu_{\sigma_{f_i},\sigma_f}\colon \mathbf{F}_p(\sigma_f) \to \mathbf{F}_p(\sigma_{f_i})$ from Section 2.3.1 therefore gives a well-defined boundary map $\partial\colon C^S_{p,q}(X) \to C^S_{p,q-1}(X)$ by

$$\partial(v \otimes f) := (-1)^p \sum_{i=0}^{q} (-1)^i \, \mu_{\sigma_{f_i},\sigma_f}(v) \otimes f_i. \tag{2.8}$$

The sign $(-1)^p$ is analogous to the sign in (2.3); it does not appear in the prior literature but is essential for compatibility with the superform complex under the de Rham pairing constructed in Section 2.6. Intuitively, it arises from commuting the boundary operator ∂ past the multi-vector v, which has degree p.

Proposition 2.4. *$\partial^2 = 0$, so $C^S_{p,\bullet}(X)$ is a chain complex.*

Proof. [GS23a, §4.7] (the sign $(-1)^p$ is a scalar twist that does not affect the verification). □

The q-th homology of $C^S_{p,\bullet}(X)$ is the *singular tropical homology* $H^S_{p,q}(X)$. While the chain complex depends on the choice of face structure, the homology groups do not [GS23a, Theorem 4.20].

The *singular tropical cochain complex* is

$$C_S^{p,\bullet}(X) := \mathrm{Hom}_{\mathbb{R}}(C^S_{p,\bullet}(X), \mathbb{R}),$$

with coboundary ∂^* dual to ∂. Its q-th cohomology is the *singular tropical cohomology*

$$H^{p,q}_S(X) := H^q(C^{p,\bullet}_S(X),\, \partial^*)\,.$$

Given $v \otimes f \in C^S_{p,q}(X)$ and $\alpha \in C^{p,q}_S(X)$, we write $\langle v \otimes f, \alpha\rangle$ for evaluation. We let $\alpha_f \in \mathbf{F}^p(\sigma_f)$ denote the multi-covector satisfying $\langle v, \alpha_f\rangle = \langle v \otimes f, \alpha\rangle$ for all $v \in \mathbf{F}_p(\sigma_f)$.

We regard $C^{p,\bullet}_S$ as a complex of presheaves with natural restriction maps, and write $\mathrm{sh}\colon C^{p,\bullet}_S \to \underline{C}^{p,\bullet}_S$ for its sheafification.

We record one further construction, used in Section 2.7. Let $\mathcal{S}$ be a weighted simplicial complex of pure dimension k with weight function m and rational facets. For each facet $\sigma \in \mathcal{S}$ choose a generator v_σ of $\mathbf{F}_k(\sigma)$ and an affine simplex $f_\sigma\colon \Delta^k \to X$ parameterizing σ, so that the multi-vector $(e_k - e_0) \wedge \cdots \wedge (e_1 - e_0) \in \mathbf{F}_k(\Delta^k)$ maps to a positive multiple of v_σ. Then $\mathcal{S}$ determines an affine tropical chain

$$\mathrm{ch}(\mathcal{S}) := \sum_{\sigma \in \mathcal{S}^{(k)}} m_\sigma\, v_\sigma \otimes f_\sigma \;\in\; C^S_{k,k}(X), \tag{2.9}$$

and for a weighted rational polyhedral complex $\mathcal{P}$ with triangulation $\mathcal{S}$ we set $\mathrm{ch}_{\mathcal{S}}(\mathcal{P}) := \mathrm{ch}(\mathcal{S})$. One cannot always take $\mathcal{S}$ rational, but one can always ensure its facets are. If $\mathcal{P}$ is balanced, then $\partial\,\mathrm{ch}_{\mathcal{S}}(\mathcal{P}) = 0$, so $\mathrm{ch}_{\mathcal{S}}(\mathcal{P})$ determines a class $\mathrm{cyc}(\mathcal{P}) \in H^S_{k,k}(X)$ independent of $\mathcal{S}$.

For the construction of the de Rham map in Section 2.6, we also need a notion of smooth chain. A singular (p,q)-simplex $v \otimes f$ is *smooth* if f is smooth as a map $\Delta^q \to X$, where σ_f is endowed with its usual smooth structure. We define $C^{S,\infty}_{p,\bullet}(X)$ as the subcomplex of $C^S_{p,\bullet}(X)$ generated by smooth simplices, with cohomology $H^{p,q,\infty}_S(X)$, and we write $C^{p,\bullet}_{S,\infty}$ for its dual presheaf of cochain complexes.

A smooth simplex $v \otimes f$ is *affine* if f is the restriction to Δ^q of an affine map $\mathbb{R}^{q+1} \to \mathbb{R}^n$ in some chart; such a simplex is uniquely determined by the images of the vertices of Δ^q. The *affinization map* $\mathrm{aff}\colon C^S_{p,\bullet}(X) \to C^{S,\infty}_{p,\bullet}(X)$ sends $v \otimes f$ to the unique affine simplex $v \otimes f_{\mathrm{aff}}$ satisfying $f_{\mathrm{aff}}(e_i) = f(e_i)$ for all i. This is well-defined since f maps into the single face σ_f. One checksthat aff is a chain map, and hence its dual gives a chain map $\mathrm{aff}^*\colon C^{p,\bullet}_{S,\infty} \to C^{p,\bullet}_S$ of presheaf complexes.

2.5 Cup and Cap Products

We define a cup product and a cap product on singular tropical (co)homology in analogy with the classical constructions on singular (co)homology of topological spaces. These definitions are new to the literature.

2.5.1 The cup product

Let X be a rational polyhedral space with face structure $\mathcal{C}$. Write $\mathrm{fr}_q\colon \Delta^q \to \Delta^{q+q'}$ and $\mathrm{ba}_{q'}\colon \Delta^{q'} \to \Delta^{q+q'}$ for the inclusions of the *front* and *back* faces of $\Delta^{q+q'}$, i.e., the faces spanned by $e_0, \dots, e_q$ and $e_q, \dots, e_{q+q'}$, respectively.

Definition 2.5. The *cup product*

$$\smile\colon C_S^{p,q}(X) \otimes C_S^{p',q'}(X) \longrightarrow C_S^{p+p',q+q'}(X)$$

is defined on a $(p+p', q+q')$-simplex $v \otimes f$ by

$$(\alpha \smile \beta)_f := (-1)^{qp'}\, \mu^*_{\sigma_{f\circ \mathrm{fr}_q}, \sigma_f}\Big(\alpha_{f\circ \mathrm{fr}_q}\Big) \wedge \mu^*_{\sigma_{f\circ \mathrm{ba}_{q'}}, \sigma_f}\Big(\beta_{f\circ \mathrm{ba}_{q'}}\Big). \tag{2.10}$$

Following the standard argument for the cup product on singular cochains of topological spaces, one verifies that the cup product satisfies the Leibniz formula

$$\partial^*(\alpha \smile \beta) = \partial^*\alpha \smile \beta + (-1)^{p+q}\, \alpha \smile \partial^*\beta. \tag{2.11}$$

The sign here is the same as in the superform Leibniz formula (2.4), a fact that will be essential for Theorem 2.12. Consequently, the cup product descends to a well-defined product on cohomology, making $H_S^{\bullet,\bullet}(X)$ into a bigraded ring.

2.5.2 The cap product

Definition 2.6. The *cap product*

$$\frown\colon C^S_{p+p',q+q'}(X) \otimes C_S^{p,q}(X) \longrightarrow C^S_{p',q'}(X)$$

is defined by

$$(v \otimes f) \frown \alpha := (-1)^{qp'} \, \mu_{\sigma_{f\circ \mathrm{ba}_{q'}}, \sigma_f}\Big(v \frown \mu^*_{\sigma_{f\circ \mathrm{fr}_q}, \sigma_f}\big(\alpha_{f\circ \mathrm{fr}_q}\big)\Big) \otimes f \circ \mathrm{ba}_{q'}, \tag{2.12}$$

where $\frown$ inside the right-hand side is the multi-vector contraction of Section 2.3.1.

Using the cup product formula (2.10) and the adjointness property (2.1) of the multi-vector contraction, one shows that the cap product is adjoint to the cup product in the sense that

$$\langle (v \otimes f) \frown \alpha, \, \beta \rangle = \langle v \otimes f, \, \alpha \smile \beta \rangle. \tag{2.13}$$

It then follows formally from (2.11) that

$$\partial((v \otimes f) \frown \alpha) = (-1)^{p+q} (\partial(v \otimes f) \frown \alpha - (v \otimes f) \frown \partial^* \alpha) \, , \tag{2.14}$$

so the cap product descends to (co)homology.

2.6 The Tropical de Rham Map

For smooth manifolds, the classical de Rham theorem identifies de Rham cohomology with real singular cohomology via the integration pairing. We now construct an explicit analogue in the tropical setting.

2.6.1 The tropical de Rham pairing

We orient Δ^q via the parameterization

$$\phi \colon \widetilde{\Delta}^q \longrightarrow \Delta^q, \qquad (t_1, \dots, t_q) \longmapsto \Big(1 - \textstyle\sum_{i=1}^q t_i, \, t_1, \dots, t_q\Big) \, ,$$

where $\widetilde{\Delta}^q := \{(t_1, \dots, t_q) \mid t_i \geq 0, \sum t_i \leq 1\} \subset \mathbb{R}^q$ carries the standard orientation from $\frac{\partial}{\partial t_1} \wedge \cdots \wedge \frac{\partial}{\partial t_q}$.

Given a smooth $\mathcal{C}$-admissible simplex $v \otimes f \in C^{S,\infty}_{p,q}(X)$ and a primitive (p,q)-superform $h \, d'x_I \wedge d''x_J$ defined in a chart neighborhood of σ_f, we define

$$\int_{v\otimes f} h \, d'x_I \wedge d''x_J := \langle v, \, dx_I \rangle \int_f h \, dx_J, \tag{2.15}$$

where $\int_f h\,dx_J := \int_{\widetilde{\Delta}^q} (f \circ \phi)^*(h\,dx_J)$. Extending by linearity, this defines the *tropical de Rham pairing*

$$\mathcal{A}^{p,q}(X) \otimes C^{S,\infty}_{p,q}(X) \longrightarrow \mathbb{R}, \qquad \omega \otimes (v \otimes f) \longmapsto \int_{v\otimes f} \omega.$$

One must verify that $\int_{v\otimes f} \omega$ depends only on the equivalence class of ω as a superform (i.e., that it vanishes when $\omega = 0$ in $\mathcal{A}^{p,q}$); this follows from the fact that $v \in \mathbf{F}_p(\sigma_f)$ and f maps into σ_f.

2.6.2 Stokes formula and the de Rham chain map

Lemma 2.7 (Stokes formula)**.** *For all $\omega \in \mathcal{A}^{p,q-1}(X)$ and $c \in C^{S,\infty}_{p,q}(X)$,*

$$\int_c d''\omega = \int_{\partial c} \omega.$$

Proof. By bilinearity, it suffices to take $c = v \otimes f$ and $\omega = h\,d'x_I \wedge d''x_J$ in a chart. Writing $f_i := f \circ \delta_i^q$ and $v_i := \mu_{\sigma_{f_i},\sigma_f}(v)$, a direct computation using (2.3) and (2.15) gives

$$\int_{v\otimes f} d''\omega = (-1)^p \langle v,\, dx_I\rangle \int_f d(h\,dx_J),$$

while the boundary formula (2.8) yields

$$\int_{\partial(v\otimes f)} \omega = (-1)^p \sum_{i=0}^{q} (-1)^i \langle v_i,\, dx_I\rangle \int_{f_i} h\,dx_J = (-1)^p \langle v,\, dx_I\rangle \int_{\partial f} h\,dx_J.$$

These coincide by the classical Stokes theorem. □

The de Rham pairing induces a chain map

$$\mathrm{dR}\colon \mathcal{A}^{p,\bullet}(X) \longrightarrow C^{p,\bullet}_{S,\infty}(X), \qquad \omega \longmapsto \left[v \otimes f \mapsto \int_{v\otimes f} \omega \right], \tag{2.16}$$

which is a chain map by Lemma 2.7. Composing with aff^* and sheafification gives the map

$$\widetilde{\mathrm{dR}} := \mathrm{sh} \circ \mathrm{aff}^* \circ \mathrm{dR}\colon \mathcal{A}^{p,\bullet} \longrightarrow \underline{C}^{p,\bullet}_S.$$

2.6.3 $\widetilde{\mathrm{dR}}$ *is a quasi-isomorphism*

The following result upgrades the abstract isomorphism of [JSS19, Theorem 3.22] to an explicit quasi-isomorphism.

Lemma 2.8. *The map* $\widetilde{\mathrm{dR}}\colon \mathcal{A}^{p,\bullet} \to \underline{C}_S^{p,\bullet}$ *is a quasi-isomorphism.*

Proof. Each of sh, aff*, and dR is a chain map, so $\widetilde{\mathrm{dR}}$ is a chain map. By [JSS19, Corollary 3.18 and Proposition 3.15], both $\mathcal{A}^{p,\bullet}$ and $\underline{C}_S^{p,\bullet}$ are exact away from degree 0, with H^0 isomorphic to Ω_X^p in each case (see also [GS23a, Remark 2.8]). It remains to show that $\widetilde{\mathrm{dR}}$ identifies the two embeddings $\Omega_X^p \hookrightarrow \mathcal{A}^{p,0}$ and $\Omega_X^p \hookrightarrow C_S^{p,0}$ described in Section 2.6.

Fix a star neighborhood U with coordinates $x_1, \dots, x_n$, and let $\omega = \sum_I c_I \, dx_I \in \Omega_X^p(U)$. Its image in $\mathcal{A}^{p,0}(U)$ is $\sum_I c_I \, d'x_I$. For any $v \otimes x \in C_{p,0}^S(U)$, the affinization of a $(p,0)$-simplex is the simplex itself, so

$$\left\langle v \otimes x, \, \mathrm{aff}^* \circ \mathrm{dR}\left(\sum_I c_I \, d'x_I\right)\right\rangle = \int_{v \otimes x} \sum_I c_I \, d'x_I = \left\langle v, \, \sum_I c_I \, dx_I \right\rangle \cdot 1 = \langle v \otimes x, \, \omega\rangle,$$

confirming that $\widetilde{\mathrm{dR}}$ identifies both degree-0 kernels. □

Theorem 2.9. *The map on cohomology induced by* $\mathrm{aff}^* \circ \mathrm{dR}$ *is an isomorphism* $H_D^{p,q}(X) \xrightarrow{\sim} H_S^{p,q}(X)$.

Proof. The global sections functor $\Gamma(\,\cdot\,, X)$ is left-exact. By [JSS19, Lemmas 2.15 and 3.14], $\mathcal{A}^{p,q}$ is fine and $\underline{C}_S^{p,q}$ is flasque, hence both are $\Gamma(\,\cdot\,, X)$-acyclic. A standard result of homological algebra then implies that $\Gamma(\,\cdot\,, X)$ sends the quasi-isomorphism $\widetilde{\mathrm{dR}}$ to a quasi-isomorphism $\mathcal{A}^{p,\bullet}(X) \to \underline{C}_S^{p,\bullet}(X)$, giving $H_D^{p,q}(X) \cong \underline{H}_S^{p,q}(X)$. Moreover, by [JSS19, Lemma 3.14], the sheafification map sh: $C_S^{p,\bullet} \to \underline{C}_S^{p,\bullet}$ is a quasi-isomorphism of presheaves, whence $H_S^{p,q}(X) \cong \underline{H}_S^{p,q}(X)$. Composing gives the result. □

2.7 Ring Isomorphism and Poincaré Duality

2.7.1 The de Rham map is a ring isomorphism

We use the derived category to compare the ring structures on $H_D^{\bullet,\bullet}$ and $H_S^{\bullet,\bullet}$. We follow [GS23a, §4], working over $\underline{\mathbb{R}}$ in place of $\underline{\mathbb{Z}}$. Let $\mathcal{K}(\underline{\mathbb{R}})$ and $\mathcal{D}(\underline{\mathbb{R}})$ denote the homotopy and derived categories of chain complexes of $\underline{\mathbb{R}}$-modules.

A cohomology class $\alpha \in H_D^{p,q}(X)$ determines a morphism $\underline{\mathbb{R}} \to \Omega_X^p[q]$ in $\mathcal{D}(\underline{\mathbb{R}})$ via the roof diagram

$$\underline{\mathbb{R}} \xrightarrow{\alpha} \mathcal{A}^{p,\bullet}[q] \xleftarrow{\simeq} \Omega_X^p[q],$$

where the right arrow is the resolution $\Omega_X^p \hookrightarrow \mathcal{A}^{p,\bullet}$ from Section 2.6, and similarly for $\alpha \in H_S^{p,q}(X)$.

The *derived cup product* $H^q(X, \Omega_X^p) \otimes H^{q'}(X, \Omega_X^{p'}) \to H^{q+q'}(X, \Omega_X^{p+p'})$ is defined by the composition

$$\underline{\mathbb{R}} \xrightarrow{\cong} \underline{\mathbb{R}} \otimes \underline{\mathbb{R}} \xrightarrow{\alpha\otimes\beta} \Omega_X^p \otimes \Omega_X^{p'}[q+q'] \xrightarrow{\wedge} \Omega_X^{p+p'}[q+q']$$

in $\mathcal{D}(\underline{\mathbb{R}})$.

Lemma 2.10. *The wedge product* $\wedge\colon H_D^{p,q}(X) \otimes H_D^{p',q'}(X) \to H_D^{p+p',q+q'}(X)$ *coincides with the derived cup product.*

Proof. It suffices to show commutativity of the diagram

$$\begin{tikzcd}
\underline{\mathbb{R}} \otimes \underline{\mathbb{R}} \arrow[r, "\alpha\otimes\beta"] & \mathcal{A}^{p,\bullet} \otimes \mathcal{A}^{p',\bullet}[q+q'] \arrow[d, "\wedge"] & \Omega_X^p \otimes \Omega_X^{p'}[q+q'] \arrow[l, "\simeq"'] \arrow[d, "\wedge"] \\
\underline{\mathbb{R}} \arrow[u, "\cong"] \arrow[r, "\alpha\wedge\beta"] & \mathcal{A}^{p+p',\bullet}[q+q'] & \Omega_X^{p+p'}[q+q'] \arrow[l, "\simeq"']
\end{tikzcd}$$

in $\mathcal{K}(\underline{\mathbb{R}})$. Commutativity of the left square is trivial. For the right square, one checks in coordinates: the tropical form $\sum_I c_I\, dx_I$ maps to the superform $\sum_I c_I\, d'x_I$ under the embedding $\Omega_X^p \hookrightarrow \mathcal{A}^{p,0}$, and the wedge products in both settings are identical. $\square$

Lemma 2.11. *The cup product* $\smile\colon H_S^{p,q}(X) \otimes H_S^{p',q'}(X) \to H_S^{p+p',q+q'}(X)$ *coincides with the derived cup product.*

Proof. By analogy with Lemma 2.10, it suffices to verify the analogous diagram with $\underline{C}_S^{p,\bullet}$ in place of $\mathcal{A}^{p,\bullet}$ and $\smile$ in place of $\wedge$. Commutativity of the left square is again trivial. For the right square: for any $x \in X$ and $v \in \mathbf{F}_p(\sigma_x)$, the discussion in Section 2.6 gives $\langle v, (\omega \wedge \omega')_x \rangle = \langle v, \omega \wedge \omega' \rangle$, and by (2.10), $\langle v, (\omega \smile \omega')_x \rangle = \langle v, \omega_x \wedge \omega'_x \rangle$. Both agree with the wedge product of covectors in the ambient $\mathbb{R}^n$, so they coincide. $\square$

Theorem 2.12 (de Rham ring isomorphism)**.** *For any rational polyhedral space X, the map $\widetilde{\mathrm{dR}}$ induces an isomorphism of bigraded rings*

$$\widetilde{\mathrm{dR}} \colon H_D^{\bullet,\bullet}(X) \xrightarrow{\sim} H_S^{\bullet,\bullet}(X).$$

Proof. By Theorem 2.9, $\widetilde{\mathrm{dR}}$ is an isomorphism of vector spaces. It remains to show that it is a ring map, i.e., that it takes the wedge product to the cup product. By Lemmas 2.10 and 2.11, both products coincide with the derived cup product. It therefore suffices to show that the map $\widetilde{\mathrm{dR}}$ induces on the derived category is the identity morphism, i.e., that the following diagram commutes in $\mathcal{K}(\underline{\mathbb{R}})$:

$$\begin{tikzcd}
\underline{\mathbb{R}} \arrow[r] \arrow[dr] & \mathcal{A}^{p,\bullet}[q] \arrow[d, "\widetilde{\mathrm{dR}}"] & \Omega^p_X[q] \arrow[l, "\simeq"'] \arrow[dl, "\simeq"'] \\
& \underline{C}_S^{p,\bullet}[q] &
\end{tikzcd}$$

The maps from $\underline{\mathbb{R}}$ define the cohomology classes, so the left triangle commutes by construction. Commutativity of the right triangle is precisely the content of Lemma 2.8. $\square$

2.7.2 Poincaré duality

Fix a compact rational polyhedral space X of dimension n with face structure $\mathcal{C}$. Suppose that endowing every facet of $\mathcal{C}$ with weight 1 defines a tropical cycle; we then say that X *admits a fundamental cycle* and write $\mathrm{cyc}(X) \in H^S_{n,n}(X)$ for the corresponding homology class (see Section 2.4 and [GS23a, §6.1]).

Following [JSS19, Definition 4.11], the *Poincaré duality map* on superforms is

$$\mathrm{PD}\colon H_D^{p,q}(X) \longrightarrow H_D^{n-p,n-q}(X)^\vee, \qquad \mathrm{PD}(\alpha) := \varepsilon_{p,q} \int_X \alpha \wedge \cdot\,, \tag{2.17}$$

where $\varepsilon_{p,q}$ is the sign determined by $\varepsilon_{p,q} = \varepsilon_{p,q+1}(-1)^{p+q+1}$, chosen so that PD is a chain map (this follows from the Stokes and Leibniz formulas (2.4) and Lemma 2.7).

The corresponding *Poincaré duality map* on singular tropical cohomology is

$$\mathrm{PD}\colon H_S^{p,q}(X) \longrightarrow H^S_{n-p,n-q}(X), \qquad \mathrm{PD}(\alpha) := \varepsilon_{p,q}\,\mathrm{cyc}(X) \frown \alpha. \tag{2.18}$$

The same sign $\varepsilon_{p,q}$ is used in both cases; this is consistent because the Leibniz formula (2.14) for the cap product has the same sign structure as (2.4). This cap product Poincaré duality map agrees with the one in [GS23a, Corollary 6.9] defined via the derived category.

We need two preparatory lemmas.

Lemma 2.13. *Let $\mathcal{P}$ be a weighted rational polyhedral complex of pure dimension k and let $\omega \in \mathcal{A}^{k,k}(U)$ be defined in a neighborhood U of $|\mathcal{P}|$. Then for any triangulation $\mathcal{S}$ of $\mathcal{P}$,*

$$\int_{\mathcal{P}} \omega = \int_{\mathrm{ch}_{\mathcal{S}}(\mathcal{P})} \omega.$$

Proof. The integral over $\mathcal{P}$ decomposes into a sum of integrals over its facets; each is pulled back by a chart and computed via (2.7), so replacing $\mathcal{P}$ by $\mathcal{S}$ leaves the integral unchanged. Similarly, the integral over the tropical chain $\mathrm{ch}_{\mathcal{S}}(\mathcal{P})$ decomposes over facets. One then checks for each rational simplex σ of dimension k in $\mathbb{R}^k$ with generator $v_\sigma \in \mathbf{F}_k(\sigma)$ and affine parameterization $f_\sigma\colon \Delta^k \to \mathbb{R}^k$:

$$\int_\sigma h\, d'x_{[k]} \wedge d''x_{[k]} = (-1)^{\binom{k}{2}} \int_\sigma h\, dx_{[k]} = \int_{v_\sigma \otimes f_\sigma} h\, d'x_{[k]} \wedge d''x_{[k]},$$

where the orientation of $\widetilde{\Delta}^k$ accounts for the sign; see [GS23a, §3.5] for the sign analysis. □

Lemma 2.14. *Given $c \in H^S_{p+p',q+q'}(X)$ and $\alpha \in H_D^{p,q}(X)$,*

$$\mathrm{dR}^\vee \circ \mathrm{aff}(c \frown \mathrm{aff}^* \circ \mathrm{dR}(\alpha)) = \int_{\mathrm{aff}(c)} \alpha \wedge \cdot$$

in $H_D^{p',q'}(X)^\vee$.

Proof. Using (2.13) and Theorem 2.12, we compute

$$\begin{aligned}\left\langle \mathrm{dR}^\vee \circ \mathrm{aff}(c \frown \mathrm{aff}^* \circ \mathrm{dR}(\alpha)),\, \beta \right\rangle &= \langle c \frown \mathrm{aff}^* \circ \mathrm{dR}(\alpha),\, \mathrm{aff}^* \circ \mathrm{dR}(\beta)\rangle \\ &= \langle c,\, (\mathrm{aff}^* \circ \mathrm{dR}(\alpha)) \smile (\mathrm{aff}^* \circ \mathrm{dR}(\beta))\rangle \\ &= \langle c,\, \mathrm{aff}^* \circ \mathrm{dR}(\alpha \wedge \beta)\rangle \\ &= \int_{\mathrm{aff}(c)} \alpha \wedge \beta.\end{aligned}$$

□

Theorem 2.15 (Poincaré duality compatibility)**.** *For any compact rational polyhedral space X of dimension n admitting a fundamental cycle, the following diagram commutes:*

$$\begin{tikzcd} H_D^{p,q}(X) \arrow[r, "\widetilde{\mathrm{dR}}"] \arrow[d, "\mathrm{PD}"'] & H_S^{p,q}(X) \arrow[d, "\mathrm{PD}"] \\ H_D^{n-p,n-q}(X)^\vee & H^S_{n-p,n-q}(X) \arrow[l, "\mathrm{dR}^\vee \circ \mathrm{aff}"'] \end{tikzcd}$$

Proof. This follows immediately from Lemmas 2.13 and 2.14. □

Chapter 3

TROPICAL ABELIAN VARIETIES

3.1 Introduction

Over the complex numbers, a principally polarized abelian variety (A, Θ) carries a canonical $(1,1)$-form: the first Chern class of the polarization line bundle $\mathcal{O}(\Theta)$. When $A = \mathrm{Jac}(C)$ is the Jacobian of a smooth projective curve C, the restriction of this form to C (via the Abel–Jacobi map) recovers the Arakelov canonical measure on C. The purpose of this chapter is to establish a complete tropical analogue of this picture, using the formalism of Lagerberg superforms developed in Chapter 2.

Let (X, ϕ) be a principally polarized tropical abelian variety of dimension g (Definition 3.1). The main results of this chapter are as follows.

(A) The polarization $\phi\colon N \to \Lambda^\vee$ determines a canonical constant superform $\nu_\phi \in H^{1,1}_D(X)$ (Definition 3.3).

(B) When $X = \mathrm{Jac}(\Gamma)$ is the tropical Jacobian of a metric graph Γ, the pullback of ν_ϕ along the Abel–Jacobi map restricts on each edge to a constant multiple of the Lebesgue measure, with coefficient equal to the Foster coefficient of that edge (Corollary 3.12); in particular, ν_ϕ pulls back to the Zhang measure μ_{Zh} on Γ.

(C) The integral of $\frac{1}{d!}\nu_\phi^{\wedge d}$ over the d-th effective locus $\widetilde{W}_d \subset \mathrm{Jac}(\Gamma)$ satisfies

$$\int_{\widetilde{W}_d} \frac{1}{d!}\,\nu_\phi^{\wedge d} = \binom{g}{d}$$

(Theorem 3.14).

(D) The tropical Poincaré formula
$$[\widetilde{W}_d] = \frac{[\Theta]^{g-d}}{(g-d)!}$$
holds in $H_{S,d,d}(\mathrm{Jac}(\Gamma))$ (Theorem 3.15). This provides an alternative proof, via the de Rham isomorphism of Chapter 2, of the formula first established in the tropical setting by Gross–Shokrieh [GS23b].

Outline. Section 3.2 introduces tropical abelian varieties and polarizations. Section 3.3 computes their singular tropical cohomology via the Künneth theorem and identifies the cohomology class h_ϕ determined by ϕ. Section 3.4 describes constant superforms on a real torus and the map from $\mathrm{Hom}(N, \Lambda^\vee)$ to superform cohomology. Section 3.5 recalls tropical Jacobians, the Abel–Jacobi map, and the polyhedral structure on the effective loci $\widetilde{W}_d$. Section 3.6 defines the canonical form ν_ϕ, establishes its coordinate formula, and verifies compatibility with the de Rham map. Section 3.7 proves the Poincaré–Lelong formula for $\|\Psi\|$ and identifies the pullback of ν_ϕ to a metric graph with the Zhang measure. Sections 3.8 and 3.9 prove Theorems 3.14 and 3.15, respectively.

Notation. We write $\binom{[g]}{d}$ for the collection of d-element subsets of $[g] := \{1, 2, \dots, g\}$. Given $I \in \binom{[g]}{p}$, $J \in \binom{[g]}{q}$, and a $g \times g$ matrix $\mathbf{A}$, we write $\mathbf{A}_{I,J}$ for the $p \times q$ submatrix with rows indexed by I and columns by J. We abbreviate $d'x_I := d'x_{i_1} \wedge \cdots \wedge d'x_{i_p}$ for $I = \{i_1 < \cdots < i_p\}$, and similarly for $d''x_J$. For an abelian group A, we set $A^\vee := \mathrm{Hom}(A, \mathbb{Z})$ and $A_\mathbb{R} := A \otimes_\mathbb{Z} \mathbb{R}$.

3.2 Tropical Abelian Varieties and Polarizations

In this section, we establish the basic structure of tropical abelian varieties and their polarizations. We introduce the key matrices $\mathbf{E}$ and $\mathbf{G}$ that will control the local behavior of superforms. The presentation follows [GS23b, §2B].

Fix two finitely generated free abelian groups Λ and N of rank g, together with an embedding $\Lambda \subset N_\mathbb{R}$. Both N and Λ are lattices inside $N_\mathbb{R} \cong \mathbb{R}^g$. Any identification of N

with $\mathbb{Z}^g$ makes the quotient torus $X := N_{\mathbb{R}}/\Lambda$ into a rational polyhedral space in the sense of Chapter 2.

Definition 3.1. A *polarization* of $X = N_{\mathbb{R}}/\Lambda$ is a homomorphism $\phi\colon N \to \Lambda^\vee$ such that the induced bilinear form

$$[\,\cdot\,,\,\cdot\,]\colon N_{\mathbb{R}} \otimes \Lambda_{\mathbb{R}} \longrightarrow \mathbb{R}, \qquad [n, \lambda] := \phi(n)(\lambda),$$

is symmetric (via the canonical identification $N_{\mathbb{R}} \cong \Lambda_{\mathbb{R}}$) and positive-definite. If ϕ is an isomorphism, we call it *principal*. A *(principally) polarized tropical abelian variety* is a pair (X, ϕ) consisting of such a torus and a (principal) polarization.

Since ϕ is necessarily injective, the quotient $\Lambda^\vee/\phi(N)$ is finite; we write $e_1, \dots, e_g$ for the elementary divisors of ϕ and $\mathbf{E}$ for the diagonal matrix $\operatorname{diag}(e_1, \dots, e_g)$. The *degree* of ϕ is $\deg\phi := \prod_{i=1}^g e_i$; the polarization is principal if and only if $\deg\phi = 1$.

By writing ϕ in Smith normal form, we obtain bases $n_1, \dots, n_g$ for N and $\lambda_1, \dots, \lambda_g$ for Λ satisfying

$$\phi(n_i) = e_i\lambda_i^\vee, \qquad \text{equivalently,} \qquad [n_i, \lambda_j] = e_i\delta_{i,j} = \mathbf{E}_{i,j}. \tag{3.1}$$

The *Gram matrix* $\mathbf{G}$ is the symmetric, positive-definite matrix with entries

$$\mathbf{G}_{i,j} := [\lambda_i, \lambda_j]. \tag{3.2}$$

Since $\Lambda \subset N_{\mathbb{R}}$, each λ_j is a linear combination of the n_i; one verifies from (3.1) and (3.2) that

$$\lambda_j = \sum_{i=1}^g (\mathbf{E}^{-1}\mathbf{G})_{i,j}\, n_i. \tag{3.3}$$

3.3 Cohomology of Tropical Abelian Varieties

The singular tropical (co)homology of a real torus is computed by a Künneth theorem. We recall the relevant isomorphisms and use them to identify the class in $H^{1,1}_S(X)$ determined by the polarization.

Let (X, ϕ) be a polarized tropical abelian variety with the notation of Section 3.2. Following [GS23b, §2B], the singular tropical (p, q)-chain group is defined as

$$C^S_{p,q}(X) := \bigwedge^p N \otimes C_q(X),$$

where $C_q(X)$ denotes the group of singular q-chains on X (see also Section 2.4). The Künneth theorem then gives the following isomorphisms.

Proposition 3.2. *There are natural isomorphisms*

$$H_{S,p,q}(X) \cong \bigwedge^p N \otimes \bigwedge^q \Lambda \qquad \text{and} \qquad H_S^{p,q}(X) \cong \bigwedge^p N^\vee \otimes \bigwedge^q \Lambda^\vee. \tag{3.4}$$

Proof. [GST22, Theorem 8.1]. □

In particular, identifying $\mathrm{Hom}(N, \Lambda^\vee) \cong N^\vee \otimes \Lambda^\vee = H_S^{1,1}(X)$, the polarization $\phi\colon N \to \Lambda^\vee$ determines the cohomology class

$$h_\phi := \sum_{i=1}^{g} e_i\, n_i^\vee \otimes \lambda_i^\vee \;\in\; H_S^{1,1}(X), \tag{3.5}$$

where $n_i^\vee$ and $\lambda_i^\vee$ are the dual basis elements.

3.4 Superforms on Tropical Abelian Varieties

The key result of this section is that the natural map $\mathrm{Hom}(N, \Lambda^\vee) \to H_D^{1,1}(X)$ is compatible with the de Rham isomorphism, as expressed by diagram (3.9). We begin by recalling how constant superforms on $X = N_\mathbb{R}/\Lambda$ interact with the de Rham map constructed in Chapter 2.

Because the universal cover of X is $N_\mathbb{R}$, the tangent space at any point is canonically identified with $N_\mathbb{R}$. Smooth 1-forms with constant coefficients on X are therefore equivalent to linear functionals on $N_\mathbb{R}$. Consequently, there is a natural map

$$\mathrm{Hom}(N, \Lambda^\vee) \longrightarrow H_D^{1,1}(X) \tag{3.6}$$

sending a homomorphism $f\colon N \to \Lambda^\vee$ to the constant $(1,1)$-superform ω whose contraction satisfies $\langle n \otimes \lambda, \omega\rangle = f(n)(\lambda)$ for all $n \in N$ and $\lambda \in \Lambda$. Indeed, $H_D^{1,1}(X)$ may be identified with the subspace of constant superforms in $\mathscr{A}^{1,1}(X)$ (see Section 2.3), so the map is well defined.

Using (3.3), one computes in coordinates that

$$n_i^\vee \otimes \lambda_j^\vee \ \longmapsto \ \sum_{k=1}^{g} (\mathbf{G}^{-1}\mathbf{E})_{j,k}\, d'x_i \wedge d''x_k. \tag{3.7}$$

By Theorem 2.12, the de Rham map $\widetilde{\mathrm{dR}}\colon H_D^{1,1}(X) \xrightarrow{\sim} H_S^{1,1}(X)$ sends a constant $(1,1)$-superform $d'x_i \wedge d''x_j$ to the linear functional

$$n \otimes \lambda \mapsto \langle n, dx_i\rangle \cdot \langle \lambda, dx_j\rangle,$$

which, by (3.3), gives

$$\widetilde{\mathrm{dR}}(d'x_i \wedge d''x_j) = n_i^\vee \otimes \sum_{k=1}^{g} (\mathbf{E}^{-1}\mathbf{G})_{j,k} \lambda_k^\vee. \tag{3.8}$$

Combining (3.7) and (3.8), one verifies that the following diagram commutes:

$$\begin{tikzcd}
\mathrm{Hom}(N, \Lambda^\vee) \arrow[d] \arrow[dr, "\cong"] & \\
H_D^{1,1}(X) \arrow[r, "\widetilde{\mathrm{dR}}"'] & H_S^{1,1}(X)_{\mathbb{R}}
\end{tikzcd} \tag{3.9}$$

where the diagonal arrow sends $n_i^\vee \otimes \lambda_j^\vee$ to its image in $H_S^{1,1}(X)$ via the identification (3.4).

3.5 Tropical Jacobians and the Abel–Jacobi Map

We recall the construction of the tropical Jacobian and the Abel–Jacobi map, following [GST22, §2.3, §8.1–8.2] and [GS23b, §2.3].

A *tropical curve* Γ is a compact, connected, 1-dimensional rational polyhedral space without boundary. Concretely, one starts with a connected finite graph G without leaf edges, with vertex set V, edge set E, and a length function $\ell\colon E \to \mathbb{R}_{>0}$. Associating to each edge e the interval $[0, \ell(e)]$ and gluing according to the adjacency of G yields a metric graph Γ. The pair (G, ℓ) is a *model* for Γ, and the *genus* of Γ is $g := |E| - |V| + 1 = \dim H_1(\Gamma, \mathbb{R})$.

3.5.1 The tropical Jacobian

Fix a model (G, ℓ) for Γ and an orientation of the edges. The *length pairing* $[\,\cdot\,,\,\cdot\,]_G\colon C_1(G,\mathbb{Z})^{\otimes 2} \to \mathbb{R}$ is defined by

$$[e, e']_G := \begin{cases} \ell(e) & \text{if } e = e', \\ 0 & \text{otherwise,} \end{cases}$$

extended bilinearly; it descends to $H_1(G,\mathbb{Z}) \cong H_1(\Gamma,\mathbb{Z})$. Set $N := H^1(\Gamma,\mathbb{Z})$ and let Λ be the image of the embedding

$$H_1(\Gamma,\mathbb{Z}) \hookrightarrow N_{\mathbb{R}}, \qquad \gamma \mapsto [\gamma, \cdot]_G.$$

The *tropical Jacobian* of Γ is the principally polarized tropical abelian variety

$$\mathrm{Jac}(\Gamma) := N_{\mathbb{R}}/\Lambda,$$

with principal polarization $\phi\colon N \to \Lambda^\vee \cong H_1(\Gamma,\mathbb{Z})$ given as the dual of the inverse of the embedding above. Choosing a basis $\gamma_1, \dots, \gamma_g$ for $H_1(G,\mathbb{Z})$ yields the dual basis $n_1, \dots, n_g$ for N and the basis $\lambda_i := [\gamma_i, \cdot]_G$ for Λ; in these bases the Gram matrix satisfies $\mathbf{G}_{i,j} = [\gamma_i, \gamma_j]_G$, and the polarization is principal so $\mathbf{E} = \mathbf{I}$.

3.5.2 The Abel–Jacobi map and the effective loci

Fix a basepoint $q \in V$. The *d-th Abel–Jacobi map*

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

sends $(p_1, \dots, p_d)$ to $\sum_{i=1}^d [\gamma_{p_i}, \cdot]_G$, where γ_{p_i} is any path from q to p_i in a refinement of (G, ℓ) containing p_i as a vertex; see [GST22, §8.1] for a careful treatment. The map $\Phi_q^{(d)}$ factors through the symmetric power $\Gamma^d/\mathfrak{S}_d$ and is a tropical map, so pullbacks of superforms are well defined.

The *d-th effective locus* is

$$\widetilde{W}_d := \Phi_q^{(d)}(\Gamma^d) \subset \mathrm{Jac}(\Gamma).$$

By [GST22, Theorem D], $\widetilde{W}_d$ is a polyhedral subset of pure dimension d. Its top-dimensional cells are parameterized as follows [GST22, §8.2]: for each set $\mathcal{E} \in \binom{E}{d}$ of d edges of G such that $G \setminus \mathcal{E}$ is connected (i.e., $\mathcal{E}$ is a *coforest* of size d), the product $\prod \mathcal{E} \subset \Gamma^d/\mathfrak{S}_d$ maps to a d-cell of $\widetilde{W}_d$. The images of distinct cells intersect in sets of measure zero, and cells coming from $\mathcal{E}$ with $G \setminus \mathcal{E}$ disconnected have image of measure zero. Consequently, the integral of any superform ω over $\widetilde{W}_d$ decomposes as

$$\int_{\widetilde{W}_d} \omega = \int_{\Gamma^d/\mathfrak{S}_d} (\Phi_q^{(d)})^* \omega = \sum_{\mathcal{E} \in \binom{E}{d}} \int_{\prod \mathcal{E}} (\Phi_q^{(d)})^* \omega. \tag{3.10}$$

3.6 The Canonical $(1,1)$*-Form*

We now define the canonical superform ν_ϕ associated to a polarization and establish its explicit coordinate expression.

Definition 3.3. Let (X, ϕ) be a polarized tropical abelian variety with bases $n_1, \dots, n_g$ and $\lambda_1, \dots, \lambda_g$ as in Section 3.2. The *canonical* $(1,1)$*-form* $\nu_\phi \in H_D^{1,1}(X)$ is the image of $\phi \in \mathrm{Hom}(N, \Lambda^\vee)$ under the map (3.6); it is uniquely characterized by the property that

$$\langle n \otimes \lambda,\, \nu_\phi \rangle = [n, \lambda] \qquad \text{for all } n \in N,\ \lambda \in \Lambda.$$

Proposition 3.4. *In the coordinates $x_1, \dots, x_g$ on $N_\mathbb{R}$ dual to $n_1, \dots, n_g$, the form ν_ϕ has the expression*

$$\nu_\phi = \sum_{i,j=1}^{g} (\mathbf{E}\mathbf{G}^{-1}\mathbf{E})_{i,j}\, d'x_i \wedge d''x_j. \tag{3.11}$$

Moreover, $\widetilde{\mathrm{dR}}(\nu_\phi) = h_\phi$*, where h_ϕ is the cohomology class* (3.5).

Proof. We use the defining property $\langle n_i \otimes \lambda_j, \nu_\phi \rangle = [n_i, \lambda_j] = \mathbf{E}_{i,j}$ and the coordinate formula (3.7) with $\phi = \sum_i e_i n_i^\vee \otimes \lambda_i^\vee$. Substituting into (3.7) and summing over i using the fact that $\mathbf{E}$ is diagonal, one obtains the expression (3.11) after a matrix multiplication with $\mathbf{G}^{-1}$. The compatibility $\widetilde{\mathrm{dR}}(\nu_\phi) = h_\phi$ follows from the commutativity of diagram (3.9). $\square$

Remark 3.5. The form ν_ϕ parallels the classical construction for complex abelian varieties: on a complex abelian variety with polarization ϕ, the first Chern form of the polarization line bundle with respect to a translation-invariant metric is a constant $(1,1)$-form satisfying the same contraction identity; see, e.g., [Wil17, §2.1].

We record next the formula for the d-th wedge power of ν_ϕ, which will be essential in Sections 3.8 and 3.9. To state it cleanly, we set $\mathbf{A} := \mathbf{E}\mathbf{G}^{-1}\mathbf{E}$, so that $\nu_\phi = \sum_{i,j} \mathbf{A}_{i,j}\, d'x_i \wedge d''x_j$.

Lemma 3.6. *For each $d \geq 1$,*

$$\nu_\phi^{\wedge d} = (-1)^{\binom{d}{2}} d! \sum_{I,J \in \binom{[g]}{d}} \det(\mathbf{A}_{I,J})\, d'x_I \wedge d''x_J. \tag{3.12}$$

Proof. Expanding the wedge product and separating the d'- and d''-parts gives

$$\nu_\phi^{\wedge d} = (-1)^{\binom{d}{2}} \sum_{i_1, j_1, \dots, i_d, j_d} \left(\prod_{\alpha=1}^{d} \mathbf{A}_{i_\alpha, j_\alpha} \right) \left(\bigwedge_\alpha d'x_{i_\alpha} \right) \wedge \left(\bigwedge_\alpha d''x_{j_\alpha} \right).$$

Only those terms with all i_α distinct and all j_α distinct are nonzero. Indexing over $I, J \in \binom{[g]}{d}$ and reordering indices within each multiset using permutations $\sigma, \tau \in \mathfrak{S}_d$ yields

$$\begin{aligned}
\nu_\phi^{\wedge d} &= (-1)^{\binom{d}{2}} \sum_{I,J} \sum_{\sigma,\tau \in \mathfrak{S}_d} \left(\prod_\alpha \mathbf{A}_{i_{\sigma(\alpha)}, j_{\tau(\alpha)}} \right) \operatorname{sgn}(\sigma)\, d'x_I \wedge \operatorname{sgn}(\tau)\, d''x_J \\
&= (-1)^{\binom{d}{2}} \sum_{I,J} \sum_{\sigma \in \mathfrak{S}_d} \operatorname{sgn}(\sigma) \cdot \det(\mathbf{A}_{I^\sigma, J})\, d'x_I \wedge d''x_J.
\end{aligned}$$

Since permuting rows of $\mathbf{A}_{I,J}$ by σ multiplies the determinant by $\operatorname{sgn}(\sigma)$, each inner sum contributes $\sum_\sigma \det(\mathbf{A}_{I,J}) = d! \cdot \det(\mathbf{A}_{I,J})$, giving (3.12). $\square$

Proposition 3.7 (Top-degree integral)**.** *For a polarized tropical abelian variety (X, ϕ) of dimension g,*

$$\int_X \frac{1}{g!}\, \nu_\phi^{\wedge g} = \deg \phi.$$

In particular, taking $\mathbf{E} = \mathbf{I}$ and using (3.11),

$$\int_X d'x_{[g]} \wedge d''x_{[g]} = (-1)^{\binom{g}{2}} \det(\mathbf{G}). \tag{3.13}$$

Proof. Lemma 3.6 at $d = g$ gives $\nu_\phi^{\wedge g} = (-1)^{\binom{g}{2}} g! \det(\mathbf{E}\mathbf{G}^{-1}\mathbf{E})\, d'x_{[g]} \wedge d''x_{[g]}$. Let $\Delta \subset N_{\mathbb{R}}$ be the parallelepiped spanned by $\lambda_1, \dots, \lambda_g$, a fundamental domain for X. By (2.7),

$$\int_X \frac{1}{g!}\, \nu_\phi^{\wedge g} = \det(\mathbf{E}\mathbf{G}^{-1}\mathbf{E})\, (-1)^{\binom{g}{2}} \int_\Delta d'x_{[g]} \wedge d''x_{[g]} = \det(\mathbf{E}\mathbf{G}^{-1}\mathbf{E}) \int_\Delta dx_{[g]}.$$

By (3.3) the matrix $\mathbf{E}^{-1}\mathbf{G}$ is the change of basis from $\lambda_1, \dots, \lambda_g$ to $n_1, \dots, n_g$, so pulling the integral back to $[0,1]^g$ contributes a factor $\det(\mathbf{E}^{-1}\mathbf{G})$. Multiplicativity of the determinant and $\det(\mathbf{E}) = \deg \phi$ give the first claim. For the second, note that the same computation shows $\int_\Delta dx_{[g]} = \det(\mathbf{E}^{-1}\mathbf{G})$, which is $\det(\mathbf{G})$ when $\mathbf{E} = \mathbf{I}$. □

3.7 Pullback to Metric Graphs and the Zhang Measure

Let Γ be a tropical curve with model (G, ℓ), Jacobian $\mathrm{Jac}(\Gamma)$, and canonical principal polarization ϕ. This section identifies the Poincaré dual of the theta divisor $\Psi \subset \mathrm{Jac}(\Gamma)$ with ν_ϕ, and shows that the pullback of ν_ϕ to Γ via the Abel–Jacobi map recovers the Zhang measure.

3.7.1 The tropical Riemann theta function and its corner locus

Following [dJS22a, §8.1], define the *normalized tropical Riemann theta function* on the universal cover $N_{\mathbb{R}}$ by

$$\|\Psi\|(x) := \frac{1}{2} \min_{\lambda \in \Lambda} [x - \lambda,\, x - \lambda]. \tag{3.14}$$

This is a Λ-periodic piecewise smooth function, so it descends to $X = \mathrm{Jac}(\Gamma)$. The *theta divisor* $\Psi := \|\Psi\| \cdot X$ is its corner locus (see [GK17, Definition 1.10] for the general notion), regarded as a Cartier divisor whose line bundle $\mathcal{L}(\Psi)$ satisfies $c_1(\mathcal{L}(\Psi)) = h_\phi$.

Lemma 3.8. *The divisor $\Psi = \|\Psi\| \cdot X$ has constant weight 1 on each facet.*

Proof. Since X and $N_{\mathbb{R}}$ are locally isomorphic, it suffices to work in $N_{\mathbb{R}}$. The function $x \mapsto [x - \lambda, x - \lambda]$ measures the squared distance from x to λ, so the singular locus of $\|\Psi\|$ is the Voronoi decomposition $\mathcal{C}$ of $N_{\mathbb{R}}$ with respect to the lattice Λ and the inner product $[\,\cdot\,,\,\cdot\,]$.

Write σ_λ for the Voronoi cell where the minimum in (3.14) is achieved at λ. By periodicity, $\sigma_\lambda = \sigma_0 + \lambda$, so σ_0 is a fundamental domain.

A codimension-1 face τ_λ between σ_0 and σ_λ lies in the hyperplane $H_\lambda := \{x \in N_\mathbb{R} : [x, \lambda] = \frac{1}{2}[\lambda, \lambda]\}$. If $\lambda = c\lambda'$ for $c \geq 2$ then $\sigma_0 \cap H_\lambda = \emptyset$, so it suffices to consider primitive $\lambda \in \Lambda$.

Fix a primitive λ. Since ϕ is principal, there exists $n_\lambda \in N$ with $[n_\lambda, \lambda] = 1$; moreover, $[n_\lambda, \cdot]$ takes integer values on Λ, so n_λ generates N/N_{τ_λ} and we may take $n_{\sigma_0,\tau_\lambda} := n_\lambda$. Write $x = \sum_i x_i n_i$; then

$$d_x\Big(\|\Psi\|_{\sigma_0}\Big) = d_x\big(\tfrac{1}{2}[x, x]\big) = \sum_{i,j} (\mathbf{G}^{-1})_{i,j}\, x_i\, dx_j,$$

from which one derives $\langle\, \cdot\,, d_x(\|\Psi\|_{\sigma_0})\rangle = [\,\cdot\,, x]$. By Λ-periodicity, the weight on τ_λ (in the sense of [GK17, §1]) is

$$\begin{aligned} m_{\tau_\lambda}(x) &= \langle n_\lambda, d_x(\|\Psi\|_{\sigma_0})\rangle + \langle -n_\lambda, d_{x-\lambda}(\|\Psi\|_{\sigma_0})\rangle \\ &= [n_\lambda, x] + [-n_\lambda, x - \lambda] = [n_\lambda, \lambda] = 1. \end{aligned}$$ □

3.7.2 The Poincaré–Lelong formula

The space of (p, q)-*supercurrents* $\mathcal{D}_{p,q}(X)$ on a rational polyhedral space X is the topological dual of the space of compactly supported (p, q)-superforms $\mathcal{A}_c^{p,q}(X)$. Exterior derivatives d' and d'' on supercurrents are induced by adjointness; a (p, q)-superform α embeds into $\mathcal{D}_{g-p,g-q}(X)$ via $\beta \mapsto \int_X \alpha \wedge \beta$, and a tropical k-cycle Z determines a (k, k)-supercurrent δ_Z via $\beta \mapsto \int_Z \beta$ (see [Gub16, §3.4]). The *Laplacian* is the operator $dd^c := d'd'' \colon \mathcal{D}_{p,q}(X) \to \mathcal{D}_{p-1,q-1}(X)$.

Remark 3.9. Our corner locus (Lemma 3.8) uses incoming slopes (the min-convention), whereas [GK17] uses outgoing slopes; the two differ by an overall sign, which accounts for the sign in the Poincaré–Lelong formula below.

Proposition 3.10 (Poincaré–Lelong formula)**.** *In* $\mathcal{D}_{g-1,g-1}(X)$,

$$dd^c\|\Psi\| = \nu_\phi - \delta_{\|\Psi\|\cdot X}.$$

Proof. Applying [GK17, Corollary 3.19] (with the sign convention of Remark 3.9) gives

$$dd^c\|\Psi\| = d'_{\mathrm{P}}d''_{\mathrm{P}}\|\Psi\| - \delta_{\|\Psi\|\cdot N_{\mathbb{R}}}$$

in $\mathcal{D}_{g-1,g-1}(N_{\mathbb{R}})$, where d'_{P} and d''_{P} are the polyhedral derivatives of [GK17, Definition 2.3]. By [GK17, Remarks 1.7 and 2.8], the polyhedral derivatives act by summing the smooth part over facets, so

$$d'_{\mathrm{P}}d''_{\mathrm{P}}\|\Psi\| = \sum_{\lambda\in\Lambda} d'd''\left(\|\Psi\||_{\sigma_\lambda}\right) \wedge \delta_{\sigma_\lambda}.$$

On σ_0 one computes directly from (3.14) and Proposition 3.4:

$$d'd''\left(\|\Psi\||_{\sigma_0}\right) = \sum_{i,j}(\mathbf{G}^{-1})_{i,j}\, d'x_i \wedge d''x_j = \nu_\phi,$$

where the last equality uses $\mathbf{E} = \mathbf{I}$ (principal polarization) in (3.11). The statement in $N_{\mathbb{R}}$ descends to X by a standard partition-of-unity argument. $\square$

3.7.3 Pullback to a metric graph and the Zhang measure

We now compute the pullback of ν_ϕ along the Abel–Jacobi map to a single edge of Γ. Let $e \in E$ be an edge parameterized by $\iota_e\colon [0,\ell(e)] \to \Gamma$ with coordinate t_e. Write $\mathbf{M}$ for the $g \times |E|$ matrix representing the inclusion $H_1(G,\mathbb{R}) \hookrightarrow C_1(G,\mathbb{R})$ in the bases $\gamma_1,\dots,\gamma_g$ and E:

$$\gamma_i = \sum_{e\in E} \mathbf{M}_{i,e}\, e.$$

One checks that $\mathbf{M}_{i,e} = \ell(e)^{-1}[e,\gamma_i]_G$. From this and the Abel–Jacobi map formula, the pullback of dx_i satisfies

$$(\Phi_q^{(1)} \circ \iota_e)^* dx_i = \mathbf{M}_{i,e}\, dt_e. \tag{3.15}$$

Define the matrix $\mathbf{P} := \mathbf{M}^\top\mathbf{G}^{-1}\mathbf{M}\mathbf{D}$, where $\mathbf{D}$ is the diagonal matrix with $\mathbf{D}_{e,e} = \ell(e)$. The following is a standard result in electrical network theory, for which we include a proof for completeness.

Lemma 3.11. *The matrix* $\mathbf{P}$ *represents the orthogonal projection* $C_1(G;\mathbb{R}) \twoheadrightarrow H_1(G;\mathbb{R})$ *in the basis* E *with respect to the length pairing* $[\,\cdot\,,\,\cdot\,]_G$.

Proof. Using $\mathbf{G} = \mathbf{M}\mathbf{D}\mathbf{M}^\top$, one computes $\mathbf{P}^2 = \mathbf{M}^\top\mathbf{G}^{-1}\mathbf{M}\mathbf{D}\mathbf{M}^\top\mathbf{G}^{-1}\mathbf{M}\mathbf{D} = \mathbf{P}$, so $\mathbf{P}$ is a projection. The adjoint with respect to $[\,\cdot\,,\,\cdot\,]_G$ is $\mathbf{P}^* = \mathbf{D}^{-1}\mathbf{P}^\top\mathbf{D}$; since $\mathbf{G}$ and $\mathbf{D}$ are symmetric, $\mathbf{P}^* = \mathbf{M}^\top\mathbf{G}^{-1}\mathbf{M}\mathbf{D} = \mathbf{P}$, confirming self-adjointness. Finally, $\mathbf{P}\mathbf{M}^\top = \mathbf{M}^\top\mathbf{G}^{-1}\mathbf{G} = \mathbf{M}^\top$, so $\mathbf{P}$ acts as the identity on $H_1(G;\mathbb{R})$. $\square$

Corollary 3.12. *For each edge* $e \in E$,

$$(\Phi_q^{(1)} \circ \iota_e)^*\nu_\phi = \ell(e)^{-1}\mathbf{P}_{e,e}\, d't_e \wedge d''t_e. \tag{3.16}$$

In particular, $\mathbf{P}_{e,e} = F(e)$ *is the* Foster coefficient *of* e *[BF11, §6.1], and the pullback of* ν_ϕ *to* Γ *via* $\Phi_q^{(1)}$ *is the Zhang measure* μ_{Zh} *introduced in [Zha93].*

Proof. Using (3.15) and $\mathbf{E} = \mathbf{I}$, the pullback of $d'x_i \wedge d''x_j$ is $\mathbf{M}_{i,e}\mathbf{M}_{j,e}\, d't_e \wedge d''t_e$. Substituting into (3.11) with $\mathbf{A} = \mathbf{G}^{-1}$ gives $\sum_{i,j}(\mathbf{G}^{-1})_{i,j}\mathbf{M}_{i,e}\mathbf{M}_{j,e} = (\mathbf{M}^\top\mathbf{G}^{-1}\mathbf{M})_{e,e} = \ell(e)^{-1}\mathbf{P}_{e,e}$, as claimed. The identification $\mathbf{P}_{e,e} = F(e)$ follows from [BF11, §6.1]. $\square$

3.8 Volumes of the Effective Loci

We prove that the integral of $\frac{1}{d!}\nu_\phi^{\wedge d}$ over the d-th effective locus $\widetilde{W}_d$ equals $\binom{g}{d}$. The proof proceeds in two steps: first we compute the integral of basis superforms over $\widetilde{W}_d$ (Lemma 3.13), then combine with Lemma 3.6 via the Cauchy–Binet formula.

The following key lemma expresses the integral of a basis (d,d)-superform over $\widetilde{W}_d$ in terms of the Gram matrix of the length pairing.

Lemma 3.13. *For* $I, J \in \binom{[g]}{d}$,

$$\int_{\widetilde{W}_d} d'x_I \wedge d''x_J = (-1)^{\binom{d}{2}} \det(\mathbf{G}_{I,J}). \tag{3.17}$$

Proof. Let $\iota_{\mathcal{E}} \colon \prod_{e\in\mathcal{E}}[0,\ell(e)] \to \Gamma^d/\mathfrak{S}_d$ be the parameterization of the cell $\prod\mathcal{E}$. Using (3.15), for $I \in \binom{[g]}{d}$ and $\mathcal{E} \in \binom{E}{d}$ one computes

$$(\Phi_q^{(d)} \circ \iota_{\mathcal{E}})^* dx_I = \det(\mathbf{M}_{I,\mathcal{E}})\, dt_{\mathcal{E}}. \tag{3.18}$$

By the definition of the pullback of a superform,

$$(\Phi_q^{(d)} \circ \iota_{\mathcal{E}})^*(d'x_I \wedge d''x_J) = \det(\mathbf{M}_{I,\mathcal{E}})\det(\mathbf{M}_{J,\mathcal{E}})\, d't_{\mathcal{E}} \wedge d''t_{\mathcal{E}}.$$

Decomposing the integral via (3.10) and using the sign formula $(-1)^{\binom{d}{2}}$ from the superform integration convention (Section 2.3):

$$\begin{aligned}\int_{\widetilde{W}_d} d'x_I \wedge d''x_J &= \sum_{\mathcal{E}\in\binom{E}{d}} \det(\mathbf{M}_{I,\mathcal{E}})\det(\mathbf{M}_{J,\mathcal{E}}) \int_{\prod_{e\in\mathcal{E}}[0,\ell(e)]} dt_{\mathcal{E}} \cdot (-1)^{\binom{d}{2}} \\ &= (-1)^{\binom{d}{2}} \sum_{\mathcal{E}\in\binom{E}{d}} \det(\mathbf{M}_{I,\mathcal{E}})\det(\mathbf{M}_{J,\mathcal{E}})\det(\mathbf{D}_{\mathcal{E},\mathcal{E}}).\end{aligned}$$

Since $\mathbf{D}$ is diagonal, $\det(\mathbf{D}_{\mathcal{E},\mathcal{F}}) = 0$ for $\mathcal{E} \neq \mathcal{F}$; expanding the sum over both $\mathcal{E}$ and $\mathcal{F}$ and applying the Cauchy–Binet formula gives

$$\begin{aligned}&= (-1)^{\binom{d}{2}} \sum_{\mathcal{E},\mathcal{F}\in\binom{E}{d}} \det(\mathbf{M}_{I,\mathcal{E}})\det(\mathbf{D}_{\mathcal{E},\mathcal{F}})\det((\mathbf{M}^\top)_{\mathcal{F},J}) \\ &= (-1)^{\binom{d}{2}} \det\big((\mathbf{M}\mathbf{D}\mathbf{M}^\top)_{I,J}\big) = (-1)^{\binom{d}{2}}\det(\mathbf{G}_{I,J}).\end{aligned}$$ □

Theorem 3.14. *For a principally polarized tropical abelian variety (X,ϕ) of genus g,*

$$\int_{\widetilde{W}_d} \frac{1}{d!}\, \nu_\phi^{\wedge d} = \binom{g}{d}. \tag{3.19}$$

Proof. By Lemma 3.6 (with $\mathbf{A} = \mathbf{G}^{-1}$ since $\mathbf{E} = \mathbf{I}$),

$$\int_{\widetilde{W}_d} \frac{1}{d!}\, \nu_\phi^{\wedge d} = (-1)^{\binom{d}{2}} \sum_{I,J\in\binom{[g]}{d}} \det((\mathbf{G}^{-1})_{I,J}) \int_{\widetilde{W}_d} d'x_I \wedge d''x_J.$$

Applying Lemma 3.13 and using the symmetry $\det(\mathbf{G}_{I,J}) = \det(\mathbf{G}_{J,I})$:

$$
\begin{aligned}
&= \sum_{I,J\in\binom{[g]}{d}} \det((\mathbf{G}^{-1})_{I,J})\det(\mathbf{G}_{J,I}) \\
&= \sum_{I\in\binom{[g]}{d}} \det\big((\mathbf{G}^{-1}\mathbf{G})_{I,I}\big) && \text{(Cauchy–Binet)} \\
&= \sum_{I\in\binom{[g]}{d}} \det(\mathbf{I}_{I,I}) = \binom{g}{d}. && \square
\end{aligned}
$$

3.9 The Tropical Poincaré Formula

The following was first proved in the tropical setting by Gross–Shokrieh [GS23b] using intersection theory on tropical Jacobians. We give an alternative proof via the de Rham isomorphism of Chapter 2.

Theorem 3.15 (Tropical Poincaré formula)**.** *In* $H_{S,d,d}(\mathrm{Jac}(\Gamma))$,

$$[\widetilde{W}_d] = \frac{[\Theta]^{g-d}}{(g-d)!}. \tag{3.20}$$

Proof. We work with superform cohomology, using the de Rham isomorphism $\widetilde{\mathrm{dR}}$ from Theorem 2.12 and the compatibility of Poincaré duality from Theorem 2.15.

Let $\nu_Z \in H_D^{g-k,g-k}(X)$ denote the Poincaré dual of a tropical k-cycle Z, characterized uniquely in cohomology by $\int_X \nu_Z \wedge \alpha = \int_Z \alpha$ for all d''-closed (k,k)-superforms α; see [JSS19, Theorem 2]. The formula (3.20) is equivalent, via Poincaré duality, to the equality

$$\nu_{\widetilde{W}_d} = \frac{\nu_\Psi^{\wedge(g-d)}}{(g-d)!} \quad \text{in } H_D^{g-d,g-d}(X). \tag{3.21}$$

We first identify $\nu_\Psi = \nu_\phi$. By Proposition 3.4, $\widetilde{\mathrm{dR}}(\nu_\phi) = h_\phi$. By [GS23a, Proposition 5.12], $\mathrm{PD}(h_\phi) = \mathrm{cyc}(\Psi)$ in $H^S_{g-1,g-1}(X)$. Since $\widetilde{\mathrm{dR}}$ is compatible with both Poincaré duality maps (Theorem 2.15), we conclude $\nu_\Psi = \nu_\phi$.

To prove (3.21), it suffices by Poincaré duality to show that

$$\int_{\widetilde{W}_d} \alpha = \int_X \frac{\nu_\phi^{\wedge(g-d)}}{(g-d)!} \wedge \alpha$$

for all $\alpha \in H_D^{d,d}(X)$. Since the d''-closed (d,d)-superforms $\{d'x_I \wedge d''x_J\}_{I,J \in \binom{[g]}{d}}$ form a basis for $H_D^{d,d}(X)$, it suffices to check on these.

Fix $I, J \in \binom{[g]}{d}$. By Lemma 3.6 and the calculation in the proof of Theorem 3.14,

$$\begin{aligned}&\int_X \frac{\nu_\phi^{\wedge(g-d)}}{(g-d)!} \wedge d'x_I \wedge d''x_J \\ &\quad = (-1)^{\binom{g-d}{2}} \sum_{I',J' \in \binom{[g]}{g-d}} \det((\mathbf{G}^{-1})_{I',J'}) \int_X d'x_{I'} \wedge d''x_{J'} \wedge d'x_I \wedge d''x_J.\end{aligned}$$

The integral $\int_X d'x_{I'} \wedge d''x_{J'} \wedge d'x_I \wedge d''x_J$ vanishes unless $I' = [g] \setminus I$ and $J' = [g] \setminus J$, in which case, after reordering and applying the sign formula from Section 2.3,

$$\int_X d'x_{[g]\setminus I} \wedge d''x_{[g]\setminus J} \wedge d'x_I \wedge d''x_J = (-1)^{d(g-d)} \int_X d'x_{[g]\setminus I} \wedge d'x_I \wedge d''x_{[g]\setminus J} \wedge d''x_J,$$

the sign coming from (2.5) with $(p,q) = (g-d, g-d)$ and $(p',q') = (d,d)$. Sorting each of the two blocks into increasing order contributes $\operatorname{sgn}(\pi_I)\operatorname{sgn}(\pi_J)$, where π_I denotes the permutation $([g] \setminus I, I)$ and likewise for J. Combining this with the prefactor $(-1)^{\binom{g-d}{2}}$ and using $\binom{g}{2} = \binom{d}{2} + \binom{g-d}{2} + d(g-d)$, the accumulated sign is $(-1)^{\binom{g}{2}+\binom{d}{2}}$, so that

$$\int_X \frac{\nu_\phi^{\wedge(g-d)}}{(g-d)!} \wedge d'x_I \wedge d''x_J = (-1)^{\binom{g}{2}+\binom{d}{2}} \operatorname{sgn}(\pi_I)\operatorname{sgn}(\pi_J)\ \det((\mathbf{G}^{-1})_{[g]\setminus I,[g]\setminus J}) \int_X d'x_{[g]} \wedge d''x_{[g]}.$$

Evaluating the top-degree integral by (3.13):

$$\int_X \frac{\nu_\phi^{\wedge(g-d)}}{(g-d)!} \wedge d'x_I \wedge d''x_J = (-1)^{\binom{d}{2}} \operatorname{sgn}(\pi_I)\operatorname{sgn}(\pi_J) \det((\mathbf{G}^{-1})_{[g]\setminus I,[g]\setminus J}) \det(\mathbf{G}).$$

By the classical formula for complementary minors of an invertible matrix,

$$\det((\mathbf{G}^{-1})_{[g]\setminus I,[g]\setminus J}) \det(\mathbf{G}) = (-1)^{\sum I + \sum J} \det(\mathbf{G}_{J,I}).$$

It remains to show that $(-1)^{\sum I+\sum J}\mathrm{sgn}(\pi_I)\mathrm{sgn}(\pi_J)=1$, so that the right-hand side equals $(-1)^{\binom{d}{2}}\det(\mathbf{G}_{J,I})=\int_{\widetilde{W}_d} d'x_I\wedge d''x_J$ by Lemma 3.13 and symmetry of $\mathbf{G}$.

To verify the sign, consider the function $s(\pi):=\sum_{k=1}^{d} b_k$ on $(g-d,d)$-shuffles $\pi=(a_1,\dots,a_{g-d},b_1,\dots,b_d)$ with $a_1<\cdots<a_{g-d}$ and $b_1<\cdots<b_d$. For any adjacent transposition t that swaps an a-entry with a neighboring b-entry, we have $s(t\circ\pi)=s(\pi)-1$ and $\mathrm{sgn}(t\circ\pi)=-\mathrm{sgn}(\pi)$, so the quantity $(-1)^{s(\pi)}\mathrm{sgn}(\pi)$ is invariant under such transpositions. Any $(g-d,d)$-shuffle can be reached from the identity by a sequence of such transpositions, so $(-1)^{s(\pi)}\mathrm{sgn}(\pi)$ is a constant $c_d\in\{\pm1\}$ depending only on g and d, namely its value $(-1)^{s(\mathrm{id})}$ at the identity shuffle. It need not equal 1: for instance $c_1=-1$ when $g=3$. This is harmless here, because I and J both have cardinality d and therefore contribute the *same* constant:

$$(-1)^{\sum I+\sum J}\mathrm{sgn}(\pi_I)\mathrm{sgn}(\pi_J)=\Big((-1)^{s(\pi_I)}\mathrm{sgn}(\pi_I)\Big)\cdot\Big((-1)^{s(\pi_J)}\mathrm{sgn}(\pi_J)\Big)=c_d^2=1. \qquad \square$$

Chapter 4

BREAK DIVISORS AND THETA CHARACTERISTICS

4.1 Introduction

On a compact Riemann surface of genus g, a *theta characteristic* is a line bundle L satisfying $L^{\otimes 2} \cong \omega_C$, where ω_C is the canonical bundle. There are exactly 2^g such classes, indexed by the 2-torsion of the Jacobian. A tropical counterpart of this picture was developed by Zharkov [Zha10]: for a metric graph Γ of genus g, the *tropical theta characteristics* are the classes $\mathcal{K} \in \mathrm{Pic}^{g-1}(\Gamma)$ satisfying $2\mathcal{K} = [K]$, where $K = \sum_p (\mathrm{val}(p) - 2)\, p$ is the canonical divisor. There are again exactly 2^g of them, and they are naturally indexed by $H_1(G, \mathbb{Z}/2\mathbb{Z})$ for any weighted model G of Γ. Zharkov produced an explicit effective divisor $\mathcal{K}_\gamma^-$ representing each class, defined via the gradient orientation of the distance function to the support of γ.

The purpose of this chapter is to give a different construction of the tropical theta characteristics, in which the algorithm producing them is a shortest-path computation. Fix a base vertex $q \in V(G)$. For a flow $f \in H_1(G, \mathbb{R})$, define directed edge weights

$$w_f(e) := \ell(e)\Big(1 - 2f(e)\Big), \qquad w_f(\bar{e}) := \ell(e)\Big(1 + 2f(e)\Big),$$

and let G_f be the resulting directed weighted graph on $V(G)$. The first structural observation is that

$$f \in \mathrm{Vor}(G) \quad \Longleftrightarrow \quad G_f \text{ has no negative directed cycle,}$$

so the (central) Voronoi polytope of the lattice $H_1(G, \mathbb{Z})$ is exactly the parameter region on which the Bellman–Ford shortest-path problem is feasible. For $f \in \mathrm{Vor}(G)$, Bellman–Ford produces distances $d_f(v)$ from q together with a shortest-path tree T, and for each non-tree edge $e = (u, v)$ the point

$$p_e := \text{point of } e \text{ at distance } \tfrac{1}{2}\Big(\ell(e) + d_f(v) - d_f(u) + 2\ell(e)f(e)\Big) \text{ from } u$$

is well defined. Setting $B(f) := \sum_{e \notin T} p_e$ produces an effective divisor of degree g, in the sense of An–Baker–Kuperberg–Shokrieh [ABKS14]. Evaluating this construction at the 2-torsion inputs $f_\gamma := \frac{1}{2}\tilde{\gamma}$, for $\tilde{\gamma} \in H_1(G, \mathbb{Z})$ any integer lift of γ with coefficients in $\{-1, 0, 1\}$, produces exactly the 2^g tropical theta characteristics.

Theorem 4.1 (Bellman–Ford computes theta characteristics; Theorem 4.30)**.** *Let Γ be a metric graph of genus g with weighted model G, and let $q \in V(G)$. For every $\gamma \in H_1(G, \mathbb{Z}/2\mathbb{Z})$:*

1. *$f_\gamma \in \mathrm{Vor}(G)$, and Bellman–Ford on G_{f_γ} produces a break divisor $B(f_\gamma)$ that is independent both of the chosen shortest-path tree and of the chosen $\{-1, 0, 1\}$-valued lift of γ;*

2. *the class $\vartheta_\gamma := [B(f_\gamma) - q] \in \mathrm{Pic}^{g-1}(\Gamma)$ is a theta characteristic, and*

$$\vartheta_\gamma - \vartheta_0 \ = \ \left[\tfrac{1}{2}\tilde{\gamma}\right] \quad \textit{in } \mathrm{Jac}(\Gamma);$$

3. *the map $\gamma \mapsto \vartheta_\gamma$ is a bijection between $H_1(G, \mathbb{Z}/2\mathbb{Z})$ and the set of tropical theta characteristics.*

Moreover, $\vartheta_\gamma = \mathcal{K}_\gamma$ for Zharkov's classes.

The proof has two ingredients. The first is that the Abel–Jacobi class $[B(f) - q]$ is affine in f (Lemma 4.20), so once ϑ_0 has been identified as half-canonical the identification of every ϑ_γ as half-canonical follows. The second is that ϑ_0 is indeed half-canonical (Proposition 4.26), which we prove by orientation-reversal on the model refined at the interior maxima of the distance function d_q. The comparison with Zharkov's family reduces to a single divisor-level identity $B(0) - q = \mathcal{K}_q^-$, which is immediate from the construction: the orientation defining $\mathcal{K}_q^-$ is by construction the gradient orientation of the Bellman–Ford distance potential d_q.

One methodological point is worth flagging at the outset. For the 2-torsion inputs f_γ, the coefficients of $\tilde{\gamma}$ lie in $\{-1, 0, 1\}$, so

$$w_{f_\gamma}(e) = \ell(e)\Big(1 - \tilde{\gamma}(e)\Big) \in \{0,\ \ell(e),\ 2\ell(e)\}.$$

These weights are nonnegative and Dijkstra's algorithm suffices. For a general $f \in \mathrm{Vor}(G)$, however, negative edge weights do occur (Remark 4.14), and Bellman–Ford is the correct algorithm. The general theorem is proved for arbitrary $f \in \mathrm{Vor}(G)$; the specialization to 2-torsion inputs is only invoked at the very last step.

Outline. Section 4.2 fixes the Hodge decomposition of $C_1(G, \mathbb{R})$, defines the Voronoi polytope, and recalls its zonotopal subdivision by spanning trees. Section 4.3 proves that Voronoi-cell membership coincides with feasibility of the Bellman–Ford relaxation inequalities, and develops the associated stable-exchange machinery. Section 4.4 defines break divisors from optimal shortest-path trees, proves the linearity lemma, and treats the base case $f = 0$; the base characteristic subsection shows ϑ_0 is half-canonical. Section 4.5 proves Theorem 4.1 by specializing to the 2-torsion inputs, and identifies ϑ_γ with Zharkov's $\mathcal{K}_\gamma$. Section 4.6 works the construction out on the complete graph K_4 with unit edge lengths and on the diamond graph with weights $(2, 2, 1, 2, 2)$. The first exhibits the tree-path term in the Abel–Jacobi formula, together with a case where our representative $B(f_\gamma) - q$ differs at the divisor level from Zharkov's $\mathcal{K}_\gamma^-$; the second exhibits a case with three optimal spanning trees, connected by stable exchanges along reverse-oriented tight edges.

Notation. Throughout, G is a finite connected graph with no loop edges, equipped with a length function $\ell \colon E(G) \to \mathbb{R}_{>0}$; the associated metric graph is Γ. We write $\mathbb{E}(G)$ for the set of directed edges and fix a reference orientation $\mathcal{O}(G) \subset \mathbb{E}(G)$ choosing one directed representative $e = (e^-, e^+)$ of each edge, with reverse $\bar{e} = (e^+, e^-)$ and $\ell(\bar{e}) = \ell(e)$. Thus $\mathbb{E}(G) = \mathcal{O}(G) \sqcup \overline{\mathcal{O}(G)}$. The assignment symbol is $:=$.

4.2 Weighted Graphs, the Hodge Decomposition, and the Voronoi Polytope

In this section we set up the ambient homological algebra of a weighted graph and recall the two structural results of de Jong–Shokrieh [dJS23] that we build on: the Hodge decomposition of $C_1(G, \mathbb{R})$ and the zonotopal subdivision of $\mathrm{Vor}(G)$ by spanning trees.

4.2.1 Chain complex and Hodge decomposition

The *space of* 1*-chains* is $C_1(G,\mathbb{R}) := \mathbb{R}^{\mathcal{O}(G)}$, equipped with the inner product

$$[e_i, e_j] := \delta_{ij}\,\ell(e_i), \qquad e_i, e_j \in \mathcal{O}(G),$$

extended bilinearly; see [dJS23, §3.3]. Following [BF11, §2.1], define the *coboundary operator* $d\colon C_0(G,\mathbb{R}) \to C_1(G,\mathbb{R})$ by

$$(dh)(e) := \frac{h(e^+) - h(e^-)}{\ell(e)},$$

and its adjoint $d^*\colon C_1(G,\mathbb{R}) \to C_0(G,\mathbb{R})$ by $(d^*\alpha)(x) = \sum_{e^+=x} \alpha(e) - \sum_{e^-=x} \alpha(e)$. The *flow space* is $H_1(G,\mathbb{R}) := \ker(d^*)$; its integer sublattice is $H_1(G,\mathbb{Z}) := \ker(d^*) \cap C_1(G,\mathbb{Z})$, and the *genus* of G is $g := \dim_{\mathbb{R}} H_1(G,\mathbb{R}) = |E(G)| - |V(G)| + 1$.

Proposition 4.2 (Hodge decomposition, [BF11, eq. (2.1)])**.** *There is an orthogonal direct sum decomposition*

$$C_1(G,\mathbb{R}) \;=\; H_1(G,\mathbb{R}) \;\oplus\; \mathrm{im}(d).$$

Proof. [BF11, eq. (2.1)]. □

Write $\pi\colon C_1(G,\mathbb{R}) \to H_1(G,\mathbb{R})$ for the orthogonal projection, so $\ker\pi = \mathrm{im}(d)$. Two consequences are used repeatedly and we record them once.

Lemma 4.3. *For every $c \in C_1(G,\mathbb{R})$ and every $x \in H_1(G,\mathbb{R})$ we have $[\pi(c), x] = [c, x]$; in particular $[\pi(e), x] = \ell(e)\,x(e)$ for $e \in \mathcal{O}(G)$.*

Proof. $c - \pi(c) \in \mathrm{im}(d) = H_1(G,\mathbb{R})^\perp$. □

By [dJS23, eq. (4.6)] the diagonal entry $\pi(e)(e)$ equals the *Foster coefficient* $F(e) = 1 - r(e^-, e^+)/\ell(e)$, where r denotes effective resistance; we use this only in Section 4.6.

4.2.2 The Voronoi polytope

The Voronoi polytope of $H_1(G,\mathbb{Z})$ inside $H_1(G,\mathbb{R})$ plays the role of a fundamental domain for the action of the integer flows on the flow space.

Definition 4.4 (Voronoi polytope, [dJS23, §2.2])**.** The *Voronoi polytope* of G is

$$\mathrm{Vor}(G) := \big\{ f \in H_1(G,\mathbb{R}) : [f,f] \le [f-z, f-z] \text{ for all } z \in H_1(G,\mathbb{Z}) \big\}.$$

Expanding the square, the defining inequalities read

$$2[f,z] \le [z,z] \qquad \text{for all } z \in H_1(G,\mathbb{Z}). \tag{4.1}$$

Recall that the *zonotope* generated by a finite set $V \subset H_1(G,\mathbb{R})$ is $\mathcal{Z}(V) = \big\{ \sum_i \alpha_i v_i : |\alpha_i| \le 1 \big\}$.

Proposition 4.5 (Zonotope formula, [dJS23, Prop. 7.3]; cf. [OS79, Prop. 5.2])**.**

$$\mathrm{Vor}(G) \;=\; \tfrac{1}{2}\,\mathcal{Z}\Big(\{ \pi(e) : e \in \mathcal{O}(G) \} \Big).$$

Proof. [dJS23, Prop. 7.3]. □

Fix a base vertex $q \in V(G)$. For a spanning tree T of G, let $T_q \subset \mathbb{E}(G)$ denote T oriented away from q; see [dJS23, §3.1]. Following [dJS23, Def. 6.8, §7], set

$$\sigma_T := \tfrac{1}{2} \sum_{a \in T_q} \pi(a), \qquad \mathcal{C}_T := \tfrac{1}{2}\,\mathcal{Z}\Big(\{ \pi(e) : e \in \mathcal{O}(G) \setminus T \} \Big).$$

Theorem 4.6 (Zonotopal subdivision, [dJS23, Thm. 7.5])**.** *The parallelotopes* $\{\sigma_T + \mathcal{C}_T\}$*, indexed by the spanning trees* T *of* G*, form a polyhedral subdivision of* $\mathrm{Vor}(G)$*.*

Proof. [dJS23, Thm. 7.5]. □

By [dJS23, Lem. 7.4(a)], the set $\{\pi(e) : e \in \mathcal{O}(G) \setminus T\}$ is a basis of $H_1(G,\mathbb{R})$ for every spanning tree T. Hence every $f \in H_1(G,\mathbb{R})$ admits a unique representation

$$f \;=\; \sigma_T + \sum_{e \in \mathcal{O}(G) \setminus T} \alpha_e\, \pi(e), \qquad \alpha_e \in \mathbb{R}, \tag{4.2}$$

and $f \in \sigma_T + \mathcal{C}_T$ if and only if $|\alpha_e| \le \frac{1}{2}$ for every $e \notin T$.

4.2.3 Divisors

A *divisor* on Γ is a finite formal $\mathbb{Z}$-linear combination of points of Γ. We write $\mathrm{Div}^d(\Gamma)$ for the degree-d divisors and $\mathrm{Div}^d_+(\Gamma)$ for the effective ones. Two divisors are *linearly equivalent* if their difference is $\mathrm{div}(\varphi)$ for a continuous piecewise-linear function φ on Γ with integer slopes, where $\mathrm{div}(\varphi) = \sum_p \Big(\sum_\xi \partial_\xi \varphi(p)\Big)\, p$ and ξ runs over the outgoing primitive tangent directions at p. We write $\mathrm{Pic}^d(\Gamma)$ for the set of degree-d classes, $K = \sum_p (\mathrm{val}(p) - 2)\, p$ for the canonical divisor, and $\mathrm{Jac}(\Gamma) := \mathrm{Pic}^0(\Gamma) \cong H_1(G, \mathbb{R})/H_1(G, \mathbb{Z})$; see [dJS23, §3.2] and [ABKS14, §2.1].

4.3 Bellman–Ford and Cell Membership

The combinatorial heart of the chapter is the following: for a flow $f \in \mathrm{Vor}(G)$, the cell of the subdivision of Theorem 4.6 containing f is exactly the shortest-path tree produced by Bellman–Ford in a weighted digraph attached to f. This section makes the identification precise and develops the associated notion of stable exchange between optimal trees.

Definition 4.7 (Induced weights)**.** For $f \in H_1(G, \mathbb{R})$ and $e \in \mathcal{O}(G)$, the *induced weight* is

$$w_f(e) := \ell(e)\Big(1 - 2f(e)\Big), \qquad w_f(\bar{e}) := \ell(e)\Big(1 + 2f(e)\Big).$$

We write G_f for the directed weighted graph on $V(G)$ with edge set $\mathbb{E}(G)$ and weight function w_f.

Note the elementary identity

$$w_f(e) + w_f(\bar{e}) = 2\ell(e) > 0, \tag{4.3}$$

which will be used repeatedly. The following proposition is the identification of the Voronoi polytope with the parameter region on which Bellman–Ford is feasible.

Proposition 4.8 (Voronoi cell = no-negative-cycle region)**.** *For $f \in H_1(G, \mathbb{R})$, the following are equivalent:*

1. $f \in \mathrm{Vor}(G)$*;*

2. every directed cycle in G_f *has nonnegative* w_f*-weight.*

Proof. Let C be a directed simple cycle that does not immediately traverse an edge and its reverse. Its characteristic flow $z_C \in H_1(G, \mathbb{Z})$ has coefficients in $\{0, \pm 1\}$, and

$$w_f(C) = \sum_{e \in C} \ell(e) - 2 \sum_{e \in C} \ell(e)\, f(e)\, z_C(e) = [z_C, z_C] - 2[f, z_C].$$

If (1) holds, then (2) follows from (4.1); the only remaining simple directed cycle is the two-arc backtrack $e\bar{e}$, which has weight $2\ell(e) > 0$ by (4.3).

Conversely, let $z \in H_1(G, \mathbb{Z})$. Replace each underlying edge e by $|z(e)|$ copies oriented according to the sign of $z(e)$; the flow condition on z makes the resulting directed multigraph balanced at every vertex, so its arcs decompose into directed simple cycles $C_1, \dots, C_N$. By (2), each $w_f(C_i) \geq 0$, hence

$$2[f, z] \;=\; \sum_{i=1}^{N} 2[f, z_{C_i}] \;\leq\; \sum_{i=1}^{N} [z_{C_i}, z_{C_i}] \;=\; \sum_{e \in E(G)} \ell(e)\, |z(e)| \;\leq\; \sum_{e \in E(G)} \ell(e)\, z(e)^2 = [z, z],$$

which is (4.1). □

4.3.1 The coefficient formula

The next definition attaches to a spanning tree T and a flow f a vertex-labelling function; it plays the role of the Bellman–Ford shortest-path distance function, and coincides with it exactly when T is optimal (Definition 4.11 below). No optimality hypothesis is required in Definition 4.9 or Proposition 4.10.

Definition 4.9 (Tree-path labels)**.** Let T be a spanning tree of G and $f \in H_1(G, \mathbb{R})$. For $v \in V(G)$ let $P^T_{q,v}$ denote the unique path from q to v in T, regarded both as a directed path in T_q and as a 1-chain $P^T_{q,v} = \sum_a \varepsilon_a\, a \in C_1(G, \mathbb{R})$, where the sum is over the edges a of the path and $\varepsilon_a = \pm 1$ according as the away-from-q direction of a agrees with $\mathcal{O}(G)$ or not. Set

$$d_{T,f}(v) := w_f\Big(P^T_{q,v}\Big) = \sum_{a \in P^T_{q,v}} \ell(a)\Big(1 - 2\varepsilon_a\, f(a)\Big).$$

Proposition 4.10 (Coefficient formula)**.** *Let T be any spanning tree of G, let $f \in H_1(G, \mathbb{R})$, and let $e = (e^-, e^+) \in \mathcal{O}(G) \setminus T$. Then the coefficient α_e in (4.2) is*

$$\alpha_e \;=\; \frac{d_{T,f}(e^+) - d_{T,f}(e^-) + 2\ell(e)\, f(e)}{2\ell(e)}. \tag{4.4}$$

Proof. Let $C_e := P^T_{q,e^-} - P^T_{q,e^+} + e$ be the fundamental cycle of e relative to T, and put $L_e := \ell(P^T_{q,e^-}) - \ell(P^T_{q,e^+})$. The tree-path label admits the orientation-free expression

$$d_{T,f}(v) = \ell(P^T_{q,v}) - 2[f, P^T_{q,v}].$$

Substituting $P^T_{q,e^-} - P^T_{q,e^+} = C_e - e$ gives

$$d_{T,f}(e^-) - d_{T,f}(e^+) = L_e - 2[f, C_e] + 2\ell(e)\, f(e). \tag{4.5}$$

Pairing the unique representation (4.2) with C_e and using Lemma 4.3, the only non-tree generator pairing nontrivially with C_e is $\pi(e)$, with $[\pi(e), C_e] = \ell(e)$; a signed count over tree edges (using the sign convention ε_a of Definition 4.9) gives $[\sigma_T, C_e] = \frac{1}{2}L_e$. Hence

$$[f, C_e] = \tfrac{1}{2}L_e + \alpha_e\, \ell(e). \tag{4.6}$$

Eliminating $[f, C_e]$ between (4.5) and (4.6) yields (4.4). □

4.3.2 Cell membership

Definition 4.11 ((f, q)-optimal tree)**.** A spanning tree T is (f, q)-*optimal* if $d_{T,f}(v)$ is, for every $v \in V(G)$, the minimum w_f-weight of a directed q-to-v walk in G_f; equivalently, $d_{T,f}$ is the shortest-path distance function from q in G_f.

Theorem 4.12 (Voronoi cell membership is Bellman–Ford feasibility)**.** *Let $f \in H_1(G, \mathbb{R})$ and let T be a spanning tree of G. The following are equivalent:*

1. $f \in \sigma_T + \mathcal{C}_T$;

2. $|\alpha_e| \le \frac{1}{2}$ *for every* $e \in \mathcal{O}(G) \setminus T$;

3. $d_{T,f}(v) \le d_{T,f}(u) + w_f(u,v)$ *for every* $(u,v) \in \mathbb{E}(G)$*;*

4. T *is* (f,q)*-optimal.*

Proof. (1) $\Leftrightarrow$ (2) is the definition of $\mathcal{C}_T$ together with the uniqueness of the representation (4.2).

(2) $\Leftrightarrow$ (3). Fix $e \in \mathcal{O}(G) \setminus T$. By (4.4),

$$\alpha_e \le \tfrac{1}{2} \iff d_{T,f}(e^+) \le d_{T,f}(e^-) + \ell(e)\big(1 - 2f(e)\big) = d_{T,f}(e^-) + w_f(e),$$

and symmetrically

$$\alpha_e \ge -\tfrac{1}{2} \iff d_{T,f}(e^-) \le d_{T,f}(e^+) + w_f(\bar{e}).$$

So (2) is precisely the Bellman–Ford relaxation condition on $\mathbb{E}(G) \setminus T_q \setminus \overline{T_q}$. For a tree edge oriented $u \to v$ away from q, the forward relaxation $d_{T,f}(v) = d_{T,f}(u) + w_f(u,v)$ holds by Definition 4.9, and the reverse relaxation then follows strictly from (4.3). Hence the tree-edge relaxations are automatic and (2) $\Leftrightarrow$ (3).

(3) $\Rightarrow$ (4). Let $q = v_0 \to v_1 \to \cdots \to v_r = v$ be any directed walk in G_f. Summing (3) telescopes to $d_{T,f}(v) \le \sum_{i=1}^{r} w_f(v_{i-1}, v_i)$, and $d_{T,f}(v)$ is itself attained by the tree walk.

(4) $\Rightarrow$ (3) is immediate, since the tree path to v competes with the tree path to u followed by the arc (u,v). $\square$

Note that Theorem 4.12 needs no hypothesis on negative cycles; that enters only through the following existence corollary.

Corollary 4.13 (Existence)**.** *For every* $f \in \mathrm{Vor}(G)$ *and* $q \in V(G)$*, an* (f,q)*-optimal tree exists, and can be computed in time* $O(|V(G)| \cdot |E(G)|)$ *by the Bellman–Ford method.*

Proof. By Proposition 4.8 the graph G_f has no negative directed cycle, so a shortest-path tree rooted at q exists [Sch03, Thm. 8.1] and Bellman–Ford finds one in the stated time [Sch03, Thm. 8.5]. Theorem 4.12 then identifies it as (f,q)-optimal. If $w_f \ge 0$ throughout, Dijkstra's algorithm also applies [Sch03, Thm. 7.2]. $\square$

Remark 4.14 (Negative edges occur for general f)**.** Bellman–Ford is genuinely required for a general $f \in \mathrm{Vor}(G)$. On the theta graph with three parallel edges e_1, e_2, e_3 oriented $q \to v$ of lengths $1, 2, 3$, one computes $r(q, v) = 6/11$ and

$$\pi(e_1) = \tfrac{1}{11}(5, -3, -2), \quad \pi(e_2) = \tfrac{1}{11}(-6, 8, -2), \quad \pi(e_3) = \tfrac{1}{11}(-6, -3, 9).$$

Set $f := \frac{1}{2}\big(\pi(e_1) - \pi(e_2) - \pi(e_3)\big) = \frac{1}{22}(17, -8, -9)$. Then $f \in \mathrm{Vor}(G)$ by Proposition 4.5, yet $w_f(e_1) = 1 - 34/22 = -6/11 < 0$. No directed *cycle* has negative weight; for instance e_1 followed by $\bar{e}_2$ costs $-6/11 + 6/11 = 0$. Negative edges but no negative cycles is exactly the regime of Bellman–Ford.

4.3.3 Stable exchanges

Two distinct spanning trees can be (f, q)-optimal for the same f, namely when f lies on a shared face of two cells. The next definition ranges over *directed* edges rather than just $\mathcal{O}(G)$; this is essential, as Section 4.6 illustrates.

Definition 4.15 (Stable edges and exchanges)**.** Let T be (f, q)-optimal. A directed edge $a \in \mathbb{E}(G)$ not lying in T_q is *tight* if $d_{T,f}(a^-) + w_f(a) = d_{T,f}(a^+)$, and *stable* if in addition adding a to T_q creates no directed cycle. The *stable exchange* along a stable edge a inserts a into T_q and deletes the unique edge of T_q with head a^+.

Lemma 4.16 (Tight edges carry their break point at the head)**.** *Let T be (f, q)-optimal and let a be a tight directed edge lying over $e \in \mathcal{O}(G) \setminus T$. If $a = e$ then $\alpha_e = \frac{1}{2}$; if $a = \bar{e}$ then $\alpha_e = -\frac{1}{2}$. In either case the point of e at distance $\ell(e)(\frac{1}{2} + \alpha_e)$ from e^- equals a^+.*

Proof. Immediate from the two displayed equivalences in the proof of Theorem 4.12 taken with equality. If $\alpha_e = \frac{1}{2}$, the point is at distance $\ell(e)$ from e^-, hence $e^+ = a^+$; if $\alpha_e = -\frac{1}{2}$, the point is at distance 0, hence $e^- = a^+$. $\square$

Theorem 4.17 (Stable exchange connectivity)**.** *Let $f \in \mathrm{Vor}(G)$.*

1. *If T' is (f,q)-optimal and a is stable relative to T', then the tree T'' obtained by the stable exchange along a is (f,q)-optimal.*

2. *If T and T' are both (f,q)-optimal, then T is obtained from T' by at most $|V(G)|-1$ stable exchanges.*

Proof. (1) For any v, the directed path from q to v in T''_q either agrees with the one in T'_q, or is obtained from it by replacing the sub-path $q \to a^+$ by $q \to a^- \to a^+$. Tightness of a says these have equal w_f-weight, so $d_{T'',f} = d_{T',f}$ and T'' is (f,q)-optimal by Theorem 4.12.

(2) Both label functions coincide with the shortest-path distance function from q; write it simply as d_f. Say T and T' *agree at* v if T_q and T'_q contain the same directed path from q to v. They agree at q, so they disagree at at most $|V(G)|-1$ vertices; if they agree everywhere we are done.

No directed edge a can satisfy $a \in T_q$ and $\bar{a} \in T'_q$: optimality of both trees would force $w_f(a) + w_f(\bar{a}) = 0$, contradicting (4.3). Search outward from q in T_q and let $a \in T_q \setminus T'_q$ be the first edge encountered for which T and T' agree at a^-; such an a exists as long as some vertex of disagreement remains. Then $d_f(a^-) + w_f(a) = d_f(a^+)$, so a is tight relative to T'. Since T and T' agree at a^-, the vertex a^+ does not lie on the T'_q-path from q to a^-, so adding a to T'_q creates no directed cycle and a is stable. The exchange produces a tree agreeing with T at all vertices where T' did, and additionally at a^+, reducing the disagreement count by at least one. $\square$

4.4 Break Divisors, Linearity, and the Base Characteristic

We now use the Bellman–Ford machinery of Section 4.3 to attach an effective divisor to each $f \in \mathrm{Vor}(G)$. The main result of this section is that the Abel–Jacobi class of this divisor depends affinely on f (Lemma 4.20), and that its base value $[B(0)-q]$ is half-canonical (Proposition 4.26).

Definition 4.18 (Break divisor; cf. [ABKS14, §3.1], [MZ08, §4.5])**.** Let $f \in \mathrm{Vor}(G)$ and let T be (f,q)-optimal. For each $e \in \mathcal{O}(G) \setminus T$, the *break point* p_e is the point of e at distance

$\ell(e)(\frac{1}{2}+\alpha_e)$ from e^-, and the *break divisor* is

$$B(f) \;:=\; \sum_{e\in\mathcal{O}(G)\setminus T} p_e \;\in\; \mathrm{Div}^g_+(\Gamma).$$

By Theorem 4.12 the coefficient $\frac{1}{2}+\alpha_e$ lies in $[0,1]$, so p_e genuinely lies on e. Removing the g open edge interiors $\{e^\circ : e \notin T\}$ from Γ leaves the contractible subgraph T, so $B(f)$ is a break divisor in the sense of [ABKS14, §3.1].

Lemma 4.19 (Independence of the optimal tree)**.** *$B(f)$ does not depend on the choice of (f,q)-optimal tree.*

Proof. By Theorem 4.17 it suffices to verify invariance under a single stable exchange. Let a be stable relative to T' and lie over $e_1 \in \mathcal{O}(G)\setminus T'$; by Lemma 4.16, $p_{e_1} = a^+$. Let e_1' be the tree edge of T'_q deleted by the exchange, with head a^+, and let a' be its directed form in T'_q. In the new tree T'', the path from q to a^+ runs through a, so $d_{T'',f} = d_{T',f}$ by Theorem 4.17(1); and since a' was a tree edge of T'_q,

$$d_{T'',f}(a'^-) + w_f(a') = d_{T'',f}(a'^+) = d_{T'',f}(a^+),$$

so a' is tight relative to T'', and Lemma 4.16 places $p_{e_1'} = a'^+ = a^+ = p_{e_1}$. The remaining break points are unchanged, since the labels $d_{T,f}$ and the coefficient formula (4.4) are unchanged. □

4.4.1 Abel–Jacobi images and the linearity lemma

For a divisor $D = \sum_i n_i p_i \in \mathrm{Div}^0(\Gamma)$, the *Abel–Jacobi map* is

$$\mu^0([D]) \;=\; \Big[\sum_i n_i\, \pi(P_{q,p_i})\Big] \;\in\; \mathrm{Jac}(\Gamma),$$

where P_{q,p_i} is any path in Γ from q to p_i, regarded as a 1-chain in a model refining G at p_i; the class is independent of the path choices [BF11, §§3–4], [ABKS14, §2.3]. For divisors of degree d set $\mu^d([D]) := \mu^0([D - d\cdot q])$; by [BF11, Thm. 3.4], μ^d is a basepoint-dependent affine bijection $\mathrm{Pic}^d(\Gamma) \to \mathrm{Jac}(\Gamma)$, compatible with the Pic^0-torsor action. We repeatedly use the identity $\mu^g([D]) = \mu^{g-1}([D-q])$, immediate from the definitions.

For a break point p_e we take the path $P_{q,p_e} = P^T_{q,e^-} + (\frac{1}{2} + \alpha_e)\, e$, giving

$$\mu^1\big([p_e - q]\big) \;=\; \big[\pi(P^T_{q,e^-})\big] + \big[(\tfrac{1}{2} + \alpha_e)\, \pi(e)\big]. \tag{4.7}$$

Lemma 4.20 (Linearity lemma)**.** *For all $f, f' \in \mathrm{Vor}(G)$,*

$$\mu^g\big([B(f)]\big) - \mu^g\big([B(f')]\big) \;=\; [f - f'] \qquad \textit{in } \mathrm{Jac}(\Gamma).$$

Proof. *Same cell.* Suppose $f, f' \in \sigma_T + \mathcal{C}_T$, with coefficients α_e and α'_e. Summing (4.7) over the g non-tree edges and subtracting, the tree-path terms $[\pi(P^T_{q,e^-})]$ cancel because the same tree is used for both:

$$\mu^g([B(f)]) - \mu^g([B(f')]) \;=\; \Big[\sum_{e \notin T} (\alpha_e - \alpha'_e)\, \pi(e)\Big] \;=\; [f - f'],$$

the last equality by (4.2).

General case. The segment $[f', f]$ lies in the convex set $\mathrm{Vor}(G)$ and, the subdivision of Theorem 4.6 being polyhedral, is covered by finitely many closed subsegments each contained in a single cell. Let $f' = f_0, f_1, \dots, f_n = f$ be the endpoints of these subsegments; consecutive pairs share a cell. By Lemma 4.19, $B(f_i)$ is unambiguous even where f_i lies in several cells, so the same-cell computation applies to each consecutive pair and the identities telescope. □

Corollary 4.21 (Break-divisor section)**.** *For every $f \in \mathrm{Vor}(G)$, $\mu^g([B(f)]) = \mu^g([B(0)]) + [f]$ in $\mathrm{Jac}(\Gamma)$. Equivalently, the canonical break-divisor section of [ABKS14, Thm. 1.1] is represented on the Voronoi fundamental domain by the affine translation $[f] \mapsto \mu^g([B(0)]) + [f]$.*

Proof. Set $f' = 0$ in Lemma 4.20. The second statement follows from the uniqueness of the break divisor in every degree-g class [ABKS14, Thm. 1.1]. □

Corollary 4.22 ($B(f)$ depends only on the class of f)**.** *For $f, f' \in \mathrm{Vor}(G)$, if $[f] = [f']$ in $\mathrm{Jac}(\Gamma)$ then $B(f) = B(f')$ as divisors.*

Proof. Lemma 4.20 gives $\mu^g([B(f)]) = \mu^g([B(f')])$, so $[B(f)] = [B(f')]$ in $\mathrm{Pic}^g(\Gamma)$. Both $B(f)$ and $B(f')$ are break divisors, and each degree-g class contains exactly one break divisor [ABKS14, Thm. 1.1]. □

4.4.2 The base characteristic

We now specialize to $f = 0$. Then $w_0(e) = w_0(\bar{e}) = \ell(e)$, so G_0 is the metric graph itself; $d_0(v)$ is the graph-distance $\mathrm{dist}(q, v)$, and we write d_q for the corresponding function on all of Γ. Let T be a shortest-path tree from q.

Lemma 4.23 (Break points are edgewise maxima of d_q)**.** *Parametrize $e = (u, v) \in \mathcal{O}(G)$ by $t \in [0, \ell(e)]$ from u. Then*

$$d_q(t) = \min\bigl\{d_q(u) + t,\ d_q(v) + \ell(e) - t\bigr\}.$$

If $e \in T$, the restriction of d_q to e is strictly monotone in the rooted direction and has no interior maximum. If $e \notin T$, its unique edgewise maximum lies at

$$t_e = \tfrac{1}{2}\Bigl(\ell(e) + d_q(v) - d_q(u)\Bigr) \in [0, \ell(e)],$$

possibly at an endpoint, and equals the break point p_e of Definition 4.18. After refining at the interior maxima, d_q is affine of slope ± 1 on every edge segment.

Proof. Any path from q into an interior point of e enters through u or through v, giving the displayed formula. If $e \in T$, one endpoint is the parent of the other in T_q, so either $d_q(v) = d_q(u) + \ell(e)$ or $d_q(u) = d_q(v) + \ell(e)$: the minimum reduces to a single affine branch on the whole edge. If $e \notin T$, the triangle inequalities $|d_q(v) - d_q(u)| \le \ell(e)$ place the crossing of the two branches inside $[0, \ell(e)]$, and d_q is maximized precisely at that crossing. Substituting $f = 0$ into (4.4) gives $\ell(e)(\frac{1}{2} + \alpha_e) = t_e$, identifying this maximum with the break point. □

Definition 4.24 (Gradient orientation and its indegree divisor)**.** Let Γ' be the model of Γ obtained by inserting a vertex at every break point p_e that lies in the interior of e. Orient

Γ' by the direction of increasing d_q: tree edges away from q, and each half of a subdivided non-tree edge towards p_e; if p_e is an endpoint of e, the whole edge is oriented towards that endpoint and no subdivision is performed. Call this orientation $\mathcal{O}_0$ and set

$$D_{\mathcal{O}_0} := \sum_{p\in\Gamma'} \Big(\mathrm{indeg}_{\mathcal{O}_0}(p) - 1\Big)\, p.$$

Lemma 4.25. *$D_{\mathcal{O}_0} = B(0) - q$.*

Proof. Each inserted point p_e has $\mathrm{indeg}_{\mathcal{O}_0} = 2$ and contributes p_e. At a vertex $v \in V(G)$, the incident edges of Γ' split into: the unique tree edge with head v (if $v \neq q$, incoming); tree edges with tail v (outgoing); half-edges of subdivided non-tree edges at v (outgoing); and undivided non-tree edges e with $p_e = v$ (incoming). Hence

$$\mathrm{indeg}_{\mathcal{O}_0}(v) = \mathbf{1}_{\{v\neq q\}} + \#\{e \notin T : p_e = v\},$$

and uniformly over $p \in \Gamma'$ we get $\mathrm{indeg}_{\mathcal{O}_0}(p) - 1 = B(0)(p) - \mathbf{1}_{\{p=q\}}$; summing gives the claim. ☐

Proposition 4.26 (Base characteristic is half-canonical)**.** *The class $\vartheta_0 := [B(0) - q] \in \mathrm{Pic}^{g-1}(\Gamma)$ satisfies $2\vartheta_0 = [K]$.*

Proof. Let $\overline{\mathcal{O}}_0$ be the reversed orientation. At every point $p \in \Gamma'$,

$$\Big(\mathrm{indeg}_{\mathcal{O}_0}(p) - 1\Big) + \Big(\mathrm{indeg}_{\overline{\mathcal{O}}_0}(p) - 1\Big) = \mathrm{val}(p) - 2,$$

so $D_{\mathcal{O}_0} + D_{\overline{\mathcal{O}}_0} = K$ (the inserted points have valence 2 and contribute 0, so K agrees on Γ' and on Γ). By Lemma 4.23, d_q has slope $+1$ on every $\mathcal{O}_0$-outgoing direction and -1 on every $\mathcal{O}_0$-incoming direction, so

$$\mathrm{div}(d_q)(p) = \mathrm{outdeg}_{\mathcal{O}_0}(p) - \mathrm{indeg}_{\mathcal{O}_0}(p) = D_{\overline{\mathcal{O}}_0}(p) - D_{\mathcal{O}_0}(p).$$

Adding the two displays gives $K = 2D_{\mathcal{O}_0} + \mathrm{div}(d_q)$, and Lemma 4.25 finishes the proof. ☐

Proposition 4.26 recovers, in the language of the Bellman–Ford distance potential, the divisor identity $K = 2\mathcal{K}_q^- + \mathrm{div}(d_q)$ of [Zha10, Lem. 5] for the singleton $S = \{q\}$: the orientation used by Zharkov to define $\mathcal{K}_q^-$ is by construction the gradient orientation of d_q, and Lemma 4.25 thus identifies

$$B(0) - q \;=\; \mathcal{K}_q^- \tag{4.8}$$

as divisors on the kink-refined model Γ'. This identity is the bridge to Zharkov's picture used in the next section.

Proposition 4.27 (Independence of the base point)**.** *For any $q, q' \in V(G)$ we have $[B_q(0) - q] = [B_{q'}(0) - q']$ in $\mathrm{Pic}^{g-1}(\Gamma)$, where the subscript records the base point used in the computation.*

Proof. Write $D := D_{\mathcal{O}_0(q)}$ and $D' := D_{\mathcal{O}_0(q')}$. Two applications of Proposition 4.26 give $2D + \mathrm{div}(d_q) = K = 2D' + \mathrm{div}(d_{q'})$, hence

$$D - D' = \mathrm{div}\Big(\tfrac{1}{2}(d_{q'} - d_q)\Big).$$

Refine so that all break points of both computations are vertices. On each resulting segment both d_q and $d_{q'}$ are affine of slope ± 1 by Lemma 4.23, so $\frac{1}{2}(d_{q'} - d_q)$ is continuous piecewise-linear with integer slopes and defines a rational function on Γ. Therefore $D \sim D'$. $\square$

4.5 Theta Characteristics via Bellman–Ford

We now specialize the linearity lemma to the 2-torsion inputs and prove the main theorem. The tropical theta characteristics form a torsor under the 2-torsion subgroup of the Jacobian, so once a single half-canonical class has been produced — as Proposition 4.26 does — the remaining ones are exactly its translates.

Lemma 4.28 ($\Theta(\Gamma)$ is a torsor)**.** *Write $\Theta(\Gamma) := \{\vartheta \in \mathrm{Pic}^{g-1}(\Gamma) : 2\vartheta = [K]\}$. If $\vartheta_0 \in \Theta(\Gamma)$ then $\Theta(\Gamma) = \vartheta_0 + \mathrm{Jac}(\Gamma)[2]$.*

Proof. If τ is 2-torsion then $2(\vartheta_0 + \tau) = [K]$; conversely $2\vartheta = 2\vartheta_0$ forces $\vartheta - \vartheta_0 \in \mathrm{Jac}(\Gamma)[2]$. □

Lemma 4.29 (Half-cycles lie in the Voronoi polytope)**.** *Let $\tilde{\gamma} \in H_1(G, \mathbb{Z})$ have edge coefficients in $\{-1, 0, 1\}$. Then $f_\gamma := \frac{1}{2}\tilde{\gamma} \in \mathrm{Vor}(G)$.*

Proof. For $z \in H_1(G, \mathbb{Z})$ and every edge e we have $\tilde{\gamma}(e)\, z(e) \le |z(e)| \le z(e)^2$, since $|z(e)| \in \mathbb{Z}_{\ge 0}$ and $|\tilde{\gamma}(e)| \le 1$. Hence $2[f_\gamma, z] = [\tilde{\gamma}, z] \le [z, z]$, which is (4.1). □

Every class $\gamma \in H_1(G, \mathbb{Z}/2\mathbb{Z})$ admits such a lift: the support of γ is an even subgraph, each connected component of which is Eulerian; orienting each component along an Eulerian circuit produces a $\{-1, 0, 1\}$-valued integer flow reducing to γ.

Theorem 4.30 (Main theorem)**.** *For every $\gamma \in H_1(G, \mathbb{Z}/2\mathbb{Z})$, define $\vartheta_\gamma := [B(f_\gamma) - q] \in \mathrm{Pic}^{g-1}(\Gamma)$, where $f_\gamma = \frac{1}{2}\tilde{\gamma}$ for any $\{-1, 0, 1\}$-valued integer lift $\tilde{\gamma}$ of γ. Then:*

1. *ϑ_γ is well defined: it depends neither on the choice of (f_γ, q)-optimal tree nor on the choice of lift $\tilde{\gamma}$;*
2. *$\vartheta_\gamma - \vartheta_0 = [\frac{1}{2}\tilde{\gamma}]$ in $\mathrm{Jac}(\Gamma)$, and $2\vartheta_\gamma = [K]$;*
3. *the map $\gamma \mapsto \vartheta_\gamma$ is a bijection $H_1(G, \mathbb{Z}/2\mathbb{Z}) \xrightarrow{\sim} \Theta(\Gamma)$.*

Proof. By Lemma 4.29, $f_\gamma \in \mathrm{Vor}(G)$, so Corollary 4.13 and Definition 4.18 apply.

(1) Independence of the optimal tree is Lemma 4.19. If $\tilde{\gamma}$ and $\tilde{\gamma}\,'$ are two $\{-1, 0, 1\}$-valued lifts of γ, they agree modulo 2, so $\tilde{\gamma} - \tilde{\gamma}\,' = 2z$ for some $z \in H_1(G, \mathbb{Z})$ and hence $f_\gamma - f'_\gamma = z$. Thus $[f_\gamma] = [f'_\gamma]$ in $\mathrm{Jac}(\Gamma)$, and Corollary 4.22 gives $B(f_\gamma) = B(f'_\gamma)$.

(2) Since $\mu^{g-1}([D - q]) = \mu^g([D])$, Lemma 4.20 applied to f_γ and 0 gives $\mu^{g-1}(\vartheta_\gamma) - \mu^{g-1}(\vartheta_0) = [f_\gamma] = [\frac{1}{2}\tilde{\gamma}]$, and hence $\vartheta_\gamma - \vartheta_0 = [\frac{1}{2}\tilde{\gamma}]$ by affinity of μ^{g-1} for the Pic^0-torsor action. Since $\tilde{\gamma} \in H_1(G, \mathbb{Z})$, we have $2[f_\gamma] = [\tilde{\gamma}] = 0$ in $\mathrm{Jac}(\Gamma)$, so $\vartheta_\gamma - \vartheta_0$ is 2-torsion, and Proposition 4.26 gives $2\vartheta_\gamma = 2\vartheta_0 + 2(\vartheta_\gamma - \vartheta_0) = [K]$.

(3) Under the isomorphism $\mathrm{Jac}(\Gamma) = H_1(G,\mathbb{R})/H_1(G,\mathbb{Z})$, the 2-torsion subgroup is

$$\mathrm{Jac}(\Gamma)[2] \;\cong\; \frac{\frac{1}{2}H_1(G,\mathbb{Z})}{H_1(G,\mathbb{Z})} \;\cong\; \frac{H_1(G,\mathbb{Z})}{2H_1(G,\mathbb{Z})} \;\cong\; H_1(G,\mathbb{Z}/2\mathbb{Z}),$$

the last isomorphism using that $H_1(G,\mathbb{Z})$ is free; the composite sends γ to $[\frac{1}{2}\tilde{\gamma}]$. Composing with the bijection $\mathrm{Jac}(\Gamma)[2] \to \Theta(\Gamma)$, $\tau \mapsto \vartheta_0 + \tau$, of Lemma 4.28, and using (2), gives the claim. □

4.5.1 Comparison with Zharkov's construction

Zharkov [Zha10] produces, for each $\gamma \in H_1(G,\mathbb{Z}/2\mathbb{Z})$, an explicit effective divisor $\mathcal{K}_\gamma^-$ representing a half-canonical class $\mathcal{K}_\gamma$, and shows that these 2^g classes exhaust $\Theta(\Gamma)$. Concretely, for $\gamma = 0$ one takes the singleton $S = \{q\}$ and forms

$$\mathcal{K}_q^- \;:=\; \sum_{p\in\Gamma}\Big(\mathrm{val}_-(p) - 1\Big)\, p,$$

where $\mathrm{val}_-(p)$ is the indegree of p under the gradient orientation of the distance function d_q, on the model Γ refined at the interior maxima of d_q; the class of $\mathcal{K}_q^-$ is called $\mathcal{K}_0$. For $\gamma \neq 0$, one takes a simple $\{-1,0,1\}$-valued lift $\tilde{\gamma}$ and forms $\mathcal{K}_\gamma^-$ by the same indegree formula, orienting the edges of the support $|\tilde{\gamma}|$ cyclically as prescribed by $\tilde{\gamma}$ and the remaining edges by the gradient of the distance function $d_\gamma(x) := \mathrm{dist}(|\tilde{\gamma}|, x)$. Zharkov proves [Zha10, Lem. 6] that

$$\mathcal{K}_\gamma - \mathcal{K}_0 \;=\; \big[\tfrac{1}{2}\tilde{\gamma}\big] \qquad \text{in } \mathrm{Jac}(\Gamma). \tag{4.9}$$

Because both families $\{\vartheta_\gamma\}$ and $\{\mathcal{K}_\gamma\}$ are translates of their $\gamma = 0$ member by the *same* element $[\frac{1}{2}\tilde{\gamma}]$ of the Jacobian, comparing them reduces to comparing the $\gamma = 0$ members, and that comparison has already been done as an identity of divisors in (4.8).

Theorem 4.31 (Identification with Zharkov's family)**.** *For every $\gamma \in H_1(G,\mathbb{Z}/2\mathbb{Z})$, $\vartheta_\gamma = \mathcal{K}_\gamma$ in $\mathrm{Pic}^{g-1}(\Gamma)$. In particular, at the divisor level $B(0) - q = \mathcal{K}_q^-$, so the Bellman–Ford break divisor at $f = 0$ is exactly Zharkov's q-reduced representative of $\mathcal{K}_0$.*

Proof. The divisor-level identity $B(0) - q = \mathcal{K}_q^-$ is (4.8), so $\vartheta_0 = \mathcal{K}_0$. Comparing Theorem 4.30(2) with (4.9) gives $\vartheta_\gamma = \vartheta_0 + [\frac{1}{2}\tilde{\gamma}] = \mathcal{K}_0 + [\frac{1}{2}\tilde{\gamma}] = \mathcal{K}_\gamma$. □

Remark 4.32 (Different effective representatives)**.** Theorem 4.31 is an identity in $\mathrm{Pic}^{g-1}(\Gamma)$. The divisors $B(f_\gamma) - q$ and $\mathcal{K}_\gamma^-$ need not coincide as divisors: our algorithm returns the translate of the canonical break divisor, whereas Zharkov's construction returns the divisor of the orientation determined by d_γ. Example 4.33 exhibits a case where they differ.

4.6 Examples: K_4 and the Diamond Graph

We illustrate the machinery on two small graphs. The first is the complete graph K_4 with unit edge lengths, the smallest example in which the tree-path term $\pi(P^T_{q,e^-})$ in (4.7) is genuinely present, and in which the two effective representatives $B(f_\gamma) - q$ and $\mathcal{K}_\gamma^-$ can differ at the divisor level (Remark 4.32). The second is the diamond graph with weights $(2, 2, 1, 2, 2)$, the smallest example in which a single γ has three (f_γ, q)-optimal spanning trees, and in which the stable exchanges connecting them run along *reverse-oriented* tight edges.

4.6.1 The complete graph K_4

Example 4.33. Let $G = K_4$ with $V(G) = \{v_0, v_1, v_2, v_3\}$, base vertex $q = v_0$, and reference orientation

$$e_0 = (v_0, v_1), \quad e_1 = (v_0, v_2), \quad e_2 = (v_0, v_3),$$

$$e_3 = (v_1, v_2), \quad e_4 = (v_1, v_3), \quad e_5 = (v_2, v_3),$$

all of length 1. The genus is $g = 3$, so $|\Theta(\Gamma)| = 8$.

By K_4's symmetry, $r(u,v) = \frac{1}{2}$ for every pair, so every Foster coefficient is $\frac{1}{2}$ and

$$\pi \;=\; \frac{1}{4}\begin{pmatrix} 2 & -1 & -1 & 1 & 1 & 0 \\ -1 & 2 & -1 & -1 & 0 & 1 \\ -1 & -1 & 2 & 0 & -1 & -1 \\ 1 & -1 & 0 & 2 & -1 & 1 \\ 1 & 0 & -1 & -1 & 2 & -1 \\ 0 & 1 & -1 & 1 & -1 & 2 \end{pmatrix},$$

with column j giving $\pi(e_j)$ in the basis $\{e_0, \dots, e_5\}$. The lattice $H_1(G, \mathbb{Z})$ is generated by the fundamental cycles of the reference tree $T_0 = \{e_0, e_1, e_2\}$, namely

$$C_3 = e_0 - e_1 + e_3, \qquad C_4 = e_0 - e_2 + e_4, \qquad C_5 = e_1 - e_2 + e_5.$$

Each nontrivial $\gamma \in H_1(G, \mathbb{Z}/2\mathbb{Z})$ has support a triangle or a 4-cycle. Running Bellman–Ford on G_{f_γ} and applying the coefficient formula (4.4) to each non-tree edge yields the eight break divisors, where m_j denotes the midpoint of e_j:

γ	$\tilde{\gamma}$	Optimal tree T; α-values	$B(f_\gamma)$
γ_0	$(0,0,0,0,0,0)$	$\{e_0,e_1,e_2\}$; $\alpha_3 = \alpha_4 = \alpha_5 = 0$	$m_3 + m_4 + m_5$
γ_{345}	$(0,0,0,-1,1,-1)$	$\{e_0,e_1,e_2\}$; $-\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}$	$v_1 + v_2 + v_3$
γ_{125}	$(0,-1,1,0,0,-1)$	$\{e_0,e_2,e_5\}$; $\alpha_1 = \alpha_3 = \alpha_4 = -\frac{1}{2}$	$q + 2v_1$
γ_{1234}	$(0,-1,1,1,-1,0)$	$\{e_2,e_3,e_4\}$; $\alpha_0 = \alpha_5 = 0,\ \alpha_1 = -\frac{1}{2}$	$q + m_0 + m_5$
γ_{024}	$(-1,0,1,0,-1,0)$	$\{e_1,e_2,e_4\}$; $\alpha_0 = \alpha_5 = -\frac{1}{2},\ \alpha_3 = \frac{1}{2}$	$q + 2v_2$
γ_{0235}	$(-1,0,1,-1,0,-1)$	$\{e_2,e_3,e_5\}$; $\alpha_0 = -\frac{1}{2},\ \alpha_1 = \alpha_4 = 0$	$q + m_1 + m_4$
γ_{0145}	$(-1,1,0,0,-1,1)$	$\{e_1,e_4,e_5\}$; $\alpha_0 = -\frac{1}{2},\ \alpha_2 = \alpha_3 = 0$	$q + m_2 + m_3$
γ_{013}	$(-1,1,0,-1,0,0)$	$\{e_1,e_2,e_3\}$; $\alpha_0 = -\frac{1}{2},\ \alpha_4 = \alpha_5 = \frac{1}{2}$	$q + 2v_3$

We work two rows in detail.

The case γ_0. All induced weights equal 1, so $d_0(v_i) = 1$ for $i = 1, 2, 3$ and $T_0 = \{e_0, e_1, e_2\}$ is optimal. Both endpoints of every non-tree edge are at distance 1, so $\alpha_{e_j} = 0$ by (4.4) and the break point is the midpoint m_j. Hence $B(0) = m_3 + m_4 + m_5$ and $\vartheta_0 = [-q + m_3 + m_4 + m_5]$, an instance of Lemma 4.25: the gradient orientation of d_q makes q a source, each v_i ordinary, and each m_j a sink. This is also the smallest case where the tree-path term of (4.7) is nonzero: the tails of e_3, e_4, e_5 are v_1, v_1, v_2, with tree paths e_0, e_0, e_1, so

$$\mu^g\Big([B(0)]\Big) = \pi\Big(2e_0 + e_1 + \tfrac{1}{2}e_3 + \tfrac{1}{2}e_4 + \tfrac{1}{2}e_5\Big).$$

It is convenient to simplify the chosen 1-chain modulo $\ker \pi$ before projecting: here $e_0 + e_1 + e_2 \in \ker \pi$, so $2e_0 + e_1$ and $e_0 - e_2$ project to the same element. Thus

$$\mu^g\Big([B(0)]\Big) \;=\; \Big(1, 0, -1, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{1}{2}\Big) \in \mathbb{R}^6,$$

and identities in $\mathrm{Jac}(\Gamma)$ are checked modulo the lattice $H_1(G, \mathbb{Z}) = \mathbb{Z}\langle C_3, C_4, C_5\rangle$.

The case γ_{345}: a divisor-level discrepancy. The lift $\tilde{\gamma}_{345} = (0, 0, 0, -1, 1, -1)$ orients the triangle on $\{v_1, v_2, v_3\}$ cyclically as $v_2 \to v_1 \to v_3 \to v_2$. On e_0, e_1, e_2 the induced weight is 1 in both directions; on e_3, e_4, e_5 the forward/reverse weight pairs are $(2, 0)$, $(0, 2)$, $(2, 0)$. Bellman–Ford gives $d_f(v_i) = 1$ for $i = 1, 2, 3$, and $T_0 = \{e_0, e_1, e_2\}$ is one of *seven* (f, q)-optimal trees — the largest number of optima occurring in this example. With T_0, formula (4.4) yields

$$\alpha_{e_3} = -\tfrac{1}{2}, \qquad \alpha_{e_4} = \tfrac{1}{2}, \qquad \alpha_{e_5} = -\tfrac{1}{2},$$

so the break points sit at v_1, v_3, v_2 and $B(f_{\gamma_{345}}) = v_1 + v_2 + v_3$. Thus $\vartheta_{\gamma_{345}} = [v_1 + v_2 + v_3 - q]$.

Zharkov's recipe returns a different effective divisor. Here $|\tilde{\gamma}_{345}|$ is the triangle on $\{v_1, v_2, v_3\}$, the vertex q is the unique point at distance 1 from it, and the gradient of $d_{\gamma_{345}}$ sends e_0, e_1, e_2 into q: $\mathrm{val}_-(q) = 3$ and $\mathrm{val}_-(v_i) = 1$, giving $\mathcal{K}^-_{\gamma_{345}} = 2q$. The two divisors differ by

$$(v_1 + v_2 + v_3 - q) - 2q \;=\; -3q + v_1 + v_2 + v_3 \;=\; \mathrm{div}(\varphi),$$

where $\varphi(q) = 1$ and $\varphi(v_i) = 0$, extended linearly on each edge. The outgoing slopes at q are -1 along each of e_0, e_1, e_2, giving $\mathrm{div}(\varphi)(q) = -3$; at each v_i the slope towards q is $+1$ and the slopes along the triangle edges vanish, giving $\mathrm{div}(\varphi)(v_i) = 1$. The two agree in $\mathrm{Pic}^{g-1}(\Gamma)$, as Theorem 4.31 asserts, but not as divisors — the phenomenon of Remark 4.32. For the remaining six nonzero rows, and also for γ_0, the two constructions return the same divisor.

The linearity lemma is visible on the table above, though only in the form in which Lemma 4.20 states it — as an identity in $\mathrm{Jac}(\Gamma)$, not in $H_1(G, \mathbb{R})$. Taking the tree paths relative to $T_0 = \{e_0, e_1, e_2\}$,

$$\mu^g\Big([B(f_{\gamma_0})]\Big) - \mu^g\Big([B(f_{\gamma_{345}})]\Big) = \Big(1, 0, -1, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{1}{2}\Big), \qquad f_{\gamma_0} - f_{\gamma_{345}} = \Big(0, 0, 0, \tfrac{1}{2}, -\tfrac{1}{2}, \tfrac{1}{2}\Big),$$

and these differ by $(1, 0, -1, 0, 1, 0) = C_4 \in H_1(G, \mathbb{Z})$. The identity therefore holds after reduction modulo the lattice, and not before. This is not a defect of the example but a feature of the statement: μ^g takes values in $\mathrm{Jac}(\Gamma) = H_1(G, \mathbb{R})/H_1(G, \mathbb{Z})$, and a different choice of the paths P_{q,p_i} used to compute it shifts the representative by an element of $H_1(G, \mathbb{Z})$. An identity between representatives in $H_1(G, \mathbb{R})$ would not even be well posed.

4.6.2 The diamond graph

Example 4.34. Let $D = K_4 \setminus \{v_0 v_3\}$ be the diamond graph, with vertex set $V(D) = \{a, b, c, d\}$, base vertex $q = a$, and reference orientation

$$e_{ab} = (a, b), \quad e_{ac} = (a, c), \quad e_{bc} = (b, c), \quad e_{bd} = (b, d), \quad e_{cd} = (c, d),$$

with lengths $\ell(e_{ab}) = \ell(e_{ac}) = \ell(e_{bd}) = \ell(e_{cd}) = 2$ and $\ell(e_{bc}) = 1$. The genus is $g = 5 - 4 + 1 = 2$, so $|\Theta(\Gamma)| = 4$.

The cycle space $H_1(D, \mathbb{R})$ is 2-dimensional; a basis of $H_1(D, \mathbb{Z})$ is

$$z_1 = e_{ab} - e_{ac} + e_{bc} \quad \text{(triangle } a\text{–}b\text{–}c\text{)}, \qquad z_2 = e_{ab} - e_{ac} + e_{bd} - e_{cd} \quad \text{(4-cycle } a\text{–}b\text{–}d\text{–}c\text{)}.$$

The four 2-torsion classes have supports the four simple cycles of D:

γ	$\tilde{\gamma}$	$\lvert\tilde{\gamma}\rvert$	$f_\gamma = \frac{1}{2}\tilde{\gamma}$
γ_0	$(0,\ 0,\ 0,\ 0,\ 0)$	$\emptyset$	$(0,\ 0,\ 0,\ 0,\ 0)$
γ_{abc}	$(1,\ -1,\ 1,\ 0,\ 0)$	$\{e_{ab}, e_{ac}, e_{bc}\}$	$(\frac{1}{2},\ -\frac{1}{2},\ \frac{1}{2},\ 0,\ 0)$
γ_{abcd}	$(1,\ -1,\ 0,\ 1,\ -1)$	$\{e_{ab}, e_{ac}, e_{bd}, e_{cd}\}$	$(\frac{1}{2},\ -\frac{1}{2},\ 0,\ \frac{1}{2},\ -\frac{1}{2})$
γ_{bcd}	$(0,\ 0,\ -1,\ 1,\ -1)$	$\{e_{bc}, e_{bd}, e_{cd}\}$	$(0,\ 0,\ -\frac{1}{2},\ \frac{1}{2},\ -\frac{1}{2})$

Bellman–Ford on G_{f_γ} from $q = a$ produces:

γ	Optimal tree(s)	$B(f_\gamma)$
γ_0	$\{e_{ab}, e_{ac}, e_{bd}\}$, $\{e_{ab}, e_{ac}, e_{cd}\}$	$m_{bc} + d$
γ_{abc}	$\{e_{ab}, e_{bc}, e_{bd}\}$, $\{e_{ab}, e_{bc}, e_{cd}\}$	$a + d$
γ_{abcd}	$\{e_{ab}, e_{bd}, e_{cd}\}$	$a + m_{bc}$
γ_{bcd}	$\{e_{ab}, e_{ac}, e_{bd}\}$, $\{e_{ab}, e_{bd}, e_{cd}\}$, $\{e_{ac}, e_{bc}, e_{bd}\}$	$b + c$

Here m_{bc} is the midpoint of e_{bc}. We now examine the case γ_{bcd} in detail; it is the smallest instance in this chapter where a single γ has three (f_γ, q)-optimal spanning trees, and where the stable exchanges connecting them run along *reverse-oriented* tight edges.

The case γ_{bcd}. The induced weights are

$$w_f(e_{ab}) = w_f(e_{ac}) = 2, \qquad w_f(e_{bc}) = 2, \qquad w_f(\bar{e}_{bc}) = 0,$$

$$w_f(e_{bd}) = 0, \qquad w_f(\bar{e}_{bd}) = 4, \qquad w_f(e_{cd}) = 4, \qquad w_f(\bar{e}_{cd}) = 0.$$

Bellman–Ford from a yields $d(b) = d(c) = d(d) = 2$. Starting from $T_1 = \{e_{ab}, e_{ac}, e_{bd}\}$ with $T_{1,q}$-orientation $e_{ab}\colon a \to b$, $e_{ac}\colon a \to c$, $e_{bd}\colon b \to d$, the non-tree edges e_{bc}, e_{cd} are *not* tight in the reference orientation:

$$d(b) + w_f(e_{bc}) = 4 \neq 2 = d(c), \qquad d(c) + w_f(e_{cd}) = 6 \neq 2 = d(d).$$

Their reverses are, however, both tight, since $w_f(\bar{e}_{bc}) = w_f(\bar{e}_{cd}) = 0$:

$$d(c) + w_f(\bar{e}_{bc}) = 2 = d(b), \qquad d(d) + w_f(\bar{e}_{cd}) = 2 = d(c).$$

The reverse $\bar{e}_{bc}$ (directed $c \to b$) is stable: c is a leaf of $T_{1,q}$, so no cycle is created; adding it and removing the tree edge e_{ab} (the unique tree edge with head b) yields $T_1' = \{e_{ac}, e_{bc}, e_{bd}\}$. Similarly $\bar{e}_{cd}$ (directed $d \to c$) is stable and its exchange removes e_{ac}, yielding $T_1'' = \{e_{ab}, e_{bd}, e_{cd}\}$. All three trees are $(f_{\gamma_{bcd}}, q)$-optimal.

Verification of Lemma 4.19. For $T_1' = \{e_{ac}, e_{bc}, e_{bd}\}$ the non-tree edges are e_{ab}, e_{cd}; (4.4) gives $\alpha_{ab} = \frac{1}{2}$ (break point at b) and $\alpha_{cd} = -\frac{1}{2}$ (break point at c). For $T_1'' = \{e_{ab}, e_{bd}, e_{cd}\}$, $\alpha_{ac} = \frac{1}{2}$ (break point at c) and $\alpha_{bc} = -\frac{1}{2}$ (break point at b). All three trees produce $B(f_{\gamma_{bcd}}) = b + c$, confirming Lemma 4.19.

The feature of this example is that a $\{-1, 0, 1\}$-valued cycle can reverse the direction of the tight relations on non-support edges. Definition 4.15, phrased in terms of directed edges $a \in \mathbb{E}(G)$ rather than only $\mathcal{O}(G)$, is what makes the stable-exchange proof of Lemma 4.19 go through in this case.

Zharkov's family on the diamond graph agrees with $\{\vartheta_\gamma\}$ divisor-by-divisor for the three cases $\gamma_0, \gamma_{abc}, \gamma_{abcd}$, and differs only for γ_{bcd}, where $\mathcal{K}^-_{\gamma_{bcd}} = a$ while $B(f_{\gamma_{bcd}}) - a = b + c - a$. These agree in $\mathrm{Pic}^{g-1}(\Gamma)$ via the rational function φ with $\varphi(a) = 2$, $\varphi(v) = 0$ for $v \in \{b, c, d\}$, extended linearly; the computation is a smaller version of the γ_{345}-analysis for K_4, and is another instance of Remark 4.32.

Chapter 5

ENERGY MINIMIZATION AND BREAK DIVISORS ON GRAPHS

5.1 Introduction

A classical theorem of Mikhalkin–Zharkov [MZ08] asserts that every divisor class of degree g on a metric graph Γ of genus g is linearly equivalent to a unique *break divisor*; uniqueness holds among break divisors, not among all effective representatives. Break divisors were studied combinatorially by An–Baker–Kuperberg–Shokrieh [ABKS14], who proved uniqueness via the matrix-tree theorem and gave an explicit firing algorithm.

The purpose of this chapter is to characterize break divisors as the *energy minimizers* of a natural quadratic functional on the space of divisors.

Fix a finite connected simple graph $G = (V, E)$ of genus $g := |E| - |V| + 1$. The *Foster–Zhang measure* on the associated metric graph Γ (obtained by assigning unit length to each edge) is the canonical measure $\mu_F := \sum_{e \in E} F(e)\, dx_e$, where $F(e) := 1 - r(e)$ is the Foster coefficient of edge e and $r(e)$ denotes its effective resistance [Zha93, BF11]. We discretize μ_F to the vertices by setting $\tilde{\mu}_F(v) := \frac{1}{2} \sum_{e \ni v} F(e)$, and define the *Foster energy* of a divisor $D \in \mathrm{Div}^g$ as

$$E_F(D) := [D - \tilde{\mu}_F]^T L^\dagger [D - \tilde{\mu}_F],$$

where $L^\dagger$ is the Moore–Penrose pseudoinverse of the Laplacian L of G. The main result of this chapter is the following.

Theorem 5.1 (Energy minimization)**.** *Let G be a connected graph of genus g. If G has positive resistance curvature ($p_v > 0$ for every vertex v), then in every equivalence class of* Div^g*, the unique break divisor is the unique minimizer of the Foster energy E_F.*

We also prove a converse (Corollary 5.16): if G is biconnected and E_F is minimized by break divisors in every class, then G has nonnegative resistance curvature. Section 5.6 shows by example that the positivity assumption cannot be dropped.

In addition, we prove in Section 5.5 that E_F coincides with the continuous Foster–Zhang energy pairing on Γ for vertex-supported divisors, connecting the combinatorial result to the superform formalism of Chapter 2.

The chapter is organized as follows. Section 5.2 introduces the setup: graphs, Laplacians, effective resistances, divisors, break divisors, and resistance curvature. Section 5.3 defines the Foster energy. Section 5.4 proves Theorem 5.1 and Corollary 5.16. Section 5.5 establishes the discrete–continuous energy equivalence. Section 5.6 gives two explicit counterexamples.

5.2 Graphs, Divisors, and Break Divisors

Throughout this chapter, $G = (V, E)$ denotes a finite connected simple graph with $n := |V|$ vertices and genus $g := |E| - |V| + 1$.

5.2.1 Laplacian and effective resistance

The combinatorial structure of G is encoded in two canonical matrices built from the Laplacian.

Definition 5.2 (Laplacian and resistance matrices)**.** The *Laplacian matrix* L of G is the $n \times n$ symmetric matrix with entries $L_{vv} = \deg(v)$, $L_{uv} = -1$ if $uv \in E$, and $L_{uv} = 0$ otherwise. It is positive semidefinite with $\ker L = \mathrm{span}(\mathbf{1})$, where $\mathbf{1} = (1, \dots, 1)^T$. Its Moore–Penrose pseudoinverse $L^\dagger$ satisfies $\ker L^\dagger = \ker L$ and $LL^\dagger L = L$, and since L is symmetric, $L^\dagger$ is symmetric too with $LL^\dagger = L^\dagger L = I - \frac{1}{n}\mathbf{1}\mathbf{1}^T$.

The *resistance matrix* of G is

$$\Omega := \mathbf{1}\,\mathrm{diag}(L^\dagger)^T + \mathrm{diag}(L^\dagger)\mathbf{1}^T - 2L^\dagger.$$

The *effective resistance* between $u, v \in V$ is $r(u, v) := \Omega_{uv} \geq 0$.

The following three identities relating L, $L^\dagger$, and Ω are the technical engine of this chapter. We write $P := \frac{1}{n}\mathbf{1}\mathbf{1}^T$ for the orthogonal projection onto $\ker L$.

Proposition 5.3 (Resistance identities)**.** *Let $p \in \mathbb{R}^V$ be the resistance curvature (Definition 5.6). The following hold:*

$$L\Omega L = -2L, \tag{5.1}$$

$$L\Omega = -2I + 2p\mathbf{1}^T, \tag{5.2}$$

$$x^T L^\dagger x = -\tfrac{1}{2} x^T \Omega x \quad \textit{for all } x \in \mathbb{R}^V \textit{ with } \mathbf{1}^T x = 0. \tag{5.3}$$

Proof. Since $L\mathbf{1} = 0$ and $LL^\dagger = I - P$:

$$L\Omega = L\mathbf{1}\operatorname{diag}(L^\dagger)^T + L\operatorname{diag}(L^\dagger)\mathbf{1}^T - 2LL^\dagger = L\operatorname{diag}(L^\dagger)\mathbf{1}^T - 2(I-P).$$

Multiplying on the right by L and using $\mathbf{1}^T L = (L\mathbf{1})^T = 0$ and $PL = \frac{1}{n}\mathbf{1}(\mathbf{1}^T L) = 0$:

$$L\Omega L = L\operatorname{diag}(L^\dagger)\underbrace{(\mathbf{1}^T L)}_{=0} - 2(I-P)L = -2L,$$

which proves (5.1).

For (5.3), if $\mathbf{1}^T x = 0$ then $x^T\mathbf{1} = 0$, so both cross-terms in $x^T\Omega x = x^T\mathbf{1}\operatorname{diag}(L^\dagger)^T x + x^T\operatorname{diag}(L^\dagger)\mathbf{1}^T x - 2x^T L^\dagger x$ vanish, giving $x^T\Omega x = -2x^T L^\dagger x$, hence (5.3).

For (5.2), we show $L\Omega$ agrees with $-2I + 2p\mathbf{1}^T$ on every vector by splitting $\mathbb{R}^V = \operatorname{im}(L) \oplus \operatorname{span}(\mathbf{1})$. If $\mathbf{1}^T x = 0$ then $x \in \operatorname{im}(L)$ (since $\ker L = \operatorname{span}(\mathbf{1})$ and L is symmetric), so $x = Ly$ for some y, and $L\Omega x = L\Omega L y = -2Ly = -2x$ by (5.1); meanwhile $(-2I + 2p\mathbf{1}^T)x = -2x + 2p\underbrace{(\mathbf{1}^T x)}_{=0} = -2x$. If $x = \mathbf{1}$, then $(-2I + 2p\mathbf{1}^T)\mathbf{1} = -2\cdot\mathbf{1} + 2np$, while the computation above gives $L\Omega\mathbf{1} = nL\operatorname{diag}(L^\dagger)$, since $P\mathbf{1} = \mathbf{1}$ kills the second term. So the case $x = \mathbf{1}$ amounts to the relation

$$n\,L\operatorname{diag}(L^\dagger) = -2\cdot\mathbf{1} + 2np, \tag{5.4}$$

equivalently $p_v = \frac{1}{2}\Big(L\operatorname{diag}(L^\dagger)\Big)_v + \frac{1}{n}$ for every v, which is a standard consequence of the definition of resistance curvature via effective resistances (see [BF11, Lemma 2.3]). Note that both sides of (5.4) have coordinate sum 0, consistent with $\mathbf{1}^T L = 0$ and $\sum_v p_v = 1$. $\square$

5.2.2 Divisors and linear equivalence

The language of divisors provides the setting for the energy minimization problem.

Definition 5.4 (Divisors and linear equivalence)**.** A *divisor* on G is a vector $D \in \mathbb{Z}^V$. Its *degree* is $\deg(D) := \sum_v D_v$; we write Div^r for divisors of degree r. The *canonical divisor* is $K \in \mathrm{Div}^{2g-2}$ with $K(v) := \deg(v) - 2$. Two divisors $D \sim D'$ are *linearly equivalent* if $D - D' \in \mathrm{im}(L)$. We write $[D]$ for the equivalence class of D.

5.2.3 Break divisors

Break divisors are the canonical effective representatives of each degree-g class, first constructed by Mikhalkin–Zharkov [MZ08].

Definition 5.5 (Break divisor)**.** Fix a spanning tree $T \subseteq E$; its complement $T^c := E \setminus T$ has $|T^c| = g$ edges. For each $e \in T^c$, choose one endpoint vertex of e. The divisor $D_e \in \mathrm{Div}^g$ obtained by adding $+1$ at each chosen vertex is a *break divisor*.

Every equivalence class in Div^g contains a unique break divisor [MZ08, ABKS14]. We write D_e for this unique representative throughout.

5.2.4 Resistance curvature

The resistance curvature is a discrete analogue of scalar curvature, measuring how effective resistances at a vertex compare to the balanced case.

Definition 5.6 (Resistance curvature)**.** The *resistance curvature* at $v \in V$ is

$$p_v := 1 - \frac{1}{2} \sum_{\substack{w \in V \\ vw \in E}} r(v, w).$$

We say G has *positive resistance curvature* if $p_v > 0$ for all v, and *nonnegative resistance curvature* if $p_v \geq 0$ for all v.

The curvature vector satisfies $\sum_v p_v = 1$, so p is itself a probability measure on V. A characterization of graphs with nonnegative resistance curvature is given in [Dev25].

5.3 Foster Coefficients and the Foster Energy

We now define the Foster energy in terms of a discrete measure derived from the Foster–Zhang measure on the metric graph.

5.3.1 Vertex Foster coefficients

The *Foster coefficient* of an edge $e = \{u, v\}$ is $F(e) := 1 - r(u, v)$. These satisfy $\sum_{e \in E} F(e) = g$ [BF11].

Definition 5.7 (Vertex Foster coefficients)**.** For $v \in V$, the *vertex Foster coefficient* is

$$\tilde{\mu}_F(v) := \frac{1}{2} \sum_{\substack{w \in V \\ e=\{v,w\} \in E}} (1 - r(v, w)).$$

The factor $\frac{1}{2}$ accounts for each edge being shared by two vertices, so $\sum_v \tilde{\mu}_F(v) = g$.

The vertex Foster coefficients interleave the canonical divisor and resistance curvature in a clean identity that plays a key role in the proof.

Proposition 5.8. *For every vertex $v \in V$,*

$$\tilde{\mu}_F(v) = \frac{K(v)}{2} + p_v.$$

Proof. Expanding directly:

$$\tilde{\mu}_F(v) = \frac{\deg(v)}{2} - \frac{1}{2} \sum_{\substack{w \\ e=\{v,w\} \in E}} r(v, w) = \frac{K(v) + 2}{2} - \frac{1}{2} \sum_{\substack{w \\ e=\{v,w\} \in E}} r(v, w) = \frac{K(v)}{2} + p_v. \square$$

In particular, $\sum_v \tilde{\mu}_F(v) = \sum_v (K(v)/2 + p_v) = (g - 1) + 1 = g$, confirming the degree count.

5.3.2 The Foster energy

The continuous energy pairing on the metric graph Γ is $\langle \nu_1, \nu_2 \rangle_{en} = \int_{\Gamma\times\Gamma} j_q(x,y)\, d\nu_1(x)\, d\nu_2(y)$, where $j_q(x,y)$ is the potential kernel of Γ [dJS22b]. For divisors supported on V, the matrix of values $j_q(v,w)$ coincides with $L^\dagger$ up to a constant that cancels when the measures have total mass zero.

Definition 5.9 (Foster energy)**.** For $D \in \mathrm{Div}^g$, the *Foster energy* of D is

$$E_F(D) := [D - \tilde{\mu}_F]^T L^\dagger [D - \tilde{\mu}_F].$$

Since $L^\dagger$ is positive semidefinite, $E_F(D) \geq 0$ for all divisors D. Equality forces $D - \tilde{\mu}_F \in \ker L^\dagger = \mathrm{span}(\mathbf{1})$; as $\mathbf{1}^T(D - \tilde{\mu}_F) = g - g = 0$, this means $D = \tilde{\mu}_F$, which is impossible whenever $\tilde{\mu}_F \notin \mathbb{Z}^V$.

More generally, call $f \in \mathbb{R}^V$ a *measure* on G if $\sum_v f_v = 1$, and set $F := \frac{1}{2}K + f \in \mathbb{R}^V$ (so that $\sum_v F_v = (g-1) + 1 = g$, and $F = \tilde{\mu}_F$ when $f = p$, by Proposition 5.8 and $\sum_v p_v = 1$). Define

$$E_f(D) := [D - F]^T L^\dagger [D - F]. \tag{5.5}$$

The Foster energy is the case $f = p$. Since $\mathbf{1}^T(D - F) = g - g = 0$, identity (5.3) gives the equivalent expression

$$E_f(D) = -\tfrac{1}{2}[D - F]^T \Omega [D - F]. \tag{5.6}$$

5.4 Proof of the Energy Minimization Theorem

The proof is organized around computing the change in f-energy when a divisor is modified by a firing operation.

5.4.1 The energy difference formula

Linear equivalence corresponds to adding Lx for $x \in \mathbb{Z}^V$.

Definition 5.10 (f-energy difference)**.** $\Delta E_f(D, x) := E_f(D + Lx) - E_f(D)$.

Write $\bar{f} := 2\mathbf{1} - 2f \in \mathbb{R}^V$.

Lemma 5.11 (Energy difference formula)**.** *For any measure f, $D \in \mathrm{Div}^g$, and $x \in \mathbb{Z}^V$,*

$$\Delta E_f(D, x) = x^T L x + x^T(2D - \deg + \bar{f}). \tag{5.7}$$

Proof. Since $\mathbf{1}^T(D - F) = 0$, we may use (5.6):

$$\begin{aligned}\Delta E_f(D, x) &= -\tfrac{1}{2}(D + Lx - F)^T \Omega (D + Lx - F) + \tfrac{1}{2}(D - F)^T \Omega (D - F)\\ &= -x^T L\Omega(D - F) - \tfrac{1}{2} x^T (L\Omega L) x.\end{aligned}$$

By (5.1), $-\frac{1}{2}x^T(L\Omega L)x = x^T L x$. By (5.2) and $\mathbf{1}^T(D - F) = 0$:

$$L\Omega(D - F) = (-2I + 2p\mathbf{1}^T)(D - F) = -2(D - F),$$

so $-x^T L\Omega(D - F) = 2x^T(D - F)$. Substituting $2F = \deg + 2f - 2\mathbf{1} = \deg - \bar{f}$ gives $2x^T(D - F) = x^T(2D - \deg + \bar{f})$. □

Lemma 5.12 (Shift invariance)**.** $\Delta E_f(D, x) = \Delta E_f(D, x + a\mathbf{1})$ *for any $a \in \mathbb{Z}$.*

Proof. Since $L\mathbf{1} = 0$ and $\mathbf{1}^T(2D - \deg + \bar{f}) = 2g - 2g + 0 = 0$, both terms in (5.7) are unchanged by $x \mapsto x + a\mathbf{1}$. □

5.4.2 A chip-count lower bound for break divisors

Lemma 5.13. *Let D_e be a break divisor and $\emptyset \neq U \subseteq V$. Writing $E(U, U) := \{e \in E :$ both endpoints of e lie in $U\}$,*

$$\sum_{u \in U} (D_e)_u \;\geq\; |E(U, U)| - |U| + 1.$$

Proof. Let T be a spanning tree from which D_e is constructed. Each edge $e \in T^c$ with both endpoints in U contributes $+1$ to $\sum_{u \in U}(D_e)_u$, while edges with only one endpoint in U may contribute zero. Hence

$$\begin{aligned}\sum_{u \in U}(D_e)_u &\geq |E(U, U) \cap T^c| = |E(U, U)| - |E(U, U) \cap T|\\ &\geq |E(U, U)| - (|U| - c(G[U])) \geq |E(U, U)| - |U| + 1,\end{aligned}$$

where $c(G[U]) \geq 1$ is the number of connected components of the induced subgraph $G[U]$, and the second inequality uses that T restricts to a forest on $G[U]$ (at most $|U| - c(G[U])$ edges). $\square$

5.4.3 The indicator estimate

The following is the key inequality driving the proof. It bounds the energy difference when firing an indicator vector e_U.

Lemma 5.14 (Indicator estimate)**.** *Let f be a measure on G, D_e a break divisor, and $\emptyset \neq U \subseteq V$. Then*

$$\Delta E_f(D_e, e_U) \;\geq\; 2 \sum_{u \in U^c} f_u.$$

In particular, if $f \geq 0$ then $\Delta E_f(D_e, e_U) \geq 0$, with strict inequality if $f > 0$.

Proof. Apply Lemma 5.11 with $x = e_U$. Using $e_U^T L e_U = |E(U, U^c)|$ (the cut size) and collecting terms:

$$\begin{aligned} \Delta E_f(D_e, e_U) &= |E(U, U^c)| + 2 \sum_{u \in U} (D_e)_u - 2|E(U, U)| - |E(U, U^c)| + \sum_{u \in U} \bar{f}_u \\ &= 2 \sum_{u \in U} (D_e)_u - 2|E(U, U)| + \sum_{u \in U} \bar{f}_u. \end{aligned}$$

Applying Lemma 5.13 and $\bar{f}_u = 2 - 2f_u$:

$$\begin{aligned} \Delta E_f(D_e, e_U) &\geq 2\Big(|E(U, U)| - |U| + 1\Big) - 2|E(U, U)| + 2|U| - 2 \sum_{u \in U} f_u \\ &= 2 - 2 \sum_{u \in U} f_u = 2 \sum_{u \in U^c} f_u. \end{aligned}$$

$\square$

5.4.4 Proof of Theorem 5.1

We first establish a degree bound used in the converse direction.

Lemma 5.15. *If G is biconnected, then* $\deg(v) \leq g + 1$ *for every vertex v.*

Proof. Since G is biconnected, $G - v$ is connected and contains a spanning tree, so $|E(G-v)| \geq |V| - 2$. Thus $g = \deg(v) + |E(G-v)| - |V| + 1 \geq \deg(v) - 1$. $\square$

Proof of Theorem 5.1. Let D_e be any break divisor and let $x \in \mathbb{Z}^V$ with $Lx \neq 0$. Set $D := D_e + Lx \neq D_e$. By Lemma 5.12 we may assume $x \geq 0$, x is not constant, and $\min_v x_v = 0$. We show $\Delta E_p(D_e, x) > 0$ by induction on $m := \max_v x_v \geq 1$.

Base case ($m = 1$). Then $x = e_U$ for some $\emptyset \neq U \subsetneq V$. Since $p_v > 0$ for all v, Lemma 5.14 gives $\Delta E_p(D_e, e_U) \geq 2 \sum_{u \in U^c} p_u > 0$.

Inductive step ($m > 1$). Let $U := \{u \in V : x_u = m\}$ and $y := x - e_U$. Then $\max y = m-1$, $y \geq 0$, y is not constant (since $m > 1$ and $\min x = 0$), and $Ly \neq 0$, so the induction hypothesis applies to y. Decompose:

$$\Delta E_p(D_e, x) = \Delta E_p(D_e + Ly,\, e_U) + \Delta E_p(D_e,\, y).$$

Expanding the first summand using (5.7):

$$\Delta E_p(D_e + Ly, e_U) = \Delta E_p(D_e, e_U) + 2e_U^T Ly = \Delta E_p(D_e, e_U) + 2 \sum_{\substack{uv \in E \\ u \in U,\, v \in U^c}} (y_u - y_v).$$

Since U consists of the maximum entries of y (value $m-1$), $y_u \geq y_v$ for each edge uv with $u \in U$, $v \in U^c$. Therefore the sum is ≥ 0, and

$$\Delta E_p(D_e, x) \geq \underbrace{\Delta E_p(D_e, e_U)}_{>0} + \underbrace{\Delta E_p(D_e, y)}_{>0} > 0.$$

This completes the induction and proves $E_F(D_e) < E_F(D)$ for every $D \neq D_e$ in $[D_e]$. $\square$

Corollary 5.16 (Partial converse)**.** *Let G be biconnected. If E_F is minimized (not necessarily strictly) by the break divisor in every equivalence class of* Div^g*, then $p_v \geq 0$ for every vertex v.*

Proof. Fix any vertex $v \in V$. Since G is biconnected, $G - v$ is connected, so there is a spanning tree T of G in which v is a leaf; by Lemma 5.15, $\deg(v) - 1 \leq g$, so the divisor constructed next indeed has degree g. Construct D_e from T by assigning v as the endpoint for every edge in T^c incident to v, giving $(D_e)_v = \deg(v) - 1$.

By assumption D_e minimizes E_F in $[D_e]$, so $\Delta E_p(D_e, -e_v) \geq 0$. Computing using (5.7) with $f = p$ and $x = -e_v$:

$$\begin{aligned}\Delta E_p(D_e, -e_v) &= e_v^T L e_v - e_v^T(2D_e - \deg + \bar{p}) \\ &= 2\deg(v) - 2(\deg(v) - 1) - (2 - 2p_v) = 2p_v.\end{aligned}$$

Hence $p_v \geq 0$. Since v was arbitrary, $p \geq 0$. □

5.5 *Continuous Energy and the Discrete–Continuous Equivalence*

We show that E_F equals the continuous Foster–Zhang energy pairing on the metric graph Γ, for divisors supported on vertices. This connects the combinatorial result of Section 5.4 to the superform formalism of Chapter 2 and to the Foster–Zhang measure appearing in Chapter 3.

5.5.1 The continuous Foster–Zhang energy

Equip each edge of G with length $\ell(e) = 1$ to form the metric graph Γ. Let $j_q(x, y)$ be the *potential kernel* on Γ with basepoint q, the unique function satisfying $\Delta_x j_q(x, y) = \delta_y - \delta_q$ with $j_q(q, y) = 0$; see [dJS22b]. The *continuous energy pairing* on signed measures of total mass zero is

$$\langle \nu_1, \nu_2 \rangle_{en} := \int_{\Gamma \times \Gamma} j_q(x, y)\, d\nu_1(x)\, d\nu_2(y),$$

which is independent of q [dJS22b]. The *Foster–Zhang measure* is $\mu_F := \sum_{e \in E} F(e)\, dx_e$ [Zha93, BF11].

Definition 5.17 (Continuous Foster energy)**.** For $D \in \mathrm{Div}^g$ viewed as a discrete measure on Γ, the *continuous Foster energy* is $E_F^{cont}(D) := \langle D - \mu_F,\, D - \mu_F \rangle_{en}$.

5.5.2 A general energy equivalence theorem

The equivalence holds for any continuous measure μ on Γ, not just μ_F. Decompose $\mu = \mu_V + \mu_E$ with $\mu_V = \sum_v N(v)\delta_v$ and $\mu_E = \sum_e M(e)\,dx_e$. Its *natural discretization* is

$$\tilde{\mu}(v) := N(v) + \frac{1}{2} \sum_{\substack{w \in V \\ e=\{v,w\}\in E}} M(e).$$

This is the same operation used to produce $\tilde{\mu}_F$ from μ_F.

Lemma 5.18 (Handshake on oriented edges)**.** *For any $a\colon V \times V \to \mathbb{R}$ supported on oriented edges, $\sum_v \sum_{w\in N(v)} a(v,w) = \sum_{\{u,w\}\in E}(a(u,w) + a(w,u))$. When a is symmetric this equals $2\sum_{e\in E} a(e)$.*

Proof. Each oriented edge (v,w) appears exactly once on the left and once in the corresponding unordered pair on the right. □

Theorem 5.19 (General energy equivalence)**.** *Let $\mu = \mu_V + \mu_E$ be a continuous measure on Γ with natural discretization $\tilde{\mu}$, and let $D \in \mathrm{Div}^g$ be supported on $V(G)$. Then*

$$\langle D - \mu,\, D - \mu\rangle_{en} = [D - \tilde{\mu}]^T L^\dagger [D - \tilde{\mu}].$$

Proof. Write $D_{int} := D - \mu_V$ and $\tilde{\mu}_E(v) := \frac{1}{2}\sum_{w:\{v,w\}\in E} M(\{v,w\})$. Then $D - \tilde{\mu} = D_{int} - \tilde{\mu}_E$. We match the three expanded terms.

Term 1. Since D_{int} is vertex-supported and $j_q(v,w) = L^\dagger_{vw}$ for $v,w \in V$, we have $\langle D_{int}, D_{int}\rangle_{en} = D_{int}^T L^\dagger D_{int}$.

Term 2. Since $j_q(x,y)$ is linear on each edge $e = \{e^-, e^+\}$ with $\ell(e) = 1$, one has $\int_e j_q(x,y)\,dx = \frac{1}{2}(L^\dagger_{e^-,y} + L^\dagger_{e^+,y})$. Therefore

$$2\langle D_{int}, \mu_E\rangle_{en} = 2\sum_{y\in V} D_{int}(y) \sum_{e\in E} M(e)\cdot \tfrac{1}{2}(L^\dagger_{e^-,y} + L^\dagger_{e^+,y}).$$

Applying Lemma 5.18 with $a(v,w) = M(\{v,w\})L^\dagger_{v,y}$ reindexes this as $2D_{int}^T L^\dagger \tilde{\mu}_E$.

Term 3. Using the midpoint rule on each pair (e, f):

$$\langle \mu_E, \mu_E \rangle_{en} = \tfrac{1}{4} \sum_{e,f} M(e)M(f)\Big(L^\dagger_{e^-,f^-} + L^\dagger_{e^-,f^+} + L^\dagger_{e^+,f^-} + L^\dagger_{e^+,f^+}\Big).$$

Expanding $\tilde{\mu}_E^T L^\dagger \tilde{\mu}_E$ and relabeling summation variables produces the same expression. All three terms match. □

Corollary 5.20. *For any $D \in \mathrm{Div}^g$ supported on $V(G)$, $E_F(D) = E_F^{cont}(D)$.*

Proof. Take $N \equiv 0$ and $M(e) = F(e)$, so $\tilde{\mu} = \tilde{\mu}_F$. □

5.6 Counterexamples: Necessity of Positive Curvature

Theorem 5.1 requires $p_v > 0$ for all v. The following two examples show that without this assumption the break divisor need not minimize E_F. In both cases the break divisor concentrates chips on a negatively curved vertex, which is the energetically costly location.

5.6.1 Counterexample 1: Genus 4

Let G be the graph on $V = \{0, 1, 2, 3, 4, 5, 6\}$ with edges $\{1, 6\}$, $\{1, 4\}$, $\{4, 6\}$, $\{6, 3\}$, $\{6, 0\}$, $\{6, 5\}$, $\{6, 2\}$, $\{3, 0\}$, $\{0, 5\}$, $\{5, 2\}$. Here $g = 4$ and the resistance curvature is

$$p = \left[\tfrac{1}{6},\ \tfrac{1}{3},\ \tfrac{8}{21},\ \tfrac{8}{21},\ \tfrac{1}{3},\ \tfrac{1}{6},\ -\tfrac{16}{21}\right]^T,$$

so vertex 6 has $p_6 = -\frac{16}{21} < 0$.

In the class containing the break divisor $D_2 = [0, 0, 0, 0, 0, 0, 4]^T$, the three linearly equivalent divisors

$$D_1 = [1, 0, 1, 1, 0, 1, 0]^T, \quad D_2 = [0, 0, 0, 0, 0, 0, 4]^T \text{ (break divisor)}, \quad D_3 = [0, 1, 0, 0, 1, 0, 2]^T$$

have Foster energies

$$E_F(D_1) = \tfrac{170}{147} \approx 1.16, \quad E_F(D_2) = \tfrac{198}{147} \approx 1.35, \quad E_F(D_3) = \tfrac{296}{147} \approx 2.01.$$

The minimizer is D_1, not the break divisor D_2. The break divisor places all chips on vertex 6, which has the most negative curvature and hence the smallest Foster coefficient $\tilde{\mu}_F(6)$.

5.6.2 Counterexample 2: Genus 7

Let G' be the graph on $V = \{0, 1, 2, 3, 4, 5, 6\}$ with 13 edges $\{0, 4\}$, $\{0, 5\}$, $\{0, 6\}$, $\{1, 4\}$, $\{1, 5\}$, $\{1, 6\}$, $\{2, 4\}$, $\{2, 6\}$, $\{3, 5\}$, $\{3, 6\}$, $\{4, 5\}$, $\{4, 6\}$, $\{5, 6\}$. Here $g = 7$ and the resistance curvature has negative entries at vertices 4, 5, and 6:

$$p = \left[\tfrac{95}{286}, \tfrac{95}{286}, \tfrac{60}{143}, \tfrac{60}{143}, -\tfrac{71}{858}, -\tfrac{71}{858}, -\tfrac{145}{429}\right]^T.$$

For a degree-7 equivalence class the energies satisfy

$$E_F(D_1') = \tfrac{115{,}497}{61{,}347} \approx 1.883 \ < \ E_F(D_e') = \tfrac{115{,}640}{61{,}347} \approx 1.885 \ < \ E_F(D_2') = \tfrac{186{,}854}{61{,}347} \approx 3.046,$$

where D_e' is the break divisor. The gap is small but nonzero, and again the break divisor places chips on the negatively curved vertices.

5.6.3 Discussion

In both examples the energy functional penalizes placing chips on vertices where $\tilde{\mu}_F(v) = K(v)/2 + p_v$ is small relative to $K(v)/2$. When $p_v < 0$, vertex v is energetically cheap (the distance to $\tilde{\mu}_F$ is smaller there), so the energy minimizer prefers v over the break divisor's choice. A complete enumeration of all 853 connected simple graphs on 7 vertices found exactly 82 with positive resistance curvature; in all 82 cases the break divisor is the unique energy minimizer, consistent with Theorem 5.1.

Chapter 6

INTEGRAL ASPECTS OF FOURIER DUALITY FOR ABELIAN VARIETIES

6.1 Introduction

A powerful tool in the study of algebraic cycles on abelian varieties is the Fourier transform, first introduced by Mukai [Muk81] and studied on Chow groups by Beauville in [Bea83, Bea86]. The classical Fourier transform

$$\mathcal{F} \colon \mathrm{CH}(X;\mathbb{Q}) \to \mathrm{CH}(X^t;\mathbb{Q})$$

is defined as the correspondence induced by $\mathrm{ch}(\mathcal{P})$, where $\mathcal{P}$ is the Poincaré bundle on $X \times X^t$. Its main drawback is that it works only with rational coefficients: the Chern character introduces denominators, and the proof that $\mathcal{F}^t \circ \mathcal{F} = (-1)^g \cdot [-1]^*$ relies on the Grothendieck–Riemann–Roch theorem, which introduces further denominators.

Already in [Bea83], Beauville observes that for an abelian variety X with dual X^t, one can define integral lifts $\mathsf{F} \colon \mathrm{CH}(X) \to \mathrm{CH}(X^t)$ and $\mathsf{F}^t \colon \mathrm{CH}(X^t) \to \mathrm{CH}(X)$ of N times the classical Fourier transforms, for some positive integer N depending on $g = \dim X$, but without making N effective. The present chapter, based on [HHL$^+$25], uses results of Pappas [Pap07] on an integral version of Grothendieck–Riemann–Roch to produce an effective N, and to prove that all of Beauville's classical results hold over the ring $\Lambda = \mathbb{Z}[1/(2g+d+1)!]$, where $d = \dim S$.

We work throughout with abelian schemes $\pi \colon X \to S$, where S is a connected smooth quasi-projective variety over a field L. There are three main results.

Theorem 6.1 (Integral inversion, Theorem 6.5)**.** *Under mild assumptions on X/S, there exist integral transformations F, F^t lifting $(2g+d)!$ times the classical Fourier transforms,* *such that*

$$N \cdot (\mathsf{F}^t \circ \mathsf{F})(\alpha) \cdot \pi^* \left(\sum_{j=0}^{d} \frac{T_d}{T_j} \cdot \widetilde{\mathrm{Td}}_j(\mathscr{H}) \right) = (-1)^g \cdot N \cdot (2g+d)!^2 \cdot T_d \cdot [-1]^*(\alpha)$$

for all $\alpha \in \mathrm{CH}(X)$*, where* $T_m = \prod_p p^{\lfloor m/(p-1) \rfloor}$, $\mathscr{H} = R^1\pi_*\Omega^\bullet_{X/S}$ *is the de Rham bundle, and* N *is an explicitly described integer.*

Theorem B (Beauville decomposition, Theorem 6.10). *With* $\Lambda = \mathbb{Z}[1/(2g+d+1)!]$*, the Beauville decomposition*

$$\mathrm{CH}^i(X;\Lambda) = \bigoplus_s \mathrm{CH}^i_{(s)}(X;\Lambda)$$

holds, is compatible with the intersection and Pontryagin products, and is stable under pullback and pushforward along isogenies. Moreover, $\mathrm{Td}(\mathscr{H}) = 1$ *in* $\mathrm{CH}(S;\Lambda)$*.*

Theorem C ($\mathfrak{sl}_2$-action, Theorem 6.14). *If* X *has a polarization* θ *of degree* $\nu(\theta)^2$*, then over* $\Lambda_\theta = \Lambda[\nu(\theta)^{-1}]$ *the operators*

$$e(x) = \ell \cdot x, \qquad h(x) = (2i - s - g) \cdot x \text{ on } \mathrm{CH}^i_{(s)}, \qquad f(x) = \lambda \star x$$

define a representation of the Lie algebra $\mathfrak{sl}_2$ *on* $\mathrm{CH}(X;\Lambda_\theta)$*, which decomposes as a direct sum of copies of* $\mathrm{Sym}^n(\mathrm{St})$*,* $n = 0, \dots, g$*.*

As a corollary, for an abelian variety X over an algebraically closed field, $\mathrm{CH}^i_{(1)}(X;\Lambda_\theta)$ contains a copy of $X(L) \otimes \Lambda_\theta$ for every $i \in \{1, \dots, g\}$, yielding torsion classes in Chow groups.

The chapter is organized as follows. Section 6.2 recalls the integral polynomials $\widetilde{\mathrm{Td}}_m$, $\tilde{s}_m$, and $\widetilde{\mathrm{CT}}_m$ introduced by Pappas, and states the integral GRR theorem. Section 6.3 recalls the setup for abelian schemes: Hodge bundles, the de Rham bundle, the Poincaré bundle, and the key formula (Proposition 6.4) that is the engine of Theorem A. Section 6.4 constructs F and F^t, proves Theorem A, and establishes the basic properties of F. Section 6.5 proves Theorem B. Section 6.6 is an algebraic interlude on representations of $\mathfrak{sl}_2$ over commutative rings. Section 6.7 proves Theorem C and identifies the primitives. Section 6.8 gives the torsion application.

6.2 Pappas's Integral Grothendieck–Riemann–Roch

In this section we recall the integral polynomials introduced in [Pap07] and state the integral GRR theorem that underlies all our results.

6.2.1 The integers T_m

To control denominators in the Grothendieck–Riemann–Roch formulas, Pappas introduced the highly divisible integers T_m. We recall their definition and basic divisibility properties.

For $m \in \mathbb{N}$, define

$$T_m = \prod_p p^{\lfloor m/(p-1) \rfloor},$$

where the product runs over all primes p. By [Pap07, Lemma 2.1], if $m_1 + \cdots + m_r + n_1 + \cdots + n_s \leq m$ then $(m_1 + 1)! \cdots (m_r + 1)! \cdot T_{n_1} \cdots T_{n_s}$ divides T_m. In particular, $n! \cdot T_{m-n}$ divides T_m for all $0 \leq n \leq m$.

The following lemma controls which denominators we need to invert.

Lemma 6.2 (Pappas; cf. [HHL+25]). *Let h be a positive integer and define*

$$N = \begin{cases} 2 & \text{if } h = 3, \\ h+1 & \text{if } h+1 \text{ is prime}, \\ 1 & \text{otherwise.} \end{cases}$$

Then T_h divides $N \cdot h!^2$.

Proof. [Pap07, Lemma 2.1] and [HHL+25, Lemma 2.1]. □

6.2.2 Integral Todd and Chern character polynomials

The Todd class of a vector bundle, viewed as a polynomial in its Chern classes, has denominators: the degree-m term Td_m is a polynomial in $\mathbb{Q}[c_1, \ldots, c_m]$ whose denominator is exactly T_m. Following [Pap07, §2], one clears these denominators by setting

$$\widetilde{\mathrm{Td}}_m = T_m \cdot \mathrm{Td}_m \in \mathbb{Z}[c_1, \ldots, c_m].$$

Similarly, the degree-m part of the Chern character has denominator $m!$; one defines

$$\tilde{s}_m(r, c'_1, \dots, c'_m) = m! \cdot \mathrm{ch}_m(r, c'_1, \dots, c'_m) \in \mathbb{Z}[r, c'_1, \dots, c'_m].$$

Finally, one defines the combined integral polynomial

$$\widetilde{\mathrm{CT}}_m = \sum_{j=0}^{m} \frac{T_m}{j! \cdot T_{m-j}} \cdot (\tilde{s}_j \cdot \widetilde{\mathrm{Td}}_{m-j}) \in \mathbb{Z}[c_1, \dots, c_m, r, c'_1, \dots, c'_m],$$

where the coefficient $T_m/(j! \cdot T_{m-j})$ is an integer by the divisibility property of the T_m. All three polynomials $\widetilde{\mathrm{Td}}_m$, $\tilde{s}_m$, and $\widetilde{\mathrm{CT}}_m$ depend only on the K-theory class of the input bundle; see [Pap07, §2] for details.

One also needs the inverse Todd polynomial. As explained in [Pap07, §4.c], the expression

$$\widetilde{\mathrm{Td}}^{\mathrm{inv}}_n(E) = (n + \mathrm{rank}(E))! \cdot \left\{\mathrm{Td}(E)^{-1}\right\}_n$$

is an integral polynomial in the Chern classes of E.

6.2.3 Pappas's theorem

Theorem 6.3 (Pappas [Pap07], Cor. 2.3, Prop. 4.4)**.** *Let X and S be smooth quasi-projective varieties over a field L of characteristic* 0*.*

1. *Let $f\colon X \to S$ be a smooth projective morphism of relative dimension g, and $\mathcal{E}$ a coherent $\mathcal{O}_X$-module. Then in* $\mathrm{CH}^n(S)$*:*

$$f_*\Big(\widetilde{\mathrm{CT}}_{g+n}(X/S, \mathcal{E})\Big) = \frac{T_{g+n}}{n!} \cdot \tilde{s}_n(S, f_*[\mathcal{E}]).$$

2. *Let $\varepsilon\colon Z \hookrightarrow X$ be a closed immersion of codimension g with normal sheaf N, and $\mathcal{E}$ a coherent $\mathcal{O}_Z$-module. For $n \in \mathbb{N}$ with $n \geq g$:*

$$\tilde{s}_n(X, \varepsilon_*[\mathcal{E}]) = \sum_{a=0}^{n-g} \binom{n}{a} \cdot \varepsilon_*\Big(\tilde{s}_a(Z, \mathcal{E}) \cdot \widetilde{\mathrm{Td}}^{\mathrm{inv}}_{n-g-a}(Z, N)\Big),$$

   and $\tilde{s}_n(X, \varepsilon_*[\mathcal{E}]) = 0$ for $n < g$.

Proof. [Pap07, Cor. 2.3, Prop. 4.4]. □

6.3 Abelian Schemes and the Key Formula

6.3.1 Notation and conventions

Throughout, L is a field, S is a connected smooth quasi-projective L-scheme, and $\pi\colon X \to S$ is an abelian scheme of relative dimension $g > 0$ with zero section $e\colon S \to X$ and dual abelian scheme $\pi^t\colon X^t \to S$.

If X is a smooth variety over L and Λ is a commutative ring, we write $\mathrm{CH}(X;\Lambda) = \mathrm{CH}(X) \otimes \Lambda$, where $\mathrm{CH}(X)$ denotes the Chow ring of X modulo rational equivalence. For $n \in \mathbb{Z}$, we write $[n]_X$ (or simply $[n]$) for multiplication by n on X.

The *Hodge bundle* of X is $E_{X/S} = e^*\Omega^1_{X/S}$, a rank-$g$ vector bundle on S. Since $\mathcal{T}_{X/S} \cong \pi^* E^\vee_{X/S}$, for a line bundle $\mathcal{L}$ on X one has

$$\widetilde{\mathrm{CT}}_n(X/S,\mathcal{L}) = \sum_{i+j=n} \frac{T_n}{i!\cdot T_j} \cdot c_1(\mathcal{L})^i \cdot \pi^*\widetilde{\mathrm{Td}}_j(S, E^\vee_{X/S}). \tag{6.1}$$

The Hodge bundle of the dual satisfies $\det(E_{X/S}) \cong \det(E_{X^t/S})$; see [GW23, Cor. 27.230].

Let $\mathcal{P}$ be the Poincaré line bundle on $X \times_S X^t$, with $\varpi = c_1(\mathcal{P}) \in \mathrm{CH}^1(X \times_S X^t)$. By [GW23, Thm. 27.243], the pushforward of $\mathcal{P}$ along $\mathrm{pr}_1\colon X \times_S X^t \to X$ is

$$R^i\mathrm{pr}_{1,*}(\mathcal{P}) = \begin{cases} 0 & i \neq g, \\ e_*(\det(E_{X/S})^{-1}) & i = g. \end{cases} \tag{6.2}$$

The *de Rham bundle* is $\mathcal{H} = \mathcal{H}^1_{\mathrm{dR}}(X/S) = R^1\pi_*(\Omega^\bullet_{X/S})$, a rank-$2g$ vector bundle on S. By degeneration of the Hodge–de Rham spectral sequence it fits in a short exact sequence

$$0 \to E_{X/S} \to \mathcal{H} \to E^\vee_{X^t/S} \to 0. \tag{6.3}$$

6.3.2 Characteristic assumptions

The integral GRR theorem of Pappas is available in characteristic 0. To handle positive characteristic, we work under one of the following conditions [HHL$^+$25, §3.2]:

(A) $\mathrm{char}(L) = 0$;

(B) $\mathrm{char}(L) = p > 0$ and X/S can be lifted to characteristic 0 over a DVR of mixed characteristic $(0, p)$ with residue field L;

(C) $X \to S$ is obtained by base change from a case satisfying (B).

In particular, condition (C) is satisfied as soon as X admits a polarization θ and a level-n structure with $n \geq 3$ and $\mathrm{char}(L) \nmid n \cdot \deg(\theta)$. When $S = \mathrm{Spec}(L)$, Norman's theorem [Nor83] guarantees that (A) or (B) holds in almost all cases; see [HHL+25, §3.2] for a complete discussion.

6.3.3 The key formula

The following proposition is the engine of all results in this chapter. It computes the pushforward of the Poincaré class against the integral Todd class of the de Rham bundle, and shows that it is supported on the zero section.

Proposition 6.4 (Key formula [HHL+25], Prop. 3.3)**.** *Assume one of the conditions in 6.3.2. Let $\mathcal{H} = \mathcal{H}^1_{\mathrm{dR}}(X/S)$. For $m \in \mathbb{N}$:*

$$\sum_{i+j=m} \frac{T_{g+m}}{(g+i)! \cdot T_j} \cdot \mathrm{pr}_{1,*}(\varpi^{g+i}) \cdot \pi^* \widetilde{\mathrm{Td}}_j(\mathcal{H}) = \begin{cases} 0 & m \neq g, \\ (-1)^g \cdot T_{2g} \cdot [e(S)] & m = g, \end{cases} \tag{6.4}$$

in $\mathrm{CH}(X)$.

Proof. [HHL+25, Prop. 3.3]. The proof in characteristic 0 applies Theorem 6.3 twice: first to $X \times_S X^t$ over X to compute $\mathrm{pr}_{1,*}(\varpi^{g+i})$ via (6.2), then to the zero section $e\colon S \hookrightarrow X$ with normal bundle $E^\vee_{X/S}$. The Todd-inverse identity in [Pap07, §2.4] (identity (3) of loc. cit.) then kills all terms with $m \neq g$, leaving only $(-1)^g T_{2g}[e(S)]$ when $m = g$. The positive-characteristic cases follow by specialization. □

6.4 The Integral Fourier Transform (Theorem A)

This section constructs the integral Fourier–Mukai transforms F and F^t and proves Theorem A from the introduction. The key input is the integration formula of Proposition 6.4, which clears the way for establishing an equivalence of integral Chow motifs.

6.4.1 Construction

Throughout, fix an abelian scheme $\pi\colon X \to S$ satisfying one of the conditions in 6.3.2, and set $d = \dim S$. Let $\pi^t\colon X^t \to S$ be the dual, with Poincaré class $\varpi \in \mathrm{CH}^1(X \times_S X^t)$.

Define the integral correspondence

$$\gamma_X = \sum_{i=0}^{2g+d} \frac{(2g+d)!}{i!} \cdot \varpi^i = (2g+d)! \cdot \mathrm{ch}(\mathscr{P}_X) \in \mathrm{CH}(X \times_S X^t),$$

and set $\mathsf{F}_X = (\gamma_X)_*\colon \mathrm{CH}(X) \to \mathrm{CH}(X^t)$, so that

$$\mathsf{F}_X(\alpha) = \mathrm{pr}_{X^t,*}(\gamma_X \cdot \mathrm{pr}_X^*(\alpha)).$$

Concretely, $\mathsf{F}_X(\alpha)$ equals $(2g+d)!$ times the classical Fourier transform of α. Let $\gamma^t = \gamma_{X^t}$ and $\mathsf{F}^t = (\gamma^t)_*\colon \mathrm{CH}(X^t) \to \mathrm{CH}(X)$.

6.4.2 Theorem A: Integral Inversion

Define N by Lemma 6.2 with $h = 2g+d$:

$$N = \begin{cases} 2 & \text{if } 2g+d = 3, \\ 2g+d+1 & \text{if } 2g+d+1 \text{ is prime}, \\ 1 & \text{otherwise}. \end{cases}$$

By that lemma, T_{g+k} divides $N \cdot (2g+d)!^2$ for all $k \le g+d$.

Theorem 6.5 (Integral inversion [HHL+25], Thm. 4.2)**.** *Let notation be as above. For all $\alpha \in \mathrm{CH}(X)$:*

$$N \cdot (\mathsf{F}^t \circ \mathsf{F})(\alpha) \cdot \pi^*\left(\sum_{j=0}^{d} \frac{T_d}{T_j} \cdot \widetilde{\mathrm{Td}}_j(\mathscr{H})\right) = (-1)^g \cdot N \cdot (2g+d)!^2 \cdot T_d \cdot [-1]^*(\alpha). \tag{6.5}$$

Proof. The composition $\mathsf{F}^t \circ \mathsf{F}$ is computed by the correspondence $(2g+d)! \cdot \mathrm{pr}_{13,*}\mu^*(\gamma)$, where $\mu\colon X \times_S X^t \times_S X \to X \times_S X^t$ sends $(x, \xi, y) \mapsto (x+y, \xi)$. By the Theorem of the Cube, $\mathrm{pr}_{12}^*(\mathcal{P}) \otimes \mathrm{pr}_{23}^*(\mathcal{P}^t) \cong \mu^*(\mathcal{P})$ as line bundles on $X \times X^t \times X$, which gives

$$\mathrm{pr}_{12}^*(\gamma) \cdot \mathrm{pr}_{23}^*(\gamma^t) = (2g+d)! \cdot \mu^*(\gamma).$$

Since the diagram

$$\begin{tikzcd} X \times_S X^t \times_S X \arrow[r, "\mu"] \arrow[d, "\mathrm{pr}_{13}"'] & X \times_S X^t \arrow[d, "\mathrm{pr}_1"] \\ X \times_S X \arrow[r, "m"'] & X \end{tikzcd}$$

is Cartesian (where $m\colon X \times_S X \to X$ is the group law), we have $(2g+d)! \cdot \mathrm{pr}_{13,*}\mu^*(\gamma) = (2g+d)! \cdot m^*\mathrm{pr}_{1,*}(\gamma)$. The left-hand side of (6.5) is therefore the result of applying the correspondence

$$m^*\left(\sum_{i+j \le g+d} \frac{N \cdot (2g+d)!^2 \cdot T_d}{T_{g+i+j}} \cdot \frac{T_{g+i+j}}{(g+i)! \cdot T_j} \cdot \mathrm{pr}_{1,*}(\varpi^{g+i}) \cdot \pi^*\widetilde{\mathrm{Td}}_j(\mathcal{H})\right)$$

to α; by our choice of N all coefficients are integers. By Proposition 6.4, this correspondence equals $(-1)^g \cdot N \cdot (2g+d)!^2 \cdot T_d \cdot m^*[e(S)]$. Since the diagram

$$\begin{tikzcd} X \arrow[r, "{(1,-1)}"] \arrow[d, "\pi"'] & X \times_S X \arrow[d, "m"] \\ S \arrow[r, "e"'] & X \end{tikzcd}$$

is Cartesian, $m^*[e(S)]$ is the class of the graph of $[-1]_X$, giving the result. □

Remark 6.6. Formula (6.5) is an integral version of the classical identity $(\mathcal{F}^t \circ \mathcal{F})(\alpha) \cdot \pi^*\mathrm{Td}(\mathcal{H}) = (-1)^g \cdot [-1]^*(\alpha)$ in $\mathrm{CH}(X; \mathbb{Q})$ [DM91].

6.4.3 Basic Properties of F

Proposition 6.7 ([HHL+25], Prop. 4.4)**.** *Let notation be as above.*

1. *For x, $y \in \mathrm{CH}(X)$: $(2g+d)! \cdot \mathsf{F}(x \star y) = \mathsf{F}(x) \cdot \mathsf{F}(y)$.*

2. *Let $f\colon X \to Y$ be a homomorphism of abelian schemes over S with dual $f^t\colon Y^t \to X^t$, and write g_X, g_Y for the relative dimensions. If $g_X \geq g_Y$ then $f^{t,*}(\mathsf{F}_X(\alpha)) = \frac{(2g_X+d)!}{(2g_Y+d)!} \cdot \mathsf{F}_Y(f_*(\alpha))$ for all $\alpha \in \mathrm{CH}(X)$; the analogous formula holds if $g_X \leq g_Y$.*

Proof. [HHL$^+$25, Prop. 4.4]. The first part uses the Cartesian diagram computing $\mathsf{F}^t \circ \mathsf{F}$, together with the Theorem of the Cube. The second part uses the isomorphism $(f \times \mathrm{id})^* \mathcal{P}_Y \cong (\mathrm{id} \times f^t)^* \mathcal{P}_X$. □

6.5 Fourier Theory over Λ (Theorem B)

In this section, we study the Fourier transform after extending scalars to the localization Λ. We establish Theorem B, which generalizes the classical Beauville decomposition of the Chow ring from fields to arbitrary base schemes over Λ.

6.5.1 The Ring Λ and the Normalized Fourier Transform

Define

$$\Lambda = \mathbb{Z}\left[\frac{1}{(2g+d+1)!}\right].$$

Since $(2g+d)! \in \Lambda^*$, we can define the normalized transform

$$\mathcal{F} = \mathcal{F}_X\colon \mathrm{CH}(X;\Lambda) \to \mathrm{CH}(X^t;\Lambda), \quad \mathcal{F}(\alpha) = (2g+d)!^{-1} \cdot \mathsf{F}_X(\alpha),$$

which lifts the classical Fourier transform. By Theorem 6.5 and Lemma 6.2, $T_d \in \Lambda^*$, so $\mathrm{Td}(\mathcal{H}) \in \mathrm{CH}(S;\Lambda)$ is well-defined, and

$$(\mathcal{F}^t \circ \mathcal{F})(\alpha) \cdot \pi^* \mathrm{Td}(\mathcal{H}) = (-1)^g \cdot [-1]^*(\alpha) \tag{6.6}$$

holds in $\mathrm{CH}(X;\Lambda)$.

6.5.2 The Beauville decomposition

Following Beauville's classical work over fields [Bea86], we isolate the eigenspaces of the multiplication maps $[n]$ to decompose the Chow ring over Λ.

Definition 6.8 (Beauville subspaces)**.** For $i \geq 0$ and $s \in \mathbb{Z}$, define

$$\mathrm{CH}^i_{(s)}(X;\Lambda) = \big\{x \in \mathrm{CH}^i(X;\Lambda) \,\big|\, [n]^*(x) = n^{2i-s} \cdot x \text{ for all } n \in \mathbb{Z}\big\}.$$

The argument of [HHL+25, §4], following the pattern of Beauville, proceeds in two steps. First, an eigenvalue lemma identifies the Fourier degree of each eigenspace:

Lemma 6.9 ([HHL+25], Lem. 4.5)**.** *Let $x \in \mathrm{CH}^i(X;\Lambda)$ and let s be an integer with $i-2g \leq s \leq 2i$ such that $[n]^*(x) = n^{2i-s} \cdot x$ for all $n \in \mathbb{Z}$. Then $\mathcal{F}(x) \in \mathrm{CH}^{g-i+s}(X^t;\Lambda)$, and $[n]^*(\mathcal{F}(x)) = n^{2g-2i+s} \cdot \mathcal{F}(x)$; equivalently, $\mathcal{F}(x) \in \mathrm{CH}^{g-i+s}_{(s)}(X^t;\Lambda)$. Moreover, if $x \neq 0$ then necessarily*

$$\max\{i-g,\, 2i-2g\} \;\leq\; s \;\leq\; \min\{i+d,\, 2i\}.$$

Proof. [HHL+25, Lem. 4.5]. □

Second, a spanning argument using the $[n]$-eigenspace structure and formula (6.6) shows that the $\mathrm{CH}^i_{(s)}$ exhaust $\mathrm{CH}^i(X;\Lambda)$. We summarize these properties in the following theorem.

Theorem 6.10 (Theorem B: Beauville decomposition [HHL+25], Thm. 4.7)**.** *Let notation be as above.*

1. $\mathrm{Td}(\mathcal{H}) = 1$ *in* $\mathrm{CH}(S;\Lambda)$, *and* $\mathcal{F}^t \circ \mathcal{F} = (-1)^g \cdot [-1]^*$ *in* $\mathrm{CH}(X;\Lambda)$.

2. *For* $\alpha, \beta \in \mathrm{CH}(X;\Lambda)$*:*

$$\mathcal{F}(\alpha \star \beta) = \mathcal{F}(\alpha) \cdot \mathcal{F}(\beta), \qquad \mathcal{F}(\alpha \cdot \beta) = (-1)^g \cdot \mathcal{F}(\alpha) \star \mathcal{F}(\beta).$$

3. *There is a decomposition*

$$\mathrm{CH}^i(X;\Lambda) = \bigoplus_{s=\max\{i-g,\,2i-2g\}}^{\min\{i+d,\,2i\}} \mathrm{CH}^i_{(s)}(X;\Lambda),$$

and $\mathcal{F}$ restricts to isomorphisms $\mathrm{CH}^i_{(s)}(X;\Lambda) \xrightarrow{\sim} \mathrm{CH}^{g-i+s}_{(s)}(X^t;\Lambda)$.

4. *The intersection product maps* $\mathrm{CH}^i_{(s)} \times \mathrm{CH}^j_{(t)} \to \mathrm{CH}^{i+j}_{(s+t)}$*, and the Pontryagin product maps* $\mathrm{CH}_{i,(s)} \times \mathrm{CH}_{j,(t)} \to \mathrm{CH}_{i+j,(s+t)}$ *(in relative notation).*

5. *For a homomorphism* $f\colon X \to Y$ *of abelian schemes over* S *with* $h := \dim Y/S \leq g = \dim X/S$*:* $(f^t)^* \circ \mathcal{F}_X = \mathcal{F}_Y \circ f_*$ *and* $\mathcal{F}_X \circ f^* = (-1)^{g-h} \cdot f^t_* \circ \mathcal{F}_Y$*.*

Proof. [HHL+25, Thm. 4.7]. The proof of 3 and the injectivity assertion of 1 use Lemma 6.9 together with (6.6): the subspaces $\mathrm{CH}^i_{(s)}$ are linearly independent (by eigenvalue separation), and every $x \in \mathrm{CH}^i(X;\Lambda)$ lies in their span (by applying $\mathcal{F}^t \circ \mathcal{F}$ to the decomposition of $\mathcal{F}(x)$). The Todd vanishing $\mathrm{Td}(\mathcal{H}) = 1$ then follows by taking $\alpha = [X]$ in (6.6). The remaining assertions follow from Proposition 6.7 by dividing by $(2g+d)!$. □

6.5.3 The augmentation ideal

We now examine the augmentation ideal inside the Chow ring, which consists of elements mapped to zero under pushforward. Understanding its nilpotence order is crucial for bounding torsion.

Define the *augmentation ideal*

$$I = \ker(\pi_*) = \bigoplus_{s>0} \mathrm{CH}_{0,(s)}(X/S;\Lambda) \subset \mathrm{CH}_0(X/S;\Lambda),$$

where $\mathrm{CH}_i(X/S;\Lambda) = \mathrm{CH}_{d+i}(X;\Lambda)$ is graded by relative dimension. Part 4 of Theorem 6.10 gives the following nilpotence result, which is the Λ-coefficient version of a theorem of Bloch [Blo76].

Corollary 6.11 ([HHL+25], Cor. 4.8)**.** *With* $\nu = \min\{g+d, 2g\}$*, we have* $I^{\star(\nu+1)} = 0$*.*

6.6 Integral $\mathfrak{sl}_2$*-Representations*

This section is an algebraic interlude on representations of $\mathfrak{sl}_2$ over commutative rings. Proposition 6.13 will be applied to Chow groups in the next section.

6.6.1 Setup and notation

Let $g \geq 1$ and let Λ be any ring of the form $\Lambda = \mathbb{Z}[1/n]$ with $(2g)! \mid n$. By $\mathfrak{sl}_2$ we mean the Lie algebra over Λ with basis $\{e, h, f\}$ and bracket $[h, e] = 2e$, $[e, f] = h$, $[h, f] = -2f$.

We consider representations $\rho\colon \mathfrak{sl}_2 \to \mathfrak{gl}(V)$ on graded Λ-modules $V = \bigoplus_{i=-g}^{g} V_i$, where $\rho(h)$ acts as multiplication by i on V_i. The *standard representation* St has basis $\{x_{-1}, x_1\}$ with

$$e = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad h = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad f = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

Its nth symmetric power $\mathrm{Sym}^n(\mathrm{St})$ has basis $\{x_{-n}, x_{-n+2}, \dots, x_n\}$ and action

$$e(x_{-n+2i}) = (n-i) \cdot x_{-n+2i+2}, \quad h(x_i) = i \cdot x_i, \quad f(x_{-n+2i}) = i \cdot x_{-n+2i-2}.$$

The key technical lemma (proved by double induction) computes the action of $f^\ell e^k$ on lowest-weight vectors.

Lemma 6.12 ([HHL+25], Lem. 5.3)**.** *Let $v \in V_{-n}$ with $f(v) = 0$ and $0 \leq n \leq g$.*

1. *For $0 \leq k \leq n$ and $\ell \geq 0$:*

$$f^\ell e^k(v) = \begin{cases} \dfrac{(n-k+\ell)!\,k!}{(n-k)!\,(k-\ell)!} \cdot e^{k-\ell}(v) & \ell \leq k, \\ 0 & \ell > k. \end{cases}$$

2. *$e^{n+1}(v) = 0$.*

3. *For $1 \leq n \leq g$: $V_n[f] := \{v \in V_n : f(v) = 0\} = 0$.*

Proof. [HHL+25, Lem. 5.3]. □

Proposition 6.13 ([HHL+25], Prop. 5.5)**.** *Let $V = \bigoplus_{i=-g}^{g} V_i$ be an $\mathfrak{sl}_2$-module as above, and for $n \in \{0, \dots, g\}$ set $M_n = \sum_{i=0}^{n} e^i V_{-n}[f]$.*

1. *The map $v \mapsto (n!)^{-1} e^n(v)$ is an isomorphism $V_{-n}[f] \xrightarrow{\sim} V_n[e]$, with inverse $(n!)^{-1} f^n$. In particular, for every $i \leq n$ the map $e^i\colon V_{-n}[f] \to V$ is injective.*

2. *$M_n = \bigoplus_{i=0}^{n} e^i V_{-n}[f]$ is an $\mathfrak{sl}_2$-submodule of V, and the map $x_{-n+2i} \otimes v \mapsto \frac{(n-i)!}{n!} e^i(v)$ defines an isomorphism $\mathrm{Sym}^n(\mathrm{St}) \otimes_\Lambda V_{-n}[f] \xrightarrow{\sim} M_n$ of $\mathfrak{sl}_2$-modules.*

3. *The natural map $\phi\colon \bigoplus_{n=0}^{g} M_n \to V$ is an isomorphism of $\mathfrak{sl}_2$-modules.*

Proof. [HHL⁺25, Prop. 5.5]. Parts (1) and (2) follow from Lemma 6.12. For injectivity of ϕ, one shows the kernel is homogeneous (using that $(2g)! \in \Lambda^*$ makes the Vandermonde determinant invertible) and then applies Lemma 6.12(2)–(3) to obtain a contradiction. Surjectivity follows by a similar minimum-index argument. □

6.7 The $\mathfrak{sl}_2$-Action on the Chow Ring (Theorem C)

This section applies the representation theory of $\mathfrak{sl}_2$ developed in the previous section to the Chow ring. By exhibiting explicit operators on $\mathrm{CH}(X)_{\Lambda_\theta}$, we prove Theorem C and establish a weak Hard Lefschetz property.

In addition to the conditions of Section 6.5, fix a polarization $\theta\colon X \to X^t$ of degree $\nu(\theta)^2$, and set

$$\Lambda_\theta = \Lambda\left[\frac{1}{\nu(\theta)}\right] = \mathbb{Z}\left[\frac{1}{\nu(\theta)\cdot(2g+d+1)!}\right].$$

6.7.1 The Classes ℓ and λ

Define $\ell \in \mathrm{CH}^1_{(0)}(X;\Lambda)$ by

$$\ell = \tfrac{1}{2}\cdot(\mathrm{id},\theta)^*(\varpi),$$

where $\varpi = c_1(\mathscr{P})$. If θ is given by a symmetric ample line bundle L, then $\ell = c_1(L)$. By standard Riemann–Roch for abelian varieties,

$$\ell^g = \nu(\theta)\cdot g!\cdot[e(S)]. \tag{6.7}$$

Working with coefficients in Λ_θ, define

$$\lambda = \frac{\ell^{g-1}}{\nu(\theta)\cdot(g-1)!} \in \mathrm{CH}^{g-1}_{(0)}(X;\Lambda_\theta).$$

As shown in [HHL⁺25, §6.2] following Künnemann [Kü93], λ is characterized by $\mathscr{F}(\ell) = (-1)^{g-1}\cdot\theta_*(\lambda)$.

6.7.2 Theorem C: The $\mathfrak{sl}_2$-Action

Theorem 6.14 (Theorem C: $\mathfrak{sl}_2$-action [HHL+25], Thm. 6.3)**.** *Define endomorphisms of* $\mathrm{CH}(X;\Lambda_\theta)$ *by*

$$e(x) = \ell \cdot x, \quad h(x) = (2i - s - g) \cdot x \text{ for } x \in \mathrm{CH}^i_{(s)}(X;\Lambda_\theta), \quad f(x) = \lambda \star x.$$

1. *The operators e, h, f define a representation of $\mathfrak{sl}_2$ on $\mathrm{CH}(X;\Lambda_\theta)$. Each $\mathrm{CH}_{(s)}(X;\Lambda_\theta) = \bigoplus_i \mathrm{CH}^i_{(s)}(X;\Lambda_\theta)$ is a subrepresentation.*

2. *For each s, with $M_j = \{x \in \mathrm{CH}_{(s)}(X;\Lambda_\theta) : h(x) = -j \cdot x,\ f(x) = 0\}$:*

$$\mathrm{CH}_{(s)}(X;\Lambda_\theta) \cong \bigoplus_{n=0}^{g} \Big(\mathrm{Sym}^n(\mathrm{St}) \otimes_{\Lambda_\theta} M_n\Big)$$

as representations of $\mathfrak{sl}_2$.

Proof. [HHL+25, Thm. 6.3]. We note the key points of the argument.

For 1: grading $\mathrm{CH}(X;\Lambda_\theta)$ by weight $2i - s - g$ on $\mathrm{CH}^i_{(s)}$, the operators e and f are homogeneous of degrees $+2$ and -2, so $[h,e] = 2e$ and $[h,f] = -2f$ are automatic. The identity $[e,f] = h$ is the content; it follows Künnemann's argument [Kü93, Kü94] via the logarithm of the diagonal correspondence $\Gamma_{[1]} \in \mathrm{CH}_0(X^2/X;\Lambda_\theta)$. Specifically, by Corollary 6.11, $(\Gamma_{[1]} - \Gamma_{[0]})^{\star(2g+d+1)} = 0$, so one may define

$$\log \Gamma_{[1]} = \sum_{j=1}^{2g+d} \frac{(-1)^{j-1}}{j} (\Gamma_{[1]} - \Gamma_{[0]})^{\star j} = \sum_n c_n \cdot \Gamma_{[n]},$$

and then set $\pi_i = (1/(2g-i)!) \cdot (\log \Gamma_{[1]})^{\star(2g-i)}$. Künnemann shows these projectors give the Beauville decomposition over $\mathbb{Q}$; the coefficients c_n satisfy $\sum_n a_{i,n} n^j = \mathbf{1}_{j=2g-i}$, so they also give the decomposition over Λ_θ. The identity $[e,f] = h$ then follows by Künnemann's calculation.

For 2: apply Proposition 6.13 with $V = \mathrm{CH}_{(s)}(X;\Lambda_\theta)$, graded by $V_i = \bigoplus_{2k-s=g+i} \mathrm{CH}^k_{(s)}$. $\square$

6.8 Torsion in Chow Groups

We now give the torsion application, which is the main motivation for working integrally. Let X be a g-dimensional abelian variety over an algebraically closed field L (so $S = \mathrm{Spec}(L)$ and $d = 0$), with polarization θ of degree $\nu(\theta)^2$. Set $\Lambda_\theta = \mathbb{Z}[1/(\nu(\theta)\cdot(2g+1)!)]$.

Lemma 6.15 ([HHL+25], Lem. 7.1)**.** *There is an isomorphism of Λ_θ-modules*

$$\mathrm{CH}^1_{(1)}(X;\Lambda_\theta) \cong X(L) \otimes \Lambda_\theta.$$

Proof. Since $2 \in \Lambda_\theta^*$, every $\alpha \in \mathrm{Pic}(X) \otimes \Lambda_\theta = \mathrm{CH}^1(X;\Lambda_\theta)$ splits uniquely as $\alpha = \frac{\alpha+[-1]^*\alpha}{2} + \frac{\alpha-[-1]^*\alpha}{2}$ into symmetric and anti-symmetric parts. This is precisely the Beauville decomposition $\mathrm{CH}^1_{(0)} \oplus \mathrm{CH}^1_{(1)}$. The anti-symmetric classes correspond to $\mathrm{Pic}^0(X) \otimes \Lambda_\theta = X^t(L) \otimes \Lambda_\theta \cong X(L) \otimes \Lambda_\theta$. ☐

Corollary 6.16 ([HHL+25], Cor. 7.2)**.** *For every $i \in \{1, \dots, g\}$, $\mathrm{CH}^i_{(1)}(X;\Lambda_\theta)$ contains a submodule isomorphic to $X(L) \otimes \Lambda_\theta$. In particular, if $p > 2g+1$ is a prime with $p \neq \mathrm{char}(L)$ and $p \nmid \nu(\theta)$, then $\mathrm{CH}^i_{(1)}(X;\Lambda_\theta)$ contains a subgroup isomorphic to $(\mathbb{Q}_p/\mathbb{Z}_p)^{2g}$.*

Proof. By Theorem 6.142 with $s = 1$, $\mathrm{CH}^1_{(1)}$ is the space M_{1-g} (the lowest-weight space of the $\mathfrak{sl}_2$-submodule of weight index $1-g$). By Proposition 6.13(1), the map e^{i-1}, i.e. intersection with ℓ^{i-1}, is injective on $\mathrm{CH}^1_{(1)}(X;\Lambda_\theta)$, so it identifies $\mathrm{CH}^1_{(1)}(X;\Lambda_\theta)$ with a submodule of $\mathrm{CH}^i_{(1)}(X;\Lambda_\theta)$ for each $i \in \{1, \dots, g\}$. The torsion claim follows by examining the p-primary part of $X(L) \cong (\mathbb{Q}_p/\mathbb{Z}_p)^{2g}$ over an algebraically closed field. ☐

Remark 6.17. This contrasts sharply with the rational setting: if L is the algebraic closure of a finite field, then $\mathrm{CH}^i_{(1)}(X;\mathbb{Q}) = 0$ for all i; see [Kü93, Thm. 7.1]. The integral approach thus genuinely captures structure that is invisible over $\mathbb{Q}$.

BIBLIOGRAPHY

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

[Bea83] Arnaud Beauville. Quelques remarques sur la transformation de Fourier dans l'anneau de Chow d'une variété abélienne. In *Algebraic geometry (Tokyo/Kyoto, 1982)*, volume 1016 of *Lecture Notes in Math.*, pages 238–260. Springer, Berlin, 1983.

[Bea86] Arnaud Beauville. Sur l'anneau de Chow d'une variété abélienne. *Math. Ann.*, 273(4):647–651, 1986.

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

[Blo76] Spencer Bloch. Some elementary theorems about algebraic cycles on abelian varieties. *Invent. Math.*, 37(3):215–228, 1976.

[CLD12] Antoine Chambert-Loir and Antoine Ducros. Formes différentielles réelles et courants sur les espaces de Berkovich. *arXiv preprint arXiv:1204.6277*, 2012.

[Dev25] Karel Devriendt. Graphs with nonnegative resistance curvature. *Ann. Combin.*, 29:1–24, 2025.

[dJS22a] Robin de Jong and Farbod Shokrieh. Faltings height and Néron–Tate height of a theta divisor. *Compos. Math.*, 158(1):1–32, 2022.

[dJS22b] Robin de Jong and Farbod Shokrieh. Metric graphs, cross ratios, and Rayleigh's laws. *Rocky Mountain J. Math.*, 52(4):1403–1422, 2022.

[dJS23] Robin de Jong and Farbod Shokrieh. Tropical moments of tropical Jacobians. *Canad. J. Math.*, 75(4):1045–1075, 2023. Preprint arXiv:1810.02639, 2018.

[DM91] Christopher Deninger and Jacob Murre. Motivic decomposition of abelian schemes and the Fourier transform. *J. reine angew. Math.*, 422:201–219, 1991.

[GK17] Walter Gubler and Klaus Künnemann. A tropical approach to nonarchimedean Arakelov geometry. *Algebra & Number Theory*, 11(1):77–180, 2017.

[GS23a] Andreas Gross and Farbod Shokrieh. A sheaf-theoretic approach to tropical homology. *J. Algebra*, 635:577–641, 2023.

[GS23b] Andreas Gross and Farbod Shokrieh. Tautological cycles on tropical Jacobians. *Algebra Number Theory*, 17(4):885–921, 2023.

[GST22] Andreas Gross, Farbod Shokrieh, and Lilla Tóthmérész. Effective divisor classes on metric graphs. *Mathematische Zeitschrift*, 302(2):663–685, 2022.

[Gub16] Walter Gubler. Forms and currents on the analytification of an algebraic variety (after Chambert-Loir and Ducros). In *Nonarchimedean and tropical geometry*, pages 1–30. Springer, 2016.

[GW23] Ulrich Görtz and Torsten Wedhorn. *Algebraic Geometry II: Cohomology of Schemes*. Springer Studium Mathematik – Master. Springer, 2023.

[HHL$^+$25] Junaid Hasan, Hazem Hassan, Milton Lin, Marcella Manivel, Lily McBeath, and Ben Moonen. Integral aspects of Fourier duality for abelian varieties. *Manuscripta Math.*, 176(5):Paper No. 64, 26, 2025.

[IKMZ19] Ilia Itenberg, Ludmil Katzarkov, Grigory Mikhalkin, and Ilia Zharkov. Tropical homology. *Mathematische annalen*, 374(1-2):963–1006, 2019.

[JRS18] Philipp Jell, Johannes Rau, and Kristin Shaw. Lefschetz (1, 1)-theorem in tropical geometry. *Épijournal de Géométrie Algébrique*, 2, 2018.

[JSS19] Philipp Jell, Kristin Shaw, and Jascha Smacka. Superforms, tropical cohomology, and Poincaré duality. *Advances in Geometry*, 19(1):101–130, 2019.

[Kü93] Klaus Künnemann. A Lefschetz decomposition for Chow motives of abelian schemes. *Invent. Math.*, 113(1):85–102, 1993.

[Kü94] Klaus Künnemann. On the Chow motive of an abelian scheme. In *Motives (Part 1)*, volume 55 of *Proc. Sympos. Pure Math.*, pages 189–205. Amer. Math. Soc., Providence, RI, 1994.

[Lag12] Aron Lagerberg. Super currents and tropical geometry. *Mathematische Zeitschrift*, 270(3):1011–1050, 2012.

[Muk81] Shigeru Mukai. Duality between $\mathrm{D}(X)$ and $\mathrm{D}(\hat{X})$ with its application to Picard sheaves. *Nagoya Math. J.*, 81:153–175, 1981.

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

[Nor83] Peter Norman. Lifting abelian varieties. *Invent. Math.*, 74:431–443, 1983.

[OS79] Tadao Oda and C. S. Seshadri. Compactifications of the generalized Jacobian variety. *Trans. Amer. Math. Soc.*, 253:1–90, 1979.

[Pap07] Georgios Pappas. Integral Grothendieck–Riemann–Roch theorem. *Invent. Math.*, 170(3):455–481, 2007.

[Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, Berlin, 2003.

[Wil17] Robert Wilms. New explicit formulas for faltings' delta-invariant. *Inventiones mathematicae*, 209(2):481–539, 2017.

[Zha93] Shouwu Zhang. Admissible pairing on a curve. *Inventiones mathematicae*, 112:171–193, 1993.

[Zha10] Ilia Zharkov. Tropical theta characteristics. In *Mirror symmetry and tropical geometry*, volume 527 of *Contemp. Math.*, pages 165–168. Amer. Math. Soc., Providence, RI, 2010.

VITA

Junaid Hasan received his Bachelor of Science in Mathematics from the Indian Institute of Science (IISc), Bangalore, in 2019. He pursued doctoral studies in Mathematics at the University of Washington under the supervision of Professor Farbod Shokrieh. His research is in algebraic and tropical geometry, with a focus on metric graphs, tropical cohomology theories, and abelian varieties.